



Complete Software Guide for Junos[®] OS for the EX9200 Switches, Release 12.3

Release
12.3



Published: 2013-10-17

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

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

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

Complete Software Guide for Junos® OS for the EX9200 Switches, Release 12.3
Release 12.3
Copyright © 2013, Juniper Networks, Inc.
All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

	About the Documentation	lvii
	Documentation and Release Notes	lvii
	Supported Platforms	lvii
	Using the Examples in This Manual	lvii
	Merging a Full Example	lviii
	Merging a Snippet	lviii
	Documentation Conventions	lix
	Documentation Feedback	lx
	Requesting Technical Support	lxi
	Self-Help Online Tools and Resources	lxi
	Opening a Case with JTAC	lxii
Chapter 1	Software Overview	63
	EX Series Switch Software Features Overview	63
Part 1	Supported Hardware	
Chapter 2	EX9208 Switch Overview	115
	EX9208 Switch Hardware Overview	115
	Software	115
	Chassis Physical Specifications	116
	Host Subsystem	117
	Line Cards	117
	Cooling System	118
	Power Supplies	118
Part 2	System Administration	
Chapter 3	Access Privileges	123
	Overview	123
	Introduction to Access Privileges	123
	Understanding Junos OS Access Privilege Levels	123
	Junos OS Login Classes Overview	128
	Access Privilege User Permission Flags Overview	128

Specifying Access Privileges for Junos OS Configuration Mode	
Hierarchies	130
Configuration	131
Configuring Access Privileges	131
Configuring Access Privilege Levels	131
Specifying Access Privileges for Junos OS Operational Mode	
Commands	132
Specifying Access Privileges for Junos OS Configuration Mode	
Hierarchies	133
Examples	134
Example: Configuring Access Privilege Levels	134
Example: Configuring Access Privileges for Operational Mode	
Commands	134
Example: Specifying Access Privileges Using	
allow/deny-configuration-regexps Statements	135
User Permission Flags Reference	139
access	140
access-control	141
admin	142
admin-control	142
all-control	143
clear	143
configure	180
control	181
field	181
firewall	181
firewall-control	182
floppy	183
flow-tap	183
flow-tap-control	184
flow-tap-operation	184
idp-profiler-operation	184
interface	185
interface-control	185
maintenance	186
network	193
pgcp-session-mirroring	194
pgcp-session-mirroring-control	195
reset	195
rollback	196
routing	196
routing-control	201
secret	205
secret-control	206
security	207
security-control	210
shell	214
snmp	214
snmp-control	214

	system	215
	system-control	216
	trace	218
	trace-control	223
	view	228
	view-configuration	292
	Administration	292
	Operational Mode Commands	292
	show cli authorization	293
Chapter 4	System Basics	295
	Overview	295
	Installation Overview	295
	Junos OS Package Names	295
	Licenses Overview	296
	Understanding Software Licenses for EX Series Switches	296
	Configuration	301
	Junos OS Packages	301
	Downloading Software Packages from Juniper Networks	301
	Statement Hierarchies	302
	[edit accounting-options] Hierarchy Level	303
	[edit chassis] Hierarchy Level	305
	[edit class-of-service] Hierarchy Level	313
	[edit dynamic-profiles] Hierarchy Level	317
	[edit event-options] Hierarchy Level	318
	[edit firewall] Hierarchy Level	320
	[edit forwarding-options] Hierarchy Level	333
	[edit forwarding-options accounting] Hierarchy Level	334
	[edit forwarding-options dhcp-relay] Hierarchy Level	334
	[edit forwarding-options enhanced-hash-key] Hierarchy Level	338
	[edit forwarding-options family] Hierarchy Level	339
	[edit forwarding-options fast-reroute-priority] Hierarchy Level	340
	[edit forwarding-options hash-key] Hierarchy Level	340
	[edit forwarding-options helpers] Hierarchy Level	341
	[edit forwarding-options load-balance] Hierarchy Level	343
	[edit forwarding-options next-hop-group] Hierarchy Level	344
	[edit forwarding-options port-mirroring] Hierarchy Level	344
	[edit forwarding-options rpf-loose-mode-discard] Hierarchy Level	345
	[edit forwarding-options sampling] Hierarchy Level	346
	[edit groups] Hierarchy Level	348
	[edit interfaces] Hierarchy Level	348
	[edit logical-systems] Hierarchy Level	359
	[edit multi-chassis] Hierarchy Level	360
	[edit policy-options] Hierarchy Level	360
	[edit protocols] Hierarchy Level	365
	bfd	367
	[edit protocols bgp] Hierarchy Level	368
	dvmrp	375
	igmp	376

igmp-snooping	377
[edit protocols isis] Hierarchy Level	378
[edit protocols l2-learning] Hierarchy Level	380
[edit protocols lacp] Hierarchy Level	380
[edit protocols layer2-control] Hierarchy Level	381
[edit protocols ldp] Hierarchy Level	381
lldp	384
mld	385
[edit protocols mpls] Hierarchy Level	386
[edit protocols mstp] Hierarchy Level	386
neighbor-discovery	388
oam	389
[edit protocols ospf] Hierarchy Level	391
[edit protocols ospf3] Hierarchy Level	395
[edit protocols pim] Hierarchy Level	398
[edit protocols rip] Hierarchy Level	402
[edit protocols ripng] Hierarchy Level	404
[edit protocols router-advertisement] Hierarchy Level	405
[edit protocols router-discovery] Hierarchy Level	406
[edit protocols rstp] Hierarchy Level	406
[edit protocols sap] Hierarchy Level	407
[edit protocols vrrp] Hierarchy Level	407
[edit protocols vstp] Hierarchy Level	408
Layer 2 Routing Instances Configuration Hierarchy	409
[edit routing-options] Hierarchy Level	411
[edit security] Hierarchy Level	421
[edit security alarms] Hierarchy Level	421
[edit security authentication-key-chains] Hierarchy Level	421
[edit security certificates] Hierarchy Level	422
[edit security ike] Hierarchy Level	422
[edit security ipsec] Hierarchy Level	423
[edit security log] Hierarchy Level	424
[edit security pki] Hierarchy Level	424
[edit security ssh-known-hosts] Hierarchy Level	425
[edit security traceoptions] Hierarchy Level	425
[edit snmp] Hierarchy Level	426
[edit switch-options] Hierarchy Level	430
[edit system] Hierarchy Level	431
[edit vlans] Hierarchy Level	445
Administration	446
Operational Commands: CLI Interface	446
set cli complete-on-space	447
set cli directory	448
set cli idle-timeout	449
set cli prompt	450
set cli restart-on-upgrade	451
set cli screen-length	452
set cli screen-width	453
set cli terminal	454

set cli timestamp	455
show cli	456
show cli authorization	458
show cli directory	462
show cli history	463
start shell	464
Operational Commands: Software Installation	465
request system license add	466
request system license delete	467
request system license save	468
request system reboot	469
request system snapshot	474
request system software add	481
request system software delete	489
request system software rollback	492
request system software validate	496
request system zeroize	499
show system boot-messages	504
show system license	511
show system snapshot	514
Operational Commands: System Monitoring	516
show chassis alarms	517
show chassis environment	528
show chassis environment cb	578
show chassis environment fpc	595
show chassis environment pem	620
show chassis environment routing-engine	628
show chassis fan	633
show chassis hardware	644
show chassis location	750
show chassis pic	754
show chassis routing-engine	767
show chassis temperature-thresholds	787
show log	804
show system alarms	807
show system processes	808
show system statistics	835
show system uptime	870
Chapter 5 System Services	875
Overview	875
DHCP Local Server	875
Extended DHCP Local Server Overview	876
DHCPv6 Local Server Overview	880
DHCP Local Server Handling of Client Information Request Messages	881
DHCP Duplicate Client Differentiation Using Client Subinterface Overview	882
Group-Specific DHCP Local Server Options	883

Understanding Dynamic Reconfiguration of Extended DHCP Local Server Clients	883
DHCP Snooping Support	887
DHCP Auto Logout Overview	888
Address-Assignment Pools Overview	890
Use of DHCP Option 50 and DHCPv6 IA_NA Option to Request a Specific IP Address	891
Multiple Address Assignment for DHCPv6 Clients	891
Centrally Configured Opaque DHCP Options	893
Graceful Routing Engine Switchover	897
Port Number Requirements for DHCP Firewall Filters	898
DHCP Relay Agent	899
Extended DHCP Relay Agent Overview	900
DHCP Relay Proxy Overview	903
DHCPv6 Relay Agent Overview	905
DHCP Duplicate Client Differentiation Using Client Subinterface Overview	906
Group-Specific DHCP Relay Options	906
DHCP Snooping Support	907
DHCP Auto Logout Overview	908
Graceful Routing Engine Switchover	910
Port Number Requirements for DHCP Firewall Filters	911
Configuration	911
DHCP Local Server Examples	912
Example: Minimum Extended DHCP Local Server Configuration	912
Example: Extended DHCP Local Server Configuration with Optional Pool Matching	912
Example: Configuring Group Liveness Detection for DHCP Local Server Clients	913
DHCP Relay Agent Examples	916
Example: Minimum DHCP Relay Agent Configuration	916
Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings	917
Example: Configuring DHCP Snooping Support for DHCP Relay Agent	921
Configuration Tasks for DHCP Local Server	923
Using External AAA Authentication Services with DHCP	923
Guidelines for Configuring Support for DHCP Duplicate Clients	925
Configuring DHCP Duplicate Client Support	925
Grouping Interfaces with Common DHCP Configurations	926
Guidelines for Configuring Interface Ranges	927
Overriding Default DHCP Local Server Configuration Settings	928
Specifying the Maximum Number of DHCP Clients Per Interface	929
Disabling ARP Table Population	930
Automatically Logging Out DHCP Clients	932
Enabling Processing of Client Information Requests	933
Specifying the Delegated Address Pool for IPv6 Prefix Assignment	934
Enabling DHCPv6 Rapid Commit Support	935
Deleting DHCP Local Server and DHCP Relay Override Settings	935

Configuring Extended DHCP Local Server Dynamic Client Reconfiguration	936
Configuring Dynamic Reconfiguration Attempts for DHCP Clients	937
Configuring Deletion of the Client When Dynamic Reconfiguration Fails	938
Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect	938
Configuring a Token for DHCP Local Server Authentication	939
Preventing Binding of Clients That Do Not Support Reconfigure Messages	939
Requesting DHCP Local Server to Initiate Reconfiguration of Client Bindings	940
Configuring Detection of DHCP Local Server Client Connectivity	941
Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces	942
Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server	944
Configuring Passwords for Usernames	945
Creating Unique Usernames for DHCP Clients	946
Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use	949
Configuration Tasks for DHCP Relay Agent	950
Using External AAA Authentication Services with DHCP	951
Configuring DHCP Duplicate Client Support	952
Grouping Interfaces with Common DHCP Configurations	953
Guidelines for Configuring Interface Ranges	954
Overriding the Default DHCP Relay Configuration Settings	955
Overwriting giaddr Information	957
Replacing the DHCP Relay Request and Release Packet Source Address	957
Overriding Option 82 Information	957
Using Layer 2 Unicast Transmission for DHCP Packets	958
Trusting Option 82 Information	958
Disabling ARP Table Population	959
Specifying the Maximum Number of DHCP Clients Per Interface	960
Automatically Logging Out DHCP Clients	961
DHCP Relay Agent Option 82 Value for Auto Logout	962
Configuring DHCP Snooping for DHCP Relay Agent	963
Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent	964
Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent	969
Sending Release Messages When Clients Are Deleted	971
Disabling Automatic Binding of Stray DHCP Requests	972
Enabling and Disabling Insertion of Option 82 Information	973
Configuring Server Groups	976
Configuring Active Server Groups	977

Enabling DHCP Relay Proxy Mode	977
Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces	978
Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets	979
DHCP Liveness Detection Overview	980
Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity	981
Disabling DHCP Relay	983
DHCP Local Server Configuration Statements	983
[edit system] Hierarchy Level	983
attempts (DHCP Local Server)	999
authentication (DHCP Local Server)	1000
bfd	1001
circuit-type (DHCP Local Server)	1002
clear-on-abort (DHCP Local Server)	1003
client-discover-match (DHCP Local Server)	1004
client-id (DHCP Local Server)	1005
delegated-pool (DHCP Local Server)	1006
delimiter (DHCP Local Server)	1007
detection-time	1008
dhcp-local-server	1009
dhcpv6 (DHCP Local Server)	1014
domain-name (DHCP Local Server)	1017
duplicate-clients-on-interface (DHCP Local Server)	1018
dynamic-profile (DHCP Local Server)	1019
external-authority	1020
failure-action	1021
forward-snooped-clients (DHCP Local Server)	1022
group (DHCP Local Server)	1023
holddown-interval	1025
interface (DHCP Local Server)	1026
interface-client-limit (DHCP Local Server)	1028
interface-delete (Subscriber Management or DHCP Client Management)	1029
interface-name (DHCP Local Server)	1030
ip-address-first	1031
liveness-detection	1032
logical-system-name (DHCP Local Server)	1033
mac-address (DHCP Local Server)	1034
method	1035
minimum-interval	1036
minimum-receive-interval	1037
multiplier	1038
no-adaptation	1039
no-arp (DHCP Local Server)	1040
option-60 (DHCP Local Server)	1041
option-82 (DHCP Local Server Authentication)	1042
option-82 (DHCP Local Server Pool Matching)	1043

overrides (DHCP Local Server)	1044
password (DHCP Local Server)	1046
pool (DHCP Local Server Overrides)	1047
pool-match-order	1048
process-inform	1049
radius-disconnect (DHCP Local Server)	1051
rapid-commit (DHCPv6 Local Server)	1052
reconfigure (DHCP Local Server)	1053
relay-agent-interface-id (DHCP Local Server)	1054
relay-agent-remote-id (DHCP Local Server)	1055
routing-instance-name (DHCP Local Server)	1056
service-profile (DHCP Local Server)	1057
session-mode	1058
strict (DHCP Local Server)	1059
threshold (detection-time)	1060
threshold (transmit-interval)	1061
timeout (DHCP Local Server)	1062
token (DHCP Local Server)	1063
transmit-interval	1064
trigger (DHCP Local Server)	1065
use-primary (DHCP Local Server)	1066
user-prefix (DHCP Local Server)	1067
username-include (DHCP Local Server)	1068
version (bfd)	1069
DHCP Relay Agent Configuration Statements	1069
[edit forwarding-options dhcp-relay] Hierarchy Level	1069
access (Dynamic Access Routes)	1073
access-internal (Dynamic Access-Internal Routes)	1074
active-server-group	1075
allow-snooped-clients	1076
always-write-giaddr	1077
always-write-option-82	1078
authentication (DHCP Relay Agent)	1079
bfd	1080
circuit-id (DHCP Relay Agent)	1081
circuit-type (DHCP Relay Agent)	1082
client-discover-match (DHCP Relay Agent)	1083
client-id (DHCP Relay Agent)	1084
delimiter (DHCP Relay Agent)	1085
detection-time	1086
dhcp-relay	1087
dhcpv6 (DHCP Relay Agent)	1093
disable-relay	1096
domain-name (DHCP Relay Agent)	1097
drop (DHCP Relay Agent Option)	1098
duplicate-clients-on-interface (DHCP Relay Agent)	1099
dynamic-profile (DHCP Relay Agent)	1100
failure-action	1101
forward-snooped-clients (DHCP Relay Agent)	1102

group (DHCP Relay Agent)	1103
holddown-interval	1106
interface (DHCP Relay Agent)	1107
interface-client-limit (DHCP Relay Agent)	1109
interface-delete (Subscriber Management or DHCP Client Management)	1110
interface-name (DHCP Relay Agent)	1111
layer2-unicast-replies	1112
liveness-detection	1113
local-server-group (DHCP Relay Agent Option)	1114
logical-system-name (DHCP Relay Agent)	1115
mac-address (DHCP Relay Agent)	1116
method	1117
minimum-interval	1118
minimum-receive-interval	1119
multiplier	1120
next-hop (Dynamic Access-Internal Routes)	1121
no-adaptation	1122
no-allow-snooped-clients	1123
no-bind-on-request (DHCP Relay Agent)	1124
no-arp (DHCP Relay Agent)	1125
option-60 (DHCP Relay Agent)	1126
option-82 (DHCP Relay Agent)	1127
option-number (DHCP Relay Agent Option)	1128
overrides (DHCP Relay Agent)	1129
password (DHCP Relay Agent)	1131
preference (Subscriber Management)	1132
prefix (DHCP Relay Agent)	1133
proxy-mode	1135
relay-agent-interface-id (DHCPv6 Relay Agent)	1136
relay-agent-remote-id (DHCPv6 Relay Agent)	1137
relay-option (DHCP Relay Agent)	1138
relay-option-82	1139
relay-server-group (DHCP Relay Agent Option)	1140
replace-ip-source-with	1141
routing-instance-name (DHCP Relay Agent)	1142
send-release-on-delete (DHCP Relay Agent)	1143
server-group	1144
service-profile (DHCP Relay Agent)	1145
session-mode	1146
threshold (detection-time)	1147
threshold (transmit-interval)	1148
transmit-interval	1149
trust-option-82	1150
use-interface-description	1151
use-primary (DHCP Relay Agent)	1153
user-prefix (DHCP Relay Agent)	1155
username-include (DHCP Relay Agent)	1156
version (bfd)	1157

Administration	1157
Verifying and Managing DHCP Local Server Configurations	1158
Verifying and Managing DHCP Local Server Configuration	1158
Verifying and Managing DHCPv6 Local Server Configuration	1158
Verifying and Managing DHCP Relay Agent Configurations	1159
Verifying and Managing DHCP Relay Configuration	1159
Verifying and Managing DHCPv6 Relay Configuration	1159
DHCP Local Server Monitoring Commands	1159
clear dhcp server binding	1160
clear dhcp server statistics	1163
clear dhcpv6 server binding	1165
clear dhcpv6 server statistics	1167
request dhcp server reconfigure	1168
request dhcpv6 server reconfigure	1170
request system reboot	1172
show dhcp server binding	1177
show dhcp server statistics	1182
show dhcpv6 server binding	1185
show dhcpv6 server statistics	1190
DHCP Relay Agent Monitoring Commands	1192
clear dhcp relay binding	1193
clear dhcp relay statistics	1195
clear dhcpv6 relay binding	1198
clear dhcpv6 relay statistics	1201
show dhcp relay binding	1203
show dhcp relay statistics	1208
show dhcpv6 relay binding	1211
show dhcpv6 relay statistics	1216
show route extensive	1219
show route protocol	1234
Troubleshooting	1245
Acquiring Troubleshooting Information	1245
Tracing Extended DHCP Operations	1245
Tracing Extended DHCP Operations for Specific Interfaces	1251
Troubleshooting Configuration Statements	1253
interface-traceoptions (DHCP)	1254
trace (DHCP Local Server)	1256
trace (DHCP Relay Agent)	1257
traceoptions (DHCP)	1258

Part 3

Chapter 6

Features

Class of Service	1263
Class of Service Overview and Examples	1263
Overview	1263
CoS Overview	1263
CoS Input and Output Configuration	1272

Packet Flow Through the CoS Process	1273
Configuration	1275
Configuration Statements	1275
Class of Service on Ethernet Interfaces	1311
Overview	1311
Class of Service for Ethernet	1311
Configuration	1312
Configuration Task	1312
Class of Service for MPLS	1317
Overview	1317
CoS for MPLS	1317
Configuration	1318
Configuration Task	1318
Classifying Packets by Behavior Aggregate	1319
Overview	1319
Behavior Aggregate Classifier	1320
BA Classifier Default Values	1324
Configuration	1330
Configuration Tasks for Classifiers	1330
Configuration Tasks for BA Classifiers	1338
Configuration Task for DSCP IPv6 Classifiers	1341
Configuration Tasks for MPLS EXP Classifiers	1341
Configuration Task for IEEE 802.1ad Classifiers	1346
Configuration Statements	1347
Defining Code-Point Aliases	1371
Overview	1371
Code-Point Alias	1372
Configuration	1375
Configuration Task	1375
Configuration Statements	1376
Classifying Packets Based on Various Packet Header Fields	1386
Overview	1386
Overview	1386
Configuration	1390
Configuration Tasks	1390
Examples	1393
Configuration Statements	1402
Tricolor Marking Policers	1440
Overview	1440
Tricolor Marking Policers	1440
Configuration	1447
Configuration Tasks for Tricolor Marking Policers	1447
Configuration Tasks for Packet Loss Priority	1460
Configuration Statements for Tricolor Marking Policers	1463

Forwarding Classes	1525
Overview	1525
Forwarding Classes	1525
Configuration	1530
Configuration Tasks	1530
Configuration Statements	1542
Forwarding Policy Options	1556
Overview	1556
Forwarding Policy	1557
Configuration	1557
Configuration Tasks	1557
Examples	1561
Configuration Statements	1565
Schedulers	1573
Overview	1573
Schedulers	1573
Configuration	1578
Configuration Tasks for Schedulers	1578
Configuration Tasks for Scheduler Maps	1613
Configuration Statements for Schedulers	1624
Queue-Level Bandwidth Sharing	1666
Overview	1666
Bandwidth Sharing	1666
Configuration	1667
Configuration Task	1667
Example	1668
RED Drop Profiles	1673
Overview	1674
RED Drop Profiles	1674
Configuration	1677
Configuration Tasks	1677
Examples	1679
Configuration Statements	1681
Rewriting Packet Header Information	1689
Overview	1689
Rewriting Packet Header Information	1689
Configuration	1692
Configuration Tasks for Applying Rewrite Rules	1692
Configuration Tasks for Rewriting Packet Header Information	1702
Example	1713
Configuration Statements	1714
Routing Engine Protocol Queue Assignments	1745
Overview	1745
Routing Engine Protocol Queue Assignments	1745
Configuration	1748
Configuration Statements	1748

	CoS for Tunnels	1765
	Overview	1765
	CoS for Tunnels	1765
	Configuration	1766
	Configuration Task	1766
	Examples	1767
	Configuration Statements	1770
Chapter 7	Device Security	1803
	Overview	1803
	Rate Limiting	1803
	Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages	1803
	Configuring the Junos OS ICMPv6 Rate Limit for ICMPv6 Routing Engine Messages	1803
	Configuration	1804
	Configuration Statements	1804
	[edit system] Hierarchy Level	1804
	icmpv4-rate-limit	1818
	icmpv6-rate-limit	1819
Chapter 8	Ethernet Switching	1821
	Overview	1821
	Bridging	1821
	Layer 2 VLANs Overview	1821
	Layer 2 Learning and Forwarding Overview	1822
	Layer 2 Learning and Forwarding for VLANs Overview	1822
	Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port	1823
	Guidelines for Configuring VLAN Identifiers for VLANs and VPLS Routing Instances	1823
	MVRP	1824
	Understanding Multiple VLAN Registration Protocol (MVRP)	1824
	Proxy ARP	1826
	Restricted and Unrestricted Proxy ARP Overview	1827
	Configuration	1829
	Configuration Tasks	1829
	Configuring a VLAN	1829
	Configuring VLAN Identifiers for VLANs and VPLS Routing Instances	1830
	Configuring Integrated Routing and Bridging for VLANs	1835
	Configuring a Set of VLANs to Act as a Switch for a Layer 2 Trunk Port	1836
	Disabling MAC Learning for a VLAN or Logical Interface	1837
	Disabling MAC Learning for a Set of VLANs	1838
	Enabling MAC Accounting	1838
	Enabling MAC Accounting for a VLAN	1839
	Enabling MAC Accounting for a Set of VLANs	1839
	Configuring Inner and Outer TPIDs and VLAN IDs	1839
	Stacking a VLAN Tag	1843

Configuring the Size of the MAC Address Table	1843
Configuring Static MAC Addresses for Logical Interfaces in a VLAN . .	1844
Disabling Layer 2 Learning and Forwarding	1845
Configuring Multiple VLAN Registration Protocol (MVRP)	1845
Configuring a Layer 2 Virtual Switch	1847
Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port	1848
Configuring VLAN Encapsulation	1849
Configuring Restricted and Unrestricted Proxy ARP	1850
Rewriting a VLAN Tag and Adding a New Tag	1851
Configuring VLAN Translation with a VLAN ID List	1852
Configuring a Logical Interface for Access Mode	1853
Configuration Statements	1853
[edit dynamic-profiles] Hierarchy Level	1853
[edit interfaces] Hierarchy Level	1854
[edit logical-systems] Hierarchy Level	1865
[edit protocols l2-learning] Hierarchy Level	1870
Layer 2 Routing Instances Configuration Hierarchy	1870
[edit switch-options] Hierarchy Level	1873
[edit vlans] Hierarchy Level	1873
bpdu-destination-mac-address	1875
bridge-priority	1876
domain-type	1877
encapsulation (Physical Interface)	1878
family	1883
fast-aps-switch	1887
global-mac-limit	1888
global-mac-move	1889
global-mac-statistics	1889
global-mac-table-aging-time	1890
global-no-mac-learning	1890
inner-tag-protocol-id	1891
inner-vlan-id	1892
input-vlan-map (Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP, and 10-Gigabit Ethernet SFP)	1893
interface	1894
interface (MVRP)	1895
interface (Spanning Tree)	1896
interface-mac-limit	1897
interface-mode	1899
interfaces	1900
isis	1900
isis-list	1901
join-timer (MVRP)	1902
l2-learning	1903
l3-interface	1903
leaveall-timer (MVRP)	1904
leave-timer (MVRP)	1905
mac-statistics	1906
mac-table-size	1907

mvrp	1908
native-vlan-id	1909
no-mac-learning	1910
output-vlan-map (Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP, and 10-Gigabit Ethernet SFP)	1911
packet-action	1912
point-to-point (MVRP)	1914
pop	1915
pop-pop	1916
pop-swap	1917
proxy-arp	1918
push	1919
push-push	1920
registration	1921
rstp	1922
service-id	1923
static-mac	1924
swap	1925
swap-push	1926
swap-swap	1927
switch-options	1928
tag-protocol-id (TPID to Rewrite)	1929
traceoptions (MVRP)	1930
vlan-id (Bridge Domain or VLAN)	1932
vlan-id (VLAN ID to Rewrite)	1934
vlan-id-list	1935
vlan-rewrite	1936
vlan-tags	1937
Administration	1937
Routine Monitoring	1937
Verifying That MVRP Is Working Correctly	1937
Operational Commands	1939
show ethernet-switching flood	1940
show ethernet-switching interface	1944
show ethernet-switching statistics	1946
show ethernet-switching table	1949
show interfaces (10-Gigabit Ethernet)	1956
show interfaces (Gigabit Ethernet)	1981
show interfaces irb	2004
show interfaces queue	2011
show mvrp	2047
show mvrp applicant-state	2049
show mvrp dynamic-vlan-memberships	2051
show mvrp interface	2052
show mvrp registration-state	2053
show mvrp statistics	2055
show vlans	2057
traceroute ethernet	2059

Chapter 9	High Availability	2061
	Overview	2061
	Graceful Routing Engine Switchover (GRES)	2061
	Understanding Graceful Routing Engine Switchover in the Junos OS	2061
	Graceful Routing Engine Switchover System Requirements	2065
	Requirements for Routers with a Backup Router Configuration	2068
	Nonstop Bridging (NSB)	2069
	Nonstop Bridging Concepts	2069
	Nonstop Bridging System Requirements	2071
	Nonstop Active Routing (NSR)	2072
	Nonstop Active Routing Concepts	2072
	Nonstop Active Routing System Requirements	2075
	Graceful Restart	2085
	Graceful Restart Concepts	2085
	Graceful Restart System Requirements	2086
	Aggregate and Static Routes	2086
	Graceful Restart and Routing Protocols	2087
	Graceful Restart and MPLS-Related Protocols	2089
	Graceful Restart and Layer 2 and Layer 3 VPNs	2091
	Graceful Restart on Logical Systems	2092
	Unified ISSU	2092
	Upgrading Routers Using ISSU	2092
	Unified ISSU Concepts	2092
	Unified ISSU System Requirements	2097
	VRRP	2109
	Understanding VRRP	2110
	Junos OS Support for VRRPv3	2111
	Improving the Convergence Time for VRRP	2114
	Configuration	2116
	Configuration: GRES	2116
	Configuring Graceful Routing Engine Switchover	2117
	Resetting Local Statistics	2118
	Configuration Statements: GRES	2118
	[edit chassis] Hierarchy Level	2118
	graceful-switchover	2127
	Configuration: Graceful Restart	2127
	Enabling Graceful Restart	2127
	Configuring Routing Protocols Graceful Restart	2128
	Configuring Graceful Restart for MPLS-Related Protocols	2134
	Configuring VPN Graceful Restart	2136
	Configuring Logical System Graceful Restart	2137
	Example: Configuring Graceful Restart	2139
	Example: Managing Helper Modes for OSPF Graceful Restart	2164
	Configuration Statements: Graceful Restart	2167
	disable	2167
	graceful-restart (Enabling Globally)	2168
	helper-disable (Multiple Protocols)	2169
	maximum-helper-recovery-time	2169
	maximum-helper-restart-time (RSVP)	2170

maximum-neighbor-reconnect-time	2170
maximum-neighbor-recovery-time	2171
no-strict-lsa-checking	2172
notify-duration	2173
reconnect-time	2174
recovery-time	2175
restart-duration	2176
restart-time (BGP Graceful Restart)	2177
stale-routes-time	2178
traceoptions (Protocols)	2179
Configuration: NSB	2180
Configuring Nonstop Bridging	2180
Resetting Local Statistics	2181
Configuration Statements: NSB	2181
[edit protocols layer2-control] Hierarchy Level	2181
nonstop-bridging	2182
Configuration: NSR	2182
Configuring Nonstop Active Routing	2183
Tracing Nonstop Active Routing Synchronization Events	2185
traceoptions (Routing Options)	2187
Example: Configuring Nonstop Active Routing	2189
Configuration Statements: NSR	2191
[edit protocols layer2-control] Hierarchy Level	2191
commit synchronize	2193
nonstop-routing	2194
traceoptions (Routing Options)	2195
Configuration: Unified ISSU	2197
Best Practices	2197
Before You Begin	2197
Performing a Unified ISSU	2200
Verifying a Unified ISSU	2212
Managing and Tracing BFD Sessions During Unified ISSU	
Procedures	2212
Configuration Statements: Unified ISSU	2214
bfd	2215
no-issu-timer-negotiation	2217
traceoptions (Protocols BFD)	2218
Configuration: VRRP	2219
Configuring the Startup Period for VRRP Operations	2219
Configuring Basic VRRP Support	2220
Configuring VRRP Authentication (IPv4 Only)	2222
Configuring the Advertisement Interval for the VRRP Master Router	2224
Configuring a Backup Router to Preempt the Master Router	2226
Modifying the Preemption Hold-Time Value	2227
Configuring Asymmetric Hold Time for VRRP Routers	2227
Configuring an Interface to Accept Packets Destined for the Virtual IP	
Address	2228
Configuring a Logical Interface to Be Tracked	2229
Configuring a Route to Be Tracked	2231

Configuring Inheritance for a VRRP Group	2232
Configuring the Silent Period	2233
Configuring Passive ARP Learning for Backup VRRP Routers	2234
Enabling the Distributed Periodic Packet Management Process for VRRP	2235
Configuring VRRP to Improve Convergence Time	2236
Example: Configuring VRRP	2237
Example: Configuring VRRP for IPv6	2239
Example: Configuring VRRP Route Tracking	2240
Tracing VRRP Operations	2241
Configuration Statements: VRRP	2242
[edit protocols vrrp] Hierarchy Level	2242
accept-data	2244
advertise-interval	2245
asymmetric-hold-time	2246
authentication-key	2247
authentication-type	2248
bandwidth-threshold	2249
delegate-processing (VRRP)	2250
failover-delay	2250
fast-interval	2251
global-advertisements-threshold	2252
hold-time (VRRP)	2253
inet6-advertise-interval	2254
interface (VRRP Group)	2255
preempt (VRRP)	2256
priority (Protocols VRRP)	2257
priority-cost (VRRP)	2258
priority-hold-time	2259
route (Interfaces)	2260
skew-timer-disable	2261
startup-silent-period	2261
traceoptions (Protocols VRRP)	2262
track (VRRP)	2264
version-3	2265
virtual-address	2266
virtual-inet6-address	2266
virtual-link-local-address	2267
vrrp-group	2268
vrrp-inet6-group	2269
Administration	2270
Routine Monitoring	2270
Resetting Local Statistics	2270
Tracing Restart Signaling-Based Helper Mode Events for OSPF Graceful Restart	2270
Verifying Graceful Restart Operation	2271
Operational Commands	2274
show system switchover	2275

	Troubleshooting	2278
	Troubleshooting Unified ISSU	2278
	Troubleshooting Unified ISSU Problems	2278
Chapter 10	Interfaces	2279
	Overview	2279
	Aggregated Ethernet Overview	2279
	Aggregated Ethernet Interfaces Overview	2279
	Load Balancing and Ethernet Link Aggregation Overview	2281
	IP-Directed Broadcast Overview	2281
	Understanding Targeted Broadcast	2282
	Reverse Path Forwarding	2283
	Understanding Unicast Reverse Path Forwarding	2283
	Understanding Multicast Reverse Path Forwarding	2283
	Configuration	2285
	Configuration	2285
	Configuring a Layer 2 Virtual Switch	2286
	Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port	2287
	Understanding Layer 2 Virtual Switches Instances	2287
	Configuring VLAN Encapsulation	2288
	Rewriting the Inner and Outer VLAN Tags	2289
	Rewriting the VLAN Tag on Tagged Frames	2290
	Binding VLAN IDs to Logical Interfaces	2291
	Configuring a Logical Interface for Access Mode	2296
	Configuring Junos OS for Supporting Aggregated Devices	2296
	Configuring an Aggregated Ethernet Interface	2299
	Deleting an Aggregated Ethernet Interface	2300
	Configuring the Number of Aggregated Ethernet Interfaces on the Device (Enhanced Layer 2 Software CLI Procedure)	2300
	Example: Configuring Aggregated Ethernet Interfaces	2301
	Configuring Tagged Aggregated Ethernet Interfaces	2302
	Configuring Untagged Aggregated Ethernet Interfaces	2303
	Configuring Aggregated Ethernet Minimum Links	2303
	Configuring Load Balancing on a LAG Link	2304
	Example: Configuring Load Balancing on a LAG Link	2305
	Configuring Multichassis Link Aggregation	2305
	Configuring Aggregated Ethernet LACP	2307
	Configuring Targeted Broadcast	2313
	Configuring Unicast RPF	2314
	Configuration Statements	2320
	[edit chassis] Hierarchy Level	2320
	[edit dynamic-profiles] Hierarchy Level	2328
	[edit forwarding-options rpf-loose-mode-discard] Hierarchy Level	2329
	[edit interfaces] Hierarchy Level	2329
	[edit multi-chassis] Hierarchy Level	2340
	[edit protocols isis] Hierarchy Level	2340
	Layer 2 Routing Instances Configuration Hierarchy	2343
	bandwidth (Interfaces)	2346
	chassis-id	2347

filter	2348
flow-control	2349
forward-and-send-to-re	2350
forward-only	2350
gratuitous-arp-reply	2351
group (RPF Selection)	2352
l2-domain-id-for-l3	2352
lacp (Aggregated Ethernet)	2353
layer3-domain-identifier	2354
link-protection	2355
mc-ae-id	2356
mode (QFX Series)	2356
mode (Interfaces)	2357
multicast-rpf-routes	2357
next-hop (PIM RPF Selection)	2358
no-gratuitous-arp-request	2358
no-local-switching	2359
policer (MAC)	2360
prefix-list (PIM RPF Selection)	2361
redundancy-group	2362
rpf-check (Dynamic Profiles)	2362
rpf-check (interfaces)	2363
rpf-check-policy (Routing Options RPF)	2364
rpf-loose-mode-discard	2365
rpf-selection	2366
source (PIM RPF Selection)	2367
status-control	2367
targeted-broadcast	2368
traceoptions (Individual Interfaces)	2369
traps	2370
unicast-reverse-path	2371
unidirectional	2372
vlan-tagging	2373
wildcard-source (PIM RPF Selection)	2374
Administration	2374
Routine Monitoring	2374
Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface	2374
Operational Commands	2376
Common Output Fields Description	2376
clear interfaces statistics	2383
show interfaces (10-Gigabit Ethernet)	2384
show interfaces (Discard)	2409
show interfaces (Gigabit Ethernet)	2414
show interfaces (Serial)	2437
show interfaces extensive	2450
show interfaces queue	2469

Chapter 11	Layer 3 Protocols	2505
	BGP	2505
	Overview	2505
	Feature Support	2505
	Configuration	2553
	Configuration Statements	2553
	Administration	2640
	Operational Commands	2640
	IS-IS	2675
	Configuration	2675
	Configuration Statements	2675
	Administration	2735
	Operational Commands	2735
	OSPF	2784
	Configuration	2784
	Configuration Statements	2784
	Administration	2855
	Operational Commands	2855
	RIP and RIPvng	2915
	Overview	2915
	RIP	2915
	Configuration	2921
	Configuration Examples	2921
	Configuration Statements	2996
	Administration	3030
	RIP Monitoring	3031
	Operational Commands	3039
	Routing Options	3193
	Overview	3193
	Routing Properties Overview	3193
	Configuration	3193
	Configuration Tasks	3193
	Configuration Statements	3216
	Administration	3232
	Operational Commands	3232
Chapter 12	LLDP	3243
	Overview	3243
	LLDP	3243
	LLDP Overview	3243
	Configuration	3244
	Configuration Examples	3244
	Example: Configuring LLDP	3244
	Configuration Tasks	3245
	Configuring LLDP	3246
	Configuration Statements	3249
	lldp	3249
	advertisement-interval	3250
	disable	3251

	hold-multiplier	3251
	interface	3252
	lldp-configuration-notification-interval	3253
	ptopo-configuration-maximum-hold-time	3253
	ptopo-configuration-trap-interval	3254
	traceoptions	3255
	transmit-delay	3257
	Administration	3257
	Routine Monitoring	3257
	Tracing LLDP Operations	3257
Chapter 13	Logical Systems	3259
	Overview	3259
	Logical Systems	3259
	Introduction to Logical Systems	3259
	Junos OS Features That Are Supported on Logical Systems	3262
	Logical Systems Operations and Restrictions	3263
	Comparing Junos OS Device Virtualization Technologies	3265
	Logical Systems Applications	3266
	Logical Systems Requirements	3267
	Logical Systems Terms and Acronyms	3267
	Configuration	3268
	Configuration Examples	3268
	Example: Configuring User Access for Logical Systems	3268
	Examples: Configuring Logical System Interfaces	3274
	Example: Configuring Static Routing on Logical Systems	3287
	Examples: Configuring Standard Firewall Filters on Logical Systems ..	3293
	Example: Configuring OSPF on Logical Systems	3307
	Examples: Configuring OSPF Routing Policy on Logical Systems ...	3320
	Examples: Configuring IS-IS on Logical Systems	3343
	Examples: Configuring BGP on Logical Systems	3362
	Example: Configuring RSVP-Signaled Point-to-Multipoint LSPs on	
	Logical Systems	3407
	Example: Configuring VPNs and VPLS on Logical Systems	3431
	Example: Configuring a Virtualized Data Center	3478
	Configuration Statements	3521
	[edit logical-systems] Hierarchy Level	3521
	Administration	3523
	Operational Commands	3523
	Operational-Mode Commands	3523
Chapter 14	MPLS	3535
	Overview	3535
	LDP	3535
	LDP Introduction	3535
	Junos OS LDP Protocol Implementation	3536
	LDP Operation	3536
	Label Operations	3536
	LDP Message Types	3538
	Discovery Messages	3538

Session Messages	3538
Advertisement Messages	3538
Notification Messages	3539
LDP Session Protection	3539
LDP Graceful Restart	3539
Configuration	3540
Configuration Tasks	3540
Minimum LDP Configuration	3541
Enabling and Disabling LDP	3541
Configuring the LDP Timer for Hello Messages	3542
Configuring the Delay Before LDP Neighbors Are Considered Down	3543
Enabling Strict Targeted Hello Messages for LDP	3544
Configuring the Interval for LDP Keepalive Messages	3544
Configuring the LDP Keepalive Timeout	3545
Configuring LDP Route Preferences	3545
Configuring LDP Graceful Restart	3545
Filtering Inbound LDP Label Bindings	3548
Filtering Outbound LDP Label Bindings	3550
Specifying the Transport Address Used by LDP	3552
Configuring the Prefixes Advertised into LDP from the Routing Table	3552
Configuring FEC Deaggregation	3553
Configuring Policers for LDP FECs	3554
Configuring LDP IPv4 FEC Filtering	3554
Configuring BFD for LDP LSPs	3555
Configuring ECMP-Aware BFD for LDP LSPs	3558
Configuring a Failure Action for the BFD Session on an LDP LSP	3558
Configuring the Holddown Interval for the BFD Session	3559
Configuring OAM Ingress Policies for LDP	3559
Configuring LDP LSP Traceroute	3560
Collecting LDP Statistics	3561
Tracing LDP Protocol Traffic	3563
Standard Firewall Filter Match Conditions for MPLS Traffic	3565
Configuration Statements	3567
[edit protocols bgp] Hierarchy Level	3567
[edit protocols ldp] Hierarchy Level	3573
[edit protocols mpls] Hierarchy Level	3575
allow-subnet-mismatch	3576
authentication-algorithm	3577
authentication-key (Protocols LDP)	3578
bfd-liveness-detection (Protocols LDP)	3579
deaggregate	3580
disable (Protocols LDP)	3581
dod-request-policy	3582
downstream-on-demand	3582
ecmp	3583
egress-policy	3583
explicit-null (Protocols LDP)	3584
export (Protocols LDP)	3584

failure-action (Protocols LDP)	3585
fec	3586
graceful-restart (Protocols LDP)	3587
hello-interval (Protocols LDP)	3588
helper-disable (LDP)	3589
holddown-interval	3589
hold-time (Protocols LDP)	3590
igp-synchronization	3591
import (Protocols LDP)	3591
ingress-policy	3592
interface (Protocols LDP)	3593
keepalive-interval	3594
keepalive-timeout	3594
l2-smart-policy	3595
label-withdrawal-delay	3595
ldp	3596
ldp-p2mp	3599
log-updown (Protocols LDP)	3600
make-before-break (LDP)	3601
maximum-neighbor-recovery-time	3602
no-forwarding	3603
oam (Protocols LDP)	3604
p2mp (Protocols LDP)	3605
periodic-traceroute	3606
policing (Protocols LDP)	3608
preference (Protocols LDP)	3609
reconnect-time	3609
recovery-time	3610
session (ldp)	3610
session-protection	3611
strict-targeted-hellos	3611
targeted-hello	3612
traceoptions (Protocols LDP)	3613
track-igp-metric	3615
traffic-statistics (Protocols LDP)	3616
transport-address	3617
Administration	3617
Operational Commands	3617
ping mpls ldp	3618
show ldp database	3621
show ldp session	3625
show ldp traffic-statistics	3630
show ldp session	3632

Chapter 15	Multicast	3637
	Overview	3637
	Multicast Overview	3637
	Multicast Overview	3637
	Multicast Protocols Overview	3649
	Supported IP Multicast Protocol Standards	3649
	Understanding MLD	3650
	PIM Overview	3653
	Configuration	3656
	Configuration	3656
	Configuring IGMP	3657
	Configuring Multiple Instances of MSDP	3679
	Configuring Basic PIM Settings	3680
	Configuring Multiple Instances of PIM	3692
	Configuring a Designated Router for PIM	3692
	Configuring Static RP	3694
	Configuring PIM Bootstrap Router	3701
	Configuring PIM Auto-RP	3705
	Configuring Embedded RP	3709
	Configuring PIM Filtering	3712
	Configuring PIM and the Bidirectional Forwarding Detection (BFD)	
	Protocol	3725
	Configuring PIM Dense Mode	3737
	Configuring PIM Sparse-Dense Mode	3740
	PIM Join Load Balancing on Multipath MVPN Routes Overview	3741
	PIM Snooping for VPLS	3745
	Configuration Statements: IGMP	3757
	igmp	3758
	accounting (Protocols IGMP Interface)	3759
	accounting (Protocols IGMP)	3759
	disable (Protocols IGMP)	3760
	exclude (Protocols IGMP)	3760
	group (Protocols IGMP)	3761
	group-count (Protocols IGMP)	3762
	group-increment (Protocols IGMP)	3762
	group-limit	3763
	group-policy (Protocols IGMP)	3764
	group-threshold (Protocols IGMP Interface)	3765
	immediate-leave (Protocols IGMP)	3766
	interface (Protocols IGMP)	3767
	log-interval (Protocols IGMP Interface)	3768
	maximum-transmit-rate (Protocols IGMP)	3769
	oif-map	3769
	passive (IGMP)	3770
	promiscuous-mode (Protocols IGMP)	3771
	query-interval (Protocols IGMP)	3771
	query-last-member-interval (Protocols IGMP)	3772
	query-response-interval (Protocols IGMP)	3773
	robust-count (Protocols IGMP)	3774

source (Protocols IGMP)	3775
source-count (Protocols IGMP)	3776
source-increment (Protocols IGMP)	3776
ssm-map (Protocols IGMP)	3777
ssm-map-policy (IGMP)	3777
static (Protocols IGMP)	3778
traceoptions (Protocols IGMP)	3779
version (Protocols IGMP)	3781
Configuration Statements: IGMP Snooping	3782
igmp-snooping	3782
group (Bridge Domains)	3783
group-limit	3784
host-only-interface	3785
immediate-leave (Bridge Domains)	3786
interface (Bridge Domains)	3787
multicast-router-interface (IGMP Snooping)	3788
proxy (Bridge Domains)	3789
query-interval (Bridge Domains)	3790
query-last-member-interval (Bridge Domains)	3791
query-response-interval (Bridge Domains)	3792
robust-count (Bridge Domains)	3793
source (Bridge Domains)	3794
source-address	3794
static (Bridge Domains)	3795
traceoptions (Protocols IGMP Snooping)	3796
vlan (Bridge Domains)	3798
Configuration Statements: MLD	3799
mld	3800
accounting (Protocols MLD Interface)	3801
accounting (Protocols MLD)	3801
disable (Protocols MLD)	3802
exclude (Protocols MLD)	3802
group (Protocols MLD)	3803
group-count (Protocols MLD)	3804
group-increment (Protocols MLD)	3804
group-limit	3805
group-policy (Protocols MLD)	3805
group-threshold (Protocols MLD Interface)	3806
immediate-leave (Protocols MLD)	3807
interface (Protocols MLD)	3808
log-interval (Protocols MLD Interface)	3809
maximum-transmit-rate (Protocols MLD)	3810
oif-map	3810
passive (MLD)	3811
query-interval (Protocols MLD)	3812
query-last-member-interval (Protocols MLD)	3812
query-response-interval (Protocols MLD)	3813
robust-count (Protocols MLD)	3813
source (Protocols MLD)	3814

source-count (Protocols MLD)	3814
source-increment (Protocols MLD)	3815
ssm-map (Protocols MLD)	3815
ssm-map-policy (MLD)	3816
static (Protocols MLD)	3817
traceoptions (Protocols MLD)	3818
version (Protocols MLD)	3820
Configuration Statements: MSDP	3821
msdp	3822
active-source-limit	3824
authentication-key	3825
data-encapsulation	3826
default-peer	3827
disable (Protocols MSDP)	3828
export (Protocols MSDP)	3829
group	3830
hold-time (Protocols MSDP)	3831
import (Protocols MSDP)	3832
keep-alive (Protocols MSDP)	3833
local-address	3834
log-interval (Protocols MSDP)	3835
log-warning (Protocols MSDP)	3836
maximum	3837
mode (Protocols MSDP)	3838
peer (Protocols MSDP)	3839
rib-group (Protocols MSDP)	3840
sa-hold-time (Protocols MSDP)	3841
source	3842
threshold	3843
traceoptions (Protocols MSDP)	3844
Configuration Statements: PIM	3846
[edit protocols pim] Hierarchy Level	3846
accept-remote-source	3851
address (Anycast RPs)	3852
address (Bidirectional Rendezvous Points)	3853
address (Local RPs)	3854
address (Static RPs)	3855
algorithm	3856
anycast-pim	3857
assert-timeout	3858
authentication (Protocols PIM)	3859
auto-rp	3860
backoff-period	3861
bfd-liveness-detection (Protocols PIM)	3862
bidirectional (Interface)	3863
bidirectional (RP)	3864
bootstrap	3865
bootstrap-export	3866
bootstrap-import	3867

bootstrap-priority	3868
dense-groups	3869
detection-time (BFD for PIM)	3870
df-election	3871
disable (PIM Graceful Restart)	3871
disable (PIM)	3872
dr-election-on-p2p	3873
dr-register-policy	3873
embedded-rp	3874
export (Protocols PIM Bootstrap)	3875
export (Protocols PIM)	3875
family (Bootstrap)	3876
family (Protocols PIM)	3877
family (Protocols PIM Interface)	3878
family (Local RP)	3879
graceful-restart (Protocols PIM)	3880
group (RPF Selection)	3881
group-ranges	3882
group-rp-mapping	3883
hello-interval (Protocols PIM)	3884
hold-time (Protocols PIM)	3885
idle-standby-path-switchover-delay	3886
import (Protocols PIM Bootstrap)	3887
import (Protocols PIM)	3888
infinity	3889
interface (Protocols PIM)	3890
join-load-balance	3892
join-prune-timeout	3893
key-chain (Protocols PIM)	3894
local	3895
local-address (Protocols PIM)	3896
log-interval (PIM Entries)	3897
loose-check	3898
mapping-agent-election	3899
maximum (PIM Entries)	3900
maximum-rps	3901
minimum-interval (PIM BFD Liveness Detection)	3902
minimum-interval (PIM BFD Transmit Interval)	3903
minimum-receive-interval	3904
mode (Protocols PIM)	3905
multiplier	3906
neighbor-policy	3906
next-hop (PIM RPF Selection)	3907
no-adaptation (PIM BFD Liveness Detection)	3907
no-bidirectional-mode	3908
no-dr-flood (PIM Snooping)	3909
offer-period	3910
override (PIM static RP)	3911
override-interval	3912

pim	3913
pim-snooping	3917
prefix-list (PIM RPF Selection)	3918
priority (Bootstrap)	3919
priority (PIM Interfaces)	3920
priority (PIM RPs)	3921
propagation-delay	3922
register-limit	3923
reset-tracking-bit	3924
restart-duration (Protocols PIM)	3925
rib-group (Protocols PIM)	3926
robustness-count	3927
rp	3928
rp-register-policy	3930
rp-set	3931
rpf-selection	3932
sglimit	3933
source (PIM RPF Selection)	3934
spt-threshold	3935
standby-path-creation-delay	3936
static (Protocols PIM)	3937
threshold (PIM BFD Detection Time)	3938
threshold (PIM BFD Transmit Interval)	3939
threshold (PIM Entries)	3940
traceoptions (Protocols PIM)	3942
traceoptions (PIM Snooping)	3945
transmit-interval (PIM BFD Liveness Detection)	3947
tunnel-devices	3948
version (BFD)	3949
version (PIM)	3950
vlan (PIM Snooping)	3951
vpn-group-address	3951
wildcard-source (PIM RPF Selection)	3952
Administration	3952
Operational Commands: IGMP	3952
clear igmp statistics	3953
show igmp group	3955
show igmp interface	3959
show multicast pim-to-igmp-proxy	3963
Operational Commands: IGMP Snooping	3964
clear igmp snooping membership	3965
clear igmp snooping statistics	3966
show igmp snooping interface	3967
show igmp snooping membership	3970
show igmp snooping statistics	3974
Operational Commands: MLD	3976
clear mld membership	3977
clear mld statistics	3978
show mld group	3979

show mld interface	3983
show mld statistics	3987
show multicast pim-to-mld-proxy	3990
Operational Commands: MSDP	3991
show msdp	3992
show msdp source	3994
show msdp source-active	3996
show msdp statistics	3999
show multicast usage	4003
show route table	4006
Operational Commands: PIM	4016
clear pim join	4017
clear pim join-distribution	4018
clear pim register	4020
clear pim statistics	4022
request pim multicast-tunnel rebalance	4024
show pim bidirectional df-election	4025
show pim bidirectional df-election interface	4028
show pim bootstrap	4031
show pim interfaces	4033
show pim join	4036
show pim neighbors	4045
show pim rps	4049
show pim source	4056
show pim statistics	4059
Chapter 16	Network Management and Monitoring 4073
Overview	4073
Ethernet OAM Link Fault Management	4073
IEEE 802.3ah OAM Link-Fault Management Overview	4073
Configuration	4074
Configuration: Ethernet OAM Link Fault Management	4074
Configuring IEEE 802.3ah OAM Link-Fault Management	4075
Enabling IEEE 802.3ah OAM Support	4076
Configuring Link Discovery	4077
Configuring the OAM PDU Interval	4078
Configuring the OAM PDU Threshold	4079
Configuring Threshold Values for Local Fault Events on an Interface	4079
Disabling the Sending of Link Event TLVs	4080
Detecting Remote Faults	4081
Configuring an OAM Action Profile	4082
Specifying the Actions to Be Taken for Link-Fault Management	
Events	4083
Monitoring the Loss of Link Adjacency	4085
Monitoring Protocol Status	4086
Configuring Threshold Values for Fault Events in an Action Profile	4087
Applying an Action Profile	4088
Setting a Remote Interface into Loopback Mode	4088

Enabling Remote Loopback Support on the Local Interface	4089
Example: Configuring IEEE 802.3ah OAM Support on an Interface . .	4090
Configuration Statements	4092
link-fault-management	4093
[edit services flow-monitoring] Hierarchy Level	4094
action (OAM)	4095
action-profile (Defining for LFM)	4096
allow-remote-loopback	4097
apply-action-profile	4097
ethernet (Protocols OAM)	4098
event (LFM)	4101
event-thresholds	4101
fast-aps-switch	4102
frame-error	4103
frame-period	4104
frame-period-summary	4105
interface (OAM Link-Fault Management)	4106
link-adjacency-loss	4107
link-down	4107
link-discovery	4108
link-event-rate	4108
negotiation-options	4109
no-allow-link-events	4109
oam	4110
pdu-interval	4112
pdu-threshold	4113
protocol-down	4113
remote-loopback	4114
send-critical-event	4114
symbol-period	4115
syslog (OAM Action)	4115
version-ipfix (Services)	4116
Administration	4116
Routine Monitoring	4116
Monitoring Traffic Through the Router or Switch	4116
Operational Commands: General	4119
monitor traffic	4120
ping	4130
traceroute	4134
traceroute monitor	4138
Operational Commands: Port Mirroring	4139
show analyzer	4140
Operational Commands: Ethernet OAM Link Fault Management	4141
show oam ethernet link-fault-management	4142
show interfaces (Fast Ethernet)	4148
show interfaces (10-Gigabit Ethernet)	4165

Chapter 17	Port Mirroring	4191
	Overview	4191
	Port Mirroring Overview	4191
	Layer 2 Port Mirroring Overview	4192
	Layer 2 Port Mirroring Properties	4192
	Layer 2 Port Mirroring Global Instance	4194
	Layer 2 Port Mirroring Named Instances	4194
	Layer 2 Port Mirroring of PE Router Logical Interfaces	4196
	Layer 2 Port Mirroring Firewall Filters	4198
	Layer 2 Port Mirroring to Multiple Destinations Using Next-Hop Groups	4200
	Configuration Guidelines for Layer 2 Port Mirroring	4201
	Application of Layer 2 Port Mirroring Types	4201
	Restrictions on Layer 2 Port Mirroring	4203
	Behavior of Layer 2 Port Mirroring at Physical Interfaces	4204
	Precedence of Multiple Levels of Layer 2 Port Mirroring on a Physical Interface	4204
	Behavior of Layer 2 Port Mirroring on PE Routers and PE Switches	4205
	Layer 2 Port Mirroring of PE Router or PE Switch Logical Interfaces	4205
	Layer 2 Port Mirroring of PE Router or PE Switch Aggregated Ethernet Interfaces	4206
	Configuration	4207
	Configuration Tasks for Layer 2 Port Mirroring at Physical Interfaces	4207
	Configuring the Global Instance of Layer 2 Port Mirroring	4208
	Defining a Named Instance of Layer 2 Port Mirroring	4210
	Binding Layer 2 Port Mirroring to Ports Grouped at the FPC Level	4214
	Binding Layer 2 Port Mirroring to Ports Grouped at the PIC Level	4216
	Disabling Layer 2 Port Mirroring Instances	4217
	Examples for Layer 2 Port Mirroring at Physical Interfaces	4218
	Examples: Layer 2 Port-Mirroring at Multiple Levels of the Chassis	4218
	Example: Layer 2 Port Mirroring with Multiple Instances	4220
	Configuration Tasks for Layer 2 Port Mirroring at Logical Interfaces	4224
	Defining a Layer 2 Port-Mirroring Firewall Filter	4224
	Applying Layer 2 Port Mirroring to a Logical Interface	4228
	Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a VLAN	4231
	Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a VPLS Routing Instance	4233
	Examples for Layer 2 Port Mirroring at Logical Interfaces	4234
	Example: Layer 2 Port Mirroring at a Logical Interface	4235
	Example: Layer 2 Port Mirroring for a Layer 2 VPN	4237
	Example: Layer 2 Port Mirroring for a Layer 2 VPN with LAG Links	4239
	Configuration Tasks for Layer 2 Port Mirroring at Multiple Destinations	4242
	Defining a Layer 2 Port-Mirroring Firewall Filter	4242
	Defining a Next-Hop Group for Layer 2 Port Mirroring	4245
	Applying Layer 2 Port Mirroring to a Logical Interface	4247
	Examples of Layer 2 Port Mirroring at Multiple Destinations	4249
	Example: Layer 2 Port Mirroring to Multiple Destinations	4249

Configuring Inline Port Mirroring	4252
Configuring Inline Port Mirroring	4252
Configuration Statements	4253
[edit forwarding-options port-mirroring] Hierarchy Level	4253
disable (Forwarding Options)	4255
disable-all-instances	4255
forwarding-options	4256
family (Port Mirroring)	4256
input (Port Mirroring)	4257
instance	4258
interface (Port Mirroring)	4259
interface (Next-Hop Group)	4259
maximum-packet-length	4260
mirror-once	4260
next-hop-group (Port Mirroring)	4261
no-filter-check	4261
output (Port Mirroring)	4262
rate	4262
run-length	4263
Administration	4263
Displaying Information	4263
Displaying Layer 2 Port-Mirroring Instance Settings and Status	4263
Displaying Next-Hop Group Settings and Status	4264
Operational Mode Commands for	
Packet Forwarding Engine Components	4264
show chassis fabric fpcs	4265
show chassis fpc	4298
show chassis hardware	4327
show chassis pic	4433
Operational Mode Commands for Layer 2 Port-Mirroring Instances	4445
show forwarding-options port-mirroring	4446
Operational Mode Commands for Firewall Filter Statistics and Logs	4447
clear firewall	4448
show firewall	4450
show firewall log	4457
Operational Mode Commands for	
Next-Hop Groups for Layer 2 Port Mirroring	4459
show forwarding-options next-hop-group	4460
Chapter 18 Routing Policy and Packet Filtering	4463
Firewall Filters	4463
Overview	4463
Introduction to Stateless Firewall Filters	4463
Standard Firewall Filters Overview	4474
Standard Firewall Filter Match Conditions Overview	4488
Introduction to Standard Firewall Filters for Fragment Handling	4503
Introduction to Standard Firewall Filters Configuration	4503
Introduction to Service Filters Configuration	4521
Introduction to Simple Filters Configuration	4529

Introduction to Firewall Filters Configuration in Logical Systems	4535
Configuration	4547
Standard Firewall Filter Configurations to Match Packets	4548
Standard Firewall Filters to Count Packets	4558
Standard Firewall Filters to Act on Packets	4569
Standard Firewall Filters for Trusted Sources	4576
Standard Firewall Filters for Flood Prevention	4600
Standard Firewall Filters for Fragment Handling	4613
Standard Firewall Filters for Setting Rate Limits	4617
Examples of Standard Firewall Filters Configuration	4620
Examples of Standard Firewall Filters Configuration Options	4669
Service Filters Configuration	4678
Simple Filters Configuration	4683
Firewall Filters Configuration in Logical Systems	4687
Configuration Statements	4691
Administration	4703
Firewall Filters Standards	4703
Firewall Filters Reference	4704
Standard Firewall Filter Match Conditions and Actions	4706
Standard Firewall Filter Match Conditions and Actions for ACX Series Routers	4751
Service Filter Match Conditions and Actions	4760
Reference Information for Firewall Filters in Logical Systems	4767
Firewall Filters Statement Hierarchies	4773
Summary of Firewall Filters Configuration Statements	4785
Traffic Policers	4802
Overview	4802
Introduction to Traffic Policing	4802
Introduction to Configuring Policers	4811
Policer Rate Limits and Actions	4821
Policer Implementation	4828
Configuration	4838
Configuring Single-Rate Two-Color Policers	4838
Configuring Three-Color Policers	4912
Configuring Logical and Physical Interface Policers	4927
Configuring Layer 2 Policers	4947
Configuration Statements	4956
Administration	5013
Traffic Policing Standards	5013
Firewall Filter and Policer Operational Mode Commands	5013
Chapter 19	
Spanning-Tree Protocols	5031
Overview	5031
Spanning-Tree	5031
Spanning-Tree Protocols Supported	5032
BPDU Overview	5033
Loop Protection for Spanning-Tree Instance Interfaces Overview . . .	5034
Root Protection for Spanning-Tree Instance Interfaces Overview . . .	5035
BPDU Protection for Spanning-Tree Instance Interfaces Overview . .	5035

VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Overview . .	5036
VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Topology . .	5037
Layer 2 Protocol Tunneling Through a Network Overview	5039
Spanning-Tree Protocols in Logical Systems	5041
Bridge Priority for Election of Root Bridge and Designated Bridge . . .	5041
Maximum Age for Awaiting Arrival of Hello BPDUs	5042
Hello Time for Root Bridge to Transmit Hello BPDUs	5042
Forward Delay Before Ports Transition to Forwarding State	5042
Spanning-Tree Instance Interface	5042
Spanning-Tree Instance Interface Priority	5043
Spanning-Tree Instance Interface Cost	5044
Spanning-Tree Instance Interface Point-to-Point Link Mode	5044
Spanning-Tree Instance Interface Configured as an Edge Port	5045
Spanning-Tree Protocol Trace Options	5045
MAC Address Rewriting Enabled for Layer 2 Protocol Tunneling . . .	5046
Layer 2 Protocol Tunnel Interface	5046
Layer 2 Protocol to be Tunneled	5047
Loop Protection for a Spanning-Tree Instance Interface	5047
Root Protect for a Spanning-Tree Instance Interface	5048
BPDUs Protection for Individual Spanning-Tree Instance Interfaces . .	5048
BPDUs Protection on All Edge Ports of the Bridge	5049
VPLS Multihoming: Priority of the Backup Bridge	5049
VPLS Multihoming: Hold Time Before Switching to Primary Priority . .	5050
VPLS Multihoming: System Identifier for Bridges in the Ring	5050
VPLS Multihoming: Bridge Flush of MAC Cache on Topology Change	5051
Configuration	5052
Configuration Examples	5052
Example: Blocking BPDUs on Aggregated Ethernet for 600 Seconds	5052
Example: Enabling Loop Protection for Spanning-Tree Protocols . . .	5052
Example: Configuring VPLS Root Topology Change Actions	5053
Example: Configuring VSTP on a Trunk Port with Tagged Traffic . . .	5053
Example: Tracing Spanning-Tree Protocol Operations	5065
Configuration Tasks	5066
Configuring Multiple Spanning-Tree Protocol	5067
Configuring MST Instances on a Physical Interface	5070
Disabling MSTP	5071
Provider Bridge Participation in RSTP or MSTP Instances	5072
RSTP or VSTP Forced to Run as IEEE 802.1D STP	5072
Configuring Rapid Spanning-Tree Protocol	5073
Reverting RSTP or VSTP Back From Forced IEEE 802.1D STP	5075
Provider Bridge Participation in RSTP or MSTP Instances	5076
System Identifier for Bridges in STP or RSTP Instances	5076
Configuring VLAN Spanning-Tree Protocol	5077
Configuring VPLS Root Protection Topology Change Actions to Control Global Spanning Tree Behavior	5080
BPDUs Protection for Individual Spanning-Tree Instance Interfaces . .	5081
Configuring Loop Protection for a Spanning-Tree Instance Interface . .	5082

Enabling Root Protect for a Spanning-Tree Instance Interface	5083
Configuring BPDU Protection on Individual Interfaces	5084
Configuring BPDU Protection on All Edge Ports	5085
Configuring VPLS Root Protection Topology Change Actions to Control VLAN Spanning Tree Behavior	5085
Configuring Layer 2 Protocol Tunneling	5086
Configuration Statements	5087
[edit protocols mstp] Hierarchy Level	5087
[edit protocols rstp] Hierarchy Level	5088
[edit protocols vstp] Hierarchy Level	5089
protocols (STP Type)	5091
access-trunk	5091
bpdu-destination-mac-address (Spanning Tree)	5092
configuration-name	5093
cost	5094
disable	5095
edge	5096
extended-system-id	5097
force-version (IEEE 802.1D STP)	5097
forward-delay	5098
hello-time	5099
interface (Spanning Tree)	5100
max-age	5101
max-hops	5102
mode (Protocols STP)	5103
msti	5104
mstp	5105
priority (Protocols STP)	5106
revision-level	5107
rstp	5108
traceoptions (Spanning Tree)	5109
vlan (MSTP)	5112
vlan (VSTP)	5113
vstp	5114
[edit protocols layer2-control] Hierarchy Level	5115
interface (Layer 2 Protocol Tunneling)	5116
layer2-control	5117
mac-rewrite	5118
protocol	5118
no-root-port	5119
priority-hold-time	5120
system-id	5121
vpls-flush-on-topology-change	5122
[edit protocols layer2-control] Hierarchy Level	5122
bpdu-block	5123
disable-timeout	5124
interface (BPDU Blocking)	5124
interface (Layer 2 Protocol Tunneling)	5125

	Administration	5125
	Routine Monitoring	5125
	Checking the Status of Spanning-Tree Instance Interfaces	5125
	Clearing the Blocked Status of a Spanning-Tree Instance Interface ..	5126
	Checking for a MAC Rewrite Error Condition Blocking a Spanning-Tree	
	Instance Interface	5127
	Clearing a MAC Rewrite Error Condition Blocking a Spanning-Tree	
	Instance Interface	5127
	Tracing Spanning-Tree Operations	5128
	Operational Commands	5130
	show bridge mac-table	5131
	show mac-rewrite interface	5135
	show spanning-tree bridge	5136
	show spanning-tree interface	5141
	show spanning-tree mstp configuration	5147
	show spanning-tree statistics	5149
Chapter 20	VPNs	5151
	Multicast VPNs	5151
	Overview	5151
	Multicast VPNs	5151
	Layer 3 VPNs	5158
	Configuration	5162
	Configuration Examples	5162
	Configuration Tasks	5221
	Configuration Statements	5250
	Troubleshooting	5298
	Troubleshooting Multicast VPNs	5299
	VPLS	5299
	Overview	5299
	VPLS	5300
	Configuration	5322
	Configuration Tasks	5322
	Configuration Statements	5376
	Administration	5431
	VPLS Reference	5432

List of Figures

Part 1	Supported Hardware	
Chapter 2	EX9208 Switch Overview	115
	Figure 1: Front View of an EX9208 Switch	116
	Figure 2: Rear View of an AC-Powered EX9208 Switch	116
	Figure 3: Rear View of a DC-Powered EX9208 Switch	117
Part 2	System Administration	
Chapter 5	System Services	875
	Figure 4: DHCP Options Data Flow	895
Part 3	Features	
Chapter 6	Class of Service	1263
	Figure 5: Packet Flow Across the Network	1265
	Figure 6: CoS Classifier, Queues, and Scheduler	1274
	Figure 7: Packet Flow Through CoS Configurable Components	1274
	Figure 8: Topology to Verify Link Redundancy Support for L2TP LNS CoS	1315
	Figure 9: Network Traffic and Burst Rates	1442
	Figure 10: Flow of Tricolor Marking Policer Operation	1445
	Figure 11: Customer-Facing and Core-Facing Forwarding Classes	1536
	Figure 12: Sample CoS-Based Forwarding	1561
	Figure 13: Segmented and Interpolated Drop Profiles	1675
	Figure 14: Segmented and Interpolated Drop Profiles	1679
	Figure 15: Packet Flow Across the Network	1690
	Figure 16: CoS with a Tunnel Configuration	1767
Chapter 8	Ethernet Switching	1821
	Figure 17: Edge Device Case for Unrestricted Proxy ARP	1828
	Figure 18: Core Device Case for Unrestricted Proxy ARP	1828
Chapter 9	High Availability	2061
	Figure 19: Preparing for a Graceful Routing Engine Switchover	2063
	Figure 20: Graceful Routing Engine Switchover Process	2064
	Figure 21: Nonstop Bridging Switchover Preparation Process	2070
	Figure 22: Nonstop Bridging During a Switchover	2071
	Figure 23: Nonstop Active Routing Switchover Preparation Process	2073
	Figure 24: Nonstop Active Routing During a Switchover	2074
	Figure 25: Basic VRRP	2110
	Figure 26: Layer 3 VPN Graceful Restart Topology	2139
Chapter 10	Interfaces	2279

	Figure 27: Multicast Routers and the RPF Check	2284
	Figure 28: Unicast RPF with Routing Asymmetry	2318
Chapter 11	Layer 3 Protocols	2505
	Figure 29: Distance-Vector Protocol	2916
	Figure 30: Split Horizon Example	2918
	Figure 31: Poison Reverse Example	2919
	Figure 32: Limitations of Unidirectional Connectivity	2919
	Figure 33: Sample RIP Network Topology	2922
	Figure 34: RIP Authentication Network Topology	2929
	Figure 35: RIP BFD Network Topology	2936
	Figure 36: RIP BFD Authentication Network Topology	2942
	Figure 37: Configuring a Point-to-Multipoint RIP Network	2949
	Figure 38: RIP Import Policy Network Topology	2955
	Figure 39: Controlling Traffic in a RIP Network with the Incoming Metric	2961
	Figure 40: Controlling Traffic in a RIP Network with the Outgoing Metric	2962
	Figure 41: RIP Incoming Metrics Network Topology	2964
	Figure 42: Sending and Receiving RIPv1 and RIPv2 Packets Network Topology	2968
	Figure 43: Redistributing Routes Between RIP Instances Network Topology	2973
	Figure 44: RIP Timers Network Topology	2979
	Figure 45: RIP Trace Operations Network Topology	2991
	Figure 46: Customer Routes Connected to a Service Provider	3198
	Figure 47: BFD Enabled on Qualified Next Hops	3203
	Figure 48: Customer Routes Connected to a Service Provider	3211
Chapter 13	Logical Systems	3259
	Figure 49: Logical Systems Concepts	3260
	Figure 50: Junos OS Without Logical Systems	3261
	Figure 51: Junos OS With Logical Systems	3261
	Figure 52: Applications of Logical Systems	3266
	Figure 53: Logical System Administrators	3273
	Figure 54: Logical System Connected to a Physical Router	3278
	Figure 55: Connecting Two Logical Systems	3280
	Figure 56: Connecting Two Logical Systems	3284
	Figure 57: Static Routes Between Logical Systems	3288
	Figure 58: Logical System with a Stateless Firewall	3295
	Figure 59: Logical Systems with Filter-Based Forwarding	3300
	Figure 60: OSPF Three-Way Handshake	3311
	Figure 61: OSPF on Logical Systems	3313
	Figure 62: OSPF with a Default Route to an ISP	3323
	Figure 63: OSPF with a Conditional Default Route to an ISP	3328
	Figure 64: OSPF Import Policy on Logical Systems	3335
	Figure 65: Install Default Route to Nearest Routing Device That Operates at Both Level 1 and Level 2	3347
	Figure 66: IS-IS on Logical Systems	3348
	Figure 67: IS-IS with a Default Route to an ISP	3357
	Figure 68: ASs, EBGp, and IBGP	3364
	Figure 69: Typical Network with IBGP Sessions	3366
	Figure 70: Typical Network with BGP Peer Sessions	3376

	Figure 71: Typical Network with IBGP Sessions	3391
	Figure 72: Typical Network with EBGP Multihop Sessions	3399
	Figure 73: Point-to-Multipoint LSPs	3408
	Figure 74: RSVP-Signaled Point-to-Multipoint LSP on Logical Systems	3411
	Figure 75: Provider Edge and Provider Logical System Topology Diagram	3434
	Figure 76: Virtualized Data Center Physical Topology	3482
	Figure 77: Virtualized Data Center Logical Topology	3483
Chapter 14	MPLS	3535
	Figure 78: Swap and Push When LDP LSPs Are Tunneled Through RSVP LSPs	3537
	Figure 79: Double Push When LDP LSPs Are Tunneled Through RSVP LSPs	3537
Chapter 15	Multicast	3637
	Figure 80: Multicast Terminology in an IP Network	3641
	Figure 81: Converting MAC Addresses to Multicast Addresses	3644
	Figure 82: Routing Devices Start Up on a Subnet	3651
	Figure 83: Querier Routing Device Is Determined	3652
	Figure 84: General Query Message Is Issued	3652
	Figure 85: Reports Are Received by the Querier Routing Device	3652
	Figure 86: Host Has No Interested Receivers and Sends a Done Message to Routing Device	3653
	Figure 87: Host Address Timer Expires and Address Is Removed from Multicast Address List	3653
	Figure 88: Extracting the Embedded RP IPv6 Address	3710
	Figure 89: BFD Liveness Detection for PIM IPv6 Topology	3733
	Figure 90: Multicast Traffic Flooded from the Source Using PIM Dense Mode	3738
	Figure 91: Prune Messages Sent Back to the Source to Stop Unwanted Multicast Traffic	3739
	Figure 92: PIM Join Load Balancing	3743
	Figure 93: PIM Snooping for VPLS	3747
Chapter 18	Routing Policy and Packet Filtering	4463
	Figure 94: Typical Network with BGP Peer Sessions	4595
	Figure 95: Firewall Filter to Protect Against TCP and ICMP Floods	4601
	Figure 96: Filter-Based Forwarding	4646
	Figure 97: Logical Systems with Filter-Based Forwarding	4655
	Figure 98: Filter-Based Forwarding to Specified Outgoing Interfaces	4665
	Figure 99: Logical System with a Stateless Firewall	4688
	Figure 100: Network Traffic and Burst Rates	4804
	Figure 101: Incoming and Outgoing Policers and Firewall Filters	4810
	Figure 102: Token Bucket Algorithm	4832
	Figure 103: Traffic Behavior Using Policer and Burst Size	4832
	Figure 104: Policer Behavior With a Single TCP Connection	4833
	Figure 105: Policer Behavior With Background Traffic (Multiple TCP Connections)	4834
	Figure 106: Bursty Traffic Without Configured Burst Size (Excessive Unused Bandwidth)	4835
	Figure 107: Bursty Traffic With Configured Burst Size (Less Unused Bandwidth)	4836

	Figure 108: Comparing Burst Size Configuration Methods: 5 ms V.S. 10 MTU . . .	4837
	Figure 109: Firewall Filter to Protect Against TCP and ICMP Floods	4865
Chapter 19	Spanning-Tree Protocols	5031
	Figure 110: VPLS Multihoming Configuration	5038
	Figure 111: Topology for VSTP Configured on a Trunk Port with Tagged Traffic . .	5055
Chapter 20	VPNs	5151
	Figure 112: Simple MVPN Topology	5154
	Figure 113: Multicast Over Layer 3 VPN Example Topology	5163
	Figure 114: PIM Join Load Balancing on Draft-Rosen MVPN	5186
	Figure 115: PIM Join Load Balancing on Next-Generation MVPN	5193
	Figure 116: PIM State Limits Topology	5202
	Figure 117: Multiple VT Interfaces in MBGP MVPN Topology	5212
	Figure 118: Internet Multicast Topology	5244
	Figure 119: GRE Tunnel Configured Between the Local CE Router and the PE Router	5246
	Figure 120: GRE Tunnel Configured Between the Remote CE Router and the PE Router	5247
	Figure 121: Flooding a Packet with an Unknown Destination to All PE Routers in the VPLS Instance	5301
	Figure 122: CE Device Multihomed to Two PE Routers	5305
	Figure 123: BGP and LDP Signaling for a VPLS Routing Instance	5315
	Figure 124: VPLS Label Block Structure	5319
	Figure 125: Label Mapping Example	5320
	Figure 126: Flooding Unknown VPLS Traffic Using Ingress Replication	5366
	Figure 127: Flooding Unknown VPLS Traffic Using a Point-to-Multipoint LSP . .	5366

List of Tables

Chapter 1

About the Documentation	lvii
Table 1: Notice Icons	lix
Table 2: Text and Syntax Conventions	lix
Software Overview	63
Table 3: First Junos OS Release for Each EX Series Switch	64
Table 4: Access Control Features on Switches by Junos OS Release	65
Table 5: Administration Features on Switches by Junos OS Release	66
Table 6: Class-of-Service (CoS) Features on Switches by Junos OS Release . . .	67
Table 7: Class-of-Service (CoS) Features on EX9200 Switches by Junos OS Release	68
Table 8: Converged Networks (LAN and SAN) Features on Switches by Junos OS Release	70
Table 9: Device Security Features on Switches by Junos OS Release	70
Table 10: High Availability and Resiliency Features on Switches by Junos OS Release	71
Table 11: High Availability and Resiliency Features on EX9200 Switches by Junos OS Release	73
Table 12: Interfaces Features on Switches by Junos OS Release	75
Table 13: Interfaces Features on EX9200 Switches by Junos OS Release	76
Table 14: IP Address Management Features on Switches by Junos OS Release	77
Table 15: IP Address Management Features on EX9200 Switches by Junos OS Release	77
Table 16: IPv6 Features on Switches by Junos OS Release	78
Table 17: Layer 2 Network Protocols Features on Switches by Junos OS Release	81
Table 18: Layer 2 Networking Features on EX9200 Switches by Junos OS Release	82
Table 19: Layer 3 Protocols Features on Switches by Junos OS Release	84
Table 20: Layer 3 Protocols Features on EX9200 Switches by Junos OS Release	86
Table 21: Logical Systems Features on EX9200 Switches by Junos OS Release	89
Table 22: MPLS Features on Switches by Junos OS Release	89
Table 23: MPLS Features on EX9200 Switches by Junos OS Release	92
Table 24: Multicast Features on Switches by Junos OS Release	94
Table 25: Multicast Features on EX9200 Switches by Junos OS Release	95
Table 26: Network Management and Monitoring Features on Switches by Junos OS Release	97

	Table 27: Network Management and Monitoring Features on EX9200 Switches by Junos OS Release	99
	Table 28: Port Security Features on Switches by Junos OS Release	100
	Table 29: Power over Ethernet (PoE) Features on Switches by Junos OS Release	102
	Table 30: Routing Policy and Packet Filtering Features on Switches by Junos OS Release	103
	Table 31: Routing Policy and Firewall Filters on EX9200 Switches by Junos OS Release	103
	Table 32: Spanning-Tree Protocols Features on Switches by Junos OS Release	105
	Table 33: Spanning-Tree Protocols Features on EX9200 Switches by Junos OS Release	106
	Table 34: System Management Features on Switches by Junos OS Release	107
	Table 35: System Management Features on EX9200 Switches by Junos OS Release	107
	Table 36: User Interface and Configuration Features on EX9200 Switches by Junos OS Release	108
	Table 37: VPN Features on EX9200 Switches by Junos OS Release	108
Part 1	Supported Hardware	
Chapter 2	EX9208 Switch Overview	115
	Table 38: Line Cards Available for EX9208 Switches	117
	Table 39: Power Supplies Supported on EX9208 Switches	118
Part 2	System Administration	
Chapter 3	Access Privileges	123
	Table 40: Login Class Permission Flags	124
	Table 41: Predefined System Login Classes	128
Chapter 4	System Basics	295
	Table 42: Junos OS EFL Part Number on EX2200 Switches	297
	Table 43: Junos OS EFL Part Number on EX3300 Switches	298
	Table 44: Junos OS AFL Part Number on EX3300 Switches	299
	Table 45: Junos OS AFL Part Number on EX3200, EX4200, EX4500, EX4550, EX6200, EX8200, and EX9200 Switches	300
	Table 46: show cli Output Fields	456
	Table 47: show cli authorization Output Fields	458
	Table 48: show cli directory Output Fields	462
	Table 49: show cli history Output Fields	463
	Table 50: show system license Output Fields	511
	Table 51: show system snapshot Output Fields	515
	Table 52: show chassis alarms Output Fields	523
	Table 53: show chassis environment Output Fields	534
	Table 54: show chassis environment cb Output Fields	580
	Table 55: show chassis environment fpc Output Fields	598
	Table 56: show chassis environment pem Output Fields	622
	Table 57: show chassis environment routing-engine Output Fields	630

	Table 58: show chassis fan Output Fields	634
	Table 59: Routing Engines Displaying DIMM Information	646
	Table 60: show chassis hardware Output Fields	649
	Table 61: show chassis location Output Fields	751
	Table 62: show chassis pic Output Fields	757
	Table 63: show chassis routing-engine Output Fields	769
	Table 64: show chassis temperature-thresholds Output Fields	788
	Table 65: show system processes Output Fields	815
	Table 66: show system uptime Output Fields	872
Chapter 5	System Services	875
	Table 67: Information in Authentication Grant	878
	Table 68: RADIUS Attributes and VSAs for DHCPv6 Local Server	880
	Table 69: Action Taken for Events That Occur During a Reconfiguration	887
	Table 70: Unsupported Opaque DHCP Options	896
	Table 71: ARP Table in Trusted Environment	931
	Table 72: ARP Table in Distrusted Environment	931
	Table 73: Actions for DHCP Local Server Snooped Packets	945
	Table 74: ARP Table in Trusted Environment	959
	Table 75: ARP Table in Distrusted Environment	959
	Table 76: DHCP Relay Agent Option 82 Value for Auto Logout	962
	Table 77: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Enabled	969
	Table 78: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Disabled	970
	Table 79: Actions for Snooped BOOTREPLY Packets	970
	Table 80: show dhcp server binding Output Fields	1178
	Table 81: show dhcp server statistics Output Fields	1183
	Table 82: show dhcpv6 server binding Output Fields	1186
	Table 83: show dhcpv6 server statistics Output Fields	1191
	Table 84: clear dhcp relay statistics Output Fields	1196
	Table 85: show dhcp relay binding Output Fields	1204
	Table 86: show dhcp relay statistics Output Fields	1209
	Table 87: show dhcpv6 relay binding Output Fields	1212
	Table 88: show dhcpv6 relay statistics Output Fields	1216
	Table 89: show route extensive Output Fields	1219
Part 3	Features	
Chapter 6	Class of Service	1263
	Table 90: Default VPLS Classifiers	1271
	Table 91: CoS Mappings—Inputs and Outputs	1272
	Table 92: LSR Default Classification	1318
	Table 93: Default IP Precedence Classifier	1325
	Table 94: Default MPLS Classifier	1325
	Table 95: Default DSCP Classifier	1326
	Table 96: Default IEEE 802.1p Classifier	1327
	Table 97: Default IEEE 802.1ad Classifier	1328
	Table 98: Default IP Precedence (ipprec-default) Classifier	1329
	Table 99: Logical Interface Classifier Combinations	1332

Table 100: Default MPLS EXP Classification Table	1341
Table 101: Default CoS Values	1372
Table 102: Color-Blind Mode TCM Color-to-PLP Mapping	1449
Table 103: Color-Aware Mode TCM PLP Mapping	1450
Table 104: Color-Blind Mode TCM Color-to-PLP Mapping	1452
Table 105: Color-Aware Mode TCM Mapping	1453
Table 106: Tricolor Marking Policer Statements	1456
Table 107: Default Forwarding Classes	1528
Table 108: Sample Forwarding Class-to-Queue Mapping	1536
Table 109: Buffer Size Temporal Value Ranges by Routing Device Type	1580
Table 110: Recommended Delay Buffer Sizes	1581
Table 111: Maximum Delay Buffer with q-pic-large-buffer Enabled by Interface	1582
Table 112: Delay-Buffer Calculations	1584
Table 113: NxDSO Transmission Rates and Delay Buffers	1585
Table 114: Bandwidth and Delay Buffer Allocations by Configuration Scenario	1603
Table 115: Bandwidth and Delay Buffer Allocations by Configuration Scenario . .	1611
Table 116: Transmission Scheduling Support by Interfaces Type	1616
Table 117: Current Behavior with Multiple Priority Levels	1669
Table 118: Current Behavior with Same Priority Levels	1669
Table 119: Current Behavior with Strict-High Priority	1670
Table 120: Strict-High Priority with Higher Load	1670
Table 121: Sharing with Multiple Priority Levels	1671
Table 122: Sharing with the Same Priority Levels	1671
Table 123: Sharing with at Least One Strict-High Priority	1671
Table 124: Sharing with at Least One Strict-High Priority and Higher Load . . .	1672
Table 125: Sharing with at Least One Strict-High Priority and Rate Limit	1672
Table 126: Default Packet Header Rewrite Mappings	1693
Table 127: Default MPLS EXP Rewrite Table	1703
Table 128: Routing Engine Protocol Queue Assignments	1746
Chapter 8 Ethernet Switching	1821
Table 129: Statement Usage and Input Rewrite Operations for VLAN Identifiers for a VLAN	1834
Table 130: Statement Usage and Output Rewrite Operations for VLAN Identifiers for a VLAN	1835
Table 131: Rewrite Operations and Statement Usage for Input VLAN Maps . . .	1840
Table 132: Rewrite Operations and Statement Usage for Output VLAN Maps . .	1841
Table 133: show ethernet-switching interface Output Fields	1944
Table 134: show ethernet-switching table Output fields	1950
Table 135: show interfaces Gigabit Ethernet Output Fields	1957
Table 136: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type	1971
Table 137: show interfaces Gigabit Ethernet Output Fields	1982
Table 138: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type	1995
Table 139: show interfaces irb Output Fields	2004
Table 140: Layer 2 Overhead, Transmitted Packets/Bytes	2012

	Table 141: show interfaces queue Output Fields	2014
	Table 142: Byte Count by PIC Type	2017
	Table 143: show mvrp Output Fields	2047
	Table 144: show mvrp applicant-state Output Fields	2049
	Table 145: show mvrp dynamic-vlan-memberships Output Fields	2051
	Table 146: show mvrp interface Output Fields	2052
	Table 147: show mvrp registration-state Output Fields	2053
	Table 148: show mvrp statistics Output Fields	2055
	Table 149: traceroute ethernet Output Fields	2059
Chapter 9	High Availability	2061
	Table 150: Effects of a Routing Engine Switchover	2065
	Table 151: Graceful Routing Engine Switchover Feature Support	2066
	Table 152: Nonstop Active Routing Platform Support	2075
	Table 153: Nonstop Active Routing Protocol and Feature Support	2076
	Table 154: Unified ISSU Platform Support	2098
	Table 155: Unified ISSU Protocol Support	2098
	Table 156: Unified ISSU PIC Support: SONET/SDH	2102
	Table 157: Unified ISSU PIC Support: Fast Ethernet and Gigabit Ethernet	2103
	Table 158: Unified ISSU PIC Support: Channelized	2105
	Table 159: Unified ISSU PIC Support: Tunnel Services	2105
	Table 160: Unified ISSU PIC Support: ATM	2106
	Table 161: Unified ISSU Support: Enhanced IQ2 Ethernet Services Engine (ESE) PIC	2107
	Table 162: Unified ISSU Support: MX Series 3D Universal Edge Routers	2108
	Table 163: Unified ISSU Support: MX Series 3D Universal Edge Routers	2109
	Table 164: Example: VRRPv2 to VRRPv3 Transition Steps and Events	2114
	Table 165: Interface State and Priority Cost Usage	2230
	Table 166: show system switchover Output Fields	2277
Chapter 10	Interfaces	2279
	Table 167: Configuration Statements Used to Bind VLAN IDs to Logical Interfaces	2291
	Table 168: show interfaces Gigabit Ethernet Output Fields	2385
	Table 169: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type	2399
	Table 170: Discard show interfaces Output Fields	2409
	Table 171: show interfaces Gigabit Ethernet Output Fields	2415
	Table 172: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type	2428
	Table 173: show interfaces (Serial) Output Fields	2437
	Table 174: Layer 2 Overhead, Transmitted Packets/Bytes	2470
	Table 175: show interfaces queue Output Fields	2472
	Table 176: Byte Count by PIC Type	2475
Chapter 11	Layer 3 Protocols	2505
	Table 177: First Junos OS Release for Each EX Series Switch	2507
	Table 178: Access Control Features on Switches by Junos OS Release	2507
	Table 179: Administration Features on Switches by Junos OS Release	2508

Table 180: Class-of-Service (CoS) Features on Switches by Junos OS Release	2509
Table 181: Class-of-Service (CoS) Features on EX9200 Switches by Junos OS Release	2510
Table 182: Converged Networks (LAN and SAN) Features on Switches by Junos OS Release	2512
Table 183: Device Security Features on Switches by Junos OS Release	2512
Table 184: High Availability and Resiliency Features on Switches by Junos OS Release	2513
Table 185: High Availability and Resiliency Features on EX9200 Switches by Junos OS Release	2515
Table 186: Interfaces Features on Switches by Junos OS Release	2517
Table 187: Interfaces Features on EX9200 Switches by Junos OS Release	2518
Table 188: IP Address Management Features on Switches by Junos OS Release	2519
Table 189: IP Address Management Features on EX9200 Switches by Junos OS Release	2519
Table 190: IPv6 Features on Switches by Junos OS Release	2520
Table 191: Layer 2 Network Protocols Features on Switches by Junos OS Release	2523
Table 192: Layer 2 Networking Features on EX9200 Switches by Junos OS Release	2524
Table 193: Layer 3 Protocols Features on Switches by Junos OS Release	2526
Table 194: Layer 3 Protocols Features on EX9200 Switches by Junos OS Release	2528
Table 195: Logical Systems Features on EX9200 Switches by Junos OS Release	2531
Table 196: MPLS Features on Switches by Junos OS Release	2531
Table 197: MPLS Features on EX9200 Switches by Junos OS Release	2534
Table 198: Multicast Features on Switches by Junos OS Release	2536
Table 199: Multicast Features on EX9200 Switches by Junos OS Release	2537
Table 200: Network Management and Monitoring Features on Switches by Junos OS Release	2539
Table 201: Network Management and Monitoring Features on EX9200 Switches by Junos OS Release	2541
Table 202: Port Security Features on Switches by Junos OS Release	2542
Table 203: Power over Ethernet (PoE) Features on Switches by Junos OS Release	2544
Table 204: Routing Policy and Packet Filtering Features on Switches by Junos OS Release	2545
Table 205: Routing Policy and Firewall Filters on EX9200 Switches by Junos OS Release	2545
Table 206: Spanning-Tree Protocols Features on Switches by Junos OS Release	2547
Table 207: Spanning-Tree Protocols Features on EX9200 Switches by Junos OS Release	2548
Table 208: System Management Features on Switches by Junos OS Release	2549
Table 209: System Management Features on EX9200 Switches by Junos OS Release	2549

Table 210: User Interface and Configuration Features on EX9200 Switches by Junos OS Release	2550
Table 211: VPN Features on EX9200 Switches by Junos OS Release	2550
Table 212: show bgp bmp Output Fields	2646
Table 213: show bgp group Output Fields	2649
Table 214: show bgp neighbor Output Fields	2656
Table 215: show bgp summary Output Fields	2669
Table 216: show policy damping Output Fields	2674
Table 217: Default Metric Values for Routes Exported into IS-IS	2710
Table 218: show isis adjacency Output Fields	2744
Table 219: show isis authentication Output Fields	2748
Table 220: show isis backup coverage Output Fields	2750
Table 221: show isis backup label-switched-path Output Fields	2752
Table 222: show isis backup spf results Output Fields	2755
Table 223: show isis database Output Fields	2759
Table 224: show isis hostname Output Fields	2765
Table 225: show isis interface Output Fields	2767
Table 226: show isis overview Output Fields	2770
Table 227: show isis route Output Fields	2773
Table 228: show isis spf Output Fields	2777
Table 229: show isis statistics Output Fields	2783
Table 230: show (ospf ospf3) interface Output Fields	2866
Table 231: show (ospf ospf3) io-statistics Output Fields	2871
Table 232: show (ospf ospf3) log Output Fields	2873
Table 233: show (ospf ospf3) neighbor Output Fields	2877
Table 234: show ospf overview Output Fields	2882
Table 235: show (ospf ospf3) route Output Fields	2888
Table 236: show (ospf ospf3) statistics Output Fields	2893
Table 237: show ospf database Output Fields	2898
Table 238: show ospf3 database Output Fields	2906
Table 239: Configuring Simple RIP Authentication	2933
Table 240: Configuring MD5 RIP Authentication	2934
Table 241: RIP Demand Circuit Packet Types	2985
Table 242: show system processes extensive Output Fields	3034
Table 243: show task memory Output Fields	3035
Table 244: Summary of Key RIP Routing Output Fields	3038
Table 245: show policy Output Fields	3052
Table 246: show policy conditions Output Fields	3054
Table 247: show rip general-statistics Output Fields	3056
Table 248: show rip neighbor Output Fields	3058
Table 249: show rip statistics Output Fields	3061
Table 250: show route Output Fields	3064
Table 251: show route advertising-protocol Output Fields	3074
Table 252: show route detail Output Fields	3085
Table 253: Next-hop Types Output Field Values	3089
Table 254: State Output Field Values	3091
Table 255: Communities Output Field Values	3093
Table 256: show route export Output Fields	3102
Table 257: show route extensive Output Fields	3105

	Table 258: show route forwarding-table Output Fields	3122
	Table 259: show route instance Output Fields	3140
	Table 260: show route receive-protocol Output Fields	3170
	Table 261: show route terse Output Fields	3189
	Table 262: BFD Operational Mode Commands	3232
	Table 263: show bfd session Output Fields	3237
Chapter 13	Logical Systems	3259
	Table 264: Benefits of Virtual Routers, VRF-Lite, and Logical Systems	3265
	Table 265: Login Class Permission Flags	3269
	Table 266: Default Route Preference Values for OSPF	3310
	Table 267: Commonly Used Operational Mode Commands	3524
Chapter 14	MPLS	3535
	Table 268: from Operators That Apply to LDP Received-Label Filtering	3548
	Table 269: to Operators for LDP Outbound-Label Filtering	3550
	Table 270: Standard Firewall Filter Match Conditions for MPLS Traffic	3566
	Table 271: show ldp database Output Fields	3621
	Table 272: show ldp session Output Fields	3625
	Table 273: show ldp traffic-statistics Output Fields	3630
	Table 274: show ldp session Output Fields	3632
Chapter 15	Multicast	3637
	Table 275: Multicast Routing Protocols Compared	3648
	Table 276: IGMP Event Messages	3675
	Table 277: Local RP and Auto-RP Message Types	3706
	Table 278: PIM Join Filter Match Conditions	3720
	Table 279: show igmp group Output Fields	3955
	Table 280: show igmp interface Output Fields	3959
	Table 281: show multicast pim-to-igmp-proxy Output Fields	3963
	Table 282: show igmp snooping interface Output Fields	3967
	Table 283: show igmp snooping membership Output Fields	3970
	Table 284: show igmp snooping statistics Output Fields	3974
	Table 285: show mld group Output Fields	3979
	Table 286: show mld interface Output Fields	3983
	Table 287: show mld statistics Output Fields	3987
	Table 288: show multicast pim-to-mld-proxy Output Fields	3990
	Table 289: show msdp Output Fields	3992
	Table 290: show msdp source Output Fields	3995
	Table 291: show msdp source-active Output Fields	3997
	Table 292: show msdp statistics Output Fields	3999
	Table 293: show multicast usage Output Fields	4003
	Table 294: show pim bidirectional df-election Output Fields	4025
	Table 295: show pim bidirectional df-election interface Output Fields	4028
	Table 296: show pim bootstrap Output Fields	4031
	Table 297: show pim interfaces Output Fields	4033
	Table 298: show pim join Output Fields	4037
	Table 299: show pim neighbors Output Fields	4046
	Table 300: show pim rps Output Fields	4050
	Table 301: show pim source Output Fields	4057

	Table 302: show pim statistics Output Fields	4060
Chapter 16	Network Management and Monitoring	4073
	Table 303: Output Control Keys for the monitor interface Command	4118
	Table 304: Match Conditions for the monitor traffic Command	4122
	Table 305: Logical Operators for the monitor traffic Command	4123
	Table 306: Arithmetic and Relational Operators for the monitor traffic Command	4125
	Table 307: traceroute Output Fields	4136
	Table 308: traceroute monitor Output Fields	4139
	Table 309: show analyzer Output Fields	4140
	Table 310: show oam ethernet link-fault-management Output Fields	4142
	Table 311: show interfaces Fast Ethernet Output Fields	4148
	Table 312: show interfaces Gigabit Ethernet Output Fields	4166
	Table 313: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type	4180
Chapter 17	Port Mirroring	4191
	Table 314: Application of Layer 2 Port Mirroring Firewall Filters on PE Routers and PE Switches	4197
	Table 315: Application of Layer 2 Port Mirroring Types	4201
	Table 316: Application of Layer 2 Port Mirroring Firewall Filters on PE Devices . .	4205
	Table 317: show chassis fabric fpcs Output Fields	4267
	Table 318: show chassis fpc Output Fields	4304
	Table 319: Routing Engines Displaying DIMM Information	4329
	Table 320: show chassis hardware Output Fields	4332
	Table 321: show chassis pic Output Fields	4436
	Table 322: show forwarding-options port-mirroring Output Fields	4446
	Table 323: show firewall Output Fields	4451
	Table 324: show firewall log Output Fields	4457
	Table 325: show forwarding-options next-hop-group Output Fields	4460
Chapter 18	Routing Policy and Packet Filtering	4463
	Table 326: Firewall Filter Protocol Families	4466
	Table 327: Filter Types	4467
	Table 328: Stateless Firewall Filter Configuration and Application Summary . .	4472
	Table 329: Packet Evaluation at a Single Firewall Filter	4476
	Table 330: Standard Firewall Filter Match Conditions by Protocol Family . . .	4481
	Table 331: Standard Firewall Filter Action Categories	4482
	Table 332: Firewall Filter Behavior by Filter Attachment Point	4484
	Table 333: Binary and Bit-Field Match Conditions for Firewall Filters	4489
	Table 334: Bit-Field Match Conditions for Common Combinations	4490
	Table 335: Bit-Field Logical Operators	4491
	Table 336: Firewall Filter List Behavior	4508
	Table 337: Syslog Message Destinations for the Firewall Facility	4519
	Table 338: Packet-Header Logs for Stateless Firewall Filter Terms	4521
	Table 339: Simple Filter Match Conditions	4533
	Table 340: Standard Firewall Filter Match Conditions for Protocol-Independent Traffic	4707
	Table 341: Standard Firewall Filter Match Conditions for IPv4 Traffic	4707

Table 342: Standard Firewall Filter Match Conditions for IPv6 Traffic	4717
Table 343: Standard Firewall Filter Match Conditions for MPLS Traffic	4723
Table 344: IP Address-Specific Firewall Filter Match Conditions for MPLS Traffic	4726
Table 345: IP Port-Specific Firewall Filter Match Conditions for MPLS Traffic . .	4727
Table 346: Standard Firewall Filter Match Conditions for VPLS Traffic	4728
Table 347: Standard Firewall Filter Match Conditions for Layer 2 CCC Traffic . .	4735
Table 348: Standard Firewall Filter Match Conditions for Layer 2 Bridging (MX Series 3D Universal Edge Routers and EX Series switches Only)	4737
Table 349: Terminating Actions for Standard Firewall Filters	4743
Table 350: Nonterminating Actions for Standard Firewall Filters	4745
Table 351: Standard Firewall Filter Match Conditions by Protocol Family for ACX Series Routers	4751
Table 352: Standard Firewall Filter Action Categories for ACX Series Routers . .	4752
Table 353: Standard Firewall Filter Match Conditions for IPv4 Traffic on ACX Series Routers	4752
Table 354: Standard Firewall Filter Match Conditions for MPLS Traffic on ACX Series Routers	4756
Table 355: Terminating Actions for Standard Firewall Filters on ACX Series Routers	4757
Table 356: Nonterminating Actions for Standard Firewall Filters on ACX Series Routers	4758
Table 357: Service Filter Match Conditions for IPv4 or IPv6 Traffic	4760
Table 358: Terminating Actions for Service Filters	4766
Table 359: Nonterminating Actions for Service Filters	4767
Table 360: Unsupported Firewall Statements for Logical Systems	4768
Table 361: Unsupported Actions for Firewall Filters in Logical Systems	4769
Table 362: Packet Lengths Considered for Traffic Policers	4811
Table 363: Two-Color Policer Configuration and Application Overview	4813
Table 364: Three-Color Policer Configuration and Application Overview	4818
Table 365: Policer Bandwidth Limits and Burst-Size Limits	4822
Table 366: Implicit and Configurable Policer Actions Based on Color Marking . .	4823
Table 367: Examples of Counter and Policer Set Size and Indexing	4876
Table 368: Summary of Prefix-Specific Action Scenarios	4884
Table 369: Recommended Naming Convention for Policers	4915
Table 370: Bandwidth Limits and Token Rates	4976
Table 371: show firewall Output Fields	5017
Table 372: show firewall filter version Output Fields	5023
Table 373: show firewall log Output Fields	5024
Table 374: show firewall prefix-action-stats Output Fields	5027
Table 375: show policer Output Fields	5029
Chapter 19 Spanning-Tree Protocols	5031
Table 376: MAC Rewrite and VPLS Configurations	5039
Table 377: DPCs Supported for Layer 2 Protocol Tunneling	5040
Table 378: show bridge mac-table Output fields	5132
Table 379: show mac-rewrite interface Output Fields	5135
Table 380: show spanning-tree bridge Output Fields	5136
Table 381: show spanning-tree Interface Output Fields	5141

Chapter 20

Table 382: show spanning-tree mstp configuration Output Fields	5147
Table 383: show spanning-tree statistics Output Fields	5149
VPNs	5151
Table 384: PIM System Log Messages	5200
Table 385: NLRI Elements	5317
Table 386: VLAN ID Range by Interface Type	5339
Table 387: Standard Firewall Filter Match Conditions for VPLS Traffic	5358

About the Documentation

- Documentation and Release Notes on page lvii
- Supported Platforms on page lvii
- Using the Examples in This Manual on page lvii
- Documentation Conventions on page lix
- Documentation Feedback on page lx
- Requesting Technical Support on page lxi

Documentation and Release Notes

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

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

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

Supported Platforms

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

- EX Series

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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


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

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

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

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

Documentation Conventions

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

Table 1: Notice Icons

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

Table 2 on page [lix](#) defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>

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

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

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at

<https://www.juniper.net/cgi-bin/docbugreport/> . If you are using e-mail, be sure to include the following information with your comments:

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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

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

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

Opening a Case with JTAC

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

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

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

CHAPTER 1

Software Overview

- [EX Series Switch Software Features Overview on page 63](#)

EX Series Switch Software Features Overview

This topic lists the Juniper Networks EX Series Ethernet Switches software features, the Juniper Networks Junos operating system (Junos OS) release in which they were introduced, and the first Junos OS release for each switch.



NOTE: For Virtual Chassis features, see *EX Series Virtual Chassis Software Features Overview*.



NOTE: In all tables in this topic, “N.S.” = “Not supported”, and “—” = “Not applicable”.

- [Table 3 on page 64](#)—First Junos OS Release for Each EX Series Switch
- [Table 4 on page 65](#)—Access Control Features
- [Table 5 on page 66](#)—Administration Features
- [Table 6 on page 67](#)—Class-of-Service (CoS) Features
- [Table 7 on page 68](#)—Class-of-Service (CoS) Features for EX9200 Switches
- [Table 8 on page 70](#)—Converged Networks (LAN and SAN) Features
- [Table 9 on page 70](#)—Device Security Features
- [Table 10 on page 71](#)—High Availability and Resiliency Features
- [Table 11 on page 73](#)—High Availability and Resiliency Features on EX9200 Switches
- [Table 12 on page 75](#)—Interfaces Features
- [Table 13 on page 76](#)—Interfaces Features on EX9200 Switches
- [Table 14 on page 77](#)—IP Address Management Features
- [Table 15 on page 77](#)—IP Address Management Features on EX9200 Switches
- [Table 16 on page 78](#)—IPv6 Features

- [Table 17 on page 81](#)—Layer 2 Network Protocols Features
- [Table 18 on page 82](#)—Layer 2 Network Protocols Features on EX9200 Switches
- [Table 19 on page 84](#)—Layer 3 Protocols Features
- [Table 20 on page 86](#)—Layer 3 Protocols Features for EX9200 Switches
- [Table 21 on page 89](#)—Logical Systems Features on EX9200 Switches
- [Table 22 on page 89](#)—MPLS Features
- [Table 23 on page 92](#)—MPLS Features on EX9200 Switches
- [Table 24 on page 94](#)—Multicast Features
- [Table 25 on page 95](#)—Multicast Features on EX9200 Switches
- [Table 26 on page 97](#)—Network Management and Monitoring Features
- [Table 27 on page 99](#)—Network Management and Monitoring Features on EX9200 Switches
- [Table 28 on page 100](#)—Port Security Features
- [Table 29 on page 102](#)—Power over Ethernet (PoE) Features
- [Table 30 on page 103](#)—Routing Policy and Packet Filtering Features
- [Table 31 on page 103](#)—Routing Policy and Packet Filtering Features on EX9200 Switches
- [Table 32 on page 105](#)—Spanning-Tree Protocols Features
- [Table 33 on page 106](#)—Spanning-Tree Protocols Features on EX9200 Switches
- [Table 34 on page 107](#)—System Management Features
- [Table 35 on page 107](#)—System Management Features on EX9200 Switches
- [Table 36 on page 108](#)—User Interface and Configuration Features on EX9200 Switches
- [Table 37 on page 108](#)—VPN Features on EX9200 Switches

The Junos OS release for software features on a switch cannot be earlier than the first Junos OS release for that switch.

Table 3: First Junos OS Release for Each EX Series Switch

Switch	Junos OS Release
EX2200 switches*	Junos OS Release 10.1R1 *EX2200-C models: Junos OS Release 11.3R1
EX3200 switches	Junos OS Release 9.0R1
EX3300 switches	Junos OS Release 11.3R1
EX4200 switches	Junos OS Release 9.0R1
EX4300 switches	Junos OS Release 13.2X50-D10

Table 3: First Junos OS Release for Each EX Series Switch (*continued*)

Switch	Junos OS Release
EX4500 switches**	Junos OS Release 10.2R1* **EX4500-C models: Junos OS Release 10.3R2
EX4550 switches	Junos OS Release 12.2R1
EX6200 switch	Junos OS Release 11.3R2
EX8208 switches	Junos OS Release 9.4R1
EX8216 switches	Junos OS Release 9.5R1
EX9200 switches	Junos OS Release 12.3R2

Table 4: Access Control Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
802.1X authentication (port-based, multiple supplicant)	10.1R1	9.0R2	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	10.2R1	N.S.
802.1X authentication with authentication bypass	10.1R1	9.0R2	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	10.2R1	N.S.
802.1X authentication with VLAN assignment, VoIP VLAN support	10.1R1	9.0R1	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	10.3R1	N.S.
802.1X user-based dynamic firewall filters	10.1R1	9.0R2	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	10.3R1	N.S.
802.1X user-based dynamic firewall filters on multiple-supplicant ports	10.1R1	9.5R2	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	10.3R1	N.S.
802.1X per-user statistics	10.1R1	9.2R1	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R1	10.3R1	N.S.
Authentication fallback	11.3R1	10.3R1	12.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	N.S.	N.S.
Captive portal authentication for Layer 3 interfaces	11.3R1	10.1R1	N.S.	13.2X50 - D10	12.1R1	12.2R1	N.S.	N.S.	N.S.
Captive portal authentication for Layer 2 interfaces	11.3R1	10.3R1	12.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	N.S.	N.S.

Table 4: Access Control Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Energy Efficient Ethernet (EEE)	N.S.	N.S.	12.2R1	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.
Infranet Controller (IC) as an external captive-portal server	12.2R1	12.2R1	12.2R1	N.S.	12.2R1	12.2R1	N.S.	N.S.	N.S.
MAC RADIUS authentication	10.1R1	9.3R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	10.3R1	N.S.
NetBIOS snooping	11.3R5	11.1R1	11.3R5	N.S.	N.S.	N.S.	11.3R5	11.1R1	N.S.
Server fail fallback	10.1R1	9.3R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	N.S.	10.2R1	N.S.
TACACS+	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2

Table 5: Administration Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
System logging (syslog) over IPv4	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2
System logging (syslog) over IPv6	10.3R1	9.3R2	11.3R1	13.2X50 - D10	10.4R1	12.2R1	11.3R2	10.1R1	12.3R2
System snapshot	N.S.	10.0R1	N.S.	13.2X50 - D10	10.2R1	12.2R1	N.S.	10.0R1	12.3R2

Table 6: Class-of-Service (CoS) Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Class of service (CoS)—Class-based queuing with prioritization, Layer 2 and Layer 3 classification, rewrite, and queuing; strict priority queuing on egress	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	See Table 6-3 for a list of EX9200 CoS features
CoS—DSCP, IEEE 802.1p, and IP precedence packet rewrites on RVIs or IRB interfaces	N.S.	9.5R1	11.3R1	N.S.	10.2R1	12.2R1	11.3R2	9.4R1	
CoS—Interface-specific classifiers on RVIs or IRB interfaces	N.S.	9.4R1	11.3R1	13.2X50 - D10	11.3R1	12.2R1	N.S.	10.2R1	
CoS—Multidestination	—	—	—	—	—	—	N.S.	9.5R1	
CoS—Per-interface classification	N.S.	9.3R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	10.2R1	
CoS support on link aggregation groups (LAGs)	10.1R1	9.2R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	N.S.	9.4R1	
CoS support on RVIs or IRB interfaces	10.3R1	9.4R1	11.3R1	13.2X50 - D10 (classifiers only; rewrites N.S.)	10.4R1	12.2R1	N.S.	9.4R1	
DSCP setting on ingress interface	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	12.3R1	
Flexible CoS-outer 802.1p marking	N.S.	9.6R1	12.3R1	N.S.	12.1R1	12.2R1	N.S.	N.S.	
Interface-specific CoS rewrite rules	10.3R1	9.4R1	N.S.		11.2R1	12.2R1	N.S.	10.2R1	

Table 6: Class-of-Service (CoS) Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
				13.2X50 - D10 (for Layer 3 interfaces; IRB interfaces and Layer 3 sub-interfaces N.S.)					
Junos EZQoS for CoS	10.1R1	9.3R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Port shaping and queue shaping	10.1R1	9.3R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	10.1R1	
Re-marking of bridged packets	11.2R1	9.4R1	N.S.	13.2X50 - D10	10.2R1	12.2R1	11.3R2	10.2R1	
Shaped-deficit weighted round-robin (SDWRR)	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Single-rate two-color marking	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	

Table 7: Class-of-Service (CoS) Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Assigning forwarding class and DSCP value for Routing Engine generated traffic	12.3R2
BA classification for VPLS based on IEEE 802.1p bits	12.3R2
Classification—Associate packets with CoS servicing levels. Types of classification: <ul style="list-style-type: none"> • Behavior aggregate (BA)—Operates on packets as they enter the switch • Multifield classification—Examines multiple fields in packets. • Fixed classification—Associate a forwarding class with a packet regardless of packet contents. 	12.3R2
Classification and DSCP marking of distributed protocol handler traffic	12.3R2
Classification of control-plane traffic	12.3R2
CoS classification and rewrite for IRB and Layer 2 interfaces and for other Layer 3 interfaces. Port-level queuing, scheduling, and shaping are supported.	12.3R2

Table 7: Class-of-Service (CoS) Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Egress GRE classification based on DSCP	12.3R2
IEEE 802.1p inheritance push and swap from transparent tags configuration	12.3R2
Elevated packet drops during oversubscription	12.3R2
Layer 2 policers for the ingress and egress interfaces. Policer types: <ul style="list-style-type: none"> • Single-rate two-color • Single-rate three-color (color-blind and color-aware) • Two-rate three-color (color-blind and color-aware) 	12.3R2
Independent values for DSCP and EXP bits	12.3R2
Apply CoS schedulers on ingress interfaces	12.3R2
Ingress DSCP bits for multicast traffic over Layer 3 VPNs	12.3R2
Layer 2 traffic policing	12.3R2
Policer support for aggregated Ethernet bundles (link aggregation groups, or LAGs)	12.3R2
Queuing support for logical tunnel interfaces	12.3R2
Rate-limit and excess rate or excess priority option	12.3R2
Re-marking of MVPN GRE encapsulation DCSP at ASBR	12.3R2
Scheduling	12.3R2
Set IPv6 DSCP and MPLS EXP independently	12.3R2
Set IPv6 DiffServ code point (DSCP) and MPLS EXP independently	12.3R2
Support for Layer 2 policers at the VLAN level	12.3R2
Support for applying a transmit rate limit to logical interfaces on Type 1, 2, or 3 Multiservices PICs	12.3R2
Support for configuring ToS rewrite rules so that DCSP bits of GRE packets are consistent with service providers' CoS policy	12.3R2
Support for copying the TOS bits to the outer IP header on GRE tunnel traffic sent by the Routing Engine	12.3R2
Support for setting the forwarding class and DSCP value for traffic generated by the Routing Engine	12.3R2
Unified command to display all CoS statistics	12.3R2

Table 8: Converged Networks (LAN and SAN) Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
NOTE: The EX4500 switch models that support Fibre Channel over Ethernet features must be Converged Enhanced Ethernet (CEE) capable. The CEE-capable EX4500 switch models have a “-C” in the hardware model number. See <i>EX4500 Switch Models</i> .									
Data Center Bridging Capability Exchange protocol (DCBX)	N.S.	N.S.	N.S.	N.S.	11.3R1	12.2R1	N.S.	N.S.	N.S.
DCBX application protocol TLV exchange	N.S.	N.S.	N.S.	N.S.	12.1R1	12.2R1	N.S.	N.S.	N.S.
FIP snooping	N.S.	N.S.	N.S.	N.S.	10.4R1	12.2R1	N.S.	N.S.	N.S.
Priority-based flow control (PFC)	N.S.	N.S.	N.S.	N.S.	10.4R1	12.2R1	N.S.	N.S.	N.S.

Table 9: Device Security Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Automatic recovery for port error disable conditions	10.1R1	9.6R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	10.0R1	N.S.
Storm control (broadcast and unicast)	10.1R1	9.1R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	N.S.
Storm control (multicast)	10.3R2	10.3R2	N.S.	13.2X50 - D10	10.3R2	12.2R1	11.3R2	10.3R2	N.S.
Unknown Layer 2 unicast forwarding	10.1R1	9.3R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	N.S.	10.0R1	12.3R2

Table 10: High Availability and Resiliency Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
NOTE: For complete lists of Virtual Chassis features, see EX Series Virtual Chassis Software Features Overview .									See Table 3 for a list of EX9200 HA features.
Graceful protocol restart for BGP	–	9.0R2	N.S.	13.2X50 - D10	–	–	11.3R2	9.4R1	
Graceful protocol restart for IS-IS	–	9.3R2	N.S.	13.2X50 - D10	–	–	11.3R2	9.4R1	
Graceful protocol restart for OSPF	–	9.0R2	N.S.	13.2X50 - D10	–	–	11.3R2	9.4R1	
Graceful protocol restart for RSVP and LDP	N.S.	N.S.	N.S.	13.2X50 - D10	N.S.	N.S.	N.S.	12.3R1	
GRES for ARP entries, forwarding database, and Layer 3 protocols	–	9.2R1 (Virtual Chassis only)	11.3R1	13.2X50 - D10	11.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	11.3R2	9.4R1	
GRES for IGMP snooping	–	11.3R1 (Virtual Chassis only)	12.1R1 (Virtual Chassis only)	13.2X50 - D10	11.4R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	N.S.	11.3R1	
GRES for LACP	N.S.	11.3R1	N.S.	13.2X50 - D10	11.3R1	12.2R1	N.S.	11.3R1	
GRES for Layer 2 and Layer 3 VPN LSPs	N.S.	N.S.	N.S.	13.2X50 - D10	N.S.	N.S.	N.S.	12.3R1	
GRES for port security (DHCP snooping, DAI, and IP source guard)	–	9.2R1 (Virtual Chassis only)	N.S.	13.2X50 - D10	–	–	11.3R2	9.6R1	
LACP support for dual-homing applications in data centers	10.1R1	10.0R1	N.S.	13.2X50 - D10	10.2R1	12.2R1	N.S.	10.0R1	
Link Aggregation Control Protocol (LACP)	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Link aggregation groups (LAGs)	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	

Table 10: High Availability and Resiliency Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Nonstop active routing (NSR) for BGP, IS-IS, IGMP with BFD, and RIP	—	11.1R1 (Virtual Chassis only)	12.1R1 (Virtual Chassis only)	N.S.	11.3R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	11.3R2	11.1R1	
Nonstop active routing (NSR) for IPv6 IS-IS, RIPng, and OSPFv3 with BFD	—	12.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	N.S.	12.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	N.S.	11.2R1	
Nonstop active routing (NSR) for OSPFv2	—	11.1R1	12.1R1 (Virtual Chassis only)	N.S.	11.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	11.3R2	10.4R1	
Nonstop active routing (NSR) for Protocol Independent Multicast (PIM)	N.S.	12.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	N.S.	12.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	N.S.	11.4R2	
Nonstop active routing (NSR) for RSVP and LDP	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	12.3R1	
Nonstop bridging (NSB) for LAGs and LACP	—	11.4R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	N.S.	11.4R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	12.1R1	11.3R1	
Nonstop bridging (NSB) for LLDP and LLDP-MED	—	11.3R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	N.S.	N.S.	N.S.	N.S.	11.3R1	
Nonstop bridging (NSB) for spanning-tree protocols	—	11.3R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	N.S.	12.1R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	12.1R1	11.3R1	
Nonstop software upgrade (NSSU)	—	12.1R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	N.S.	12.1R1 (Virtual Chassis only)	12.1R1 (Virtual Chassis only)	12.2R1	10.4R1	
Power budget management	—	—	N.S.	13.2X50 - D10	—	—	11.3R2	10.2R1	
	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	11.4R1	

Table 10: High Availability and Resiliency Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Virtual Router—Network Time Protocol (NTP), system logging, Simple Network Management Protocol (SNMP), RADIUS, and TACACS support in a virtual router									
Virtual Router Redundancy Protocol (VRRP)	12.3R1	9.0R2	12.1R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Virtual Router Redundancy Protocol (VRRP)—Support for multiple VRRP owners per physical interface	12.3R1	12.3R1	12.3R1	N.S.	12.3R1	12.3R1	N.S.	12.3R1	
Virtual Router Redundancy Protocol (VRRP) for IPv6 (except authentication type and authentication key)	N.S.	10.2R1	12.3R1	13.2X50 - D10	11.2R1	12.2R1	12.1R1	10.1R1	

Table 11: High Availability and Resiliency Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Graceful Routing Engine switchover (GRES)	12.3R2
Nonstop active routing (NSR) support for Protocol Independent Multicast (PIM) for IPv4 and IPv6	12.3R2
Nonstop active routing (NSR) support for VPLS and for LDP-based VPLS	12.3R2
Nonstop active routing (NSR) support for LDP OAM features	12.3R2
Nonstop active routing (NSR) support for Layer 2 VPNs	12.3R2
Unified ISSU (requires EX9200-40T or EX9200-40F line cards)	12.3R3

Table 11: High Availability and Resiliency Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Virtual Router Redundancy Protocol version 3 (VRRPv3)	12.3R2

Table 12: Interfaces Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Digital optical monitoring (DOM)	N.S.	10.0R1	11.3R1	13.2X50 - D10	N.S.	N.S.	N.S.	10.0R1	See Table 12-1 for a list of EX9200 interfaces features.
Interface ranges	10.1R1	10.0R1	11.3R1	13.2X50 - D10	10.2R	12.2R1	11.3R2	10.1R1	
IPv4 over generic routing encapsulation (GRE) tunnels—encapsulation support	N.S.	12.1R1	N.S.	N.S.	N.S.	N.S.	N.S.	12.1R1	
IPv4 over generic routing encapsulation (GRE) tunnels—de-encapsulation support	N.S.	12.1R1	N.S.	N.S.	N.S.	N.S.	N.S.	12.3R1	
IPv6 over generic routing encapsulation (GRE) tunnels using IPv4 transport—encapsulation support	N.S.	12.1R1	N.S.	N.S.	N.S.	N.S.	N.S.	12.1R1	
IPv6 over generic routing encapsulation (GRE) tunnels using IPv4 transport—de-encapsulation support	N.S.	12.1R1	N.S.	N.S.	N.S.	N.S.	N.S.	12.3R1	
IP directed broadcast	11.3R1	9.4R1	12.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Time domain reflectometry (TDR)	10.1R1	9.0R2	11.3R1	13.2X50 - D10	N.S.	N.S.	11.3R2	9.4R1	
Unicast reverse-path forwarding (RPF)	N.S.	9.3R2	12.3R1	13.2X50 - D10	11.2R1	12.2R1	11.3R2	10.1R1	
	N.S.	9.2R1	N.S.	13.2X50 - D10	11.2R1	12.2R1	11.3R2	9.4R1	

Table 12: Interfaces Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
VLAN-tagged Layer 3 subinterfaces									

Table 13: Interfaces Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
ICMP redirect	12.3R2
Clear MAC address information	12.3R2
IPv6 subnet support on loopback interfaces	12.3R2
IPv6 support for unnumbered Ethernet interfaces	12.3R2
Multichassis link aggregation (MC-LAG)	12.3R2
Nonstop active routing (NSR) support for Bidirectional Forwarding Detection (BFD)	12.3R2
Protection against distributed denial-of-service (DDOS) attacks	12.3R2
Software support for IPv4 to IPv6 transition	12.3R2
Static mapping for port forwarding	12.3R2
Support for active monitoring on logical systems	12.3R2
Support for VRF in Routing Engine-based sampling	12.3R2
Support for integrated routing and bridging (IRB) MAC synchronization in MC-LAG for aggregated Ethernet	12.3R2
Targeted broadcast for virtual routing and forwarding (VRF) instances	12.3R2
Trunk interface enhancements: <ul style="list-style-type: none"> Configure a single logical trunk interface to support a list of VLANs or to accept packets with no VLAN tag. Configure multiple logical trunk interfaces on a single physical interface. 	12.3R2
Unicast reverse-path forwarding (RPF) loose mode, with ability to discard packets with source addresses pointing to the discard interface	12.3R2
Unnumbered Ethernet—Configure IPv4 processing on interfaces without assigning explicit IP addresses to the interfaces.	12.3R2
VLAN rewrite operations on incoming and outgoing frames	12.3R2

Table 14: IP Address Management Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
DHCP server and relay with option 82 for Layer 2 VLANs	10.1R1	9.3R2	11.3R1	N.S.	10.2R1	12.2R1	11.3R2	9.4R1	See Table 14-1 for a list of EX9200 IP address management features.
DHCP server and relay with option 82 for Layer 3 interfaces	10.1R1	9.0R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
DNS for IPv6	N.S.	9.3R2	N.S.	13.2X50 - D10	N.S.	N.S.	N.S.	N.S.	
Local DHCP server	10.1R1	9.3R2	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	9.4R1	
Virtual router aware DHCP (VR-aware DHCP)	N.S.	12.3R1	N.S.	13.2X50 - D10	12.3R1	12.3R1	N.S.	12.3R1	
Virtual router aware DHCPv6 (VR-aware DHCPv6)	N.S.	N.S.	N.S.	13.2X50 - D10	12.3R1	12.3R1	12.3R1	N.S.	
Static addresses	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	N.S.	9.4R1	

Table 15: IP Address Management Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
DHCP server and relay	12.3R2
DHCPv6 local server	12.3R2
DHCPv6 support	12.3R2
Distinguishing DHCP duplicate clients by subinterface	12.3R2
Dynamic reconfiguration of extended DHCP and DHCPv6 local server clients	12.3R2
Dynamic IPv6 filters	12.3R2
Expression support for dynamic profiles	12.3R2
Extended DHCP relay proxy	12.3R2

Table 15: IP Address Management Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Optional disabling of automatic ARP table population	12.3R2
IPv6 address assignment pools	12.3R2
Overriding DHCP settings on specific interfaces	12.3R2
Per-interface DHCP tracing operations	12.3R2
S-VLAN-based shaping for dynamic profiles	12.3R2
Sending a DHCP relay and relay proxy release message	12.3R2
Specifying the DHCP source address used for IP packets	12.3R2
Support for MAC address validation	12.3R2
Support for address pool threshold traps	12.3R2
Address assignment pools	12.3R2
Per-interface DHCP lease limits	12.3R2

Table 16: IPv6 Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
NOTE: A separate software license is required for IPv6. See Understanding Software Licenses for EX Series Switches .									
Application identification (APPID) for IPv6 packets	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	12.3R2
BFD for IPv6	N.S.	12.3R1	N.S.	13.2X50 - D10	12.3R1	12.3R1	N.S.	12.3R1	12.3R2 (also for static routes)
BGP for IPv6	N.S.	9.4R1	12.3R1	13.2X50 - D10	11.1R1	12.2R1	12.1R1	10.1R1	12.3R2
IPv6 CoS (multifield classification and rewrite)	N.S.	10.2R1	12.3R1	13.2X50 - D10	12.1R1	12.2R1	12.1R1	10.4R1	12.3R2

Table 16: IPv6 Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
IPv6 management	10.3R1	9.3R2 (using loopback addresses only)	11.3R1	13.2X50 - D10	10.4R1	12.2R1	N.S.	10.1R1	12.3R2
IPv6 multicast protocols (PIM, MLDv1/v2)	N.S.	10.1R1	12.3R1	13.2X50 - D10	11.2R1	12.2R1	12.1R1	10.2R1	12.3R2
IPv6 neighbor redirect compliance with RFC 4861	12.3R1	12.3R1	12.3R1	13.2X50 - D10	12.3R1	12.3R1	12.3R1	12.3R1	12.3R2
IPv6 path MTU discovery	10.3R1	9.3R1	12.3R1	13.2X50 - D10	10.4R1	12.2R1	N.S.	10.2R1	12.3R2
IS-IS for IPv6	N.S.	9.4R1	N.S.	13.2X50 - D10	11.2R1	12.2R1	12.1R1	10.1R1	12.3R2
MBGP for IPv6	N.S.	9.3R1	12.3R1	13.2X50 - D10	N.S.	N.S.	12.1R1	10.1R1	12.3R2
OSPFv3	N.S.	9.3R1	12.3R1	13.2X50 - D10	11.1R1	12.2R1	12.1R1	10.1R1	12.3R2
RFC 4291 compliance	12.3R1 See note at end of table	12.3R1 See note at end of table	12.3R1 See note at end of table	N.S.	12.3R1 See note at end of table	12.3R1 See note at end of table	12.3R1 See note at end of table	12.3R1 See note at end of table	12.3R2 See note at end of table
RIPng	N.S.	9.3R1	12.3R1	13.2X50 - D10	11.1R1	12.2R1	12.1R1	10.1R1	12.3R2
VRRPv3 (RFC 5798 compliance, ability to send SNMP traps)	N.S.	12.3R1	N.S.	13.2X50 - D10	12.3R1	12.3R1	N.S.	12.3R1	12.3R2



NOTE: Compliance with RFC 4291

EX Series switches drop the following types of illegal IPv6 packets:

- Packets that have a link-local source or destination address. Because link-local addresses are intended to be used for addressing only on a single link, EX Series switches do not forward any packets with such addresses to other links.
- Packets with the IPv6 unspecified source address 0:0:0:0:0:0:0:0.
- Packets that are to be sent outside a node but have the IPv6 loopback address 0:0:0:0:0:0:0:1 as the source address. When IPv6 packets are received on an interface, EX Series switches drop packets that have the loopback address as the destination address.

EX Series switches do not support Subnet-Router Anycast address.

Table 17: Layer 2 Network Protocols Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
802.1Q VLAN tagging	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	See Table 82 for a list of EX9200 Layer 2 networking protocols features.
Edge virtual bridging (EVB) support with virtual Ethernet port aggregator (VEPA)	N.S.	N.S.	N.S.	N.S.	12.1R1	12.2R1	N.S.	12.1R1	
Ethernet ring protection switching (ERPS, G.8032/Y.1344)	12.1R1	12.1R1	12.3R1	13.2X50 - D10	12.3R1	12.3R1	N.S.	12.3R1	
Layer 2 protocol tunneling (L2PT)	11.1R1	10.0R1	12.3R1	N.S.	11.2R1	12.2R1	12.1R1	N.S.	
Link Layer Discovery Protocol (LLDP)	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Link Layer Discovery Protocol—Media Endpoint Discovery (LLDP-MED) with voice over IP (VoIP) integration	10.1R1	9.0R2	11.3R1	13.2X50 - D10	N.S.	N.S.	N.S.	N.S.	
MAC-based VLANs	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Multiple VLAN Registration Protocol (MVRP, IEEE 802.1ak)	11.3R1	10.0R1	12.3R1	13.2X50 - D10	11.2R1	12.2R1	12.1R1	10.0R1	
Private VLANs (PVLANS)	11.1R1	9.3R2	12.1R1	N.S.	11.2R1	12.2R1	11.3R2	10.1R1	
Private VLANs (PVLANS) support across switches	11.1R1	10.4R1	12.1R1	N.S.	11.2R1	12.2R1	11.3R2	11.2R1	
Proxy ARP—Restricted	10.1R1	10.0R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	N.S.	10.0R1	

Table 17: Layer 2 Network Protocols Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Proxy ARP—Unrestricted	10.1R1	9.6R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	12.1R1	10.0R1	
Proxy ARP per VLAN	10.1R1	10.1R1	N.S.	13.2X50 - D10	10.2R1	12.2R1	N.S.	10.1R1	
Q-in-Q tunneling	11.1R1	9.3R2	11.4R1	N.S.	11.2R1	12.2R1	12.1R1	11.1R1	
Q-in-Q VLAN extended support for multiple S-VLANs per access interface, firewall-filter-based VLAN assignment, and RVIs or IRB interfaces	N.S.	9.6R1	12.3R1	N.S.	11.2R1	12.2R1	12.1R1	11.1R1	
Redundant trunk groups	10.1R1	9.0R2	11.3R1	13.2X50 - D10	11.2R1	12.2R1	11.3R2	9.4R1	
Routed VLAN interfaces (RVIs)—Also known as integrated routing and bridging (IRB) interfaces	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
VLAN ID translation	11.1R1	10.0R1	N.S.	N.S.	11.2R1	12.2R1	N.S.	11.1R1	
VLAN ranges	10.1R1	9.2R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	

Table 18: Layer 2 Networking Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
VLANs and virtual switching	12.3R2
DHCP support for integrated routing and bridging (IRB)	12.3R2
MC-LAG support for IGMP snooping in IRB	12.3R2
Hash-key load-balancing support for Layer 3 and Layer 4 fields	12.3R2

Table 18: Layer 2 Networking Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
IP multicast over Layer 2 trunk port support	12.3R2
Integrated routing and bridging (IRB)	12.3R2
Layer 2 Ethernet OAM: <ul style="list-style-type: none"> Distributed periodic packet management process (ppmd) for improved scaling Graceful Routing Engine switchover (GRES) Remote defect indication (RDI) Configuration of action profiles 	12.3R2
Layer 2 address learning in logical systems	12.3R2
Layer 2 forwarding support for bridging and VPLS	12.3R2
Layer 2 policer statistics MIB	12.3R2
Firewall filter match conditions for Layer 2 bridging and VPLS	12.3R2
Next-hop groups using either IP addresses or Layer 2 addresses for the next hop	12.3R2
Unicast reverse-path forwarding (RPF) loose mode, with ability to discard packets with source addresses pointing to the discard interface	12.3R2
Spanning-tree protocols support for Layer 2 bridging and VPLS	12.3R2
VLAN rewrite operations on incoming and outgoing frames	12.3R2
STP root guard (root protection)	12.3R2
Support for Layer 2 and Layer 2.5 features: <ul style="list-style-type: none"> Extensive set of Layer 2 label-manipulation capabilities, Q-in-Q support MC-LAG active / standby and MC-LAG active / active xSTP protocol support Integrated Routing and Bridging (IRB) interface support IGMP snooping for multichassis link aggregation group (MC-LAG) interfaces Configurable label block sizes for VPLS Connectivity fault management process flooding to interfaces based on mesh groups Layer 2 address learning in logical systems Virtual switch support, providing virtual Layer 2 switch instances with separate Layer 2 learning domains, isolated 4K VLAN ID spaces, and STP instances Ethernet Ring Protocol (ERP) for multiple ring instances on the same physical ring Transit and bypass static label-switched paths (LSPs) Layer 2 Gigabit Ethernet logical interface policing Static LSP statistics Multiple VLAN Registration Protocol (MVRP)—IEEE 802.1ak-2007 	12.3R2

Table 18: Layer 2 Networking Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
VPLS root protection topology change-action control	12.3R2
VLAN ranges	12.3R2
Q-in-Q VLAN extended support for multiple S-VLANs per access interface, firewall-filter-based VLAN assignment, and RVIs or IRB interfaces	12.3R2
Q-in-Q tunneling	12.3R2
Proxy ARP—Unrestricted and restricted	12.3R2
MAC-based VLANs	12.3R2
Link Layer Discovery Protocol (LLDP)	12.3R2
Layer 2 protocol tunneling (L2PT)	12.3R2

Table 19: Layer 3 Protocols Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Bidirectional Forwarding Detection (BFD)	11.3R1	9.0R2	12.1R1	13.2X50 - D10	10.2R1	12.2R1	12.1R1	9.4R1	See Table 18 for a list of EX9200 Layer 3 protocols features.
Border Gateway Protocol (BGP)	N.S.	9.0R2	12.1R1	13.2X50 - D10	11.1R1	12.2R1	11.3R2	9.4R1	
Multiprotocol Border Gateway Protocol (MBGP)	N.S.	9.3R1	12.3R1	13.2X50 - D10	11.2R1	12.2R1	12.1R1	9.4R1	

A separate software license is required for BGP and MBGP. See [Understanding Software Licenses for EX Series Switches](#).

Table 19: Layer 3 Protocols Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Distributed periodic packet management (PPM) with BFD	N.S.	10.4R1	N.S.	13.2X50 - D10	N.S.	N.S.	12.1R1	10.4R1	
Distributed periodic packet management (PPM) with LACP	N.S.	10.2R1	N.S.	13.2X50 - D10	11.1R1	12.2R1	11.3R2	10.2R1	
Filter-based forwarding	N.S.	9.4R1	12.3R1	13.2X50 - D10	11.2R1	12.2R1	11.3R2	9.6R1	
Filter-based forwarding over IPv6	N.S.	10.1R1	N.S.	N.S.	N.S.	N.S.	N.S.	10.3R1	
Intermediate System-to-Intermediate System (IS-IS)	N.S.	9.0R2	N.S.	13.2X50 - D10	11.1R1	12.2R1	11.3R2	9.4R1	
A separate software license is required for IS-IS. See Understanding Software Licenses for EX Series Switches .									
IPv6 Layer 3 multicast protocols	N.S.	10.1R1	N.S.	13.2X50 - D10	N.S.	N.S.	N.S.	10.2R1	
Jumbo frames on RVIs or IRB interfaces	N.S.	9.4R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	N.S.	9.4R1	
OSPF Multitopology Routing (MT-OSPF)	N.S.	9.5R1	N.S.	13.2X50 - D10	N.S.	N.S.	N.S.	N.S.	
See the Junos OS Routing Protocols Configuration Guide .									

Table 19: Layer 3 Protocols Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
OSPFv2	11.1R1	9.0R2	11.4R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
OSPFv3 IPsec support	N.S.	10.3R1	N.S.	13.2X50 - D10	N.S.	N.S.	N.S.	N.S.	
Routing Information Protocol version 1 (RIPv1) and RIPv2	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Static routes	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Virtual routing and forwarding (VRF) with IPv4—Virtual routing instances	12.3R1	9.2R1	12.3R1	13.2X50 - D10	11.1R1	12.2R1	11.3R2	9.6R1	
VRF with IPv4—Virtual routing instances for PIM and IGMP	N.S.	10.0R1	N.S.	13.2X50 - D10	11.1R1	12.2R1	11.3R2	10.0R1	
VRF with IPv4—Virtual routing instances for IGMP snooping	N.S.	11.4R1	N.S.	13.2X50 - D10	12.1R1	12.2R1	N.S.	11.3R1	
VRF with IPv6—Virtual routing instances for multicast traffic	N.S.	10.1R1	N.S.	13.2X50 - D10	N.S.	N.S.	N.S.	10.1R1	
VRF with IPv6—Virtual routing instances for unicast traffic	N.S.	10.1R1	12.3R1	13.2X50 - D10	N.S.	N.S.	N.S.	10.1R1	

Table 20: Layer 3 Protocols Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Accumulated IGP attribute for BGP	12.3R2
Advertisement of the best external BGP path to internal peers	12.3R2

Table 20: Layer 3 Protocols Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Alias support for local autonomous system numbers for BGP	12.3R2
BFD liveness detection	12.3R2
BFD protocol support for OSPFv3	12.3R2
BGP remote next-hop support for single-hop EBGP peers	12.3R2
BGP support for 4-byte autonomous system numbers	12.3R2
BGP support for MDT-SAFI updates without a route target	12.3R2
Behavior change for BGP-independent autonomous system (AS) domains	12.3R2
Bidirectional Forwarding Detection (BFD) hold-down timer	12.3R2
Distributed periodic packet management support for aggregate interfaces	12.3R2
Egress filtering PIMv4/v6 join messages	12.3R2
For internal BGP (IBGP), advertise multiple paths to a destination	12.3R2
Frequent BGP keepalive messages and short BGP hold time	12.3R2
Hitless authentication key rollover for IS-IS	12.3R2
Hub-and-spoke support for multiprotocol BGP-based multicast VPNs with PIM-SSM GRE S-PMSI transport	12.3R2
IPv4 subnet support on loopback interfaces	12.3R2
IS-IS hold-down timer for subsequent SPF calculations	12.3R2
Keepalive support for GRE interfaces	12.3R2
Multitopology routing (MTR)	12.3R2
Nonstop active routing (NSR) support for the Routing Information Protocol (RIP) and RIP next generation (RIPng)	12.3R2
Nonstop active routing (NSR) support	12.3R2
OSPF graceful restart enhancement	12.3R2
OSPF hold-down timer for subsequent SPF calculations	12.3R2
Only the system log notes failure to add routes	12.3R2

Table 20: Layer 3 Protocols Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Origin validation for BGP	12.3R2
PIM join suppression support	12.3R2
Priority assignment for prefixes in OSPF import policies	12.3R2
Reduction in flooding of self-originated OSPF LSAs	12.3R2
Support for BFD over multihop static routes	12.3R2
Support for BFD on logical switches	12.3R2
Support for IPSec authentication for OSPFv2	12.3R2
Support for OSPF database protection for OSPF and OSPFv3	12.3R2
Support for OSPF export and import policies for network-summary LSAs	12.3R2
Support for alternate loop-free routes for IS-IS and OSPF	12.3R2
Support for clearing the VPN tag	12.3R2
Support for disabling the attribute set messages on independent AS domains for BGP loop detection	12.3R2
Support for disabling traps for passive OSPFv2 interfaces	12.3R2
Support for display of flood next-hop branch overflow condition	12.3R2
Support for dropping and ignoring path attributes during BGP neighbor updates	12.3R2
Support for the algorithm that determines the single best path to skip the step that evaluates an AS path	12.3R2
Support for limiting the number of prefixes accepted from a BGP peer	12.3R2
Support for multiarea adjacency in OSPFv2	12.3R2
Support for multiple address families in OSPFv3	12.3R2
Support for route leaking when the switch is in overload mode	12.3R2
Support for route-filter-based BGP outbound route filtering	12.3R2
Support for the BGP Monitoring Protocol	12.3R2
Support to hold down BGP peering sessions after a nonstop active routing (NSR) switchover Timer to delay MED updates for routes advertised by BGP groups or peers configured with the metric-out igp statement Virtual Router Redundancy Protocol (VRRP)	12.3R2

Table 20: Layer 3 Protocols Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Timer to delay MED updates for routes advertised by BGP groups or peers configured with the metric-out igp statement	12.3R2
Virtual Router Redundancy Protocol (VRRP)	12.3R2

Table 21: Logical Systems Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
A separate software license is required for logical systems. See Understanding Software Licenses for EX Series Switches .	
Logical systems	
Layer 2 address learning in logical systems	12.3R2

Table 22: MPLS Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
A separate software license is required for MPLS. See Understanding Software Licenses for EX Series Switches .									

Table 22: MPLS Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Aggregated Ethernet interfaces (LAGs) on circuit cross-connects (CCCs)	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	See Table 22-1 for a list of EX9200 MPLS features.
BFD for an LDP-based LSP	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	12.2R1	
BFD for an RSVP-based LSP	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	12.2R1	
CCC between 2 interfaces in the same switch	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
Interior gateway protocol (IGP) IS-IS and OSPF shortcuts	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
IP over MPLS	N.S.	10.1R1	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
					See Note at end of table	See Note at end of table			
IPv6 over MPLS label-switched paths (LSPs)	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	12.1R1	
					See Note at end of table	See Note at end of table			
LDP-based MPLS	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
LDP tunneling (LDP over RSVP)	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
MPLS-based circuit cross-connects (CCC)	N.S.	9.5R1	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
MPLS label-switched router (LSR) support	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
	N.S.	9.5R1	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	

Table 22: MPLS Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
MPLS Layer 2 CCC on Ethernet- encapsulated interfaces (RFC 6624)									
MPLS Layer 2 CCC on VLAN-encapsulated interfaces (RFC 4905)	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
MPLS Layer 2 VLAN CCC on Ethernet- encapsulated interfaces (RFC 6624)	N.S.	9.5R1	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.3R1	
MPLS Layer 2 VLAN CCC on VLAN-encapsulated interfaces (RFC 4905)	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.3R1	
MPLS Layer 2 VPN over CCC	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
MPLS Layer 2 VPN over VLAN CCC	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.3R1	
MPLS OAM-LSP ping	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	11.1R1	
MPLS over untagged Layer 3 interfaces	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
MPLS with class of service (CoS)	N.S.	9.5R1	N.S.	N.S.	12.2R5	12.2R5	N.S.	12.1R1	
MPLS Layer 3 VPNs	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
MPLS with RSVP-based label-switched paths (LSPs)	N.S.	9.5R1	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
	N.S.	N.S.	N.S.	N.S.			N.S.	12.1R1	

Table 22: MPLS Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Layer 3 subinterfaces as MPLS core interfaces					12.2R1 See Note at end of table.	12.2R1 See Note at end of table.			
Routed VLAN interfaces (RVIs) as MPLS core interfaces	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	12.1R1	
Path maximum transmission unit (MTU) and unicast reverse-path forwarding (RPF) checks for VPNs	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
Resource Reservation Protocol—traffic engineering (RSVP-TE)	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
Standby secondary path protection	N.S.	12.1R1	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
Static label-switched paths (LSPs)	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	12.1R1	



NOTE: Layer 3 subinterfaces as MPLS core interfaces—For EX4500 and EX4550 switches to support Layer 3 subinterfaces as MPLS core interfaces, the peer switch that the Layer 3 subinterfaces connect to, must be an EX8200 switch.

IP over MPLS—The EX4500 and EX4550 switches do not support IP over MPLS (single MPLS label in the packet) when the switch is positioned as a non-penultimate-hop popping (non-PHP) switch.

Table 23: MPLS Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
---------	------------------

A separate software license is required for MPLS. See [Understanding Software Licenses for EX Series Switches](#).

Table 23: MPLS Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Bypass static LSPs	12.3R2
LDP LSP action based on a BFD failure event	12.3R2
LDP downstream on demand	12.3R2
LDP, BGP, and VPLS interworking	12.3R2
P2MP LSP traceroute	12.3R2
Static LSP: <ul style="list-style-type: none"> • Revert timer • Statistics • Traceoptions • At the ingress switch • At the transit switch 	12.3R2
Statistics for P2MP LSPs	12.3R2

Table 24: Multicast Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
IGMP (Internet Group Management Protocol) version 1 (IGMPv1) and IGMPv2	11.1R1	9.0R2	12.1R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	See Table 24-5 for a list of EX9200 multicast features
IGMP filtering	11.3R1	9.5R1	12.3R1	13.2X50 - D10	11.3R1	11.3R1	11.3R1	9.5R1	
IGMP snooping with RVIs or IRB interfaces	10.1R1	9.2R1	12.1R1	13.2X50 - D10	10.2R1	12.2R1	N.S.	9.4R1	
IGMPv3	11.1R1	9.3R2	12.1R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.6R1	
IGMPv1 and IGMPv2 snooping	10.1R1	9.1R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
IGMPv3 snooping	10.1R1	9.6R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.6R1	
Multicast Listener Discovery version 1 and 2 (MLDv1 and MLDv2)	N.S.	10.1R1	N.S.	13.2X50 - D10	11.2R1	12.2R1	12.1R1	10.2R1	
Multicast Listener Discovery version 1 (MLDv1) snooping (MLDv1 snooping)	12.1R1	12.1R1	12.1R1	N.S.	12.1R1	12.2R1	12.1R1	12.1R1	
Multicast Listener Discovery version 2 (MLDv2) snooping (MLDv2 snooping)	12.1R1	12.1R1	12.1R1	N.S.	12.1R1	12.2R1	12.1R1	12.1R1	
Multicast Source Discovery Protocol (MSDP)	N.S.	9.4R1	12.3R1	13.2X50 - D10	10.2R1	12.2R1	12.1R1	9.4R1	
See the Junos OS Multicast Protocols Configuration Guide .									
Multicast VLAN registration (MVR)	11.3R1	9.6R1	12.1R1	N.S.	N.S.	N.S.	N.S.	N.S.	
	11.1R1	9.2R1	12.1R1		11.2R1	12.2R1	11.3R2	9.4R1	

Table 24: Multicast Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Protocol Independent Multicast dense mode (PIM DM)				13.2X50 - D10					
See the Junos OS Multicast Protocols Configuration Guide .									
Protocol Independent Multicast sparse mode (PIM SM)	11.1R1	9.0R2	12.1R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
See the Junos OS Multicast Protocols Configuration Guide .									
Protocol Independent Multicast source-specific multicast (PIM SSM)	11.1R1	9.3R1	12.1R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
See the Junos OS Multicast Protocols Configuration Guide .									
Single-source multicast	N.S.	9.0R2	N.S.	13.2X50 - D10	N.S.	N.S.	N.S.	9.4R1	

Table 25: Multicast Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
BFD for PIM—IPv6	12.3R2
BFD support for ECMP LSPs signaled using LDP	12.3R2
Bidirectional PIM (RFC 5015)	12.3R2
Control of PIM resources for multicast VPNs	12.3R2
Disable PIM for IPv6 only	12.3R2
Dynamic reuse of data multicast distribution tree (MDT) group addresses	12.3R2
Flexible configuration for IGMP or MLD static-join	12.3R2
IGMPv3 and MLDv2 full support	12.3R2

Table 25: Multicast Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
IGMP and MLD enhancements— <ul style="list-style-type: none"> • immediate-leave (IGMP and MLD) • promiscuous-mode (IGMP only) 	12.3R2
IGMP and PIM support for unnumbered interfaces	12.3R2
IGMP join and leave recording for system or for specific interfaces	12.3R2
IGMP and MLD source or group access lists and MLD join and leave recording	12.3R2
IGMP and MLD support for dynamic interfaces	12.3R2
Independently configurable loopback addresses for VRF VPNs	12.3R2
Internet multicast using ingress replication provider tunnels	12.3R2
Software support for configuring accept any-source multicast (ASM) join messages (*;G) for group addresses	12.3R2
Software support for configuring a provider network to operate in source-specific multicast (SSM) mode	12.3R2
LDP signaling for point-to-multipoint LSPs in next-generation MBGP multicast VPNs	12.3R2
Load-balancing PIM join messages on multicast VPNs	12.3R2
Multicast flow maps	12.3R2
Nonstop active routing (NSR) PIM for Draft-Rosen VPNs	12.3R2
PIM automatic make-before-break (MBB) join load balancing	12.3R2
PIM join load balancing	12.3R2
Source-specific multicast (SSM)-map definition for different groups to different sources	12.3R2
Support for filtering unwanted PIM neighbors	12.3R2
Support for multicast output interface (OIF) mapping	12.3R2
Translation of PIM join/prune messages to IGMP or MLD report/leave messages	12.3R2
Turn off spanning-tree interface state in multicast snooping	12.3R2

Table 26: Network Management and Monitoring Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
802.1ag Ethernet OAM connectivity fault management (CFM)	11.2R1	10.2R1	12.3R1	13.2X50 - D10	12.2R1	12.2R1	N.S.	11.4R1	See Table 26-99 for a list of EX9200 network management and monitoring features.
Ethernet frame delay measurement (ETH-DM, Y.1731)	N.S.	11.4R1 (EX4200 only)	N.S.	13.2X50 - D10	11.4R1	12.2R1	N.S.	11.4R1	
Ethernet OAM link fault management (LFM—also known as Ethernet in the First Mile, EFM)	11.1R1	9.4R1	12.2R1	13.2X50 - D10	12.2R1	12.2R1	N.S.	10.0R1	
Port mirroring	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Port mirroring enhancements <ul style="list-style-type: none"> • Layer 3 interface support • Multiple VLAN support 	N.S.	9.5R1	N.S.	—	N.S.	N.S.	N.S.	9.5R1	
Port mirroring enhancements <ul style="list-style-type: none"> • For remote port mirroring, ingress and egress options on VLAN member interfaces on the intermediate (transit) switch to avoid flooding mirrored traffic to those interfaces 	N.S.	10.0R1	N.S.	—	N.S.	N.S.	N.S.	N.S.	
	N.S.	N.S.	N.S.	13.2X50 - D10	11.2R1	12.2R1	N.S.	N.S.	

Table 26: Network Management and Monitoring Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Port mirroring support for multiple analyzers per session									
Real-time performance monitoring (RPM)	10.1R1	9.3R2	12.2R1	13.2X50 - D10	10.2R1	12.2R1	12.1R1	10.1R1	
Real-time performance monitoring (RPM)—hardware timestamps with RVIs or IRB interfaces	N.S.	10.3R1	12.2R1	13.2X50 - D10	10.2R1	12.2R1	12.1R1	10.3R1	
Real-time performance monitoring (RPM)—client and server on same interface	10.3R1	10.3R1	12.2R1	13.2X50 - D10	11.1R1	12.2R1	N.S.	10.3R1	
Routing Engine Software Development Kit (SDK)	N.S.	12.2R1	12.2R1 (EX4200 only)	N.S.	12.2R1	12.2R1	N.S.	12.2R1	
RMON	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
sFlow monitoring technology	11.1R1	9.3R2	12.1R1	13.2X50 - D10	11.2R1	12.2R1	12.1R1	10.0R1	
sFlow monitoring technology—Persistent IP addresses for agent IDs and use in datagrams	11.1R1	10.2R1	N.S.	13.2X50 - D10	N.S.	N.S.	12.1R1	10.2R1	
Simple Network Management Protocol version 1 (SNMPv1), SNMPv2, and SNMPv3	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
	11.1R1	11.1R1	12.1R2		11.1R1	12.2R1	N.S.	12.1R1	

Table 26: Network Management and Monitoring Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Uplink failure detection				13.2X50 - D10					

Table 27: Network Management and Monitoring Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Junos OS XML API and scripting—NETCONF Java toolkit for rapid development of Java applications to manage devices running Junos OS	12.3R2
Junos OS XML API and scripting—NETCONF Perl client installation—Supports loading prerequisites from Comprehensive Perl Archive Network (CPAN) global repository	12.3R2
Junos OS XML API and scripting—NETCONF tracing operations	12.3R2
Junos OS XML API and scripting: <ul style="list-style-type: none"> • Dedicated directory for user script library • Global variable provided to Junos OS automation scripts • References to a correlating event in a policy action • Trigger a policy based on the event count • Unique filenames for uploaded files • Upload files created by event scripts • XML schemata for Junos OS XML operational tag elements • jcs:open() extension function support for routing instances 	12.3R2
Configuration options to filter out interfaces from SNMP Get and GetNext operations	12.3R2
Enhanced SNMP support for logical switches and routing instances	12.3R2
Generating SNMP traps when MAC address table is full	12.3R2
Junos OS MIB support for VPLS	12.3R2
MIB support for VRF route entries	12.3R2
Proxy SNMP agent	12.3R2
SNMP MIB support for OSPFv3	12.3R2
SNMP poll and trap support for DHCP leases	12.3R2
SNMP support for the DHCP bindings table	12.3R2
SNMP support for the authd daemon and for radius-acc-server-mib and radius-auth-server-mib	12.3R2

Table 27: Network Management and Monitoring Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
SNMP support for spanning-tree protocols	12.3R2
Support for Internet draft draft-ietf-bfd-mib-02.txt—MIB for BFD liveness detection	12.3R2
Support for MIB objects in accounting profiles	12.3R2
Support for an enterprise-specific event MIB (mib-jnx-event.txt)	12.3R2
Support for sending traps over routing instances	12.3R2
Support for adding lists of clients to the SNMP community	12.3R2
Support for the enterprise-specific Packet Forwarding Engine MIB (mib-jnx-pfe.txt)	12.3R2
Support for the pimNeighborLoss trap	12.3R2
Support for trap spoofing	12.3R2
IEEE 802.3ah link fault management (LFM) for Ethernet OAM (also known as Ethernet in the First Mile, or EFM)	12.3R2
Port mirroring of Layer 2 VLAN and VPLS traffic	12.3R2
Fast update filters for dynamic profiles	12.3R2
Flow aggregation to multiple collectors	12.3R2
IPv6 flow aggregation templates	12.3R2
Inline flow monitoring	12.3R2

Table 28: Port Security Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Automatic recovery for port error disable conditions	10.1R1	9.6R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	10.0R1	N.S.
DHCP option 82	10.1R1	9.3R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	N.S.	9.4R1	12.3R2
DHCP snooping	10.1R1	9.0R2	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	10.3R1	N.S.

Table 28: Port Security Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Dynamic ARP inspection (DAI)	10.1R1	9.0R2	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	10.3R1	N.S.
IP source guard	10.1R1	9.2R1	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	10.3R1	N.S.
Layer 3 virtual private network (VPN) for IPv4 (RFC 2547 and 4364)	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	12.3R2
Layer 3 virtual private network (VPN) for IPv6 through IPv4 MPLS	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	12.3R2
MAC limiting	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2
MAC address limit per port	10.1R1	9.0R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	10.3R1	12.3R2
MAC limiting per port and per VLAN (VLAN membership MAC limit)	12.3R1	12.3R1	12.3R1	13.2X50 - D10	12.3R1	12.3R1	12.3R1	12.3R1	12.3R2
MAC move limiting	10.1R1	9.0R2	11.3R1	13.2X50 - D10	N.S.	N.S.	11.3R2	N.S.	N.S.
Persistent MAC learning (sticky MAC)	11.4R1	11.4R1	12.3R1	13.2X50 - D10	11.4R1	12.2R1	11.4R1	11.4R1	N.S.
Persistent storage for DHCP snooping	10.1R1	9.4R1	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	10.3R1	N.S.
Self-signed digital certificates for enabling SSL services	11.1R1	11.1R1	N.S.	13.2X50 - D10	11.1R1	12.2R1	12.1R1	11.1R1	N.S.
Static ARP support	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2

Table 29: Power over Ethernet (PoE) Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Link Layer Discovery Protocol (LLDP) with granular Power over Ethernet (PoE) management	12.2R1	12.2R1 (EX4200-24PX and EX4200-48PX models only)	12.2R1	13.2X50 - D10	N.S.	N.S.	12.2R1	12.2R1	N.S.
NOTE: EX4200 switches must be running PoE controller software firmware version 4.04 or later to support the Link Layer Discovery Protocol (LLDP) with granular Power over Ethernet (PoE) management feature. See show chassis firmware detail and request system firmware upgrade poe to check or upgrade this firmware.									N.S.
Power over Ethernet (PoE)	10.1R1	9.0R2	11.3R1	13.2X50 - D10	—	—	11.3R2	11.2R1	N.S.
Power over Ethernet Plus (PoE+)	10.3R1	11.2R1 (EX4200-24PX and EX4200-48PX models only)	11.3R1	13.2X50 - D10	—	—	11.3R2	11.2R1	N.S.
Power over Ethernet (PoE) power management mode	10.1R1	9.3R2	11.3R1	13.2X50 - D10	—	—	11.3R2	11.2R1	N.S.

Table 30: Routing Policy and Packet Filtering Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Dynamic allocation of TCAM memory to firewall filters	10.1R1	10.0R1	11.3R1	–	10.2R1	12.2R1	N.S.	10.3R1	See Table 30 for a list of EX9200 routing policy and firewall filter features.
Firewall filters and rate limiting	10.1R1	9.0R2	11.3R1	13.2X50 – D10	10.2R1	12.2R1	11.3R2	9.4R1	
For a list of supported firewall filter match conditions and actions, see <i>Platform Support for Firewall Filter Match Conditions, Actions, and Action Modifiers on EX Series Switches</i> .									
Firewall filters on LAGs	10.1R1	9.0R2	11.3R1	13.2X50 – D10	10.2R1	12.2R1	N.S.	10.0R1	
Firewall filters on the loopback interface	10.1R1	9.2R1	11.3R1	13.2X50 – D10	11.1R1	12.2R1	12.1R1	9.6R1	
For a list of supported firewall filter match conditions and actions on a loopback interface, see <i>Support for Match Conditions and Actions for Loopback Firewall Filters on Switches</i> .									
Firewall filters on the management interface	11.3R1	10.4R1	N.S.	13.2X50 – D10	10.4R1	12.2R1	12.1R1	10.4R1	
Firewall filters on the virtual management interface	–	10.4R1 (EX4200 Virtual Chassis only)	N.S.	13.2X50 – D10	–	–	–	–	
Firewall filters with IPv6	11.3R1	10.1R1	12.3R1	13.2X50 – D10	12.1R1	12.2R1	12.1R1	10.3R1	
Policing	10.1R1	9.0R2	11.3R1	13.2X50 – D10	10.2R1	12.2R1	11.3R2	9.4R1	
Tricolor marking policers	11.2R1	11.2R1	11.4R8	N.S.	12.1R1	12.2R1	11.3R1	N.S.	

Table 31: Routing Policy and Firewall Filters on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Access and access-internal routes	12.3R2
Extension of numeric-range match conditions in firewall filters	12.3R2
Aggregate policer support for different family address types configured on the same interface	12.3R2

Table 31: Routing Policy and Firewall Filters on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Authentication for BFD (MD5/SHA1)	12.3R2
BGP multipath link-bandwidth attribute	12.3R2
DHCP state persistence for DHCP relay agent	12.3R2
Dynamic configuration support for routing policies	12.3R2
Extended DHCP relay agent	12.3R2
Filter-based forwarding to a specific outgoing interface or destination IP address	12.3R2
Firewall filters within logical systems	12.3R2
IEEE 802.1p priority match conditions for Layer 2 VPN firewall filters	12.3R2
Filter-based forwarding to a specific outgoing interface or destination IP address	12.3R2
Layer 2 Gigabit Ethernet logical interface extended policing support	12.3R2
Layer 2 support for firewall filter match conditions	12.3R2
Load balancing of VPLS traffic	12.3R2
Option 60 support for extended DHCP relay agents	12.3R2
Policers on physical interfaces	12.3R2
Firewall filters feature support	12.3R2
Support for policers that limit traffic on logical interfaces in ingress or egress directions	12.3R2
Support for policers that rate-limit based on a percentage of physical port speed on an interface	12.3R2
Support for the discard action for tricolor marking policers applied to firewall filters	12.3R2
Support for the prefix-list match condition for firewall filters for the VPLS protocol family	12.3R2
Support for enhanced policer statistics	12.3R2
Support for MAC address validation	12.3R2
Tricolor marking policers	12.3R2

Table 32: Spanning-Tree Protocols Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
BPDU protection for spanning-tree protocols	10.1R1	9.1R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	See Table 33 for a list of EX9200 spanning-tree protocols features.
BPDU filter	12.2R1	12.2R1	12.2R1	13.2X50 - D10	12.2R1	12.2R1	12.2R1	12.2R1	
Distributed periodic packet management (PPM) for Spanning Tree Protocols	N.S.	12.3R1	N.S.	N.S.	N.S.	N.S.	N.S.	12.3R1	
Loop protection for spanning-tree protocols	10.1R1	9.1R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Root protection for spanning-tree protocols	10.1R1	9.1R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Spanning tree: <ul style="list-style-type: none"> RSTP and VSTP concurrent configuration 	N.S.	10.2R1	12.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	10.2R1	
Spanning tree: <ul style="list-style-type: none"> Spanning Tree Protocol (STP) Rapid Spanning Tree Protocol (RSTP) Multiple Spanning Tree Protocol (MSTP) 	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Spanning tree: <ul style="list-style-type: none"> VLAN Spanning Tree Protocol (VSTP) 	10.1R1	9.4R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.6R1	

Table 33: Spanning-Tree Protocols Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Spanning tree:	12.3R2
<ul style="list-style-type: none">• Spanning Tree Protocol (STP)• Rapid Spanning Tree Protocol (RSTP)• Multiple Spanning Tree Protocol (MSTP)• VLAN Spanning Tree Protocol (VSTP)	
Root protection for spanning-tree protocols	12.3R2
Loop protection for spanning-tree protocols	12.3R2
Distributed periodic packet management (PPM) for Spanning Tree Protocols	12.3R2
BPDU filter	12.3R2

Table 34: System Management Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Autoinstallation of configuration files	10.1R1	9.4R1	11.3R1	N.S.	10.2R1	12.2R1	11.3R2	N.S.	See Table 34-1 for a list of EX9200 system management features.
Automatic software download	10.1R1	9.6R1	11.3R1	N.S.	10.2R1	12.2R1	11.3R2	9.6R1	
Automatic repair of corrupted partition when booting from alternate partition	12.3R1	12.3R1	12.3R1	13.2X50 - D10	12.3R1	12.3R1	12.3R1	12.3R1	
Configuration rollback	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Zero Touch Provisioning (EZ Touchless Provisioning using DHCP)	12.2R1	12.2R1	12.2R5	N.S.	12.2R1	12.2R1	N.S.	N.S.	
J-Web interface, for switch configuration and management	10.1R1 (12.1R1 for EX2200-C switches)	9.0R2	12.1R1	13.2X50 - D10	10.2R1	12.2R1	12.1R1	9.4R1	
Junos Space Service Now support	12.3R1	12.3R1	12.3R1	N.S.	12.3R1	12.3R1	12.3R1	12.3R1	
LCD panel management support	—	9.0R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R1	9.4R1	
Online insertion and removal (OIR) of uplink modules	—	10.0R1	—	13.2X50 - D10	11.1R1	12.2R1	—	—	

Table 35: System Management Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Configuration rollback	12.3R2

Table 36: User Interface and Configuration Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Device-initiated SSH connection (outbound SSH)	12.3R2
Dynamic IPv6 filters	12.3R2
Dynamic configuration of the switch advertisement protocol	12.3R2
Dynamic profiles support by extended DHCP local server and extended DHCP relay agent	12.3R2
Enhanced IPv6 statistics	12.3R2
Extended DHCP local server	12.3R2
IGMP dynamic profiles	12.3R2
Extended DHCP local server	12.3R2
Protection for device configuration	12.3R2
RADIUS MSCHAPv2 protocol support for administrator authentication, password aging, and update	12.3R2
Limit configuration command output	12.3R2
Remote tracing	12.3R2
Support for CLI edit mode wildcard range	12.3R2
Support for configuring ARP aging time for a logical interface	12.3R2
Support for configuring a proxy server for downloading licenses	12.3R2
Support for configuring time-based user access	12.3R2
Support for logical router system administrators	12.3R2

Table 37: VPN Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Aggregated Ethernet interfaces for VPLS routing instances	12.3R2
BGP autodiscovery for LDP VPLS (FEC 129)	12.3R2
Clearing MAC addresses for better convergence	12.3R2
Configurable label block sizes for VPLS	12.3R2
Disable TTL propagation behavior for the routes in a VRF routing instance	12.3R2

Table 37: VPN Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
EXP-based traffic classification for VPLS	12.3R2
Enhanced show interface command for Layer 3 VPN functionality	12.3R2
Expanded interface support for the vrf-table-label statement	12.3R2
Extranet next-generation MVPN GRE tunnels for Layer 3 VPNs	12.3R2
GRE tunnels for Layer 3 VPNs Ignore MTU mismatch on Layer 2 circuits Integrated routing and bridging support for inter-AS VPLS between BGP-signaled VPLS and LDP-signaled VPLS LDP-based VPLS Label allocation and substitution policy	12.3R2
Ignore MTU mismatch on Layer 2 circuits	12.3R2
Integrated routing and bridging support for inter-AS VPLS between BGP-signaled VPLS and LDP-signaled VPLS	12.3R2
LDP-based VPLS	12.3R2
Label allocation and substitution policy	12.3R2
Layer 2 VPN multihoming	12.3R2
Layer 3 VPN BGP routes and labels	12.3R2
Layer 3 VPN localization	12.3R2
Load balancing and IP header filtering for Layer 3 VPNs	12.3R2
Local switching support for the ignore-encapsulation-mismatch statement	12.3R2
Multipath load balancing for EBGp and IBGP VPNs	12.3R2
Multiple logical trunk interfaces per physical interface	12.3R2
Multiprotocol BGP-based multicast VPN	12.3R2
NTP support for IPv4 VRF and IPv6 VRF	12.3R2
Nonstop active routing support for Layer 3 VPNs	12.3R2
PIM source-specific multicast (PIM-SSM) provider tunnel support added to Multiprotocol BGP-based multicast VPNs	12.3R2
Point-to-multipoint LSP support for VPLS	12.3R2
Point-to-multipoint LSP support for multicast VPNs	12.3R2

Table 37: VPN Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Proxy BGP route target filtering	12.3R2
Static VPLS	12.3R2
Static route target filtering	12.3R2
Support for autorp, BSR, PIM dense mode and mtrace for next-generation multicast VPNs	12.3R2
VLAN range for Layer 2 VPN	12.3R2
VPLS automatic site ID	12.3R2
VPLS automatic site ID for nonstop active routing	12.3R2
VPLS ping	12.3R2
VPLS trunk interfaces	12.3R2
eBGP and iBGP load-balancing support for MVPN and PIM	12.3R2

**Related
Documentation**

- *EX Series Virtual Chassis Software Features Overview*
- *EX2200 Switches Hardware Overview*
- *EX3200 Switches Hardware Overview*
- *EX3300 Switches Hardware Overview*
- *EX4200 Switches Hardware Overview*
- <will add topic-ref to EX4300 HW overview topic>
- *EX4500 Switches Hardware Overview*
- *EX4550 Switches Hardware Overview*
- *EX6210 Switch Hardware Overview*
- *EX8208 Switch Hardware Overview*
- *EX8216 Switch Hardware Overview*
- *EX9204 Switch Hardware Overview*
- [EX9208 Switch Hardware Overview on page 115](#)
- *EX9214 Switch Hardware Overview*
- *Line Card Model and Version Compatibility in an EX6200 Switch*
- *Line Card Model and Version Compatibility in an EX8200 Switch*
- *Line Card Model and Version Compatibility in an EX9200 Switch*

- *XRE200 External Routing Engine Hardware Overview*
- *Layer 3 Protocols Supported on EX Series Switches*
- *Layer 3 Protocols Not Supported on EX Series Switches*

PART 1

Supported Hardware

- [EX9208 Switch Overview on page 115](#)

CHAPTER 2

EX9208 Switch Overview

- [EX9208 Switch Hardware Overview on page 115](#)

EX9208 Switch Hardware Overview

Juniper Networks EX9208 Ethernet Switches provide high performance, scalable connectivity, and carrier-class reliability for high-density environments such as campus-aggregation and data-center networks. The EX9208 switch has a capacity of up to 4.8 terabits per second (Tbps) or up to 240 gigabit per second (Gbps) per slot full duplex. The EX9208 switch is a modular system that provides high availability and redundancy for all major hardware components, including Routing Engine modules (RE modules), Switch Fabric modules (SF modules), fan tray (redundant fans), and power supplies.

A fully populated EX9208 switch provides a maximum port density of 240 1-Gigabit Ethernet ports, 160 10-Gigabit Ethernet ports, or 24 40-Gigabit Ethernet ports in a redundant configuration.

You can manage EX9208 switches by using the same interfaces that you use for managing other devices running the Juniper Networks Junos operating system (Junos OS)—the command-line interface (CLI) and the Network and Security Manager (NSM).

- [Software on page 115](#)
- [Chassis Physical Specifications on page 116](#)
- [Host Subsystem on page 117](#)
- [Line Cards on page 117](#)
- [Cooling System on page 118](#)
- [Power Supplies on page 118](#)

Software

The Juniper Networks EX Series Ethernet Switches run Junos OS, which provides Layer 2 and Layer 3 switching, routing, and security services. The same Junos OS code base that runs on EX Series switches also runs on all Juniper Networks J Series, M Series, MX Series, and T Series routers and SRX Series Services Gateways.

Chassis Physical Specifications

The EX9208 switch is eight rack units (8 U) in size. Five EX9208 switches can fit in a standard 48 U rack. Each EX9208 switch is designed to optimize rack space and cabling. See [Figure 1 on page 116](#), [Figure 2 on page 116](#), and [Figure 3 on page 117](#).

Figure 1: Front View of an EX9208 Switch




Figure 2: Rear View of an AC-Powered EX9208 Switch





Figure 3: Rear View of a DC-Powered EX9208 Switch



Host Subsystem

Switching and routing functionality, system management, and system control functions of an EX9208 switch are performed by the host subsystem. The host subsystem consists of a Routing Engine functioning together with a Switch Fabric.

You can install either one or two host subsystems in the slots labeled **0** and **1** in the front panel of the chassis. A base configuration EX9208 switch has one host subsystem. A redundant configuration EX9208 switch has a second host subsystem. See *EX9208 Switch Configurations*.

Line Cards

The EX9208 switch has six horizontal line card slots and supports line rate for each line card. The line cards in EX9208 switches combine a Packet Forwarding Engine and Ethernet interfaces in a single assembly. Line cards are field-replaceable units (FRUs) that can be installed in the line card slots labeled **0** through **5** on the front of the switch chassis. All line cards are hot-removable and hot-insertable.

Four line cards are available for EX9208 switches:

Table 38: Line Cards Available for EX9208 Switches

Model	Description	Additional Information
EX9200-32XS	32-port SFP+ line card	<i>EX9200-32XS Line Card</i>
EX9200-40T	40-port 10/100/1000BASE-T RJ-45 line card	<i>EX9200-40T Line Card</i>

Table 38: Line Cards Available for EX9208 Switches (*continued*)

Model	Description	Additional Information
EX9200-40F	40-port 100FX/1000BASE-X SFP line card	<i>EX9200-40F Line Card</i>
EX9200-4QS	4-port 40-Gigabit Ethernet QSFP+ line card	<i>EX9200-4QS Line Card</i>

Cooling System

The cooling system in an EX9208 switch is a field-replaceable unit (FRU). It consists of a hot-removable and hot-insertable fan tray. The fan tray contains six fans. The fan tray installs vertically on the right back of the chassis and provides side-to-side chassis cooling. See *Cooling System and Airflow in an EX9208 Switch*.

Power Supplies

Power supplies for the EX9208 switch are fully redundant, load-sharing, and hot-removable and hot-insertable FRUs. Each EX9208 switch chassis can hold up to four AC or DC power supplies.

[Table 39 on page 118](#) shows the details of the power supplies available for EX9208 switches.

Table 39: Power Supplies Supported on EX9208 Switches

Power Supply	Input Voltage	Output Power
2520 W AC	Low-voltage line (100–120 VAC)	1167 W
	High-voltage line (220–240 VAC)	2050 W
2400 W DC	–40 VDC through –70 VDC	2400 W to 2600 W

An AC base configuration EX9208 switch ships with three low-line (110–120 VAC) AC power supplies or two high-line (220–240 VAC) AC power supplies. A redundant AC configuration EX9208 switch ships with four low-line (110–120 VAC) or four high-line (220–240 VAC) AC power supplies. See *AC Power Supply in an EX9208 Switch*.

A redundant DC configuration EX9208 switch ships with four DC power supplies. See *DC Power Supply in an EX9208 Switch*.



CAUTION: Mixing different types of power supplies (AC and DC) in the same chassis is not supported.

Related Documentation

- *Connecting and Configuring an EX9200 Switch (CLI Procedure)*
- *Field-Replaceable Units in an EX9200 Switch*

- *Line Card Model and Version Compatibility in an EX9200 Switch*

PART 2

System Administration

- [Access Privileges on page 123](#)
- [System Basics on page 295](#)
- [System Services on page 875](#)

CHAPTER 3

Access Privileges

- [Overview on page 123](#)
- [Configuration on page 131](#)
- [Administration on page 292](#)

Overview

- [Introduction to Access Privileges on page 123](#)

Introduction to Access Privileges

- [Understanding Junos OS Access Privilege Levels on page 123](#)
- [Junos OS Login Classes Overview on page 128](#)
- [Access Privilege User Permission Flags Overview on page 128](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)

Understanding Junos OS Access Privilege Levels

Each top-level command-line interface (CLI) command and each configuration statement have an access privilege level associated with them. Users can execute only those commands and configure and view only those statements for which they have access privileges. The access privileges for each login class are defined by one or more *permission flags*.

For each login class, you can explicitly deny or allow the use of operational and configuration mode commands that would otherwise be permitted or not allowed by a privilege level specified in the **permissions** statement.

The following sections provide additional information about permissions:

- [Junos OS Login Class Permission Flags on page 123](#)
- [Allowing or Denying Individual Commands for Junos OS Login Classes on page 127](#)

Junos OS Login Class Permission Flags

The **permissions** statement specifies one or more of the permission flags listed in [Table 40 on page 124](#). Permission flags are not cumulative, so for each class you must list all the permission flags needed, including **view** to display information and **configure**

to enter configuration mode. Two forms of permissions control for individual parts of the configuration are:

- "Plain" form—Provides read-only capability for that permission type. An example is **interface**.
- Form that ends in **-control**—Provides read and write capability for that permission type. An example is **interface-control**.

Table 40 on page 124 lists the Junos® operating system (Junos OS) login class permission flags that you can configure by including the **permissions** statement at the **[edit system login class class-name]** hierarchy level.

Table 40: Login Class Permission Flags

Permission Flag	Description
access	Can view the access configuration in configuration mode and with the show configuration operational mode command.
access-control	Can view and configure access information at the [edit access] hierarchy level.
admin	Can view user account information in configuration mode and with the show configuration operational mode command.
admin-control	Can view user accounts and configure them at the [edit system login] hierarchy level.
all-control	Can access all operational mode commands and configuration mode commands. Can modify configuration in all the configuration hierarchy levels.
clear	Can clear (delete) information learned from the network that is stored in various network databases by using the clear commands.
configure	Can enter configuration mode by using the configure command.
control	Can perform all control-level operations—all operations configured with the -control permission flags.
field	Can view field debug commands. Reserved for debugging support.
firewall	Can view the firewall filter configuration in configuration mode.
firewall-control	Can view and configure firewall filter information at the [edit firewall] hierarchy level.
floppy	Can read from and write to the removable media.
flow-tap	Can view the flow-tap configuration in configuration mode.

Table 40: Login Class Permission Flags (*continued*)

Permission Flag	Description
flow-tap-control	Can view the flow-tap configuration in configuration mode and can configure flow-tap configuration information at the [edit services flow-tap] hierarchy level.
flow-tap-operation	Can make flow-tap requests to the router or switch. For example, a Dynamic Tasking Control Protocol (DTCP) client must authenticate itself to the Junos OS as an administrative user. That account must have flow-tap-operation permission. NOTE: The flow-tap-operation option is not included in the all-control permissions flag.
idp-profiler-operation	Can view profiler data.
interface	Can view the interface configuration in configuration mode and with the show configuration operational mode command.
interface-control	Can view chassis, class of service (CoS), groups, forwarding options, and interfaces configuration information. Can edit configuration at the following hierarchy levels: <ul style="list-style-type: none"> • [edit chassis] • [edit class-of-service] • [edit groups] • [edit forwarding-options] • [edit interfaces]
maintenance	Can perform system maintenance, including starting a local shell on the router or switch and becoming the superuser in the shell by using the su root command, and can halt and reboot the router or switch by using the request system commands.
network	Can access the network by using the ping , ssh , telnet , and traceroute commands.
pgcp-session-mirroring	Can view the pgcp session mirroring configuration.
pgcp-session-mirroring-control	Can modify the pgcp session mirroring configuration.
reset	Can restart software processes by using the restart command and can configure whether software processes are enabled or disabled at the [edit system processes] hierarchy level.
rollback	Can use the rollback command to return to a previously committed configuration other than the most recently committed one.
routing	Can view general routing, routing protocol, and routing policy configuration information in configuration and operational modes.

Table 40: Login Class Permission Flags (*continued*)

Permission Flag	Description
routing-control	Can view general routing, routing protocol, and routing policy configuration information and can configure general routing at the [edit routing-options] hierarchy level, routing protocols at the [edit protocols] hierarchy level, and routing policy at the [edit policy-options] hierarchy level.
secret	Can view passwords and other authentication keys in the configuration.
secret-control	Can view passwords and other authentication keys in the configuration and can modify them in configuration mode.
security	Can view security configuration in configuration mode and with the show configuration operational mode command.
security-control	Can view and configure security information at the [edit security] hierarchy level.
shell	Can start a local shell on the router or switch by using the start shell command.
snmp	Can view Simple Network Management Protocol (SNMP) configuration information in configuration and operational modes.
snmp-control	Can view SNMP configuration information and can modify SNMP configuration at the [edit snmp] hierarchy level.
system	Can view system-level information in configuration and operational modes.
system-control	Can view system-level configuration information and configure it at the [edit system] hierarchy level.
trace	Can view trace file settings and configure trace file properties.
trace-control	Can modify trace file settings and configure trace file properties.
view	Can use various commands to display current system-wide, routing table, and protocol-specific values and statistics. Cannot view the secret configuration.
view-configuration	Can view all of the configuration excluding secrets, system scripts, and event options. NOTE: Only users with the maintenance permission can view commit script, op script, or event script configuration.

Allowing or Denying Individual Commands for Junos OS Login Classes

By default, all top-level CLI commands have associated access privilege levels. Users can execute only those commands and view only those statements for which they have access privileges. For each login class, you can explicitly deny or allow the use of operational and configuration mode commands that would otherwise be permitted or not allowed by a privilege level specified in the **permissions** statement.

- The **all** login class permission bits take precedence over extended regular expressions when a user with **rollback** permission issues the **rollback** command.
- Expressions used to allow and deny commands for users on RADIUS and TACACS+ servers have been simplified. Instead of a single, long expression with multiple commands (**allow-commands=cmd1 cmd2 ... cmdn**), you can specify each command as a separate expression. This new syntax is valid for **allow-configuration** and **deny-configuration**, **allow-commands** and **deny-commands**, and all user permission bits.
- Users cannot issue the **load override** command when specifying an extended regular expression. Users can only issue the **merge**, **replace**, and **patch** configuration commands.
- If you allow and deny the same commands, the **allow-commands** permissions take precedence over the permissions specified by the **deny-commands**. For example, if you include **allow-commands "request system software add"** and **deny-commands "request system software add"**, the login class user is allowed to install software using the **request system software add** command.
- Regular expressions for **allow-commands** and **deny-commands** can also include the **commit**, **load**, **rollback**, **save**, **status**, and **update** commands.
- If you specify a regular expression for **allow-commands** and **deny-commands** with two different variants of a command, the longest match is always executed.

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

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

Related Documentation

- [Configuring Access Privilege Levels on page 131](#)

Junos OS Login Classes Overview

All users who can log in to the router or switch must be in a login class. With login classes, you define the following:

- Access privileges that users have when they are logged in to the router or switch
- Commands and statements that users can and cannot specify
- How long a login session can be idle before it times out and the user is logged out

You can define any number of login classes and then apply one login class to an individual user account.

The Junos OS contains a few predefined login classes, which are listed in [Table 41 on page 128](#). The predefined login classes cannot be modified.

Table 41: Predefined System Login Classes

Login Class	Permission Flag Set
operator	clear, network, reset, trace, and view
read-only	view
superuser or super-user	all
unauthorized	None



NOTE:

- You cannot modify a predefined login class name. If you issue the **set** command on a predefined class name, the Junos OS appends **-local** to the login class name. The following message also appears:

warning: '<class-name>' is a predefined class name; changing to '<class-name>-local'

- You cannot issue the **rename** or **copy** command on a predefined login class. Doing so results in the following error message:

error: target '<class-name>' is a predefined class

Related Documentation

- *Defining Junos OS Login Classes*
- *Defining Junos OS Login Classes*
- *Understanding QFabric System Login Classes*

Access Privilege User Permission Flags Overview

Permission flags are used to grant a user access to operational mode commands and configuration hierarchy levels and statements. By specifying a specific permission flag

on the user's login class at the **[edit system login class]** hierarchy level, you grant the user access to the corresponding commands and configuration hierarchy levels and statements. To grant access to all commands and configuration statements, use the **all** permissions flag.

For permission flags that grant access to configuration hierarchy levels and statements, the flags grant read-only privilege to that configuration. For example, the **interface** permissions flag grants read-only access to the **[edit interfaces]** hierarchy level. The **-control** form of the flag grants read-write access to that configuration. Using the preceding example, **interface-control** grants read-write access to the **[edit interfaces]** hierarchy level.

The permission flags listed in "Related Documentation" grant a specific set of access privileges. Each permission flag is listed with the operational mode commands and configuration hierarchy levels and statements for which that flag grants access.



NOTE: Each command listed represents that command and all subcommands with that command as a prefix. Each configuration statement listed represents the top of the configuration hierarchy to which that flag grants access.

Related Documentation

- [Understanding Junos OS Access Privilege Levels on page 123](#)
- [access on page 140](#)
- [access-control on page 141](#)
- [admin on page 142](#)
- [admin-control on page 142](#)
- [all-control on page 143](#)
- [clear on page 143](#)
- [configure on page 180](#)
- [control on page 181](#)
- [field on page 181](#)
- [firewall on page 181](#)
- [firewall-control on page 182](#)
- [floppy on page 183](#)
- [flow-tap on page 183](#)
- [flow-tap-operation on page 184](#)
- [idp-profiler-operation on page 184](#)
- [interface on page 185](#)
- [interface-control on page 185](#)
- [maintenance on page 186](#)

- [network on page 193](#)
- [pgcp-session-mirroring on page 194](#)
- [pgcp-session-mirroring-control on page 195](#)
- [reset on page 195](#)
- [rollback on page 196](#)
- [routing on page 196](#)
- [routing-control on page 201](#)
- [secret on page 205](#)
- [secret-control on page 206](#)
- [security on page 207](#)
- [security-control on page 210](#)
- [shell on page 214](#)
- [snmp on page 214](#)
- [system on page 215](#)
- [system-control on page 216](#)
- [trace on page 218](#)
- [trace-control on page 223](#)
- [view on page 228](#)
- [view-configuration on page 292](#)

Specifying Access Privileges for Junos OS Configuration Mode Hierarchies

The **allow/deny-configuration** and **allow/deny-configuration-regexps** statements let you explicitly allow or deny users access privileges to portions of the configuration hierarchy. Each of these statements is added to named login classes and configured with one or more regular expressions to be allowed or denied. Each login class is assigned to specific users or user IDs.

The search and match methods differ in the two forms of these statements. You must select which form to use within a login class—you cannot configure **allow-configuration** and **allow-configuration-regexps** together in the same login class. You must select just one. If you have existing configurations using the **allow/deny-configuration** form of the statements, using the same configuration options with the **allow/deny-configuration-regexps** form of the statements might not produce the same results.

- **Allow/deny-configuration** statements perform slower matching, with more flexibility, especially in wildcard matching. However, it can take a very long time to evaluate all of the possible statements if a great number of full path regular expressions or wildcard

expressions are configured, possibly impacting performance. These statements were introduced before Junos OS Release 7.4.

- **Allow/deny-configuration-regexps** statements perform faster matching, with less flexibility. You configure a set of strings in which each string is a regular expression, with spaces between the terms of the string. This provides very fast matching. However, it is more tedious to use wildcard expressions in this form of the statement, because you must set up wildcards for each token (term) of the space-delimited string you want to match. These statements were introduced in Junos OS Release 11.2.

Related Documentation

- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- [Example: Configuring Access Privilege Levels on page 134](#)
- [*Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies*](#)
- [Understanding Junos OS Access Privilege Levels on page 123](#)

Configuration

- [Configuring Access Privileges on page 131](#)
- [Examples on page 134](#)
- [User Permission Flags Reference on page 139](#)

Configuring Access Privileges

- [Configuring Access Privilege Levels on page 131](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 133](#)

Configuring Access Privilege Levels

Each top-level command-line interface (CLI) command and each configuration statement have an access privilege level associated with it. Users can execute only those commands and configure and view only those statements for which they have access privileges.

To configure access privilege levels, include the **permissions** statement at the **[edit system login class *class-name*]** hierarchy level:

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

Related Documentation

- [Example: Configuring Access Privilege Levels on page 134](#)
- [Understanding Junos OS Access Privilege Levels on page 123](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- *permissions*

Specifying Access Privileges for Junos OS Operational Mode Commands

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

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

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

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

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

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

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



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

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

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

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

If you allow and deny the same commands, the **allow-commands** permissions take precedence over the permissions specified by **deny-commands**. For example, if you include **allow-commands "request system software add"** and **deny-commands "request system**

software add", the login class user is allowed to install software using the **request system software add** command.

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

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

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

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

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

Related Documentation

- [Example: Configuring Access Privileges for Operational Mode Commands on page 134](#)
- *Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands*
- *allow-commands*
- *deny-commands*

Specifying Access Privileges for Junos OS Configuration Mode Hierarchies

The **allow/deny-configuration** and **allow/deny-configuration-regexps** statements let you explicitly allow or deny users access privileges to portions of the configuration hierarchy. Each of these statements is added to named login classes and configured with one or more regular expressions to be allowed or denied. Each login class is assigned to specific users or user IDs.

The search and match methods differ in the two forms of these statements. You must select which form to use within a login class—you cannot configure **allow-configuration** and **allow-configuration-regexps** together in the same login class. You must select just one. If you have existing configurations using the **allow/deny-configuration** form of the statements, using the same configuration options with the **allow/deny-configuration-regexps** form of the statements might not produce the same results.

- **Allow/deny-configuration** statements perform slower matching, with more flexibility, especially in wildcard matching. However, it can take a very long time to evaluate all of the possible statements if a great number of full path regular expressions or wildcard expressions are configured, possibly impacting performance. These statements were introduced before Junos OS Release 7.4.
- **Allow/deny-configuration-regexps** statements perform faster matching, with less flexibility. You configure a set of strings in which each string is a regular expression, with spaces between the terms of the string. This provides very fast matching. However, it is more tedious to use wildcard expressions in this form of the statement, because you must set up wildcards for each token (term) of the space-delimited string you want to match. These statements were introduced in Junos OS Release 11.2.

Related Documentation

- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- [Example: Configuring Access Privilege Levels on page 134](#)
- [*Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies*](#)
- [Understanding Junos OS Access Privilege Levels on page 123](#)

Examples

- [Example: Configuring Access Privilege Levels on page 134](#)
- [Example: Configuring Access Privileges for Operational Mode Commands on page 134](#)
- [Example: Specifying Access Privileges Using allow/deny-configuration-regexps Statements on page 135](#)

Example: Configuring Access Privilege Levels

Create two access privilege classes on the router or switch, one for configuring and viewing user accounts only and the second for configuring and viewing SNMP parameters only:

```
[edit]
system {
  login {
    class user-accounts {
      permissions [ configure admin admin-control ];
    }
    class network-mgmt {
      permissions [ configure snmp snmp-control ];
    }
  }
}
```

Related Documentation

- [Configuring Access Privilege Levels on page 131](#)

Example: Configuring Access Privileges for Operational Mode Commands

The following example shows how to configure access privileges for different login classes for individual operational mode commands:

```
[edit]
```

```

system {
  # This login class has operator privileges and the additional ability
  # to reboot the router.
  login {
    # This login class has operator privileges and the additional ability to reboot the
    # router or switch.
    class operator-and-boot {
      permissions [ clear network reset trace view ];
      allow-commands "request system reboot";
    }
    # This login class has operator privileges but can't use any commands beginning
    # with "set".
    # This login class has operator privileges
    # but cannot use any commands beginning with "set"
    class operator-no-set {
      permissions [ clear network reset trace view ];
      deny-commands "^set";
    }
    # This login class has operator privileges and can install software but not view
    # BGP information, and can issue the show route command, without specifying
    # commands or arguments under it.
    class operator-and-install-but-no-bgp {
      permissions [ clear network reset trace view ];
      allow-commands "(request system software add)|(show route$)";
      deny-commands "show bgp";
    }
  }
}

```

**Related
Documentation**

- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)

Example: Specifying Access Privileges Using allow/deny-configuration-regexps Statements

This example shows how to set up configuration access privileges using the **allow-configuration-regexps** and **deny-configuration-regexps** statements.

- [Requirements on page 135](#)
- [Overview on page 136](#)
- [Configuration on page 136](#)
- [Examples on page 136](#)

Requirements

This example uses the following hardware and software components:

- One Juniper Networks J Series, M Series, MX Series, or T Series device
- Junos OS Release 11.2 or later
 - There must be at least one user assigned to a login class.
 - There can be more than one login class, each with varying permission configurations, and more than one user on the device.

Overview

The **allow-configuration-regexps** and **deny-configuration-regexps** statements let you explicitly allow or deny users assigned to named user classes access privileges to portions of the configuration hierarchy, giving the system administrator precision control over who can change specific configurations in the system.



NOTE: The statements **allow-configuration-regexps** and **deny-configuration-regexps** perform similar functions as the statements **allow-configuration** and **deny-configuration**, except you can configure sets of strings in which the strings include spaces when using the first set of statements. You cannot use the two kinds of statements together.

Configuration

To set up configuration access privileges:

1. To explicitly allow one or more individual configuration mode hierarchies that would otherwise be denied, include the **allow-configuration-regexps** statement at the **[edit system login class class-name]** hierarchy level, configured with the regular expressions to be allowed.

```
[edit system login class class-name]
user@host# set allow-configuration-regexps "regular expression 1" "regular expression
2" "regular expression 3" "regular expression 4" ...
```

2. To explicitly deny one or more individual configuration hierarchies that would otherwise be allowed, include the **deny-configuration-regexps** statement at the **[edit system login class class-name]** hierarchy level, configured with the regular expressions to be denied.

```
[edit system login class class-name]
user@host# set deny-configuration-regexps "regular expression 1" "regular-expression
2" "regular expression 3" "regular expression 4"...
```

3. Assign the login class to one or more users.

```
[edit system login]
user@host# set user username class class-name
```

4. Commit your changes.

Users assigned this login class have the permissions you have set for the class.

Examples

Using Allow or Deny Configurations with Regular Expressions

Purpose This section provides examples of access privilege configurations to give you ideas for creating configurations appropriate for your system. You can use combinations of privilege statements for configuration access and for operational mode commands to give precise control over classes of access privileges.

Allow Configuration Changes The following example login class lets the user make changes at the **[edit system services]** hierarchy level and issue configuration mode commands (such as **commit**), in addition to the permissions specified by the **configure** permissions flag, which allows the user to enter configuration mode using the **configure** command.

```
[edit system login class class-name]
user@host# set permissions configure view view-configuration
user@host# set allow-configuration-regexps "system services"
```

Deny Configuration Changes The following example login class lets the user perform all operations allowed by the **all** permissions flag. However, it denies modifying the configuration at the **[edit system services]** hierarchy level.

```
[edit system login class class-name]
user@host# set permissions all configure view view-configuration
user@host# set deny-configuration-regexps "system services"
```

If the following statement is included in the configuration and the user's login class permission bit is set to **all**, the user cannot configure telnet parameters:

```
[edit system login class class-name]
user@host# set deny-configuration "system services telnet"
```

If the following statement is included in the configuration and the user's login class permission bit is set to **all**, the user cannot issue login class commands within any login class whose name begins with "m":

```
[edit system login class class-name]
user@host# set deny-configuration "system login class m.*"
```

If the following statement is included in the configuration and the user's login class permission bit is set to **all**, the user cannot edit the configuration or issue commands (such as **commit**) at the **[edit system login class]** or the **[edit system services]** hierarchy levels:

```
[edit system login class class-name]
user@host# set deny-configuration "system login class" "system services"
```

Allow and Deny Configuration Changes The following example login class lets the user perform all operations allowed by the **all** permissions flag, and explicitly grants configuration access to **[system "interfaces .* unit .* family inet address .*" protocols]**. However, the user is denied configuration access to the SNMP hierarchy level.



NOTE: You can use the ***** wildcard character when denoting regular expressions. However, it must be used as a portion of a regular expression. You cannot use **[*]** or **[.*]** alone.

```
[edit system login class class-name]
user@host# set permissions all configure view view-configuration
user@host# set allow-configuration-regexps system "interfaces .* unit .* family inet
address .*" protocols
user@host# set deny-configuration-regexps snmp
```

Allow and Deny Multiple Configuration Changes

The following example login class lets the user perform all operations allowed by the **all** permissions flag, and explicitly grants configuration access to multiple hierarchy levels for interfaces. It denies configuration access to the **[edit system]** and **[edit protocols]** hierarchy levels.



NOTE: You can configure as many regular expressions as needed to be allowed or denied. Regular expressions to be denied take precedence over configurations to be allowed.

```
[edit system login class class-name]
user@host# set permissions all configure view view-configuration
user@host# set allow-configuration-regexps "interfaces .* description .*" "interfaces .*
unit .* description .*" "interfaces .* unit .* family inet address .*" "interfaces .* disable"
user@host# set deny-configuration-regexps "system" "protocols"
```

Allow Configuration Changes and Deny Operations Commands

You can combine allow and deny configuration statements with allow and deny operational commands statements to fine-tune access privileges. The following example login class uses a combination of the **deny-commands** operational permissions statement and the **allow-configuration-regexps** configuration permissions statement to let the user configure and commit changes to the OSPF and BGP protocols. However, this class of user cannot issue the **show system statistics** or the **show bgp summary** commands.

```
[edit system login class class-name]
user@host# set permissions all configure view view-configuration
user@host# set deny-commands "(show system statistics)|(show bgp summary)"
user@host# set allow-configuration-regexps "protocols ospf|bgp"
```

The following shows permissions set for individual configuration mode hierarchies:

```
[edit]
system {
  login { # This login class has operator privileges and the additional ability to edit
    # configuration at the system services hierarchy level.
    class only-system-services {
      permissions [ configure ];
      allow-configuration "system services";
    }
    # services commands.
    class all-except-system-services { # This login class has operator privileges but
      # cannot edit any system services configuration.
      permissions [ all ];
      deny-configuration "system services";
    }
  }
}
```


Verification To verify that you have set the access privileges correctly:

1. Configure a login class and commit the changes.
2. Assign the login class to a *username*.
3. Log in as the *username* assigned with the new login class.
4. Attempt to perform the configurations that have been allowed or denied.
 - You should be able to perform configuration changes to hierarchy levels and regular expressions that have been allowed.
 - You should not be able to perform configuration changes to hierarchy levels and regular expressions that have been denied.
 - Denied expressions should take precedence over allowed expressions.
 - Any allowed or denied expressions should take precedence over any permissions granted with the **permissions** statement.

- Related Documentation**
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
 - [Example: Configuring Access Privilege Levels on page 134](#)
 - [Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies](#)
 - [Understanding Junos OS Access Privilege Levels on page 123](#)

User Permission Flags Reference

- [access on page 140](#)
- [access-control on page 141](#)
- [admin on page 142](#)
- [admin-control on page 142](#)
- [all-control on page 143](#)
- [clear on page 143](#)
- [configure on page 180](#)
- [control on page 181](#)
- [field on page 181](#)
- [firewall on page 181](#)
- [firewall-control on page 182](#)
- [floppy on page 183](#)
- [flow-tap on page 183](#)
- [flow-tap-control on page 184](#)
- [flow-tap-operation on page 184](#)
- [idp-profiler-operation on page 184](#)
- [interface on page 185](#)
- [interface-control on page 185](#)

- [maintenance on page 186](#)
- [network on page 193](#)
- [pgcp-session-mirroring on page 194](#)
- [pgcp-session-mirroring-control on page 195](#)
- [reset on page 195](#)
- [rollback on page 196](#)
- [routing on page 196](#)
- [routing-control on page 201](#)
- [secret on page 205](#)
- [secret-control on page 206](#)
- [security on page 207](#)
- [security-control on page 210](#)
- [shell on page 214](#)
- [snmp on page 214](#)
- [snmp-control on page 214](#)
- [system on page 215](#)
- [system-control on page 216](#)
- [trace on page 218](#)
- [trace-control on page 223](#)
- [view on page 228](#)
- [view-configuration on page 292](#)

access

Can view the access configuration in configuration mode.

Commands No associated CLI commands.

Configuration Hierarchy Levels

```
[edit access]
[edit access ppp-options]
[edit dynamic-profile]
[edit logical-systems access]
[edit logical-systems routing-instances instance system services static-subscribers
access-profile]
[edit logical-systems routing-instances instance system services static-subscribers
dynamic-profile]
[edit logical-systems routing-instances instance system services static-subscribers group
access-profile]
[edit logical-systems routing-instances instance system services static-subscribers group
dynamic-profile]
[edit logical-systems system services static-subscribers access-profile]
[edit logical-systems system services static-subscribers dynamic-profile]
[edit logical-systems system services static-subscribers group access-profile]
[edit logical-systems system services static-subscribers group dynamic-profile]
[edit routing-instances instance system services static-subscribers access-profile]
```

[edit routing-instances instance system services static-subscribers dynamic-profile]
 [edit routing-instances instance system services static-subscribers group access-profile]
 [edit routing-instances instance system services static-subscribers group dynamic-profile]
 [edit system services static-subscribers access-profile]
 [edit system services static-subscribers dynamic-profile]
 [edit system services static-subscribers group access-profile]
 [edit system services static-subscribers group dynamic-profile]

**Related
Documentation**

- [Access Privilege User Permission Flags Overview on page 128](#)
- [Understanding Junos OS Access Privilege Levels on page 123](#)
- [Configuring Access Privilege Levels on page 131](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)
- [access-control on page 141](#)

access-control

Can view access configuration information. Can edit access configuration at the **[edit access]**, **[edit logical-systems]**, **[edit routing-instances]**, and **[edit system services]** hierarchy levels.

**Configuration
Hierarchy Levels**

[edit access]
 [edit access ppp-options]
 [edit dynamic-profile]
 [edit logical-systems access]
 [edit logical-systems routing-instances instance system services static-subscribers access-profile]
 [edit logical-systems routing-instances instance system services static-subscribers dynamic-profile]
 [edit logical-systems routing-instances instance system services static-subscribers group access-profile]
 [edit logical-systems routing-instances instance system services static-subscribers group dynamic-profile]
 [edit logical-systems system services static-subscribers access-profile]
 [edit logical-systems system services static-subscribers dynamic-profile]
 [edit logical-systems system services static-subscribers group access-profile]
 [edit logical-systems system services static-subscribers group dynamic-profile]
 [edit routing-instances instance system services static-subscribers dynamic-profile]
 [edit routing-instances instance system services static-subscribers group access-profile]
 [edit routing-instances instance system services static-subscribers group dynamic-profile]
 [edit system services static-subscribers access-profile]
 [edit system services static-subscribers dynamic-profile]
 [edit system services static-subscribers group access-profile]
 [edit system services static-subscribers group dynamic-profile]

**Related
Documentation**

- [Access Privilege User Permission Flags Overview on page 128](#)
- [Understanding Junos OS Access Privilege Levels on page 123](#)

- [Configuring Access Privilege Levels on page 131](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)
- [access on page 140](#)

admin

Can view user account information in configuration mode.

Commands	show system audit
Configuration Hierarchy Levels	[edit protocols uplink-failure-detection] [edit system] [edit system accounting] [edit system diag-port-authentication] [edit system extensions] [edit system login] [edit system pic-console-authentication] [edit system root-authentication] [edit system services ssh ciphers] [edit system services ssh client-alive-count-max] [edit system services ssh client-alive-interval]] [edit system services ssh hostkey-algorithm] [edit system services ssh key-exchange] [edit system services ssh macs] [edit system services ssh max-sessions-per-connection] [edit system services ssh no-tcp-forwarding] [edit system services ssh protocol-version] [edit system services ssh root-login] [edit system services ssh tcp-forwarding]
Related Documentation	<ul style="list-style-type: none">• Access Privilege User Permission Flags Overview on page 128• Understanding Junos OS Access Privilege Levels on page 123• Configuring Access Privilege Levels on page 131• Specifying Access Privileges for Junos OS Operational Mode Commands on page 132• Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130• admin-control on page 142

admin-control

Can view user account information and configure it at the **[edit system]** hierarchy level.

Commands	show system audit
Configuration Hierarchy Levels	[edit protocols uplink-failure-detection] [edit system] [edit system accounting]

```
[edit system diag-port-authentication]
[edit system extensions]
[edit system login]
[edit system pic-console-authentication]
[edit system root-authentication]
[edit system services ssh ciphers]
[edit system services ssh hostkey-algorithm]
[edit system services ssh key-exchange]
[edit system services ssh macs]
[edit system services ssh protocol-version]
[edit system services ssh root-login]
```

- Related Documentation**
- [Access Privilege User Permission Flags Overview on page 128](#)
 - [Understanding Junos OS Access Privilege Levels on page 123](#)
 - [Configuring Access Privilege Levels on page 131](#)
 - [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
 - [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)
 - [admin on page 142](#)

all-control

Can access all operational mode commands and configuration mode commands. Can modify configuration in all the configuration hierarchy levels.

Commands All CLI commands.

Configuration Hierarchy Levels All CLI configuration hierarchy levels and statements.

- Related Documentation**
- [Access Privilege User Permission Flags Overview on page 128](#)
 - [Understanding Junos OS Access Privilege Levels on page 123](#)
 - [Configuring Access Privilege Levels on page 131](#)
 - [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
 - [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)

clear

Can clear (delete) information learned from the network that is stored in various network databases.

Commands

```
clear
clear amt
clear amt statistics
<clear-amt-statistics>
clear amt tunnel
clear-amt-tunnel
clear amt tunnel gateway-address
```

```
<clear amt tunnel gateway-address>
clear amt tunnel statistics
<clear-amt-tunnel-statistics>
clear amt tunnel statistics gateway-address
<clear-amt-tunnel-gateway-address-statistics>
clear amt tunnel statistics tunnel-interface
<clear-amt-tunnel-interface-statistics>
clear amt tunnel tunnel-interface
<clear-amt-tunnel-interface<>
clear ancp
clear ancp neighbor
<clear-ancp-neighbor-connection>
clear ancp subscriber
<clear-ancp-subscriber-connection>
clear arp
<clear-arp-table>
clear auto-configuration
clear auto-configuration interfaces
<clear-auto-configuration-interfaces>
clear bfd
clear bfd adaptation
<clear-bfd-adaptation-information>
clear bfd adaptation address
<clear-bfd-adaptation-address>
clear bfd adaptation discriminator
<clear-bfd-adaptation-discriminator>
clear bfd session
<clear-bfd-session-information>
clear bfd session address
<clear-bfd-session-address>
clear bfd session discriminator
<clear-bfd-session-discriminator>
clear bgp
clear bgp damping
<clear-bgp-damping>
clear bgp neighbor
<clear-bgp-neighbor>
clear bgp table
<clear-bgp-table>
clear bridge
clear bridge mac-table
<clear-bridge-mac-table>
clear bridge mac-table interface
<clear-bridge-interface-mac-table>
clear appqos-counter
clear appqos-rate-limiter-statistics
clear appqos-rule-statistics
clear cli
clear cli logical-system
<clear-cli-logical-system>
clear database-replication
clear database-replication statistics
<clear-database-replication-statistics-information>
clear ddos-protection
clear ddos-protection protocols
clear ddos-protection protocols amtv4
```

```

clear ddos-protection protocols amtv4 aggregate
clear ddos-protection protocols amtv4 aggregate culprit-flows
clear ddos-protection protocols amtv4 aggregate states
clear ddos-protection protocols amtv4 aggregate statistics
clear ddos-protection protocols amtv4 culprit-flows
clear ddos-protection protocols amtv4 states
clear ddos-protection protocols amtv4 statistics
clear ddos-protection protocols amtv6
clear ddos-protection protocols amtv6 aggregate
clear ddos-protection protocols amtv6 aggregate culprit-flows
<clear-ddos-amtv6-aggregate-flows>
clear ddos-protection protocols amtv6 aggregate states
<clear-ddos-amtv6-aggregate-states>
clear ddos-protection protocols amtv6 aggregate statistics
<clear-ddos-amtv6-aggregate-statistics>
clear ddos-protection protocols amtv6 culprit-flows
<clear-ddos-amtv6-flows>
clear ddos-protection protocols amtv6 states
<clear-ddos-amtv6-states>
clear ddos-protection protocols amtv6 statistics
<clear-ddos-amtv6-statistics>
clear ddos-protection protocols ancp aggregate culprit-flows
<clear-ddos-ancp-aggregate-flows>
clear ddos-protection protocols ancp culprit-flows

clear ddos-protection protocols ancp
clear ddos-protection protocols ancp aggregate
clear ddos-protection protocols ancp aggregate states
clear ddos-protection protocols ancp aggregate statistics
<clear-ddos-ancp-aggregate-statistics>
clear ddos-protection protocols ancp states
<clear-ddos-ancp-states>
clear ddos-protection protocols ancp statistics
<clear-ddos-ancp-statistics>
clear ddos-protection protocols ancpv6
clear ddos-protection protocols ancpv6 aggregate
clear ddos-protection protocols ancpv6 aggregate states

clear ddos-protection protocols ancpv6 aggregate culprit-flows
clear ddos-protection protocols arp aggregate statistics
clear-ddos-arp-aggregate-statistics
clear ddos-protection protocols arp aggregate culprit-flows
clear ddos-protection protocols arp states
clear-ddos-arp-states
clear ddos-protection protocols arp statistics
<clear-ddos-arp-statistics>
clear ddos-protection protocols arp culprit-flows
clear ddos-protection protocols atm
clear ddos-protection protocols atm aggregate
clear ddos-protection protocols atm aggregate culprit-flows
clear ddos-protection protocols atm aggregate states
<clear-ddos-atm-aggregate-states>
clear ddos-protection protocols atm aggregate statistics
<clear-ddos-atm-aggregate-statistics>
clear ddos-protection protocols atm culprit-flows
clear ddos-protection protocols bfd aggregate culprit-flows

```

```
clear ddos-protection protocols atm states
clear-ddos-atm-states
clear ddos-protection protocols atm statistics
clear-ddos-atm-statistics
clear ddos-protection protocols bfd
clear ddos-protection protocols bfd aggregate
clear ddos-protection protocols bfd culprit-flows
clear ddos-protection protocols bfd aggregate states
clear-ddos-bfd-aggregate-states
clear ddos-protection protocols bfd aggregate statistics
clear-ddos-bfd-aggregate-statistics
clear ddos-protection protocols bfd states
clear-ddos-bfd-states
clear ddos-protection protocols bfd statistics
clear-ddos-bfd-statistics
clear ddos-protection protocols bfdv6
clear ddos-protection protocols bfdv6 aggregate
clear ddos-protection protocols bfdv6 culprit-flows
clear ddos-protection protocols bfdv6 aggregate states
clear-ddos-bfdv6-aggregate-states
clear ddos-protection protocols bfdv6 aggregate statistics
clear-ddos-bfdv6-aggregate-statistics
clear ddos-protection protocols bfdv6 states
clear-ddos-bfdv6-states
clear ddos-protection protocols bfdv6 statistics
clear-ddos-bfdv6-statistics
clear ddos-protection protocols bgp
clear ddos-protection protocols bgp aggregate
clear ddos-protection protocols bgp aggregate culprit-flows
clear ddos-protection protocols bgp aggregate states
clear-ddos-bgp-aggregate-states
clear ddos-protection protocols bgp aggregate statistics
clear-ddos-bgp-aggregate-statistics
clear ddos-protection protocols bgp culprit-flows
clear-ddos-bgp-aggregate-statistics
clear ddos-protection protocols bgp states
clear-ddos-bgp-states
clear ddos-protection protocols bgp statistics
clear-ddos-bgp-statistics
clear ddos-protection protocols bgpv6
clear ddos-protection protocols bgpv6 aggregate
clear ddos-protection protocols bgpv6 aggregate culprit-flows
clear ddos-protection protocols bgpv6 aggregate states
clear-ddos-bgpv6-aggregate-states
clear ddos-protection protocols bgpv6 aggregate statistics
clear-ddos-bgpv6-aggregate-statistics
clear ddos-protection protocols bgpv6 culprit-flows
clear ddos-protection protocols bgpv6 states
clear-ddos-bgpv6-states
clear ddos-protection protocols bgpv6 statistics
<clear-ddos-bgpv6-statistics>
clear ddos-protection protocols culprit-flows
clear ddos-protection protocols demux-autosense
clear ddos-protection protocols demux-autosense aggregate
clear ddos-protection protocols demux-autosense aggregate culprit-flows
clear ddos-protection protocols demux-autosense culprit-flows
clear ddos-protection protocols demux-autosense aggregate states
```



```
clear-ddos-demuxauto-aggregate-states
clear ddos-protection protocols demux-autosense aggregate statistics
clear-ddos-demuxauto-aggregate-statistics
clear ddos-protection protocols demux-autosense states
clear-ddos-demuxauto-states
clear ddos-protection protocols demux-autosense statistics
clear-ddos-demuxauto-statistics
clear ddos-protection protocols dhcpv4
clear ddos-protection protocols dhcpv4 ack
clear ddos-protection protocols dhcpv4 ack culprit-flows
clear ddos-protection protocols dhcpv4 ack states
clear-ddos-dhcpv4-ack-states
clear ddos-protection protocols dhcpv4 ack statistics
clear-ddos-dhcpv4-ack-statistics
clear ddos-protection protocols dhcpv4 aggregate
clear ddos-protection protocols dhcpv4 aggregate culprit-flows
clear ddos-protection protocols dhcpv4 aggregate states
clear-ddos-dhcpv4-aggregate-states
clear ddos-protection protocols dhcpv4 aggregate statistics
clear-ddos-dhcpv4-aggregate-statistics
clear ddos-protection protocols dhcpv4 bad-packets
clear ddos-protection protocols dhcpv4 bad-packets culprit-flows
clear ddos-protection protocols dhcpv4 bad-packets states
clear-ddos-dhcpv4-bad-pack-states
clear ddos-protection protocols dhcpv4 bad-packets statistics
clear-ddos-dhcpv4-bad-pack-statistics
clear ddos-protection protocols dhcpv4 bootp
clear ddos-protection protocols dhcpv4 bootp culprit-flows
clear ddos-protection protocols dhcpv4 bootp states
clear-ddos-dhcpv4-bootp-states
clear ddos-protection protocols dhcpv4 bootp statistics
clear ddos-protection protocols dhcpv4 culprit-flows
<clear-ddos-dhcpv4-bootp-statistics>
clear ddos-protection protocols dhcpv4 culprit-flows
clear ddos-protection protocols dhcpv4 decline
clear ddos-protection protocols dhcpv4 decline culprit-flows
clear ddos-protection protocols dhcpv4 decline states
clear-ddos-dhcpv4-decline-states
clear ddos-protection protocols dhcpv4 decline statistics
clear-ddos-dhcpv4-decline-statistics
clear ddos-protection protocols dhcpv4 discover
clear ddos-protection protocols dhcpv4 discover states
clear-ddos-dhcpv4-discover-states
clear ddos-protection protocols dhcpv4 discover statistics
clear-ddos-dhcpv4-discover-statistics
clear ddos-protection protocols dhcpv4 force-renew
clear ddos-protection protocols dhcpv4 force-renew culprit-flows
clear ddos-protection protocols dhcpv4 force-renew states
clear-ddos-dhcpv4-forcerenew-states
clear ddos-protection protocols dhcpv4 force-renew statistics
clear-ddos-dhcpv4-forcerenew-statistics
clear ddos-protection protocols dhcpv4 inform
clear ddos-protection protocols dhcpv4 inform culprit-flows
clear ddos-protection protocols dhcpv4 inform states
clear-ddos-dhcpv4-inform-states
clear ddos-protection protocols dhcpv4 inform statistics
```

```
clear-ddos-dhcpv4-inform-statistics
clear ddos-protection protocols dhcpv4 lease-active
clear ddos-protection protocols dhcpv4 lease-active culprit-flows
clear ddos-protection protocols dhcpv4 lease-active states
clear-ddos-dhcpv4-leaseact-states
clear ddos-protection protocols dhcpv4 lease-active statistics
clear-ddos-dhcpv4-leaseact-statistics
clear ddos-protection protocols dhcpv4 lease-query
clear ddos-protection protocols dhcpv4 lease-query culprit-flows
clear ddos-protection protocols dhcpv4 lease-query states
clear-ddos-dhcpv4-leasequery-states
clear ddos-protection protocols dhcpv4 lease-query statistics
clear-ddos-dhcpv4-leasequery-statistics
clear ddos-protection protocols dhcpv4 lease-unassigned
clear ddos-protection protocols dhcpv4 lease-unassigned culprit-flows
clear ddos-protection protocols dhcpv4 lease-unassigned states
clear-ddos-dhcpv4-leaseuna-states
clear ddos-protection protocols dhcpv4 lease-unassigned statistics
clear-ddos-dhcpv4-leaseuna-statistics
clear ddos-protection protocols dhcpv4 lease-unknown
clear ddos-protection protocols dhcpv4 lease-unknown culprit-flows
clear ddos-protection protocols dhcpv4 lease-unknown states
clear-ddos-dhcpv4-leaseunk-states
clear ddos-protection protocols dhcpv4 lease-unknown statistics
clear-ddos-dhcpv4-leaseunk-statistics
clear ddos-protection protocols dhcpv4 nak
clear ddos-protection protocols dhcpv4 nak culprit-flows
clear ddos-protection protocols dhcpv4 nak states
clear-ddos-dhcpv4-nak-states
clear ddos-protection protocols dhcpv4 nak statistics
clear-ddos-dhcpv4-nak-statistics
clear ddos-protection protocols dhcpv4 no-message-type
clear ddos-protection protocols dhcpv4 no-message-type culprit-flows
clear ddos-protection protocols dhcpv4 no-message-type states
clear-ddos-dhcpv4-no-msgtype-states
clear ddos-protection protocols dhcpv4 no-message-type statistics
clear-ddos-dhcpv4-no-msgtype-statistics
clear ddos-protection protocols dhcpv4 offer
clear ddos-protection protocols dhcpv4 offer culprit-flows
clear ddos-protection protocols dhcpv4 offer states
clear-ddos-dhcpv4-offer-states
clear ddos-protection protocols dhcpv4 offer statistics
clear-ddos-dhcpv4-offer-statistics
clear ddos-protection protocols dhcpv4 release
clear ddos-protection protocols dhcpv4 release culprit-flows
clear ddos-protection protocols dhcpv4 release states
clear-ddos-dhcpv4-release-states
clear ddos-protection protocols dhcpv4 release statistics
clear-ddos-dhcpv4-release-statistics
clear ddos-protection protocols dhcpv4 renew
clear ddos-protection protocols dhcpv4 renew culprit-flows
clear ddos-protection protocols dhcpv4 renew states
clear-ddos-dhcpv4-renew-states
clear ddos-protection protocols dhcpv4 renew statistics
clear-ddos-dhcpv4-renew-statistics
clear ddos-protection protocols dhcpv4 request
```

```
clear ddos-protection protocols dhcpv4 request culprit-flows
clear ddos-protection protocols dhcpv4 request states
clear-ddos-dhcpv4-request-states
clear ddos-protection protocols dhcpv4 request statistics
clear-ddos-dhcpv4-request-statistics
clear ddos-protection protocols dhcpv4 states
clear-ddos-dhcpv4-states
clear ddos-protection protocols dhcpv4 statistics
clear-ddos-dhcpv4-statistics
clear ddos-protection protocols dhcpv4 unclassified
clear ddos-protection protocols dhcpv4 unclassified culprit-flows
clear ddos-protection protocols dhcpv4 unclassified states
clear-ddos-dhcpv4-unclass-states
clear ddos-protection protocols dhcpv4 unclassified statistics
clear-ddos-dhcpv4-unclass-statistics
clear ddos-protection protocols dhcpv6
clear ddos-protection protocols dhcpv6 advertise
clear ddos-protection protocols dhcpv6 advertise culprit-flows
clear ddos-protection protocols dhcpv6 advertise states
clear-ddos-dhcpv6-advertise-states
clear ddos-protection protocols dhcpv6 advertise statistics
clear-ddos-dhcpv6-advertise-statistics
clear ddos-protection protocols dhcpv6 aggregate
clear ddos-protection protocols dhcpv6 aggregate states
clear-ddos-dhcpv6-aggregate-states
clear ddos-protection protocols dhcpv6 aggregate statistics
clear-ddos-dhcpv6-aggregate-statistics
clear ddos-protection protocols dhcpv6 confirm
clear ddos-protection protocols dhcpv6 confirm culprit-flows
clear ddos-protection protocols dhcpv6 confirm states
clear-ddos-dhcpv6-confirm-states
clear ddos-protection protocols dhcpv6 confirm statistics
clear-ddos-dhcpv6-confirm-statistics
clear ddos-protection protocols dhcpv6 decline
clear ddos-protection protocols dhcpv6 decline states
clear-ddos-dhcpv6-decline-states
clear ddos-protection protocols dhcpv6 decline statistics
clear-ddos-dhcpv6-decline-statistics
clear ddos-protection protocols dhcpv6 information-request
clear ddos-protection protocols dhcpv6 information-request states
clear-ddos-dhcpv6-info-req-states
clear ddos-protection protocols dhcpv6 information-request statistics
clear-ddos-dhcpv6-info-req-statistics
clear ddos-protection protocols dhcpv6 leasequery
clear ddos-protection protocols dhcpv6 leasequery states
clear-ddos-dhcpv6-leasequery-states
clear ddos-protection protocols dhcpv6 leasequery statistics
clear-ddos-dhcpv6-leasequery-statistics
clear ddos-protection protocols dhcpv6 leasequery-data
clear ddos-protection protocols dhcpv6 leasequery-data states
clear-ddos-dhcpv6-leaseq-da-states
clear ddos-protection protocols dhcpv6 leasequery-data statistics
clear-ddos-dhcpv6-leaseq-da-statistics
clear ddos-protection protocols dhcpv6 leasequery-done
clear ddos-protection protocols dhcpv6 leasequery-done states
clear-ddos-dhcpv6-leaseq-do-states
```

```
clear ddos-protection protocols dhcpv6 leasequery-done statistics
clear-ddos-dhcpv6-leaseq-do-statistics
clear ddos-protection protocols dhcpv6 leasequery-reply
clear ddos-protection protocols dhcpv6 leasequery-reply states
clear-ddos-dhcpv6-leaseq-re-states
clear ddos-protection protocols dhcpv6 leasequery-reply statistics
clear-ddos-dhcpv6-leaseq-re-statistics
clear ddos-protection protocols dhcpv6 rebind
clear ddos-protection protocols dhcpv6 rebind states
clear-ddos-dhcpv6-rebind-states
clear ddos-protection protocols dhcpv6 rebind statistics
clear-ddos-dhcpv6-rebind-statistics
clear ddos-protection protocols dhcpv6 reconfigure
clear ddos-protection protocols dhcpv6 reconfigure states
clear-ddos-dhcpv6-reconfig-states
clear ddos-protection protocols dhcpv6 reconfigure statistics
clear-ddos-dhcpv6-reconfig-statistics
clear ddos-protection protocols dhcpv6 relay-forward
clear ddos-protection protocols dhcpv6 relay-forward states
clear-ddos-dhcpv6-relay-for-states
clear ddos-protection protocols dhcpv6 relay-forward statistics
clear-ddos-dhcpv6-relay-for-statistics
clear ddos-protection protocols dhcpv6 relay-reply
clear ddos-protection protocols dhcpv6 relay-reply states
clear-ddos-dhcpv6-relay-rep-states
clear ddos-protection protocols dhcpv6 relay-reply statistics
clear-ddos-dhcpv6-relay-rep-statistics
clear ddos-protection protocols dhcpv6 release
clear ddos-protection protocols dhcpv6 release states
clear-ddos-dhcpv6-release-states
clear ddos-protection protocols dhcpv6 release statistics
clear-ddos-dhcpv6-release-statistics
clear ddos-protection protocols dhcpv6 renew
clear ddos-protection protocols dhcpv6 renew states
clear-ddos-dhcpv6-renew-states
clear ddos-protection protocols dhcpv6 renew statistics
clear-ddos-dhcpv6-renew-statistics
clear ddos-protection protocols dhcpv6 reply
clear ddos-protection protocols dhcpv6 reply states
clear-ddos-dhcpv6-reply-states
clear ddos-protection protocols dhcpv6 reply statistics
clear-ddos-dhcpv6-reply-statistics
clear ddos-protection protocols dhcpv6 request
clear ddos-protection protocols dhcpv6 request culprit-flows
clear ddos-protection protocols dhcpv6 request states
clear-ddos-dhcpv6-request-states
clear ddos-protection protocols dhcpv6 request statistics
clear-ddos-dhcpv6-request-statistics
clear ddos-protection protocols dhcpv6 solicit
clear ddos-protection protocols dhcpv6 solicit culprit-flows
clear ddos-protection protocols dhcpv6 solicit states
clear-ddos-dhcpv6-solicit-states
clear ddos-protection protocols dhcpv6 solicit statistics
clear-ddos-dhcpv6-solicit-statistics
clear ddos-protection protocols dhcpv6 states
clear-ddos-dhcpv6-states
```

```
clear ddos-protection protocols dhcpv6 statistics
clear-ddos-dhcpv6-statistics
clear ddos-protection protocols dhcpv6 unclassified
clear ddos-protection protocols dhcpv6 unclassified culprit-flows
clear ddos-protection protocols dhcpv6 unclassified states
clear-ddos-dhcpv6-unclass-states
clear ddos-protection protocols dhcpv6 unclassified statistics
clear-ddos-dhcpv6-unclass-statistics
clear ddos-protection protocols diameter
clear ddos-protection protocols diameter aggregate
clear ddos-protection protocols diameter aggregate culprit-flows
clear ddos-protection protocols diameter aggregate states
clear-ddos-diameter-aggregate-states
clear ddos-protection protocols diameter aggregate statistics
<clear-ddos-diameter-aggregate-statistics>
clear ddos-protection protocols diameter culprit-flows
clear ddos-protection protocols diameter states
<clear-ddos-diameter-states>
clear ddos-protection protocols diameter statistics
clear-ddos-diameter-statistics
clear ddos-protection protocols dns
clear ddos-protection protocols dns aggregate
clear ddos-protection protocols dns aggregate states
clear-ddos-dns-aggregate-states
clear ddos-protection protocols dns aggregate statistics
clear-ddos-dns-aggregate-statistics
clear ddos-protection protocols dns states
clear-ddos-dns-states
clear ddos-protection protocols dns statistics
clear-ddos-dns-statistics
clear ddos-protection protocols dtcp
clear ddos-protection protocols dtcp aggregate
clear ddos-protection protocols dtcp aggregate culprit-flows
clear ddos-protection protocols dtcp aggregate states
clear-ddos-dtcp-aggregate-states
clear ddos-protection protocols dtcp aggregate statistics
clear ddos-protection protocols dtcp culprit-flows
clear-ddos-dtcp-aggregate-statistics
clear ddos-protection protocols dtcp states
clear-ddos-dtcp-states
clear ddos-protection protocols dtcp statistics
clear-ddos-dtcp-statistics
clear ddos-protection protocols dynamic-vlan
clear ddos-protection protocols dynamic-vlan aggregate
clear ddos-protection protocols dynamic-vlan aggregate culprit-flows
clear ddos-protection protocols dynamic-vlan aggregate states
clear-ddos-dynvlan-aggregate-states
clear ddos-protection protocols dynamic-vlan aggregate statistics
clear-ddos-dynvlan-aggregate-statistics
clear ddos-protection protocols dynamic-vlan states
clear-ddos-dynvlan-states
clear ddos-protection protocols dynamic-vlan statistics
clear-ddos-dynvlan-statistics
clear ddos-protection protocols egpv6
clear ddos-protection protocols egpv6 aggregate
clear ddos-protection protocols egpv6 aggregate culprit-flows
```

```
clear ddos-protection protocols egpv6 aggregate states
clear-ddos-egpv6-aggregate-states
clear ddos-protection protocols egpv6 aggregate statistics
clear-ddos-egpv6-aggregate-statistics
clear ddos-protection protocols egpv6 states
clear-ddos-egpv6-states
clear ddos-protection protocols egpv6 statistics
clear-ddos-egpv6-statistics
clear ddos-protection protocols eoam
clear ddos-protection protocols eoam aggregate
clear ddos-protection protocols eoam aggregate culprit-flows
clear ddos-protection protocols eoam aggregate states
clear-ddos-eoam-aggregate-states
clear ddos-protection protocols eoam aggregate statistics
clear-ddos-eoam-aggregate-statistics
clear ddos-protection protocols eoam states
clear-ddos-eoam-states
clear ddos-protection protocols eoam statistics
clear-ddos-eoam-statistics
clear ddos-protection protocols esmc
clear ddos-protection protocols esmc aggregate
clear ddos-protection protocols esmc aggregate culprit-flows
clear ddos-protection protocols esmc aggregate states
clear-ddos-esmc-aggregate-states
clear ddos-protection protocols esmc aggregate statistics
clear-ddos-esmc-aggregate-statistics
clear ddos-protection protocols esmc culprit-flows
clear ddos-protection protocols esmc states
clear-ddos-esmc-states
clear ddos-protection protocols esmc statistics
clear ddos-protection protocols fab-probe
clear ddos-protection protocols fab-probe aggregate
clear ddos-protection protocols fab-probe aggregate states
clear ddos-protection protocols fab-probe aggregate statistics
<clear-ddos-fab-probe-aggregate-statistics>
clear ddos-protection protocols fab-probe states
<clear-ddos-fab-probe-states>
clear ddos-protection protocols fab-probe statistics
<clear-ddos-fab-probe-statistics>
clear-ddos-esmc-statistics
clear ddos-protection protocols firewall-host
clear ddos-protection protocols firewall-host aggregate
clear ddos-protection protocols firewall-host aggregate culprit-flows
clear ddos-protection protocols firewall-host aggregate states
clear-ddos-fw-host-aggregate-states
clear ddos-protection protocols firewall-host aggregate statistics
clear-ddos-fw-host-aggregate-statistics
clear ddos-protection protocols firewall-host states
clear-ddos-fw-host-states
clear ddos-protection protocols firewall-host statistics
<clear-ddos-fw-host-statistics>
clear ddos-protection protocols frame-relay
clear ddos-protection protocols frame-relay aggregate
clear ddos-protection protocols frame-relay aggregate culprit-flows
clear ddos-protection protocols frame-relay aggregate states
clear ddos-protection protocols frame-relay aggregate statistics
```

```
clear ddos-protection protocols frame-relay culprit-flows
clear ddos-protection protocols frame-relay frf15
clear ddos-protection protocols frame-relay frf15 culprit-flows
clear ddos-protection protocols frame-relay frf15 states
clear ddos-protection protocols frame-relay frf15 statistics
clear ddos-protection protocols frame-relay frf16
clear ddos-protection protocols frame-relay frf16 culprit-flows
clear ddos-protection protocols frame-relay frf16 states
clear ddos-protection protocols frame-relay frf16 statistics
clear ddos-protection protocols frame-relay states
clear ddos-protection protocols frame-relay statistics
clear ddos-protection protocols ftp
clear ddos-protection protocols ftp aggregate
clear ddos-protection protocols ftp aggregate culprit-flows
clear ddos-protection protocols ftp aggregate states
clear-ddos-ftp-aggregate-states
clear ddos-protection protocols ftp aggregate statistics
clear-ddos-ftp-aggregate-statistics
clear ddos-protection protocols ftp states
<clear-ddos-ftp-states>
clear ddos-protection protocols ftp statistics
clear-ddos-ftp-statistics
clear ddos-protection protocols ftpv6
clear ddos-protection protocols ftpv6 aggregate
clear ddos-protection protocols ftpv6 aggregate culprit-flows
clear ddos-protection protocols ftpv6 aggregate states
clear-ddos-ftp6-aggregate-states
clear ddos-protection protocols ftpv6 aggregate statistics
clear-ddos-ftp6-aggregate-statistics
clear ddos-protection protocols ftpv6 culprit-flows
clear ddos-protection protocols ftpv6 states
clear-ddos-ftp6-states
clear ddos-protection protocols ftpv6 statistics
clear-ddos-ftp6-statistics
clear ddos-protection protocols gre
clear ddos-protection protocols gre culprit-flows
clear ddos-protection protocols gre aggregate
clear ddos-protection protocols gre aggregate culprit-flow
clear ddos-protection protocols gre aggregate states
clear-ddos-gre-aggregate-states
clear ddos-protection protocols gre aggregate statistics
clear-ddos-gre-aggregate-statistics
clear ddos-protection protocols gre states
clear-ddos-gre-states
clear ddos-protection protocols gre statistics
clear-ddos-gre-statistics
clear ddos-protection protocols icmp
clear ddos-protection protocols icmp aggregate
clear ddos-protection protocols icmp aggregate culprit-flows
clear ddos-protection protocols icmp aggregate states
clear-ddos-icmp-aggregate-states
clear ddos-protection protocols icmp aggregate statistics
<clear-ddos-icmp-aggregate-statistics>
clear ddos-protection protocols icmp culprit-flows
clear ddos-protection protocols icmp states
clear-ddos-icmp-states
```

```
clear ddos-protection protocols icmp statistics
clear-ddos-icmp-statistics
clear ddos-protection protocols icmpv6
clear ddos-protection protocols icmpv6 aggregate
clear ddos-protection protocols icmpv6 aggregate culprit-flows
clear ddos-protection protocols icmpv6 aggregate states
<clear-ddos-icmpv6-aggregate-states>
clear ddos-protection protocols icmpv6 aggregate statistics
<clear-ddos-icmpv6-aggregate-statistics>
clear ddos-protection protocols icmpv6 states
<clear-ddos-icmpv6-states>
clear ddos-protection protocols icmpv6 statistics
<clear-ddos-icmpv6-statistics>
clear ddos-protection protocols igmp
clear ddos-protection protocols igmp aggregate
clear ddos-protection protocols igmp aggregate culprit-flows
clear ddos-protection protocols igmp aggregate states
clear-ddos-igmp-aggregate-states
clear ddos-protection protocols igmp aggregate statistics
clear-ddos-igmp-aggregate-statistics
clear ddos-protection protocols igmp states
clear-ddos-igmp-states
clear ddos-protection protocols igmp statistics
clear-ddos-igmp-statistics
clear ddos-protection protocols igmp-snoop
clear ddos-protection protocols igmp-snoop aggregate
clear ddos-protection protocols igmp-snoop aggregate culprit-flows
clear ddos-protection protocols igmp-snoop aggregate states
clear-ddos-igmp-snoop-aggregate-states
clear ddos-protection protocols igmp-snoop aggregate statistics
clear-ddos-igmp-snoop-aggregate-statistics
clear ddos-protection protocols igmp-snoop states
clear-ddos-igmp-snoop-states
clear ddos-protection protocols igmp-snoop statistics
clear-ddos-igmp-snoop-statistics
clear ddos-protection protocols igmpv4v6
clear ddos-protection protocols igmpv4v6 culprit-flows
clear ddos-protection protocols igmpv4v6 aggregate
clear ddos-protection protocols igmpv4v6 aggregate states
clear-ddos-igmpv4v6-aggregate-states
clear ddos-protection protocols igmpv4v6 aggregate statistics
clear-ddos-igmpv4v6-aggregate-statistics
clear ddos-protection protocols igmpv4v6 states
clear-ddos-igmpv4v6-states
clear ddos-protection protocols igmpv4v6 statistics
clear-ddos-igmpv4v6-statistics
clear ddos-protection protocols igmpv6
clear ddos-protection protocols igmpv6 aggregate
clear ddos-protection protocols igmpv6 aggregate culprit-flows
clear ddos-protection protocols igmpv6 aggregate states
clear-ddos-igmpv6-aggregate-states
clear ddos-protection protocols igmpv6 aggregate statistics
clear-ddos-igmpv6-aggregate-statistics
clear ddos-protection protocols igmpv6 states
clear-ddos-igmpv6-states
clear ddos-protection protocols igmpv6 statistics
```



```
<clear-ddos-igmpv6-statistics>
clear ddos-protection protocols inline-ka
clear ddos-protection protocols inline-ka aggregate
clear ddos-protection protocols inline-ka aggregate culprit-flows
clear ddos-protection protocols inline-ka aggregate states
clear ddos-protection protocols inline-ka aggregate statistics
clear ddos-protection protocols inline-ka culprit-flows
clear ddos-protection protocols inline-ka states
clear ddos-protection protocols inline-ka statistics
clear ddos-protection protocols inline-svcs
clear ddos-protection protocols inline-svcs aggregate
clear ddos-protection protocols inline-svcs aggregate culprit-flows
clear ddos-protection protocols inline-svcs aggregate states
clear ddos-protection protocols inline-svcs aggregate statistics
clear ddos-protection protocols inline-svcs culprit-flows
clear ddos-protection protocols inline-svcs states
clear ddos-protection protocols inline-svcs statistics
clear ddos-protection protocols ip-fragments
clear ddos-protection protocols ip-fragments aggregate
clear ddos-protection protocols ip-fragments culprit-flows
clear ddos-protection protocols ip-fragments aggregate states
clear-ddos-ip-frag-aggregate-states
clear ddos-protection protocols ip-fragments aggregate statistics
clear-ddos-ip-frag-aggregate-statistics
clear ddos-protection protocols ip-fragments first-fragment
clear ddos-protection protocols ip-fragments first-fragment states
clear-ddos-ip-frag-first-frag-states
clear ddos-protection protocols ip-fragments first-fragment statistics
clear-ddos-ip-frag-first-frag-statistics
clear ddos-protection protocols ip-fragments states
clear-ddos-ip-frag-states
clear ddos-protection protocols ip-fragments statistics
clear-ddos-ip-frag-statistics
clear ddos-protection protocols ip-fragments trail-fragment
clear ddos-protection protocols ip-fragments trail-fragment culprit-flows
clear ddos-protection protocols ip-fragments trail-fragment states
clear-ddos-ip-frag-trail-frag-states
clear ddos-protection protocols ip-fragments trail-fragment statistics
clear-ddos-ip-frag-trail-frag-statistics
clear ddos-protection protocols ip-options
clear ddos-protection protocols ip-options aggregate
clear ddos-protection protocols ip-options aggregate states
clear-ddos-ip-opt-aggregate-states
clear ddos-protection protocols ip-options aggregate statistics
clear-ddos-ip-opt-aggregate-statistics
clear ddos-protection protocols ip-options non-v4v6
clear ddos-protection protocols ip-options non-v4v6 states
<clear-ddos-ip-opt-non-v4v6-states>
clear ddos-protection protocols ip-options non-v4v6 statistics
<clear-ddos-ip-opt-non-v4v6-statistics>
clear ddos-protection protocols ip-options router-alert
clear ddos-protection protocols ip-options router-alert culprit-flows
clear ddos-protection protocols ip-options router-alert states
clear-ddos-ip-opt-rt-alert-states
clear ddos-protection protocols ip-options router-alert statistics
clear-ddos-ip-opt-rt-alert-statistics
```

```
clear ddos-protection protocols ip-options states
clear-ddos-ip-opt-states
clear ddos-protection protocols ip-options statistics
clear-ddos-ip-opt-statistics
clear ddos-protection protocols ip-options unclassified
clear ddos-protection protocols ip-options unclassified culprit-flows
clear ddos-protection protocols ip-options unclassified states
clear-ddos-ip-opt-unclass-states
clear ddos-protection protocols ip-options unclassified statistics
clear-ddos-ip-opt-unclass-statistics
clear ddos-protection protocols ipv4-unclassified
clear ddos-protection protocols ipv4-unclassified aggregate
clear ddos-protection protocols ipv4-unclassified culprit-flows
clear ddos-protection protocols ipv4-unclassified aggregate states
clear-ddos-ipv4-uncls-aggregate-states
clear ddos-protection protocols ipv4-unclassified aggregate statistics
clear-ddos-ipv4-uncls-aggregate-statistics
clear ddos-protection protocols ipv4-unclassified states
clear-ddos-ipv4-uncls-states
clear ddos-protection protocols ipv4-unclassified statistics
clear-ddos-ipv4-uncls-statistics
clear ddos-protection protocols ipv6-unclassified
clear ddos-protection protocols ipv6-unclassified aggregate
clear ddos-protection protocols ipv6-unclassified aggregate states
clear-ddos-ipv6-uncls-aggregate-states
clear ddos-protection protocols ipv6-unclassified aggregate statistics
clear-ddos-ipv6-uncls-aggregate-statistics
clear ddos-protection protocols ipv6-unclassified states
clear-ddos-ipv6-uncls-states
clear ddos-protection protocols ipv6-unclassified statistics
clear-ddos-ipv6-uncls-statistics
clear ddos-protection protocols isis
clear ddos-protection protocols isis aggregate
clear ddos-protection protocols isis aggregate culprit-flows
clear ddos-protection protocols isis aggregate states
clear-ddos-isis-aggregate-states
clear ddos-protection protocols isis aggregate statistics
<clear-ddos-isis-aggregate-statistics>
clear ddos-protection protocols isis culprit-flows
clear ddos-protection protocols isis states
clear-ddos-isis-states
clear ddos-protection protocols isis statistics
clear-ddos-isis-statistics
clear ddos-protection protocols jfm
clear ddos-protection protocols jfm aggregate
clear ddos-protection protocols jfm aggregate culprit-flows
clear ddos-protection protocols jfm aggregate states
clear-ddos-jfm-aggregate-states
clear ddos-protection protocols jfm aggregate statistics
clear-ddos-jfm-aggregate-statistics
clear ddos-protection protocols jfm states
clear-ddos-jfm-states
clear ddos-protection protocols jfm statistics
<clear-ddos-jfm-statistics>
clear ddos-protection protocols keepalive
clear ddos-protection protocols keepalive aggregate
```

```
clear ddos-protection protocols keepalive aggregate culprit-flows
clear ddos-protection protocols keepalive aggregate states
clear ddos-protection protocols keepalive aggregate statistics
clear ddos-protection protocols keepalive culprit-flows
clear ddos-protection protocols keepalive states
clear ddos-protection protocols keepalive statistics
clear ddos-protection protocols l2pt
clear ddos-protection protocols l2pt aggregate
clear ddos-protection protocols l2pt aggregate states
clear ddos-protection protocols l2pt aggregate statistics
clear ddos-protection protocols l2pt culprit-flows
clear ddos-protection protocols l2pt states
clear ddos-protection protocols l2pt statistics
clear ddos-protection protocols l2tp
clear ddos-protection protocols l2tp aggregate
clear ddos-protection protocols l2tp aggregate culprit-flows
clear ddos-protection protocols l2tp aggregate states
clear-ddos-l2tp-aggregate-states
clear ddos-protection protocols l2tp aggregate statistics
clear-ddos-l2tp-aggregate-statistics
clear ddos-protection protocols l2tp states
clear-ddos-l2tp-states
clear ddos-protection protocols l2tp statistics
clear-ddos-l2tp-statistics
clear ddos-protection protocols lacp
clear ddos-protection protocols lacp aggregate
clear ddos-protection protocols lacp aggregate culprit-flows
clear ddos-protection protocols lacp aggregate states
clear-ddos-lacp-aggregate-states
clear ddos-protection protocols lacp aggregate statistics
clear-ddos-lacp-aggregate-statistics
clear ddos-protection protocols lacp states
clear-ddos-lacp-states
clear ddos-protection protocols lacp statistics
clear-ddos-lacp-statistics
clear ddos-protection protocols ldp
clear ddos-protection protocols ldp aggregate
clear ddos-protection protocols ldp aggregate culprit-flows
clear ddos-protection protocols ldp aggregate states
clear-ddos-ldp-aggregate-states
clear ddos-protection protocols ldp aggregate statistics
<clear-ddos-ldp-aggregate-statistics>
clear ddos-protection protocols ldp culprit-flows
clear ddos-protection protocols ldp states
clear-ddos-ldp-states
clear ddos-protection protocols ldp statistics
clear-ddos-ldp-statistics
clear ddos-protection protocols ldpv6
clear ddos-protection protocols ldpv6 aggregate
clear ddos-protection protocols ldpv6 aggregate culprit-flows
clear ddos-protection protocols ldpv6 aggregate states
clear-ddos-ldpv6-aggregate-states
clear ddos-protection protocols ldpv6 aggregate statistics
clear-ddos-ldpv6-aggregate-statistics
clear ddos-protection protocols ldpv6 states
clear-ddos-ldpv6-states
```

```

clear ddos-protection protocols ldpv6 statistics
clear-ddos-ldpv6-statistics
clear ddos-protection protocols lldp
clear ddos-protection protocols lldp aggregate
clear ddos-protection protocols lldp aggregate culprit-flows

clear ddos-protection protocols lldp aggregate states
clear-ddos-lddp-aggregate-states
clear ddos-protection protocols lldp aggregate statistics
clear-ddos-lddp-aggregate-statistics
clear ddos-protection protocols lldp states
clear-ddos-lddp-states
clear ddos-protection protocols lldp statistics
clear-ddos-lddp-statistics
clear ddos-protection protocols lmp
clear ddos-protection protocols lmp aggregate
clear ddos-protection protocols lmp aggregate culprit-flows
clear ddos-protection protocols lmp aggregate states
clear-ddos-lmp-aggregate-states
clear ddos-protection protocols lmp aggregate statistics
clear-ddos-lmp-aggregate-statistics
clear ddos-protection protocols lmp states
clear-ddos-lmp-states
clear ddos-protection protocols lmp statistics
clear-ddos-lmp-statistics
clear ddos-protection protocols lmpv6
clear ddos-protection protocols lmpv6 aggregate
clear ddos-protection protocols lmpv6 aggregate culprit-flows
clear ddos-protection protocols lmpv6 aggregate states
clear-ddos-lmpv6-aggregate-states
clear ddos-protection protocols lmpv6 aggregate statistics
<clear-ddos-lmpv6-aggregate-statistics>
clear ddos-protection protocols lmpv6 culprit-flows
clear ddos-protection protocols lmpv6 states
clear-ddos-lmpv6-states
clear ddos-protection protocols lmpv6 statistics
clear-ddos-lmpv6-statistics
clear ddos-protection protocols mac-host
clear ddos-protection protocols mac-host aggregate
clear ddos-protection protocols mac-host aggregate culprit-flows
clear ddos-protection protocols mac-host aggregate states
clear-ddos-mac-host-aggregate-states
clear ddos-protection protocols mac-host aggregate statistics
clear-ddos-mac-host-aggregate-statistics
clear ddos-protection protocols mac-host states
clear-ddos-mac-host-states
clear ddos-protection protocols mac-host statistics
clear-ddos-mac-host-statistics
clear ddos-protection protocols mlp
clear ddos-protection protocols mlp aggregate
clear ddos-protection protocols mlp aggregate culprit-flows
clear ddos-protection protocols mlp aggregate states
clear-ddos-mlp-aggregate-states
clear ddos-protection protocols mlp aggregate statistics
clear-ddos-mlp-aggregate-statistics
clear ddos-protection protocols mlp aging-exception

```

```
clear ddos-protection protocols mlp aging-exception culprit-flows
clear ddos-protection protocols mlp aging-exception states
clear-ddos-mlp-aging-exc-states
clear ddos-protection protocols mlp aging-exception statistics
clear-ddos-mlp-aging-exc-statistics
clear ddos-protection protocols mlp packets
clear ddos-protection protocols mlp packets states
clear-ddos-mlp-packets-states
clear ddos-protection protocols mlp packets statistics
clear-ddos-mlp-packets-statistics
clear ddos-protection protocols mlp states
clear-ddos-mlp-states
clear ddos-protection protocols mlp statistics
clear-ddos-mlp-statistics
clear ddos-protection protocols mlp unclassified
clear ddos-protection protocols mlp unclassified states
clear-ddos-mlp-unclass-states
clear ddos-protection protocols mlp unclassified statistics
clear-ddos-mlp-unclass-statistics
clear ddos-protection protocols msdp
clear ddos-protection protocols msdp aggregate
clear ddos-protection protocols msdp culprit-flows
clear ddos-protection protocols msdp aggregate states
clear-ddos-msdp-aggregate-states
clear ddos-protection protocols msdp aggregate statistics
clear-ddos-msdp-aggregate-statistics
clear ddos-protection protocols msdp states
clear-ddos-msdp-states
clear ddos-protection protocols msdp statistics
clear-ddos-msdp-statistics
clear ddos-protection protocols msdpv6
clear ddos-protection protocols msdpv6 aggregate
clear ddos-protection protocols msdpv6 aggregate culprit-flows
clear ddos-protection protocols msdpv6 aggregate states
clear-ddos-msdpv6-aggregate-states
clear ddos-protection protocols msdpv6 aggregate statistics
clear-ddos-msdpv6-aggregate-statistics
clear ddos-protection protocols msdpv6 states
clear-ddos-msdpv6-states
clear ddos-protection protocols msdpv6 statistics
clear-ddos-msdpv6-statistics
clear ddos-protection protocols multicast-copy
clear ddos-protection protocols multicast-copy aggregate
clear ddos-protection protocols multicast-copy aggregate states
clear-ddos-mcast-copy-aggregate-states
clear ddos-protection protocols multicast-copy aggregate statistics
clear-ddos-mcast-copy-aggregate-statistics
clear ddos-protection protocols multicast-copy states
clear-ddos-mcast-copy-states
clear ddos-protection protocols multicast-copy statistics
clear-ddos-mcast-copy-statistics
clear ddos-protection protocols mvrp
clear ddos-protection protocols mvrp aggregate
clear ddos-protection protocols mvrp culprit-flows
clear ddos-protection protocols mvrp aggregate states
clear-ddos-mvrp-aggregate-states
```

```
clear ddos-protection protocols mvrp aggregate statistics
clear-ddos-mvrp-aggregate-statistics
clear ddos-protection protocols mvrp states
clear-ddos-mvrp-states
clear ddos-protection protocols mvrp statistics
clear-ddos-mvrp-statistics
clear ddos-protection protocols ndpv6
clear ddos-protection protocols ndpv6 aggregate
clear ddos-protection protocols ndpv6 aggregate states
clear ddos-protection protocols ndpv6 aggregate statistics
clear ddos-protection protocols ndpv6 states
clear ddos-protection protocols ndpv6 statistics
<clear-ddos-ndpv6-statistics>clear ddos-protection protocols ntp
clear ddos-protection protocols ntp aggregate
clear ddos-protection protocols ntp culprit-flows
clear ddos-protection protocols ntp aggregate states
clear-ddos-ntp-aggregate-states
clear ddos-protection protocols ntp aggregate statistics
clear-ddos-ntp-aggregate-statistics
clear ddos-protection protocols ntp states
clear-ddos-ntp-states
clear ddos-protection protocols ntp statistics
clear-ddos-ntp-statistics
clear ddos-protection protocols oam-lfm
clear ddos-protection protocols oam-lfm aggregate
clear ddos-protection protocols oam-lfm aggregate states
clear-ddos-oam-lfm-aggregate-states
clear ddos-protection protocols oam-lfm aggregate statistics
clear-ddos-oam-lfm-aggregate-statistics
clear ddos-protection protocols oam-lfm states
clear-ddos-oam-lfm-states
clear ddos-protection protocols oam-lfm statistics
clear-ddos-oam-lfm-statistics
clear ddos-protection protocols ospf
clear ddos-protection protocols ospf aggregate
clear ddos-protection protocols ospf aggregate culprit-flows
clear ddos-protection protocols ospf aggregate states
clear-ddos-ospf-aggregate-states
clear ddos-protection protocols ospf aggregate statistics
clear-ddos-ospf-aggregate-statistics
clear ddos-protection protocols ospf states
clear-ddos-ospf-states
clear ddos-protection protocols ospf statistics
clear-ddos-ospf-statistics
clear ddos-protection protocols ospfv3v6
clear ddos-protection protocols ospfv3v6 aggregate
clear ddos-protection protocols ospfv3v6 aggregate culprit-flows
clear ddos-protection protocols ospfv3v6 aggregate states
clear-ddos-ospfv3v6-aggregate-states
clear ddos-protection protocols ospfv3v6 aggregate statistics
clear-ddos-ospfv3v6-aggregate-statistics
clear ddos-protection protocols ospfv3v6 states
clear-ddos-ospfv3v6-states
clear ddos-protection protocols ospfv3v6 statistics
clear-ddos-ospfv3v6-statistics
clear ddos-protection protocols pfe-alive
```

```
clear ddos-protection protocols pfe-alive aggregate
clear ddos-protection protocols pfe-alive culprit-flows
clear ddos-protection protocols pfe-alive aggregate states
clear-ddos-pfe-alive-aggregate-states
clear ddos-protection protocols pfe-alive aggregate statistics
clear-ddos-pfe-alive-aggregate-statistics
clear ddos-protection protocols pfe-alive states
clear-ddos-pfe-alive-states
clear ddos-protection protocols pfe-alive statistics
clear-ddos-pfe-alive-statistics
clear ddos-protection protocols pim
clear ddos-protection protocols pim aggregate
clear ddos-protection protocols pim culprit-flows
clear ddos-protection protocols pim aggregate states
clear-ddos-pim-aggregate-states
clear ddos-protection protocols pim aggregate statistics
clear-ddos-pim-aggregate-statistics
clear ddos-protection protocols pim states
clear-ddos-pim-states
clear ddos-protection protocols pim statistics
clear-ddos-pim-statistics
clear ddos-protection protocols pimv6
clear ddos-protection protocols pimv6 aggregate
clear ddos-protection protocols pimv6 aggregate culprit-flows
clear ddos-protection protocols pimv6 aggregate states
clear ddos-protection protocols pimv6 aggregate statistics
clear ddos-protection protocols pimv6 states
clear ddos-protection protocols pimv6 statistics
clear ddos-protection protocols pmvrp
clear ddos-protection protocols pmvrp aggregate
clear ddos-protection protocols pmvrp culprit-flows
clear ddos-protection protocols pmvrp aggregate states
clear-ddos-pmvrp-aggregate-states
clear ddos-protection protocols pmvrp aggregate statistics
clear-ddos-pmvrp-aggregate-statistics
clear ddos-protection protocols pmvrp states
clear-ddos-pmvrp-states
clear ddos-protection protocols pmvrp statistics
clear-ddos-pmvrp-statistics
clear ddos-protection protocols pos
clear ddos-protection protocols pos aggregate

clear ddos-protection protocols pos aggregate states
clear-ddos-pos-aggregate-states
clear ddos-protection protocols pos aggregate statistics
clear-ddos-pos-aggregate-statistics
clear ddos-protection protocols pos states
clear-ddos-pos-states
clear ddos-protection protocols pos statistics
clear-ddos-pos-statistics
clear ddos-protection protocols ppp
clear ddos-protection protocols ppp aggregate
clear ddos-protection protocols ppp aggregate states
clear-ddos-ppp-aggregate-states
clear ddos-protection protocols ppp aggregate statistics
clear-ddos-ppp-aggregate-statistics
```

```
clear ddos-protection protocols ppp authentication
clear ddos-protection protocols ppp authentication states
clear-ddos-ppp-auth-states
clear ddos-protection protocols ppp authentication statistics
clear ddos-protection protocols pppvp culprit-flows
<clear-ddos-ppp-auth-statistics>
clear ddos-protection protocols ppp ipcp
clear ddos-protection protocols ppp ipcp states
clear-ddos-ppp-ipcp-states
clear ddos-protection protocols ppp ipcp statistics
clear-ddos-ppp-ipcp-statistics
clear ddos-protection protocols ppp ipv6cp
clear ddos-protection protocols ppp ipv6cp states
clear-ddos-ppp-ipv6cp-states
clear ddos-protection protocols ppp ipv6cp statistics
clear-ddos-ppp-ipv6cp-statistics
clear ddos-protection protocols ppp isis
clear ddos-protection protocols ppp isis states
clear-ddos-ppp-isis-states
clear ddos-protection protocols ppp isis statistics
clear-ddos-ppp-isis-statistics
clear ddos-protection protocols ppp lcp
clear ddos-protection protocols ppp lcp states
clear-ddos-ppp-lcp-states
clear ddos-protection protocols ppp lcp statistics
clear-ddos-ppp-lcp-statistics
clear ddos-protection protocols ppp mplsdp
clear ddos-protection protocols pppvp culprit-flows
clear ddos-protection protocols ppp mplsdp states
clear-ddos-ppp-mplsdp-states
clear ddos-protection protocols ppp mplsdp statistics
clear-ddos-ppp-mplsdp-statistics
clear ddos-protection protocols ppp states
clear-ddos-ppp-states
clear ddos-protection protocols ppp statistics
clear-ddos-ppp-statistics
clear ddos-protection protocols ppp unclassified
clear ddos-protection protocols ppp unclassified states
clear ddos-protection protocols ppp unclassified statistics
<clear-ddos-ppp-unclass-statistics>
clear ddos-protection protocols pppoe
clear ddos-protection protocols pppoe aggregate
clear ddos-protection protocols pppoe aggregate states
clear-ddos-pppoe-aggregate-states
clear ddos-protection protocols pppoe aggregate statistics
clear-ddos-pppoe-aggregate-statistics
clear ddos-protection protocols pppoe padi
clear ddos-protection protocols pppvp culprit-flows
clear ddos-protection protocols pppoe padi states
clear-ddos-pppoe-padi-states
clear ddos-protection protocols pppoe padi statistics
clear-ddos-pppoe-padi-statistics
clear ddos-protection protocols pppoe padm
clear ddos-protection protocols pppoe padm states
clear-ddos-pppoe-padm-states
clear ddos-protection protocols pppoe padm statistics
```


clear-ddos-pppoe-padm-statistics
clear ddos-protection protocols pppoe padn
clear ddos-protection protocols pppoe padn states
clear-ddos-pppoe-padm-states
clear ddos-protection protocols pppoe padn statistics
clear-ddos-pppoe-padm-statistics
clear ddos-protection protocols pppoe pado
clear ddos-protection protocols pppoe pado states
clear-ddos-pppoe-pado-states
clear ddos-protection protocols pppoe pado statistics
clear-ddos-pppoe-pado-statistics
clear ddos-protection protocols pppoe padr
clear ddos-protection protocols pppoe padr states
clear-ddos-pppoe-padr-states
clear ddos-protection protocols pppoe padr statistics
clear-ddos-pppoe-padr-statistics
clear ddos-protection protocols pppoe pads
clear ddos-protection protocols pmvrp culprit-flows
clear ddos-protection protocols pppoe pads states
clear-ddos-pppoe-pads-states
clear ddos-protection protocols pppoe pads statistics
clear-ddos-pppoe-pads-statistics
clear ddos-protection protocols pppoe padt
clear ddos-protection protocols pppoe padt states
clear-ddos-pppoe-padt-states
clear ddos-protection protocols pppoe padt statistics
clear-ddos-pppoe-padt-statistics
clear ddos-protection protocols pppoe states
clear-ddos-pppoe-states
clear ddos-protection protocols pppoe statistics
clear-ddos-pppoe-statistics
clear ddos-protection protocols ptp
clear ddos-protection protocols ptp aggregate
clear ddos-protection protocols ptp aggregate states
clear-ddos-ntp-aggregate-states
clear ddos-protection protocols ptp aggregate statistics
clear-ddos-ntp-aggregate-statistics
clear ddos-protection protocols ptp states
clear-ddos-ntp-states
clear ddos-protection protocols ptp statistics
clear-ddos-ntp-statistics
clear ddos-protection protocols pvstp
clear ddos-protection protocols pvstp aggregate
clear ddos-protection protocols pmvrp culprit-flows
clear ddos-protection protocols pvstp aggregate states
clear-ddos-pvstp-aggregate-states
clear ddos-protection protocols pvstp aggregate statistics
clear-ddos-pvstp-aggregate-statistics
clear ddos-protection protocols pvstp states
clear-ddos-pvstp-states
clear ddos-protection protocols pvstp statistics
clear-ddos-pvstp-statistics
clear ddos-protection protocols radius
clear ddos-protection protocols radius accounting
clear ddos-protection protocols radius accounting states
clear-ddos-radius-account-states

```
clear ddos-protection protocols radius accounting statistics
clear-ddos-radius-account-statistics
clear ddos-protection protocols radius aggregate
clear ddos-protection protocols radius aggregate states
clear-ddos-radius-aggregate-states
clear ddos-protection protocols radius aggregate statistics
clear-ddos-radius-aggregate-statistics
clear ddos-protection protocols radius authorization
clear ddos-protection protocols pmvrp culprit-flows
clear ddos-protection protocols radius authorization states
clear-ddos-radius-auth-states
clear ddos-protection protocols radius authorization statistics
clear-ddos-radius-auth-statistics
clear ddos-protection protocols radius server
clear ddos-protection protocols radius server states
clear-ddos-radius-server-states
clear ddos-protection protocols radius server statistics
clear-ddos-radius-server-statistics
clear ddos-protection protocols radius states
clear-ddos-radius-states
clear ddos-protection protocols radius statistics
clear-ddos-radius-statistics
clear ddos-protection protocols redirect
clear ddos-protection protocols redirect aggregate
clear ddos-protection protocols redirect aggregate states
clear-ddos-redirect-aggregate-states
clear ddos-protection protocols redirect aggregate statistics
clear ddos-protection protocols pmvrp culprit-flows
<clear-ddos-redirect-aggregate-statistics>
clear ddos-protection protocols redirect states
clear-ddos-redirect-states
clear ddos-protection protocols redirect statistics
clear-ddos-redirect-statistics
clear ddos-protection protocols reject
clear ddos-protection protocols reject aggregate
clear ddos-protection protocols reject aggregate states
clear ddos-protection protocols reject aggregate statistics
clear ddos-protection protocols reject states
clear ddos-protection protocols reject statistics
clear ddos-protection protocols rip
clear ddos-protection protocols rip aggregate
clear ddos-protection protocols rip aggregate states
clear-ddos-rip-aggregate-states
clear ddos-protection protocols rip aggregate statistics
clear-ddos-rip-aggregate-statistics
clear ddos-protection protocols rip states
clear-ddos-rip-states
clear ddos-protection protocols rip statistics
clear-ddos-rip-statistics
clear ddos-protection protocols ripv6
clear ddos-protection protocols ripv6 aggregate
clear ddos-protection protocols ripv6 aggregate states
clear-ddos-ripv6-aggregate-states
clear ddos-protection protocols ripv6 aggregate statistics
clear-ddos-ripv6-aggregate-statistics
clear ddos-protection protocols ripv6 states
```

```
clear-ddos-ripv6-states
clear ddos-protection protocols ripv6 statistics
clear-ddos-ripv6-statistics
clear ddos-protection protocols rsvp
clear ddos-protection protocols rsvp aggregate
clear ddos-protection protocols rsvp aggregate states
clear-ddos-rsvp-aggregate-states
clear ddos-protection protocols rsvp aggregate statistics
clear-ddos-rsvp-aggregate-statistics
clear ddos-protection protocols rsvp states
clear-ddos-rsvp-states
clear ddos-protection protocols rsvp statistics
clear-ddos-rsvp-statistics
clear ddos-protection protocols rsvpv6
clear ddos-protection protocols rsvpv6 aggregate
clear ddos-protection protocols rsvpv6 aggregate states
clear-ddos-rsvpv6-aggregate-states
clear ddos-protection protocols rsvpv6 aggregate statistics
clear-ddos-rsvpv6-aggregate-statistics
clear ddos-protection protocols rsvpv6 states
clear-ddos-rsvpv6-states
clear ddos-protection protocols rsvpv6 statistics
clear-ddos-rsvpv6-statistics
clear ddos-protection protocols sample
clear ddos-protection protocols sample aggregate
clear ddos-protection protocols sample aggregate states
<clear-ddos-sample-aggregate-states>
clear ddos-protection protocols sample aggregate statistics
<clear-ddos-sample-aggregate-statistics>
clear ddos-protection protocols sample host
clear ddos-protection protocols sample host states
<clear-ddos-sample-host-states>
clear ddos-protection protocols sample host statistics
<clear-ddos-sample-host-statistics>
clear ddos-protection protocols sample pfe
clear ddos-protection protocols sample pfe culprit-flows
clear ddos-protection protocols sample pfe states
<clear-ddos-sample-pfe-states>
clear ddos-protection protocols sample pfe statistics
<clear-ddos-sample-pfe-statistics>
clear ddos-protection protocols sample states
<clear-ddos-sample-states>
clear ddos-protection protocols sample statistics
<clear-ddos-sample-statistics>
clear ddos-protection protocols sample syslog
clear ddos-protection protocols sample syslog culprit-flows
clear ddos-protection protocols sample syslog states
<clear-ddos-sample-syslog-states>
clear ddos-protection protocols sample syslog statistics
<clear-ddos-sample-syslog-statistics>
clear ddos-protection protocols sample tap
clear ddos-protection protocols sample tap states
<clear-ddos-sample-tap-states>
clear ddos-protection protocols sample tap statistics
<clear-ddos-sample-tap-statistics>
clear ddos-protection protocols services
```

```
clear ddos-protection protocols services aggregate
clear ddos-protection protocols services aggregate states
clear-ddos-services-aggregate-states
clear ddos-protection protocols services aggregate statistics
clear-ddos-services-aggregate-statistics
clear ddos-protection protocols services states
clear-ddos-services-states
clear ddos-protection protocols services statistics
clear-ddos-services-statistics
clear ddos-protection protocols snmp
clear ddos-protection protocols snmp aggregate
clear ddos-protection protocols snmp culprit-flows
clear ddos-protection protocols snmp aggregate states
clear-ddos-snmp-aggregate-states
clear ddos-protection protocols snmp aggregate statistics
clear-ddos-snmp-aggregate-statistics
clear ddos-protection protocols snmp states
clear-ddos-snmp-states
clear ddos-protection protocols snmp statistics
clear-ddos-snmp-statistics
clear ddos-protection protocols snmpv6
clear ddos-protection protocols snmpv6 aggregate
clear ddos-protection protocols snmpv6 aggregate states
clear-ddos-snmpv6-aggregate-states
clear ddos-protection protocols snmpv6 aggregate statistics
clear-ddos-snmpv6-aggregate-statistics
clear ddos-protection protocols snmpv6 states
clear-ddos-snmpv6-states
clear ddos-protection protocols snmpv6 statistics
clear-ddos-snmpv6-statistics
clear ddos-protection protocols ssh
clear ddos-protection protocols ssh aggregate
clear ddos-protection protocols ssh aggregate states
clear-ddos-ssh-aggregate-states
clear ddos-protection protocols ssh aggregate statistics
clear-ddos-ssh-aggregate-statistics
clear ddos-protection protocols ssh states
clear-ddos-ssh-states
clear ddos-protection protocols ssh statistics
clear-ddos-ssh-statistics
clear ddos-protection protocols sshv6
clear ddos-protection protocols sshv6 aggregate
clear ddos-protection protocols sshv6 culprit-flows
clear ddos-protection protocols sshv6 aggregate states
clear-ddos-sshv6-aggregate-states
clear ddos-protection protocols sshv6 aggregate statistics
clear-ddos-sshv6-aggregate-statistics
clear ddos-protection protocols sshv6 states
clear-ddos-sshv6-states
clear ddos-protection protocols sshv6 statistics
clear-ddos-sshv6-statistics
clear ddos-protection protocols states
clear-ddos-protocols-states
clear ddos-protection protocols statistics
clear-ddos-protocols-statistics
clear ddos-protection protocols stp
```

```
clear ddos-protection protocols stp aggregate
clear ddos-protection protocols stp aggregate states
clear-ddos-stp-aggregate-states
clear ddos-protection protocols stp aggregate statistics
clear-ddos-stp-aggregate-statistics
clear ddos-protection protocols stp states
clear-ddos-stp-states
clear ddos-protection protocols stp statistics
clear-ddos-stp-statistics
clear ddos-protection protocols tacacs
clear ddos-protection protocols tacacs aggregate
clear ddos-protection protocols tacacs aggregate states
clear-ddos-tacacs-aggregate-states
clear ddos-protection protocols tacacs aggregate statistics
clear-ddos-tacacs-aggregate-statistics
clear ddos-protection protocols tacacs states
clear-ddos-tacacs-states
clear ddos-protection protocols tacacs statistics
clear-ddos-tacacs-statistics
clear ddos-protection protocols tcp-flags
clear ddos-protection protocols tcp-flags aggregate
clear ddos-protection protocols tcp-flags initial culprit-flows
clear ddos-protection protocols tcp-flags aggregate states
clear-ddos-tcp-flags-aggregate-states
clear ddos-protection protocols tcp-flags aggregate statistics
clear-ddos-tcp-flags-aggregate-statistics
clear ddos-protection protocols tcp-flags established
clear ddos-protection protocols tcp-flags established states
clear-ddos-tcp-flags-establish-states
clear ddos-protection protocols tcp-flags established statistics
clear-ddos-tcp-flags-establish-statistics
clear ddos-protection protocols tcp-flags initial
clear ddos-protection protocols tcp-flags initial states
clear-ddos-tcp-flags-initial-states
clear ddos-protection protocols tcp-flags initial statistics
clear-ddos-tcp-flags-initial-statistics
clear ddos-protection protocols tcp-flags states
clear-ddos-tcp-flags-states
clear ddos-protection protocols tcp-flags statistics
clear-ddos-tcp-flags-statistics
clear ddos-protection protocols tcp-flags unclassified
clear ddos-protection protocols tcp-flags unclassified states
clear-ddos-tcp-flags-unclass-states
clear ddos-protection protocols tcp-flags unclassified statistics
clear-ddos-tcp-flags-unclass-statistics
clear ddos-protection protocols telnet
clear ddos-protection protocols telnet aggregate
clear ddos-protection protocols telnet aggregate culprit-flows
clear ddos-protection protocols telnet aggregate states
clear-ddos-telnet-aggregate-states
clear ddos-protection protocols telnet aggregate statistics
clear-ddos-telnet-aggregate-statistics
clear ddos-protection protocols telnet states
clear-ddos-telnet-states
clear ddos-protection protocols telnet statistics
clear-ddos-telnet-statistics
```

```
clear ddos-protection protocols telnetv6
clear ddos-protection protocols telnetv6 aggregate
clear ddos-protection protocols telnetv6 aggregate states
clear-ddos-telnetv6-aggregate-states
clear ddos-protection protocols telnetv6 aggregate statistics
clear-ddos-telnetv6-aggregate-statistics
clear ddos-protection protocols telnetv6 states
clear-ddos-telnetv6-states
clear ddos-protection protocols telnetv6 statistics
clear-ddos-telnetv6-statistics
clear ddos-protection protocols ttl
clear ddos-protection protocols ttl aggregate
clear ddos-protection protocols ttl aggregate culprit-flows
clear ddos-protection protocols ttl aggregate states
clear-ddos-ttl-aggregate-states
clear ddos-protection protocols ttl aggregate statistics
clear-ddos-ttl-aggregate-statistics
clear ddos-protection protocols ttl states
clear-ddos-ttl-states
clear ddos-protection protocols ttl statistics
clear-ddos-ttl-statistics
clear ddos-protection protocols tunnel-fragment
clear ddos-protection protocols tunnel-fragment aggregate
clear ddos-protection protocols tunnel-fragment aggregate states
clear-ddos-tun-frag-aggregate-states
clear ddos-protection protocols tunnel-fragment aggregate statistics
clear-ddos-tun-frag-aggregate-statistics
clear ddos-protection protocols tunnel-fragment states
clear-ddos-tun-frag-states
clear ddos-protection protocols tunnel-fragment statistics
clear-ddos-tun-frag-statistics
clear ddos-protection protocols unclassified
clear ddos-protection protocols unclassified aggregate
clear ddos-protection protocols unclassified aggregate states
clear ddos-protection protocols unclassified aggregate statistics
clear ddos-protection protocols unclassified states
clear ddos-protection protocols unclassified statistics
<clear-ddos-uncls-statistics>
clear ddos-protection protocols virtual-chassis
clear ddos-protection protocols virtual-chassis aggregate
clear ddos-protection protocols virtual-chassis aggregate culprit-flows
clear ddos-protection protocols virtual-chassis aggregate states
clear-ddos-vchassis-aggregate-states
clear ddos-protection protocols virtual-chassis aggregate statistics
clear-ddos-vchassis-aggregate-statistics
clear ddos-protection protocols virtual-chassis control-high
clear ddos-protection protocols virtual-chassis control-high states
clear-ddos-vchassis-control-hi-states
clear ddos-protection protocols virtual-chassis control-high statistics
clear-ddos-vchassis-control-hi-statistics
clear ddos-protection protocols virtual-chassis control-low
clear ddos-protection protocols virtual-chassis control-low states
clear-ddos-vchassis-control-lo-states
clear ddos-protection protocols virtual-chassis control-low statistics
clear-ddos-vchassis-control-lo-statistics
clear ddos-protection protocols virtual-chassis states
```

```
clear-ddos-vchassis-states
clear ddos-protection protocols virtual-chassis statistics
clear-ddos-vchassis-statistics
clear ddos-protection protocols virtual-chassis unclassified
clear ddos-protection protocols virtual-chassis unclassified culprit-flows
clear ddos-protection protocols virtual-chassis unclassified states
clear-ddos-vchassis-unclass-states
clear ddos-protection protocols virtual-chassis unclassified statistics
clear-ddos-vchassis-unclass-statistics
clear ddos-protection protocols virtual-chassis vc-packets
clear ddos-protection protocols virtual-chassis vc-packets states
clear-ddos-vchassis-vc-packets-states
clear ddos-protection protocols virtual-chassis vc-packets statistics
clear-ddos-vchassis-vc-packets-statistics
clear ddos-protection protocols virtual-chassis vc-ttl-errors
clear ddos-protection protocols virtual-chassis vc-ttl-errors states
clear-ddos-vchassis-vc-ttl-err-states
clear ddos-protection protocols virtual-chassis vc-ttl-errors statistics
clear-ddos-vchassis-vc-ttl-err-statistics
clear ddos-protection protocols vrrp
clear ddos-protection protocols vrrp aggregate
clear ddos-protection protocols vrrp aggregate states
clear-ddos-vrrp-aggregate-states
clear ddos-protection protocols vrrp aggregate statistics
clear-ddos-vrrp-aggregate-statistics
<clear ddos-protection protocols vrrp states>
clear ddos-protection protocols vrrp culprit-flows
clear-ddos-vrrp-states
clear ddos-protection protocols vrrp statistics
clear-ddos-vrrp-statistics
clear ddos-protection protocols vrrpv6
clear ddos-protection protocols vrrpv6 aggregate
clear ddos-protection protocols vrrpv6 aggregate states
clear-ddos-vrrpv6-aggregate-states
clear ddos-protection protocols vrrpv6 aggregate statistics
clear-ddos-vrrpv6-aggregate-statistics
clear ddos-protection protocols vrrpv6 states
clear-ddos-vrrpv6-states
clear ddos-protection protocols vrrpv6 statistics
clear-ddos-vrrpv6-statistics
clear dhcp
clear dhcp relay
clear dhcp relay binding
<clear-dhcp-relay-binding-information>

clear dhcp relay binding interface
<clear-dhcp-interface-bindings>
clear dhcp relay statistics
<clear-dhcp-relay-statistics-information>

clear dhcp server
clear dhcp server binding
<clear-dhcp-server-binding-information>

clear dhcp server binding interface
<clear-dhcp-server-binding-interface>
```

```
clear dhcp server statistics
<clear-server-statistics-information>
clear dhcp statistics
<clear-dhcp-service-statistics-information>
clear dhcpv6
clear dhcpv6 relay
clear dhcpv6 relay binding
clear dhcpv6 relay binding interface
clear dhcpv6 relay statistics
<clear-dhcpv6-relay-statistics-information>
clear dhcpv6 server
clear dhcpv6 server binding
<clear-dhcpv6-server-binding-information>
clear dhcpv6 server binding interface
<clear-dhcpv6-server-binding-interface>
clear dhcpv6 server statistics
<clear-dhcpv6-server-statistics-information>
clear dhcpv6 statistics
<clear-dhcpv6-service-statistics-information>
clear diameter

clear diameter function
<clear-diameter-function>

clear diameter peer
<clear-diameter-peer>
<clear-dhcp-binding-information>

<clear-dhcp-conflict-information>

<clear-dhcp-statistics-information>

clear validation
clear validation database
clear validation session
clear validation statistics
clear dot1x
clear dot1x interface
<clear-dot1x-interface-session>

clear dot1x mac-address
<clear-dot1x-mac-session>

clear error
clear error bpdu
clear error bpdu interface
<clear-bpdu-error>
clear error mac-rewrite
clear error mac-rewrite interface
<clear-mac-rewrite-error>
clear esis
clear esis adjacency
<clear-esis-adjacency>
clear esis statistics
<clear-esis-statistics>
clear fabric
```



```
<clear-fabric>
clear fabric statistics
<clear-fabric-statistics>
clear firewall
<clear-firewall-counters>
clear firewall all
<clear-all-firewall-counters>
clear firewall log
<clear-firewall-log>
clear firewall policer
clear firewall policer counter
clear firewall policer counter all
clear helper
clear helper statistics
<clear-helper-statistics-information>

clear igmp
clear igmp membership
<clear-igmp-membership>
clear igmp snooping
clear igmp snooping membership
<clear-igmp-snooping-membership>
clear igmp snooping membership bridge-domain
<clear-igmp-snooping-bridge-domain-membership>
clear igmp snooping statistics
<clear-igmp-snooping-statistics>
clear igmp snooping statistics bridge-domain
<clear-igmp-snooping-bridge-domain-statistics>
clear igmp statistics
<clear-igmp-statistics>
clear ike
clear ike security-associations
<clear-ike-security-associations>
clear ilmi
clear ilmi statistics
<clear-ilmi-statistics>
clear interfaces
clear interfaces interface-set
clear interfaces interface-set statistics
<clear-interface-set-statistics>
clear interfaces interface-set statistics all
<clear-interface-set-statistics-all>
clear interfaces interval
<clear-interfaces-interval>
clear interfaces mac-database
<clear-interfaces-mac-database>
clear interfaces mac-database statistics
<clear-interface-mac-database-statistics>
clear interfaces mac-database statistics all
<clear-interface-mac-database-statistics-all>
clear interfaces statistics
<clear-interfaces-statistics>

clear interfaces statistics all
<clear-interfaces-statistics-all>
```

```
clear ipsec
clear ipsec security-associations
<clear-ipsec-security-associations>
clear ipv6
clear ipv6 neighbors
<clear-ipv6-nd-information>

clear ipv6 neighbors all
<clear-ipv6-all-neighbors>
clear isis
clear isis adjacency
<clear-isis-adjacency-information>

clear isis database
<clear-isis-database-information>

clear isis overload
<clear-isis-overload-information>

clear isis statistics
<clear-isis-statistics-information>

clear ipv6 router-advertisement
clear lacp
clear lacp statistics
clear l2-learning
clear l2-learning mac-move-buffer
<clear-l2-learning-mac-move-buffer>
clear l2-learning-redundancy-group
<clear-l2-learning-redundancy-group-statistics>
clear l2-learning remote-backbone-edge-bridges
<clear-l2-learning-remote-backbone-edge-bridges>
clear ldp
clear ldp statistics
<clear-ldp-statistics>
clear ldp statistics interface
<clear-ldp-interface-hello-statistics>
clear ldp neighbor
<clear-ldp-neighbors>
clear ldp session
<clear-ldp-sessions>
clear lldp
clear lldp neighbors
<clear-lldp-neighbors>
clear lldp neighbors interface
<clear-lldp-interface-neighbors>
clear lldp statistics
<clear-lldp-statistics>
clear lldp statistics interface
<clear-lldp-interface-statistics>
clear mld
clear mld membership
<clear-mld-membership>
clear mld statistics
<clear-mld-statistics>
clear mobile-ip
```

```
clear mobile-ip binding
clear mobile-ip binding all
<clear-binding-all>

clear mobile-ip binding ip-address
<clear-binding-ip>

clear mobile-ip binding nai
<clear-binding-nai>

clear mobile-ip visitor
clear mobile-ip visitor all
<clear-visitor-all>

clear mobile-ip visitor ip-address
<clear-visitor-ip>

clear mobile-ip visitor nai
<clear-visitor-nai>

clear mpls
clear mpls lsp
<clear-mpls-lsp-information>

clear mpls static-lsp
<clear-mpls-static-lsp-information>

clear mpls traceroute
clear mpls traceroute database
clear mpls traceroute database ldp
<clear-mpls-traceroute-database-ldp>
clear msdp
clear msdp cache
<clear-msdp-cache>
clear msdp statistics
<clear-msdp-statistics>
clear multicast
clear multicast bandwidth-admission
<clear-multicast-bandwidth-admission >
clear multicast forwarding-cache
clear multicast scope
<clear-multicast-scope-statistics>
clear multicast sessions
<clear-multicast-sessions>
clear multicast statistics
<clear-multicast-statistics>
clear mvrp
clear mvrp statistics
<clear-mvrp-interface-statistics>

clear network-access
clear network-access aaa
clear network-access aaa statistics
<clear-aaa-statistics-table>

clear network-access aaa statistics address-assignment
```

```
clear network-access aaa statistics address-assignment client
<clear-aaa-address-assignment-client-statistics>
clear network-access aaa statistics address-assignment pool
<clear-aaa-address-assignment-pool-statistics>
clear network-access aaa subscriber
<clear-aaa-subscriber-table>

clear network-access aaa subscriber statistics
<clear-aaa-subscriber-table-specific-statistics>

clear network-access requests
clear network-access requests pending
<clear-authentication-pending-table>

clear network-access requests statistics
<clear-authentication-statistics>

clear network-access securid-node-secret-file
<clear-node-secret-file>

clear oam
clear oam ethernet
clear oam ethernet connectivity-fault-management
clear oam ethernet connectivity-fault-management continuity-measurement
<clear-cfm-continuity-measurement>
clear oam ethernet connectivity-fault-management delay-statistics
<clear-cfm-delay-statistics>
clear oam ethernet connectivity-fault-management loss-statistics
<clear-cfm-loss-statistics>
clear oam ethernet connectivity-fault-management path-database
<clear-cfm-linktrace-path-database>

clear oam ethernet connectivity-fault-management policer
<clear-cfm-policer-statistics>
clear oam ethernet connectivity-fault-management sla-iterator-statistics
<clear-cfm-iterator-statistics>
clear oam ethernet connectivity-fault-management statistics
<clear-cfm-statistics>

clear oam ethernet link-fault-management
clear oam ethernet link-fault-management state
<clear-lfmd-state>
clear oam ethernet link-fault-management statistics
<clear-lfmd-statistics>
clear oam ethernet link-fault-management statistics action-profile
<clear-lfmd-action-profile-statistics>
clear oam ethernet lmi
clear oam ethernet lmi statistics
<clear-elmi-statistics>

clear ospf
clear ospf database
<clear-ospf-database-information>
clear ospf database-protection
<clear-ospf-database-protection>
```

```
clear ospf io-statistics
<clear-ospf-io-statistics-information>

clear ospf neighbor
<clear-ospf-neighbor-information>

clear ospf overload
<clear-ospf-overload-information>

clear ospf statistics
<clear-ospf-statistics-information>

clear ospf3
clear ospf3 database
<clear-ospf3-database-information>
clear ospf3 database-protection
<clear-ospf-database-protection>
clear ospf3 io-statistics
<clear-ospf3-io-statistics-information>
clear ospf3 neighbor
<clear-ospf3-neighbor-information>

clear ospf3 overload
<clear-ospf3-overload-information>

clear ospf3 statistics
<clear-ospf3-io-statistics-information>
clear pfe
clear pfe statistics
clear pfe statistics fabric
clear passive-monitoring
<clear-passive-monitoring>
clear passive-monitoring statistics
<clear-passive-monitoring-statistics>
clear pgm
clear pgm negative-acknowledgments
<clear-pgm-negative-acknowledgments>
clear pgm source-path-messages
<clear-pgm-source-path-messages>
clear pgm statistics
<clear-pgm-statistics>
clear pim
clear pim join
<clear-pim-join-state>
clear pim join-distribution
<clear-pim-join-distribution>
clear pim register
<clear-pim-register-state>
clear pim snooping
clear pim snooping join
clear pim snooping statistics
clear pim statistics
<clear-pim-statistics>
clear ppp
clear ppp statistics
<clear-ppp-statistics-information>
```

```
clear pppoe
clear pppoe lockout
<clear-pppoe-lockout-timers>
clear pppoe sessions
<clear-pppoe-sessions-information>
clear pppoe statistics
<clear-pppoe-statistics-information>
clear pppoe statistics interfaces
<clear-pppoe-statistics-interface-information>
clear protection-group
<clear-protection-group>
clear protection-group ethernet-ring
<clear-ethernet-ring-information>
clear protection-group ethernet-ring statistics
<clear-ethernet-ring-information>
clear r2cp
clear r2cp radio
<clear-r2cp-radio>
clear r2cp session
<clear-r2cp-session>
clear r2cp statistics
<clear-r2cp-statistics>
clear r2cp statistics radio
clear r2cp statistics session
clear rip
clear rip general-statistics
<clear-rip-general-statistics>
clear rip statistics
<clear-rip-statistics>
clear rip statistics peer
<clear-rip-peer-statistics>
clear ripng
clear ripng general-statistics
<clear-ripng-general-statistic>
clear ripng statistics
<clear-ripng-statistics>
clear rsvp
clear rsvp session
<clear-rsvp-session-information>
clear rsvp statistics
<clear-rsvp-counters-information>
clear services
clear services alg
clear services alg statistics
<clear-services-alg-statistics>
clear services application-aware-access-list
clear services application-aware-access-list statistics
<clear-application-aware-access-list-statistics-interface>
clear services application-aware-access-list statistics interface
<clear-application-aware-access-list-statistics-interface>
clear services application-aware-access-list statistics subscriber
<clear-application-aware-access-list-statistics-subscriber>
clear services application-identification
clear services application-identification application-system-cache
<clear-appid-application-system-cache>
```

```

clear services application-identification counter
<clear-appid-counter>
clear services application-identification counter ssl-encrypted-sessions
<clear-appid-counter-encrypted>
clear services application-identification statistics
<clear-appid-application-statistics>
clear services application-identification statistics cumulative
<clear-appid-application-statistics-cumulative>
clear services application-identification statistics interval
<clear-appid-application-statistics-interval>
clear services border-signaling-gateway
clear services border-signaling-gateway denied-messages
<clear-service-bsg-denied-messages>

clear services border-signaling-gateway name-resolution-cache

clear services border-signaling-gateway name-resolution-cache all
<clear-service-border-signaling-gateway-name-resolution-cache-all>

clear services border-signaling-gateway name-resolution-cache by-fqdn
<clear-border-signaling-gateway-name-resolution-cache-by-fqdn>
clear services border-signaling-gateway statistics
<clear-service-border-signaling-gateway-statistics>
clear services captive-portal-content-delivery
clear services captive-portal-content-delivery statistics
clear services captive-portal-content-delivery statistics interface
<clear-cpcdd-interface-statistics>
clear services cos
clear services cos statistics
<clear-services-cos-statistics>
clear services crtp
clear services crtp statistics
<clear-services-crtp-statistics>
clear services dynamic-flow-capture
clear services dynamic-flow-capture criteria
<clear-services-dynamic-flow-capture-criteria>
clear services dynamic-flow-capture sequence-number
clear services flow-collector
<clear-services-flow-collector-information>
clear services flow-collector statistics
<clear-services-flow-collector-statistics>
clear service-msp-flow-ipaction-table
clear services ids
<clear-services-ids-tables>
clear services ids destination-table
<clear-services-ids-destination-table>
clear services ids pair-table
<clear-services-ids-pair-table>
clear services ids source-table
<clear-services-ids-source-table>
clear services inline
clear services inline nat
clear services inline nat pool
<clear-inline-nat-pool-information>

```

```
clear services inline nat statistics
<clear-inline-nat-statistics>
clear services inline software
clear services inline software statistics
<clear-inline-software-statistics>
clear services ipsec-vpn
clear services ipsec-vpn ipsec
clear services ipsec-vpn ipsec security-associations
<clear-services-ipsec-vpn-security-associations>
clear services ipsec-vpn ike
clear services ipsec-vpn ike security-associations
<clear-services-ike-security-associations>
clear services pcp
clear services pcp epoch
clear services pcp statistics
clear services ipsec-vpn ipsec statistics
<clear-ipsec-vpn-statistics>
clear services l2tp
<clear-l2tp-destinations-information>
clear services l2tp disconnect-cause-summary
<clear-l2tp-disconnect-cause-summary>
clear services l2tp multilink
<clear-l2tp-multilink-information>
clear services l2tp session
<clear-l2tp-session-information>
clear services l2tp destination
<clear-l2tp-destinations-information>
clear services l2tp disconnect-cause-summary
<clear-l2tp-disconnect-cause-summary>
clear services l2tp tunnel
<clear-l2tp-tunnel-information>
clear services l2tp user
<clear-l2tp-user-session-information>
clear services local-policy-decision-function
clear services local-policy-decision-function statistics
clear services local-policy-decision-function statistics interface
<clear-local-policy-decision-function-statistics-interface>
clear services local-policy-decision-function statistics subscriber
<clear-local-policy-decision-function-statistics-subscriber>
clear services server-load-balance
clear services server-load-balance external-manager-statistics
<clear-external-manager-statistics>
clear services server-load-balance hash-table
<clear-hash-table-information>
clear services server-load-balance health-monitor-statistics
<clear-health-monitor-statistics>
clear services server-load-balance real-server-group-statistics
<clear-real-server-group-statistics>
clear services server-load-balance real-server-statistics
<clear-real-server-statistics>
clear services server-load-balance sticky
<clear-sticky-table>
clear services server-load-balance virtual-server-statistics
<clear-virtual-server-statistics>
clear services service-sets statistics syslog
<clear-service-set-syslog-statistics>
```



```
clear services stateful-firewall flow-analysis
<clear-service-flow-analysis>
clear services stateful-firewall flows
<clear-service-sfw-flow-table-information>
clear services stateful-firewall sip-call
<clear-service-sfw-sip-call-information>
clear services stateful-firewall sip-register
<clear-service-sfw-sip-register-information>
clear services stateful-firewall statistics
<clear-stateful-firewall-statistics>
clear services stateful-firewall subscriber-analysis
<clear-service-subs-analysis>
clear services subscriber
clear services subscriber sessions
<get-services-subscriber-sessions>
clear services software
clear services software statistics
<clear-services-software-statistics>
clear services stateful-firewall
clear services stateful-firewall flow-analysis
<clear-service-flow-analysis>
clear services stateful-firewall flows
<clear-service-sfw-flow-table-information>
clear services pgcp
clear services pgcp gates
<clear-service-pgcp-gates>

clear services pgcp gates gateway
<clear-service-pgcp-gates-gateway>

clear services pgcp statistics
<clear-service-pgcp-statistics>

clear services pgcp statistics gateway
<clear-service-pgcp-statistics-gateway>
clear twamp-information
clear twamp-server-information
clear twamp-server-connection-information
clear snmp
clear snmp history
<clear-snmp-history>
clear snmp statistics
<clear-snmp-statistics>
clear spanning-tree
clear spanning-tree protocol-migration
clear spanning-tree protocol-migration interface
<clear-interface-stp-protocol-migration>
clear spanning-tree statistics
<clear-stp-interface-statistics>
clear spanning-tree statistics interface
clear spanning-tree statistics routing-instance
<clear-stp-routing-instance-statistics>
clear spanning-tree topology-change-counter
<clear-stp-topology-change-counter>
clear synchronous-ethernet
clear synchronous-ethernet esmc
```

```
clear synchronous-ethernet esmc statistics
clear system
clear system login
clear system login lockout
< clear-system-login-lockout>

clear vpls
clear vpls mac-address
<clear-vpls-mac-address>
clear vpls mac-table
<clear-vpls-mac-table>

clear vpls mac-table interface
<clear-vpls-interface-mac-table>
request interface rebalance
request pppoe
request pppoe connect
request pppoe disconnect
request snmp
<request-snmp-utility-mib-clear>
<request-snmp-utility-mib-set>
clear vrrp
clear vrrp interface
request services ipsec-vpn ipsec
request services ipsec-vpn ipsec switch
request services ipsec-vpn ipsec switch tunnel
```

Configuration Hierarchy Levels No associated CLI configuration hierarchy levels and statements.

- Related Documentation**
- [Access Privilege User Permission Flags Overview on page 128](#)
 - [Understanding Junos OS Access Privilege Levels on page 123](#)
 - [Configuring Access Privilege Levels on page 131](#)
 - [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
 - [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)

configure

Can enter configuration mode.

Commands

```
configure
request snmp
request-snmp-utility-mib-clear
request-snmp-utility-mib-set
```

Configuration Hierarchy Levels No associated CLI configuration hierarchy levels and statements.

- Related Documentation**
- [Access Privilege User Permission Flags Overview on page 128](#)
 - [Understanding Junos OS Access Privilege Levels on page 123](#)

- [Configuring Access Privilege Levels on page 131](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)

control

Can perform all control-level operations; can modify any configuration.

Commands	test configuration
Configuration Hierarchy Levels	No associated CLI configuration hierarchy levels and statements.
Related Documentation	<ul style="list-style-type: none">• Access Privilege User Permission Flags Overview on page 128• Understanding Junos OS Access Privilege Levels on page 123• Configuring Access Privilege Levels on page 131• Specifying Access Privileges for Junos OS Operational Mode Commands on page 132• Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130

field

Can view field debug commands.

Commands	No associated CLI commands.
Configuration Hierarchy Levels	No associated CLI configuration hierarchy levels and statements.
Related Documentation	<ul style="list-style-type: none">• Access Privilege User Permission Flags Overview on page 128• Understanding Junos OS Access Privilege Levels on page 123• Configuring Access Privilege Levels on page 131• Specifying Access Privileges for Junos OS Operational Mode Commands on page 132• Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130

firewall

Can view the firewall filter configuration in configuration mode.

Commands	<pre>show firewall <get-firewall-information> show firewall counter <get-firewall-counter-information> show firewall filter <get-firewall-filter-information></pre>
-----------------	---

	<pre>show firewall filter version <get-filter-version> show firewall log <get-firewall-log-information> show firewall policer show firewall policer counters <get-firewall-policer-counter-information> show firewall policer counters counter-id <get-firewall-policer-per-counter-information> show firewall templates-in-use show firewall prefix-action-stats <get-firewall-prefix-action-information> show policer <get-policer-information></pre>
Configuration Hierarchy Levels	<pre>[edit dynamic-profiles firewall] [edit firewall] [edit logical-systems firewall]</pre>
Related Documentation	<ul style="list-style-type: none">• Access Privilege User Permission Flags Overview on page 128• Understanding Junos OS Access Privilege Levels on page 123• Configuring Access Privilege Levels on page 131• Specifying Access Privileges for Junos OS Operational Mode Commands on page 132• Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130• firewall-control on page 182

firewall-control

Can view and configure firewall filter information at the **[edit dynamic-profiles firewall]**, **[edit firewall]**, and **[edit logical-systems firewall]** hierarchy levels.

Commands	<pre>show firewall <get-firewall-information> show firewall counter <get-firewall-counter-information> show firewall filter <get-firewall-filter-information> show firewall filter version <get-filter-version> show firewall log <get-firewall-log-information> show firewall prefix-action-stats</pre>
-----------------	---

	<code><get-firewall-prefix-action-information></code>
	<code>show policer</code>
Configuration Hierarchy Levels	<ul style="list-style-type: none"> <code>[edit dynamic-profiles firewall]</code> <code>[edit firewall]</code> <code>[edit logical-systems firewall]</code>
Related Documentation	<ul style="list-style-type: none"> • Access Privilege User Permission Flags Overview on page 128 • Understanding Junos OS Access Privilege Levels on page 123 • Configuring Access Privilege Levels on page 131 • Specifying Access Privileges for Junos OS Operational Mode Commands on page 132 • Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130 • firewall on page 181
floppy	
	Can read from and write to the removable media.
Commands	No associated CLI commands.
Configuration Hierarchy Levels	No associated CLI configuration hierarchy levels and statements.
Related Documentation	<ul style="list-style-type: none"> • Access Privilege User Permission Flags Overview on page 128 • Understanding Junos OS Access Privilege Levels on page 123 • Configuring Access Privilege Levels on page 131 • Specifying Access Privileges for Junos OS Operational Mode Commands on page 132 • Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130
flow-tap	
	Can view the flow-tap configuration in configuration mode.
Commands	No associated CLI commands.
Configuration Hierarchy Levels	<ul style="list-style-type: none"> <code>[edit services flow-tap]</code> <code>[edit services radius-flow-tap]</code> <code>[edit system services flow-tap-dtcp]</code>
Related Documentation	<ul style="list-style-type: none"> • Access Privilege User Permission Flags Overview on page 128 • Understanding Junos OS Access Privilege Levels on page 123 • Configuring Access Privilege Levels on page 131 • Specifying Access Privileges for Junos OS Operational Mode Commands on page 132

- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)
- [flow-tap-control on page 184](#)

flow-tap-control

Can view the flow-tap configuration in configuration mode and can configure flow-tap configuration information at the **[edit services flow-tap]**, **[edit services radius-flow-tap]**, and **[edit system services flow-tap-dtcp]** hierarchy levels.

Commands No associated CLI commands.

Configuration Hierarchy Levels **[edit services flow-tap]**
[edit services radius-flow-tap]
[edit system services flow-tap-dtcp]

- Related Documentation**
- [Access Privilege User Permission Flags Overview on page 128](#)
 - [Understanding Junos OS Access Privilege Levels on page 123](#)
 - [Configuring Access Privilege Levels on page 131](#)
 - [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
 - [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)
 - [flow-tap on page 183](#)

flow-tap-operation

Can make flow-tap requests to the router.

Commands No associated CLI commands.

Configuration Hierarchy Levels No associated CLI configuration hierarchy levels and statements.

- Related Documentation**
- [Access Privilege User Permission Flags Overview on page 128](#)
 - [Understanding Junos OS Access Privilege Levels on page 123](#)
 - [Configuring Access Privilege Levels on page 131](#)
 - [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
 - [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)

idp-profiler-operation

Can view profiler data.

Commands	No associated CLI commands.
CLI Configuration Hierarchy Levels	No associated CLI configuration hierarchy levels and statements.
	interface
	Can view the interface configuration in configuration mode.
Commands	No associated CLI commands.
Configuration Hierarchy Levels	<ul style="list-style-type: none"> [edit accounting-options] [edit chassis] [edit class-of-service] [edit class-of-service interfaces] [edit dynamic-profiles class-of-service] [edit dynamic-profiles class-of-service interfaces] [edit dynamic-profiles interfaces] [edit dynamic-profiles routing-instances instance system services dhcp-local-server] [edit dynamic-profiles routing-instances instance system services static-subscribers group] [edit forwarding-options] [edit interfaces] [edit jnx-example] [edit logical-systems forwarding-options] [edit logical-systems interfaces] [edit logical-systems routing-instances instance system services dhcp-local-server] [edit logical-systems routing-instances instance system services static-subscribers group] [edit logical-systems system services dhcp-local-server] [edit logical-systems system services static-subscribers group] [edit routing-instances instance system services dhcp-local-server] [edit routing-instances instance system services static-subscribers group] [edit services logging] [edit services radius-flow-tap] [edit services radius-flow-tap interfaces] [edit system services dhcp-local-server] [edit system services static-subscribers group]
Related Documentation	<ul style="list-style-type: none"> • Access Privilege User Permission Flags Overview on page 128 • Understanding Junos OS Access Privilege Levels on page 123 • Configuring Access Privilege Levels on page 131 • Specifying Access Privileges for Junos OS Operational Mode Commands on page 132 • Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130 • interface-control on page 185

interface-control

Can view chassis, class of service (CoS), groups, forwarding options, and interfaces configuration information. Can edit configuration at the **[edit chassis]**, **[edit**

`class-of-service`], `[edit groups]`, `[edit forwarding-options]`, and `[edit interfaces]` hierarchy levels.

Commands No associated CLI commands.

Configuration Hierarchy Levels

- `[edit accounting-options]`
- `[edit chassis]`
- `[edit class-of-service]`
- `[edit class-of-service interfaces]`
- `[edit dynamic-profiles class-of-service]`
- `[edit dynamic-profiles class-of-service interfaces]`
- `[edit dynamic-profiles interfaces]`
- `[edit dynamic-profiles routing-instances instance system services dhcp-local-server]`
- `[edit dynamic-profiles routing-instances instance system services static-subscribers group]`
- `[edit forwarding-options]`
- `[edit interfaces]`
- `[edit jnx-example]`
- `[edit logical-systems forwarding-options]`
- `[edit logical-systems interfaces]`
- `[edit logical-systems routing-instances instance system services dhcp-local-server]`
- `[edit logical-systems routing-instances instance system services static-subscribers group]`
- `[edit logical-systems system services dhcp-local-server]`
- `[edit logical-systems system services static-subscribers group]`
- `[edit routing-instances instance system services dhcp-local-server]`
- `[edit routing-instances instance system services static-subscribers group]`
- `[edit services logging]`
- `[edit services radius-flow-tap]`
- `[edit services radius-flow-tap interfaces]`
- `[edit system services dhcp-local-server]`
- `[edit system services static-subscribers group]`

- Related Documentation**
- [Access Privilege User Permission Flags Overview on page 128](#)
 - [Understanding Junos OS Access Privilege Levels on page 123](#)
 - [Configuring Access Privilege Levels on page 131](#)
 - [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
 - [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)
 - [interface on page 185](#)

maintenance

Can perform system maintenance, including starting a local shell on the router and becoming the superuser in the shell, and can halt and reboot the router.

Commands

- `clear system reboot`
 - `<clear-reboot>`
- `clear-system-services-reverse-information`
- `file archive`
 - `<file-archive>`


```
monitor traffic
request chassis beacon
<request-chassis-beacon>
request chassis ccg
<request-chassis-ccg>
request chassis cb
request chassis cfeb
request chassis cfeb master
request chassis cip
request chassis fabric
request chassis fabric device
request chassis fabric plane
request chassis fabric upgrade-bandwidth
request chassis fabric upgrade-bandwidth fpc
request chassis fabric upgrade-bandwidth info
request chassis feb
  <request-feb>

request chassis fpc
request chassis mcs
request chassis mic
request chassis optics
request chassis pcg
request chassis pic
request chassis redundancy
request chassis redundancy feb
  <request-redundancy-feb>
request chassis routing-engine
request chassis routing-engine hard-disk-test
request chassis routing-engine master
request chassis scg
request chassis sfb
request chassis sfm
request chassis sfm master
request chassis sib
request chassis sib f13

request chassis sib f2s
request chassis spmb
request chassis ssb
request chassis ssb master
request chassis synchronization
request chassis synchronization force
request chassis synchronization force automatic-switching
request chassis synchronization force mark-failed
request chassis synchronization force unmark-failed
request chassis synchronization switch
request chassis tfeb
request chassis vcpu
request chassis vnpu
request l2circuit-switchover
request mpls
request mpls lsp
request mpls lsp adjust-autobandwidth
<request-mpls-lsp-autobandwidth-adjust>
request security
```

```
request security certificate
request security certificate enroll
request security datapath-debug
request security datapath-debug action-profile
request security datapath-debug action-profile reload-all
  <reload-eedebg-action-profile>

request security idp
  <request-idp-security-policy-load>

request security idp security-package
request security idp security-package download
  <request-idp-security-package-download>

request security idp security-package download version
  <request-idp-security-package-download-version>

request security idp security-package install
  <request-idp-security-package-install>

request security idp ssl-inspection
request security idp ssl-inspection key
request security idp ssl-inspection key add
  <request-idp-ssl-key-add>

request security idp ssl-inspection key delete
  <request-idp-ssl-key-delete>
request security idp storage-cleanup
  <request-idp-storage-cleanup>
request security key-pair
request security pki
request security pki ca-certificate
request security pki ca-certificate ca-profile-group
request security pki ca-certificate ca-profile-group load
request security pki ca-certificate enroll
request security pki local-certificate export
request security pki ca-certificate load
  <load-pki-ca-certificate>
request security pki ca-certificate verify
  <verify-pki-ca-certificate>
request security pki crt
request security pki crt load
  <load-pki-crt>
request security pki generate-certificate-request
  <generate-pki-certificate-request>
request security pki generate-key-pair
  <generate-pki-key-pair>
request security pki local-certificate
request security pki local-certificate enroll
request security pki local-certificate generate-self-signed
  <generate-pki-self-signed-local-certificate>
request security pki local-certificate load
  <load-pki-local-certificate>
request security pki local-certificate verify
  <verify-pki-local-certificate>
request security pki verify-integrity-status
```

```
<verify-integrity-status>
request services fips
request services fips authorize
request services fips authorize pic
request services fips zeroize
request services fips zeroize pic
request services flow-collector
request services flow-collector change-destination
  <request-services-flow-collector-destination>

request services ggsn
request services ggsn pdp
request services ggsn pdp terminate
request services ggsn pdp terminate apn
  <request-ggsn-terminate-contexts-apn>

request services ggsn pdp terminate context
  <request-ggsn-terminate-context>

request services ggsn pdp terminate context msisdn
  <request-ggsn-terminate-msisdn-context>

request services ggsn restart
request services ggsn restart interface
  <request-ggsn-restart-interface>

request services ggsn restart node
  <request-ggsn-restart-node>

request services ggsn start
request services ggsn start interface
request services ggsn stop
request services ggsn stop interface
  <request-ggsn-stop-interface>

request services ggsn stop node
  <request-ggsn-stop-node>

request services ggsn trace
request services ggsn trace software
request services ggsn trace software update
  <request-ggsn-software-update>

request services ggsn trace start
request services ggsn trace start imsi
  <request-ggsn-start-imsi-trace>

request services ggsn trace start msisdn
  <request-ggsn-start-msisdn-trace>

request services ggsn trace stop
request services ggsn trace stop all
  <request-ggsn-stop-trace-activity>

request services ggsn trace stop imsi
  <request-ggsn-stop-imsi-trace>
```

```
request services ggsn trace stop msisdn
  <request-ggsn-stop-msisdn-trace>

request support
request support information
request system
request system certificate
request system certificate add
request system commit
request system commit server
request system commit server pause
  <request-commit-server-pause>
request system commit server queue
request system commit server queue cleanup
  <request-commit-server-cleanup>
request system commit server start
  <request-commit-server-start>
request system configuration
request system configuration rescue
request system configuration rescue delete
  <request-delete-rescue-configuration>

request system configuration rescue save
  <request-save-rescue-configuration>
request system diagnostics
request system diagnostics transfer-control
  <transfer-control>
request system firmware
request system firmware downgrade
request system firmware downgrade feb
request system firmware downgrade fpc
request system firmware downgrade pic
request system firmware downgrade poe
request system firmware downgrade re
request system firmware downgrade scb
request system firmware downgrade sfm
request system firmware downgrade spmb
request system firmware downgrade ssb
request system firmware downgrade vcpu
request system firmware upgrade
request system firmware upgrade feb
request system firmware upgrade fpc
request system firmware upgrade fpga
request system firmware upgrade fpga fpc
request system firmware upgrade fpga scb
  <request-scb-fpga-upgrade>
request system firmware upgrade pic
request system firmware upgrade poe
request system firmware upgrade re
request system firmware upgrade re bios
request system firmware upgrade scb
request system firmware upgrade sfm
request system firmware upgrade spmb
request system firmware upgrade ssb
request system firmware upgrade vcpu
```

```
request system halt
  <request-halt>

request system keep-alive
request system license
request system license add
request system license delete
  <request-license-delete>

request system license save
request system license update
  <request-license-update>
request system logout
request system partition
request system partition abort
request system partition compact-flash
request system partition hard-disk
request system power-off
  <request-power-off>

request system power-on
request system reboot
  <request-reboot>

request system scripts
request system scripts add
  <request-scripts-package-add>

request system scripts convert
request system scripts convert slax-to-xslt
request system scripts convert xslt-to-slax
request system scripts delete
  <request-scripts-package-delete>

request system scripts event-scripts
request system scripts event-scripts reload
  <reload-event-scripts>

request system scripts refresh-from
  <request-script-refresh-from>

request system scripts rollback
  <request-scripts-package-rollback>

request system snapshot
  <request-snapshot>

request system software
request system software abort
request system software abort in-service-upgrade
  <abort-in-service-upgrade>

request system software add
  <request-package-add>

request system software delete
```

```

<request-package-delete>

request system software delete-backup
  <request-package-delete-backup>

request system software in-service-upgrade
  <request-package-in-service-upgrade>

request system software nonstop-upgrade
  <request-package-nonstop-upgrade>
request system software recovery-package
request system software recovery-package add
request system software recovery-package delete
request system software recovery-package extract
request system software recovery-package extract ex-8200-package
request system software recovery-package extract ex-xre200-package
request system software rollback
  <request-package-rollback>

request system software validate
  <request-package-validate>
request system software validate in-service-upgrade
  <check-in-service-upgrade>

request system storage
request system storage cleanup
  <request-system-storage-cleanup>
request system storage cleanup qfabric
  <remove-qfabric-repository-contents>
request system zeroize
request vpls-switchover
set date
set date ntp
show chassis usb
show chassis usb storage
  <get-usb-storage-status>
show services fips
start shell
start shell user
test access
test access profile
  <get-radius-profile-access-test-result>

test access radius-server
  <get-radius-server-access-test-result>
get-test-services-l2tp-tunnel-result

```

Configuration Hierarchy Levels

```

[edit event-options]
[edit security ipsec internal]
[edit security ipsec trusted-channel]
[edit services dynamic-flow-capture traceoptions]
[edit services ggsn]
[edit system fips]
[edit services ggsn rule-space]
[edit system processes daemon-process command]

```

[edit system scripts]
 [edit system scripts commit]
 [edit system scripts op]

Related Documentation

- [Access Privilege User Permission Flags Overview on page 128](#)
- [Understanding Junos OS Access Privilege Levels on page 123](#)
- [Configuring Access Privilege Levels on page 131](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)

network

Can access the network by using the **ping**, **ssh**, **telnet**, and **traceroute** commands.

Commands

```
mtrace
mtrace from-source
mtrace monitor
mtrace to-gateway
ping
  <ping>

ping atm
ping clns
ping ethernet
  <request-ping-ethernet>
ping fibre-channel
ping mpls
ping mpls bgp
  <request-ping-bgp-lsp>
ping mpls l2circuit
ping mpls l2circuit interface
  <request-ping-l2circuit-interface>

ping mpls l2circuit virtual-circuit
  <request-ping-l2circuit-virtual-circuit>

ping mpls l2vpn
ping mpls l2vpn fec129
ping mpls l2vpn instance
  <request-ping-l2vpn-instance>

ping mpls l2vpn interface
  <request-ping-l2vpn-interface>

ping mpls l3vpn
  <request-ping-l3vpn>

ping mpls ldp
  <request-ping-ldp-lsp>

ping mpls ldp p2mp
```

```
<request-ping-ldp-p2mp-lsp>

ping mpls lsp-end-point
  <request-ping-lsp-end-point>

ping mpls rsvp
  <request-ping-rsvp-lsp>

ping vpls
ping vpls instance
  <request-ping-vpls-instance>

request routing-engine
request routing-engine login
<request-routing-engine-login>
request routing-engine login other-routing-engine
<request-login-to-other-routing-engine>
request services flow-collector
request services flow-collector test-file-transfer
  <request-services-flow-collector-test-file-transfer>

show host
show interfaces level-extra descriptions
show multicast mrinfo
ssh
telnet
traceroute
  <traceroute>

traceroute clns
traceroute ethernet
  <request-traceroute-ethernet>

traceroute monitor
traceroute mpls
traceroute mpls ldp
  <traceroute-mpls-ldp>
traceroute mpls rsvp
  <traceroute-mpls-rsvp>
```

Configuration Hierarchy Levels No associated CLI configuration hierarchy levels and statements.

- Related Documentation**
- [Access Privilege User Permission Flags Overview on page 128](#)
 - [Understanding Junos OS Access Privilege Levels on page 123](#)
 - [Configuring Access Privilege Levels on page 131](#)
 - [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
 - [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)

pgcp-session-mirroring

Can view session mirroring configuration by using the **pgcp** command.

Commands	show services pgcp gates gate-way display session-mirroring
Configuration Hierarchy Levels	[edit services pgcp gateway session-mirroring] [edit services pgcp session-mirroring]
Related Documentation	<ul style="list-style-type: none"> • Access Privilege User Permission Flags Overview on page 128 • Understanding Junos OS Access Privilege Levels on page 123 • Configuring Access Privilege Levels on page 131 • Specifying Access Privileges for Junos OS Operational Mode Commands on page 132 • Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130 • pgcp-session-mirroring-control on page 195

pgcp-session-mirroring-control

Can modify the PGCP session mirroring configuration.

Commands	show services pgcp gates gate-way display session-mirroring
Configuration Hierarchy Levels	[edit services pgcp gateway session-mirroring] [edit services pgcp session-mirroring]
Related Documentation	<ul style="list-style-type: none"> • Access Privilege User Permission Flags Overview on page 128 • Understanding Junos OS Access Privilege Levels on page 123 • Configuring Access Privilege Levels on page 131 • Specifying Access Privileges for Junos OS Operational Mode Commands on page 132 • Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130 • pgcp-session-mirroring on page 194

reset

Can restart software processes by using the **restart** command and can configure whether software processes configured at the **[edit system processes]** hierarchy level are enabled or disabled.

Commands	request chassis cfeb master switch request chassis cfeb master switch no-confirm request chassis routing-engine master acquire request chassis routing-engine master acquire force request chassis routing-engine master acquire force no-confirm request chassis routing-engine master acquire no-confirm request chassis routing-engine master release request chassis routing-engine master release no-confirm request chassis routing-engine master switch request chassis routing-engine master switch no-confirm request chassis sfm master switch
-----------------	--

```
request chassis sfm master switch no-confirm
request chassis ssb master switch
request chassis ssb master switch no-confirm
restart
restart kernel-replication
  <restart-kernel-replication>
restart-named-service
restart routing
  <routing-restart>
restart services
restart services border-signaling-gateway
  <restart-border-signaling-gateway-service>
restart services pgcp
  <restart-pgcp-service>
restart web-management
  <restart-web-management>
```

Configuration Hierarchy Levels No associated CLI configuration hierarchy levels and statements.

- Related Documentation**
- [Access Privilege User Permission Flags Overview on page 128](#)
 - [Understanding Junos OS Access Privilege Levels on page 123](#)
 - [Configuring Access Privilege Levels on page 131](#)
 - [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
 - [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)

rollback

Can roll back to previous configurations.

Commands rollback

Configuration Hierarchy Levels [edit]

- Related Documentation**
- [Access Privilege User Permission Flags Overview on page 128](#)
 - [Understanding Junos OS Access Privilege Levels on page 123](#)
 - [Configuring Access Privilege Levels on page 131](#)
 - [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
 - [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)

routing

Can view general routing, routing protocol, and routing policy configuration information.

Commands	No associated CLI commands.
Configuration Hierarchy Levels	<pre> [edit bridge-domains] [edit bridge-domains domain multicast-snooping-options] [edit bridge-domains domain multicast-snooping-options traceoptions] [edit dynamic-profiles protocols igmp traceoptions] [edit dynamic-profiles protocols mld traceoptions] [edit dynamic-profiles protocols router-advertisement traceoptions] [edit dynamic-profiles routing-instances] [edit dynamic-profiles routing-instances instance bridge-domains] [edit dynamic-profiles routing-instances instance bridge-domains domain multicast-snooping-options] [edit dynamic-profiles routing-instances instance bridge-domains domain multicast-snooping-options traceoptions] [edit dynamic-profiles routing-instances instance multicast-snooping-options] [edit dynamic-profiles routing-instances instance multicast-snooping-options traceoptions] [edit dynamic-profiles routing-instances instance pbb-options] [edit dynamic-profiles routing-instances instance protocols] [edit dynamic-profiles routing-instances instance protocols bgp group neighbor traceoptions] [edit dynamic-profiles routing-instances instance protocols bgp group traceoptions] [edit dynamic-profiles routing-instances instance protocols bgp traceoptions] [edit dynamic-profiles routing-instances instance protocols esis traceoptions] [edit dynamic-profiles routing-instances instance protocols isis traceoptions] [edit dynamic-profiles routing-instances instance protocols l2vpn traceoptions] [edit dynamic-profiles routing-instances instance protocols ldp traceoptions] [edit dynamic-profiles routing-instances instance protocols msdp group peer traceoptions] [edit dynamic-profiles routing-instances instance protocols msdp group traceoptions] [edit dynamic-profiles routing-instances instance protocols msdp peer traceoptions] [edit dynamic-profiles routing-instances instance protocols msdp traceoptions] [edit dynamic-profiles routing-instances instance protocols mvpn traceoptions] [edit dynamic-profiles routing-instances instance protocols ospf traceoptions] [edit dynamic-profiles routing-instances instance protocols pim traceoptions] [edit dynamic-profiles routing-instances instance protocols rip traceoptions] [edit dynamic-profiles routing-instances instance protocols ripng traceoptions] [edit dynamic-profiles routing-instances instance protocols router-discovery traceoptions] [edit dynamic-profiles routing-instances instance protocols vpls traceoptions] [[edit dynamic-profiles routing-instances instance routing-options] [edit dynamic-profiles routing-instances instance routing-options multicast traceoptions] [edit dynamic-profiles routing-instances instance routing-options traceoptions] [edit dynamic-profiles routing-instances instance service-groups] [edit dynamic-profiles routing-instances instance switch-options] [edit dynamic-profiles routing-options multicast traceoptions] [edit jnx-example] [edit fabric protocols] [edit fabric protocols bgp group neighbor traceoptions] [edit fabric protocols bgp group traceoptions] [edit fabric protocols bgp traceoptions] [edit fabric routing-instances] [edit fabric routing-instances instance routing-options] [edit fabric routing-instances instance routing-options traceoptions] [edit fabric routing-options] [edit fabric routing-options traceoptions] </pre>

```
[edit logical-systems bridge-domains]
[edit logical-systems bridge-domains domain multicast-snooping-options]
[edit logical-systems bridge-domains domain multicast-snooping-options traceoptions]
[edit logical-systems policy-options]
[edit logical-systems protocols]
[edit logical-systems protocols bgp group neighbor traceoptions]
[edit logical-systems protocols bgp group traceoptions]
[edit logical-systems protocols bgp traceoptions]
[edit logical-systems protocols dvmrp traceoptions]
[edit logical-systems protocols esis traceoptions]
[edit logical-systems protocols igmp traceoptions]
[edit logical-systems protocols igmp-host traceoptions]
[edit logical-systems protocols isis traceoptions]
[edit logical-systems protocols l2circuit traceoptions]
[edit logical-systems protocols l2iw traceoptions]
[edit logical-systems protocols ldp traceoptions]
[edit logical-systems protocols mld traceoptions]
[edit logical-systems protocols mld-host traceoptions]
[edit logical-systems protocols msdp group peer traceoptions]
[edit logical-systems protocols msdp group traceoptions]
[edit logical-systems protocols msdp peer traceoptions]
[edit logical-systems protocols msdp traceoptions]
[edit logical-systems protocols ospf traceoptions]
[edit logical-systems protocols pim traceoptions]
[edit logical-systems protocols rip traceoptions]
[edit logical-systems protocols ripng traceoptions]
[edit logical-systems protocols router-advertisement traceoptions]
[edit logical-systems protocols router-discovery traceoptions]
[edit logical-systems protocols rsvp lsp-set]
[edit logical-systems protocols rsvp traceoptions]
[edit logical-systems routing-instances]
[edit logical-systems routing-instances instance bridge-domains]
[edit logical-systems routing-instances instance bridge-domains domain
multicast-snooping-options]
[edit logical-systems routing-instances instance bridge-domains domain
multicast-snooping-options traceoptions]
[edit logical-systems routing-instances instance multicast-snooping-options]
[edit logical-systems routing-instances instance multicast-snooping-options
traceoptions]
[edit logical-systems routing-instances instance pbb-options]
[edit logical-systems routing-instances instance protocols]
[edit logical-systems routing-instances instance protocols bgp group neighbor
traceoptions]
[edit logical-systems routing-instances instance protocols bgp group traceoptions]
[edit logical-systems routing-instances instance protocols bgp traceoptions]
[edit logical-systems routing-instances instance protocols esis traceoptions]
[edit logical-systems routing-instances instance protocols isis traceoptions]
[edit logical-systems routing-instances instance protocols l2vpn traceoptions]
[edit logical-systems routing-instances instance protocols ldp traceoptions]
[edit logical-systems routing-instances instance protocols msdp group peer traceoptions]
[edit logical-systems routing-instances instance protocols msdp group traceoptions]
[edit logical-systems routing-instances instance protocols msdp peer traceoptions]
[edit logical-systems routing-instances instance protocols msdp traceoptions]
[edit logical-systems routing-instances instance protocols mvpn traceoptions]
[edit logical-systems routing-instances instance protocols ospf traceoptions]
[edit logical-systems routing-instances instance protocols pim traceoptions]
```

```
[edit logical-systems routing-instances instance protocols rip traceoptions]
[edit logical-systems routing-instances instance protocols ripng traceoptions]
[edit logical-systems routing-instances instance protocols router-discovery traceoptions]
[edit logical-systems routing-instances instance protocols vpls traceoptions]
[edit logical-systems routing-instances instance routing-options]
[edit logical-systems routing-instances instance routing-options multicast traceoptions]
[edit logical-systems routing-instances instance routing-options validation group session
traceoptions]
[edit logical-systems routing-instances instance routing-options validation traceoptions]
[edit logical-systems routing-instances instance routing-options traceoptions]
[edit logical-systems routing-options validation group session traceoptions]
[edit logical-systems routing-instances instance service-groups]
[edit logical-systems routing-instances instance switch-options]
[edit logical-systems routing-options]
[edit logical-systems routing-options validation group session traceoptions]
[edit logical-systems routing-options validation traceoptions]
[edit logical-systems routing-options multicast traceoptions]
[edit logical-systems routing-options traceoptions]
[edit logical-systems switch-options]
[edit multicast-snooping-options]
[edit multicast-snooping-options traceoptions]
[edit policy-options]
[edit protocols]
[edit protocols amt traceoptions]
[edit protocols bgp group neighbor traceoptions]
[edit protocols bgp group traceoptions]
[edit protocols bgp traceoptions]
[edit protocols connections]
[edit protocols dot1x]
[edit protocols dvmrp traceoptions]
[edit protocols esis traceoptions]
[edit protocols igmp traceoptions]
[edit protocols igmp-host traceoptions]
[edit protocols igmp-snooping]
[edit protocols isis traceoptions]
[edit protocols l2circuit traceoptions]
[edit protocols l2iw traceoptions]
[edit protocols ldp traceoptions]
[edit protocols lldp]
[edit protocols lldp-med]
[edit protocols mld traceoptions]
[edit protocols mld-host traceoptions]
[edit protocols msdp group peer traceoptions]
[edit protocols msdp group traceoptions]
[edit protocols msdp peer traceoptions]
[edit protocols msdp traceoptions]
[edit protocols mstp]
[edit protocols mvrp]
[edit protocols oam]
[edit protocols ospf traceoptions]
[edit protocols pim traceoptions]
[edit protocols rip traceoptions]
[edit protocols ripng traceoptions]
[edit protocols router-advertisement traceoptions]
[edit protocols router-discovery traceoptions]
[edit protocols rsvp traceoptions]
```

```
[edit protocols sflow]
[edit protocols stp]
[edit protocols uplink-failure-detection]
[edit protocols vstp]
[edit routing-instances]
[edit routing-instances instance bridge-domains]
[edit routing-instances instance bridge-domains domain multicast-snooping-options]
[edit routing-instances instance bridge-domains domain multicast-snooping-options
traceoptions]
[edit routing-instances instance multicast-snooping-options]
[edit routing-instances instance multicast-snooping-options traceoptions]
[edit routing-instances instance pbb-options]
[edit routing-instances instance protocols]
[edit routing-instances instance protocols bgp group neighbor traceoptions]
[edit routing-instances instance protocols bgp group traceoptions]
[edit routing-instances instance protocols bgp traceoptions]
[edit routing-instances instance protocols esis traceoptions]
[edit routing-instances instance protocols isis traceoptions]
[edit routing-instances instance protocols l2vpn traceoptions]
[edit routing-instances instance protocols ldp traceoptions]
[edit routing-instances instance protocols msdp group peer traceoptions]
[edit routing-instances instance protocols msdp group traceoptions]
[edit routing-instances instance protocols msdp peer traceoptions]
[edit routing-instances instance protocols msdp traceoptions]
[edit routing-instances instance protocols mvpn traceoptions]
[edit routing-instances instance protocols ospf traceoptions]
[edit routing-instances instance protocols pim traceoptions]
[edit routing-instances instance protocols rip traceoptions]
[edit routing-instances instance protocols ripng traceoptions]
[edit routing-instances instance protocols router-discovery traceoptions]
[edit routing-instances instance protocols vlpls traceoptions]
[edit routing-instances instance routing-options]
[edit routing-instances instance routing-options validation group session traceoptions]
[edit routing-instances instance routing-options validation traceoptions]
[edit routing-instances instance routing-options multicast traceoptions]
[edit routing-instances instance routing-options traceoptions]
[edit routing-instances instance service-groups]
[edit routing-instances instance switch-options]
[edit routing-options]
[edit routing-options validation group session]
[edit routing-options multicast traceoptions]
[edit routing-options validation]
[edit routing-options traceoptions]
[edit switch-options]
```

**Related
Documentation**

- [Access Privilege User Permission Flags Overview on page 128](#)
- [Understanding Junos OS Access Privilege Levels on page 123](#)
- [Configuring Access Privilege Levels on page 131](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)
- [routing-control on page 201](#)

routing-control

Can view general routing, routing protocol, and routing policy configuration information and can configure general routing at the **[edit routing-options]** hierarchy level, routing protocols at the **[edit protocols]** hierarchy level, and routing policy at the **[edit policy-options]** hierarchy level.

Commands No associated CLI commands.

Configuration Hierarchy Levels	<ul style="list-style-type: none"> [edit bridge-domains] [edit bridge-domains domain multicast-snooping-options] [edit bridge-domains domain multicast-snooping-options traceoptions] [edit dynamic-profiles protocols igmp traceoptions] [edit dynamic-profiles protocols mld traceoptions] [edit dynamic-profiles protocols router-advertisement traceoptions] [edit dynamic-profiles routing-instances] [edit dynamic-profiles routing-instances instance bridge-domains] [edit dynamic-profiles routing-instances instance bridge-domains domain multicast-snooping-options] [edit dynamic-profiles routing-instances instance bridge-domains domain multicast-snooping-options traceoptions] [edit dynamic-profiles routing-instances instance multicast-snooping-options] [edit dynamic-profiles routing-instances instance multicast-snooping-options traceoptions] [edit dynamic-profiles routing-instances instance pbb-options] [edit dynamic-profiles routing-instances instance protocols] [edit dynamic-profiles routing-instances instance protocols bgp group neighbor traceoptions] [edit dynamic-profiles routing-instances instance protocols bgp group traceoptions] [edit dynamic-profiles routing-instances instance protocols bgp traceoptions] [edit dynamic-profiles routing-instances instance protocols esis traceoptions] [edit dynamic-profiles routing-instances instance protocols isis traceoptions] [edit dynamic-profiles routing-instances instance protocols l2vpn traceoptions] [edit dynamic-profiles routing-instances instance protocols ldp traceoptions] [edit dynamic-profiles routing-instances instance protocols msdp group peer traceoptions] [edit dynamic-profiles routing-instances instance protocols msdp group traceoptions] [edit dynamic-profiles routing-instances instance protocols msdp peer traceoptions] [edit dynamic-profiles routing-instances instance protocols msdp traceoptions] [edit dynamic-profiles routing-instances instance protocols mvpn traceoptions] [edit dynamic-profiles routing-instances instance protocols ospf traceoptions] [edit dynamic-profiles routing-instances instance protocols pim traceoptions] [edit dynamic-profiles routing-instances instance protocols rip traceoptions] [edit dynamic-profiles routing-instances instance protocols ripng traceoptions] [edit dynamic-profiles routing-instances instance protocols router-discovery traceoptions] [edit dynamic-profiles routing-instances instance protocols vpls traceoptions] [edit dynamic-profiles routing-instances instance routing-options] [edit dynamic-profiles routing-instances instance routing-options multicast traceoptions] [edit dynamic-profiles routing-instances instance routing-options traceoptions] [edit dynamic-profiles routing-instances instance service-groups] [edit dynamic-profiles routing-instances instance switch-options] [edit dynamic-profiles routing-options multicast traceoptions] [edit jnx-example] [edit fabric protocols]
---------------------------------------	--

```
[edit fabric protocols bgp group neighbor traceoptions]
[edit fabric protocols bgp group traceoptions]
[edit fabric protocols bgp traceoptions]
[edit fabric routing-instances]
[edit fabric routing-instances instance routing-options]
[edit fabric routing-instances instance routing-options traceoptions]
[edit fabric routing-options]
[edit fabric routing-options traceoptions]
[edit logical-systems bridge-domains]
[edit logical-systems bridge-domains domain multicast-snooping-options]
[edit logical-systems bridge-domains domain multicast-snooping-options traceoptions]
[edit logical-systems policy-options]
[edit logical-systems protocols]
[edit logical-systems protocols bgp group neighbor traceoptions]
[edit logical-systems protocols bgp group traceoptions]
[edit logical-systems protocols bgp traceoptions]
[edit logical-systems protocols dvmrp traceoptions]
[edit logical-systems protocols esis traceoptions]
[edit logical-systems protocols igmp traceoptions]
[edit logical-systems protocols igmp-host traceoptions]
[edit logical-systems protocols isis traceoptions]
[edit logical-systems protocols l2circuit traceoptions]
[edit logical-systems protocols l2iw traceoptions]
[edit logical-systems protocols ldp traceoptions]
[edit logical-systems protocols mld traceoptions]
[edit logical-systems protocols mld-host traceoptions]
[edit logical-systems protocols msdp group peer traceoptions]
[edit logical-systems protocols msdp group traceoptions]
[edit logical-systems protocols msdp peer traceoptions]
[edit logical-systems protocols msdp traceoptions]
[edit logical-systems protocols ospf traceoptions]
[edit logical-systems protocols pim traceoptions]
[edit logical-systems protocols rip traceoptions]
[edit logical-systems protocols ripng traceoptions]
[edit logical-systems protocols router-advertisement traceoptions]
[edit logical-systems protocols router-discovery traceoptions]
[edit logical-systems protocols rsvp traceoptions]
[edit logical-systems routing-instances]
[edit logical-systems routing-instances instance bridge-domains]
[edit logical-systems routing-instances instance bridge-domains domain
multicast-snooping-options]
[edit logical-systems routing-instances instance bridge-domains domain
multicast-snooping-options traceoptions]
[edit logical-systems routing-instances instance multicast-snooping-options]
[edit logical-systems routing-instances instance multicast-snooping-options
traceoptions]
[edit logical-systems routing-instances instance pbb-options]
[edit logical-systems routing-instances instance protocols]
[edit logical-systems routing-instances instance protocols bgp group neighbor
traceoptions]
[edit logical-systems routing-instances instance protocols bgp group traceoptions]
[edit logical-systems routing-instances instance protocols bgp traceoptions]
[edit logical-systems routing-instances instance protocols esis traceoptions]
[edit logical-systems routing-instances instance protocols isis traceoptions]
[edit logical-systems routing-instances instance protocols l2vpn traceoptions]
[edit logical-systems routing-instances instance protocols ldp traceoptions]
```



```
[edit logical-systems routing-instances instance protocols msdp group peer traceoptions]
[edit logical-systems routing-instances instance protocols msdp group traceoptions]
[edit logical-systems routing-instances instance protocols msdp peer traceoptions]
[edit logical-systems routing-instances instance protocols msdp traceoptions]
[edit logical-systems routing-instances instance protocols mvpn traceoptions]
[edit logical-systems routing-instances instance protocols ospf traceoptions]
[edit logical-systems routing-instances instance protocols pim traceoptions]
[edit logical-systems routing-instances instance protocols rip traceoptions]
[edit logical-systems routing-instances instance protocols ripng traceoptions]
[edit logical-systems routing-instances instance protocols router-discovery traceoptions]
[edit logical-systems routing-instances instance protocols vpls traceoptions]
[edit logical-systems routing-instances instance routing-options]
[edit logical-systems routing-instances instance routing-options multicast traceoptions]
[edit logical-systems routing-instances instance routing-options traceoptions]
[edit logical-systems routing-instances instance service-groups]
[edit logical-systems routing-instances instance switch-options]
[edit logical-systems routing-options]
[edit logical-systems routing-options multicast traceoptions]
[edit logical-systems routing-options traceoptions]
[edit logical-systems switch-options]
[edit multicast-snooping-options]
[edit multicast-snooping-options traceoptions]
[edit policy-options]
[edit protocols]
[edit protocols amt traceoptions]
[edit protocols bgp group neighbor traceoptions]
[edit protocols bgp group traceoptions]
[edit protocols bgp traceoptions]
[edit protocols connections][edit protocols dot1x]
[edit protocols dvmrp traceoptions]
[edit protocols esis traceoptions]
[edit protocols igmp traceoptions]
[edit protocols igmp-host traceoptions]
[edit protocols igmp-snooping]
[edit protocols isis traceoptions]
[edit protocols l2circuit traceoptions]
[edit protocols l2iw traceoptions]
[edit protocols ldp traceoptions]
[edit protocols lldp]
[edit protocols lldp-med]
[edit protocols mld traceoptions]
[edit protocols mld-host traceoptions]
[edit protocols msdp group peer traceoptions]
[edit protocols msdp group traceoptions]
[edit protocols msdp peer traceoptions]
[edit protocols msdp traceoptions]
[edit protocols mstp]
[edit protocols mvrp]
[edit protocols oam]
[edit protocols ospf traceoptions]
[edit protocols pim traceoptions]
[edit protocols ptp]
[edit protocols rip traceoptions]
[edit protocols ripng traceoptions]
[edit protocols router-advertisement traceoptions]
[edit protocols router-discovery traceoptions]
```

```
[edit protocols rsvp traceoptions]
[edit protocols sflow]
[edit protocols stp]
[edit protocols uplink-failure-detection]
[edit protocols vstp]
[edit routing-instances]
[edit routing-instances instance bridge-domains]
[edit routing-instances instance bridge-domains domain multicast-snooping-options]
[edit routing-instances instance bridge-domains domain multicast-snooping-options
traceoptions]
[edit routing-instances instance multicast-snooping-options]
[edit routing-instances instance multicast-snooping-options traceoptions]
[edit routing-instances instance pbb-options]
[edit routing-instances instance protocols]
[edit routing-instances instance protocols bgp group neighbor traceoptions]
[edit routing-instances instance protocols bgp group traceoptions]
[edit routing-instances instance protocols bgp traceoptions]
[edit routing-instances instance protocols esis traceoptions]
[edit routing-instances instance protocols isis traceoptions]
[edit routing-instances instance protocols l2vpn traceoptions]
[edit routing-instances instance protocols ldp traceoptions]
[edit routing-instances instance protocols msdp group peer traceoptions]
[edit routing-instances instance protocols msdp group traceoptions]
[edit routing-instances instance protocols msdp peer traceoptions]
[edit routing-instances instance protocols msdp traceoptions]
[edit routing-instances instance protocols mvpn traceoptions]
[edit routing-instances instance protocols ospf traceoptions]
[edit routing-instances instance protocols pim traceoptions]
[edit routing-instances instance protocols rip traceoptions]
[edit routing-instances instance protocols ripng traceoptions]
[edit routing-instances instance protocols router-discovery traceoptions]
[edit routing-instances instance protocols vpls traceoptions]
[edit routing-instances instance routing-options]
[edit routing-instances instance routing-options multicast traceoptions]
[edit routing-instances instance routing-options traceoptions]
[edit routing-instances instance service-groups]
[edit routing-instances instance switch-options]
[edit routing-options]
[edit routing-options multicast traceoptions]
[edit routing-options traceoptions]
[edit switch-options]
```

**Related
Documentation**

- [Access Privilege User Permission Flags Overview on page 128](#)
- [Understanding Junos OS Access Privilege Levels on page 123](#)
- [Configuring Access Privilege Levels on page 131](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)
- [routing on page 196](#)

secret

Can view passwords and other authentication keys in the configuration.

Commands No associated CLI commands.

Configuration Hierarchy Levels

```
[edit access profile client chap-secret]
[edit access profile client firewall-user password]
[edit access profile client l2tp shared-secret]
[edit access profile client pap-password]
[edit access profile radius-server secret]
[edit access radius-disconnect secret]
[edit dynamic-profiles interfaces interface ppp-options chap default-chap-secret]
[edit dynamic-profiles interfaces interface ppp-options pap default-password]
[edit dynamic-profiles interfaces interface ppp-options pap local-password]
[edit dynamic-profiles interfaces interface unit ppp-options chap default-chap-secret]
[edit dynamic-profiles interfaces interface unit ppp-options pap default-password]
[edit dynamic-profiles interfaces interface unit ppp-options pap local-password]
[edit interfaces interface ppp-options chap default-chap-secret]
[edit interfaces interface ppp-options pap default-password]
[edit interfaces interface ppp-options pap local-password]
[edit interfaces interface unit ppp-options chap default-chap-secret]
[edit interfaces interface unit ppp-options pap default-password]
[edit interfaces interface unit ppp-options pap local-password]
[edit logical-systems interfaces interface unit ppp-options chap]
[edit logical-systems interfaces interface unit ppp-options pap default-password]
[edit logical-systems interfaces interface unit ppp-options pap local-password]
[edit logical-systems routing-instances instance system services static-subscribers
authentication password]
[edit logical-systems routing-instances instance system services static-subscribers group
authentication password]
[edit logical-systems system services static-subscribers authentication password]
[edit logical-systems system services static-subscribers group authentication password]
[edit routing-instances instance system services static-subscribers authentication
password]
[edit routing-instances instance system services static-subscribers group authentication
password]
[edit services ggsn apn radius accounting server secret]
[edit services ggsn apn radius authentication server secret]
[edit services ggsn radius server secret]
[edit system accounting destination radius server secret]
[edit system accounting destination tacplus server secret]
[edit system radius-server secret]
[edit system services outbound-ssh client secret]
[edit system services packet-triggered-subscribers partition-radius
accounting-shared-secret]
[edit system services static-subscribers authentication password]
[edit system services static-subscribers group authentication password]
[edit system tacplus-server secret]
```

Related Documentation

- [Access Privilege User Permission Flags Overview on page 128](#)
- [Understanding Junos OS Access Privilege Levels on page 123](#)
- [Configuring Access Privilege Levels on page 131](#)

- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)
- [secret-control on page 206](#)

secret-control

Can view passwords and other authentication keys in the configuration and can modify them in configuration mode.

Commands No associated CLI commands.

Configuration Hierarchy Levels	<pre>[edit access profile client chap-secret] [edit access profile client firewall-user password] [edit access profile client l2tp shared-secret] [edit access profile client pap-password] [edit access profile radius-server secret] [edit access radius-disconnect secret] [edit dynamic-profiles interfaces interface ppp-options chap default-chap-secret] [edit dynamic-profiles interfaces interface ppp-options pap default-password] [edit dynamic-profiles interfaces interface ppp-options pap local-password] [edit dynamic-profiles interfaces interface unit ppp-options chap default-chap-secret] [edit dynamic-profiles interfaces interface unit ppp-options pap default-password] [edit dynamic-profiles interfaces interface unit ppp-options pap local-password] [edit interfaces interface ppp-options chap default-chap-secret] [edit interfaces interface ppp-options pap default-password] [edit interfaces interface ppp-options pap local-password] [edit interfaces interface unit ppp-options chap default-chap-secret] [edit interfaces interface unit ppp-options pap default-password] [edit interfaces interface unit ppp-options pap local-password] [edit logical-systems interfaces interface unit ppp-options chap] [edit logical-systems interfaces interface unit ppp-options pap default-password] [edit logical-systems interfaces interface unit ppp-options pap local-password] [edit logical-systems routing-instances instance system services static-subscribers authentication password] [edit logical-systems routing-instances instance system services static-subscribers group authentication password] [edit logical-systems system services static-subscribers authentication password] [edit logical-systems system services static-subscribers group authentication password] [edit routing-instances instance system services static-subscribers authentication password] [edit routing-instances instance system services static-subscribers group authentication password] [edit services ggsn apn radius accounting server secret] [edit services ggsn apn radius authentication server secret] [edit services ggsn radius server secret] [edit system accounting destination radius server secret] [edit system accounting destination tacplus server secret] [edit system radius-server secret] [edit system services outbound-ssh client secret] [edit system services packet-triggered-subscribers partition-radius accounting-shared-secret] [edit system services static-subscribers authentication password] [edit system services static-subscribers group authentication password]</pre>
---------------------------------------	---

[edit system tacplus-server secret]

Related Documentation

- [Access Privilege User Permission Flags Overview on page 128](#)
- [Understanding Junos OS Access Privilege Levels on page 123](#)
- [Configuring Access Privilege Levels on page 131](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)
- [secret on page 205](#)

security

Can view security configuration.

Commands

```
clear security
clear security alarms
  <clear-security-alarm-information>
clear security idp
clear security idp application-ddos
clear security idp application-ddos cache
  <clear-idp-appddos-cache>

clear security idp application-identification
clear security idp application-identification application-system-cache
  <clear-idp-application-system-cache>

clear security idp application-statistics
  <clear-idp-applications-information>

clear security idp attack
clear security idp attack table
  <clear-idp-attack-table>

clear security idp counters
  <clear-idp-counters-by-counter-class>

clear security idp ssl-inspection
clear security idp ssl-inspection session-id-cache
  <clear-idp-ssl-session-cache-information>
clear security idp status
  <clear-idp-status-information>
clear security log
  <clear-security-log-information>
clear security pki
clear security pki ca-certificate
  <clear-pki-ca-certificate>
clear security pki certificate-request
  <clear-pki-certificate-request>
clear security pki crl
  <clear-pki-crl>
clear security pki key-pair
```

```
<clear-pki-key-pair>
clear security pki local-certificate
  <clear-pki-local-certificate>
request security
request security certificate
request security certificate enroll
request security datapath-debug
request security datapath-debug action-profile
request security datapath-debug action-profile reload-all
request security idp
  <request-idp-policy-load>
request security idp security-package
request security idp security-package download
  <request-idp-security-package-download>

request security idp security-package download version
  <request-idp-security-package-download-version>

request security idp security-package install
  <request-idp-security-package-install>

request security idp ssl-inspection
request security idp ssl-inspection key
request security idp ssl-inspection key add
  <request-idp-ssl-key-add>

request security idp ssl-inspection key delete
  <request-idp-ssl-key-delete>
request security idp storage-cleanup
  <request-idp-storage-cleanup>
request security key-pair
request security pki
request security pki ca-certificate
request security pki ca-certificate verify
  <verify-pki-ca-certificate>
request security pki ca-certificate enroll
request security pki ca-certificate ca-profile-group
request security pki ca-certificate ca-profile-group load
request security pki ca-certificate load
  <load-pki-ca-certificate>
request security pki crt
request security pki crt load
  <request security pki crt load>
request security pki generate-certificate-request
  <generate-pki-certificate-request>
request security pki generate-key-pair
  <generate-pki-key-pair>
request security pki local-certificate
request security pki local-certificate verify
  <verify-pki-local-certificate>
request security pki verify-integrity-status
  <verify-integrity-status>
request security pki local-certificate enroll
request security pki local-certificate export
request security pki local-certificate generate-self-signed
  <generate-pki-self-signed-local-certificate>
```

```
request security pki local-certificate load
  <load-pki-local-certificate>
request system set-encryption-key
show security
show security alarms
  <get-security-alarm-information>
show security idp
show security idp application-ddos
show security idp application-ddos application
  <get-idp-addos-application-information>

show security idp application-identification
show security idp application-identification application-system-cache
  <get-idp-application-system-cache>

show security idp application-statistics
  <get-idp-applications-information>

show security idp attack
show security idp attack description
  <get-idp-attack-description-information>
show security idp attack detail
  <get-idp-attack-detail-information>
show security idp attack table
  <get-idp-attack-table-information>

show security idp counters
  <get-idp-counter-information>

show security idp logical-system
show security idp logical-system policy-association
show security idp memory
  <get-idp-memory-information>

show security idp policies
  <get-idp-subscriber-policy-list>

show security idp policy-templates-list
  <get-idp-policy-template-information>
  <get-idp-predefined-attack-groups>
  <get-idp-predefined-attack-group-filters>
  <get-idp-predefined-attacks>
  <get-idp-predefined-attack-filters>
  <get-idp-recent-security-package-information>
show security idp policy-commit-status
  <get-idp-policy-commit-status>

<get-idp-recent-security-package-information>

show security idp security-package-version
  <get-idp-security-package-information>

show security idp ssl-inspection
show security idp ssl-inspection key
  <get-idp-ssl-key-information>
```

```
show security idp ssl-inspection session-id-cache  
  <get-idp-ssl-session-cache-information>
```

```
show security idp status  
  <get-idp-status-information>
```

```
show security idp status detail  
  <get-idp-detail-status-information>
```

```
show security keychain  
  <get-hakr-keychain-information>
```

```
show security log  
  <get-security-log-information>
```

```
show security pki  
show security pki ca-certificate  
  <get-pki-ca-certificate>  
show security pki certificate-request  
  <get-pki-certificate-request>  
show security pki crl  
  <get-pki-crl>  
show security pki local-certificate  
  <get-pki-local-certificate>
```

**Configuration
Hierarchy Levels**

```
[edit security]  
[edit security alarms]  
[edit security log]
```

**Related
Documentation**

- [Access Privilege User Permission Flags Overview on page 128](#)
- [Understanding Junos OS Access Privilege Levels on page 123](#)
- [Configuring Access Privilege Levels on page 131](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)
- [security-control on page 210](#)

security-control

Can view and configure security information at the **[edit security]** hierarchy level.

Commands

```
clear security  
clear security alarms  
  <clear-security-alarm-information>  
clear security idp  
clear security idp application-ddos  
clear security idp application-ddos cache  
  <clear-idp-appddos-cache>
```

```
clear security idp application-identification  
clear security idp application-identification application-system-cache  
  <clear-idp-application-system-cache>
```



```
clear security idp application-statistics
  <clear-idp-applications-information>

clear security idp attack
clear security idp attack table
  <clear-idp-attack-table>

clear security idp counters
  <clear-idp-counters-by-counter-class>

clear security idp ssl-inspection
clear security idp ssl-inspection session-id-cache
  <clear-idp-ssl-session-cache-information>
clear security idp status
  <clear-idp-status-information>
clear security log
  <clear-security-log-information>
clear security pki
clear security pki ca-certificate
  <clear-pki-ca-certificate>
clear security pki certificate-request
  <clear-pki-certificate-request>
clear security pki crl
  <clear-pki-crl>
clear security pki key-pair
  <clear-pki-key-pair>
clear security pki local-certificate
  <clear-pki-local-certificate>
request security
request security certificate
request security certificate enroll
request security datapath-debug
request security datapath-debug action-profile
request security datapath-debug action-profile reload-all
request security idp
  <request-idp-policy-load>
request security idp security-package
request security idp security-package download
  <request-idp-security-package-download>

request security idp security-package download version
  <request-idp-security-package-download-version>

request security idp security-package install
  <request-idp-security-package-install>

request security idp ssl-inspection
request security idp ssl-inspection key
request security idp ssl-inspection key add
  <request-idp-ssl-key-add>

request security idp ssl-inspection key delete
  <request-idp-ssl-key-delete>
request security idp storage-cleanup
  <request-idp-storage-cleanup>
request security key-pair
```

```
request security pki
request security pki ca-certificate
request security pki ca-certificate verify
  <verify-pki-ca-certificate>
request security pki ca-certificate enroll
request security pki ca-certificate load
  <load-pki-ca-certificate>
request security pki crl
request security pki crl load
  <request security pki crl load>
request security pki generate-certificate-request
  <generate-pki-certificate-request>
request security pki generate-key-pair
  <generate-pki-key-pair>
request security pki local-certificate
request security pki local-certificate verify
  <verify-pki-local-certificate>
request security pki local-certificate enroll
request security pki local-certificate generate-self-signed
  <generate-pki-self-signed-local-certificate>
request security pki local-certificate load
  <load-pki-local-certificate>
request system set-encryption-key
show security
show security alarms
  <get-security-alarm-information>
show security idp
show security idp application-ddos
show security idp application-ddos application
  <get-idp-addos-application-information>

show security idp application-identification
show security idp application-identification application-system-cache
  <get-idp-application-system-cache>

show security idp application-statistics
  <get-idp-applications-information>

show security idp attack
show security idp attack description
  <get-idp-attack-description-information>
show security idp attack detail
  <get-idp-attack-detail-information>
show security idp attack table
  <get-idp-attack-table-information>

show security idp counters
  <get-idp-counter-information>

show security idp logical-system
show security idp logical-system policy-association
show security idp memory
  <get-idp-memory-information>

show security idp policies
  <get-idp-subscriber-policy-list>
```

```

show security idp policy-templates-list
  <get-idp-policy-template-information>
  <get-idp-predefined-attack-groups>
  <get-idp-predefined-attack-group-filters>
  <get-idp-predefined-attacks>
  <get-idp-predefined-attack-filters>
  <get-idp-recent-security-package-information>
show security idp policy-commit-status
  <get-idp-policy-commit-status>

<get-idp-recent-security-package-information>

show security idp security-package-version
  <get-idp-security-package-information>

show security idp ssl-inspection
show security idp ssl-inspection key
  <get-idp-ssl-key-information>

show security idp ssl-inspection session-id-cache
  <get-idp-ssl-session-cache-information>

show security idp status
  <get-idp-status-information>

show security idp status detail
  <get-idp-detail-status-information>
show security keychain
  <get-hakr-keychain-information>
show security log
  <get-security-log-information>

show security pki
show security pki ca-certificate
  <get-pki-ca-certificate>
show security pki certificate-request
  <get-pki-certificate-request>
show security pki crl
  <get-pki-crl>
show security pki local-certificate
  <get-pki-local-certificate>

```

**Configuration
Hierarchy Levels**

```

[edit security]
[edit security alarms]
[edit security log]

```

**Related
Documentation**

- [Access Privilege User Permission Flags Overview on page 128](#)
- [Understanding Junos OS Access Privilege Levels on page 123](#)
- [Configuring Access Privilege Levels on page 131](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)

- [security on page 207](#)

shell

Can start a local shell on the router.

Commands start shell
 start shell user

Configuration Hierarchy Levels No associated CLI configuration hierarchy levels and statements.

- Related Documentation**
- [Access Privilege User Permission Flags Overview on page 128](#)
 - [Understanding Junos OS Access Privilege Levels on page 123](#)
 - [Configuring Access Privilege Levels on page 131](#)
 - [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
 - [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)

snmp

Can view Simple Network Management Protocol (SNMP) configuration.

Commands No associated CLI commands.

Configuration Hierarchy Levels [edit snmp]

- Related Documentation**
- [Access Privilege User Permission Flags Overview on page 128](#)
 - [Understanding Junos OS Access Privilege Levels on page 123](#)
 - [Configuring Access Privilege Levels on page 131](#)
 - [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
 - [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)
 - [snmp-control on page 214](#)

snmp-control

Can view SNMP configuration information and can modify SNMP configuration at the [edit snmp] hierarchy level.

Commands No associated CLI commands.

Configuration Hierarchy Levels [edit snmp]

- Related Documentation**
- [Access Privilege User Permission Flags Overview on page 128](#)
 - [Understanding Junos OS Access Privilege Levels on page 123](#)

- [Configuring Access Privilege Levels on page 131](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)
- [snmp on page 214](#)

system

Can view system-level configuration information.

Commands	request chassis synchronization request chassis synchronization force request chassis synchronization force automatic-switching request chassis synchronization force mark-failed request chassis synchronization force unmark-failed request chassis synchronization switch
Configuration Hierarchy Levels	[edit applications] [edit chassis system-domains] [edit dynamic-profiles routing-instances instance forwarding-options helpers tftp] [edit dynamic-profiles routing-instances instance routing-options fate-sharing] [edit ethernet-switching-options] [edit forwarding-options helpers bootp] [edit forwarding-options helpers domain] [edit forwarding-options helpers port] [edit forwarding-options helpers tftp] [edit logical-systems] [edit logical-systems routing-instances instance forwarding-options helpers bootp] [edit logical-systems routing-instances instance forwarding-options helpers domain] [edit logical-systems routing-instances instance forwarding-options helpers port] [edit logical-systems routing-instances instance forwarding-options helpers tftp] [edit logical-systems routing-instances instance routing-options fate-sharing] [edit logical-systems routing-options fate-sharing] [edit logical-systems system] [edit logical-systems system syslog] [edit poe] [edit routing-instances instance forwarding-options helpers bootp] [edit routing-instances instance forwarding-options helpers domain] [edit routing-instances instance forwarding-options helpers port] [edit routing-instances instance forwarding-options helpers tftp] [edit routing-instances instance routing-options fate-sharing] [edit routing-options fate-sharing] [edit services] [edit services ggsn charging charging-log traceoptions] [edit system] [edit system archival] [edit system backup-router] [edit system compress-configuration-files] [edit system default-address-selection] [edit system domain-name] [edit system domain-search] [edit system encrypt-configuration-files] [edit system host-name] [edit system inet6-backup-router]

```
[edit system internet-options gre-path-mtu-discovery]
[edit system internet-options ipip-path-mtu-discovery]
[edit system internet-options ipv6-path-mtu-discovery]
[edit system internet-options ipv6-path-mtu-discovery-timeout]
[edit system internet-options ipv6-reject-zero-hop-limit]
[edit system internet-options no-tcp-reset]
[edit system internet-options no-tcp-rfc1323]
[edit system internet-options no-tcp-rfc1323-paws]
[edit system internet-options path-mtu-discovery]
[edit system internet-options source-port upper-limit]
[edit system internet-options source-quench]
[edit system internet-options tcp-drop-synfin-set]
[edit system internet-options tcp-mss]
[edit system license]
[edit system max-configuration-rollbacks]
[edit system max-configurations-on-flash]
[edit system mirror-flash-on-disk]
[edit system no-debugger-on-alt-break]
[edit system no-redirects-ipv6]
[edit system name-server]
[edit system no-multicast-echo]
[edit system no-neighbor-learn]
[edit system no-redirects]
[edit system ports auxiliary log-out-on-disconnect]
[edit system ports auxiliary port-type]
[edit system ports console log-out-on-disconnect]
[edit system ports console port-type]
[edit system processes]
[edit system proxy]
[edit system saved-core-context]
[edit system saved-core-files]
[edit system services]
[edit system services web-management]
[edit system static-host-mapping]
[edit system syslog]
[edit system time-zone]
[edit virtual-chassis]
[edit vlans]
```

**Related
Documentation**

- [Access Privilege User Permission Flags Overview on page 128](#)
- [Understanding Junos OS Access Privilege Levels on page 123](#)
- [Configuring Access Privilege Levels on page 131](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)
- [system-control on page 216](#)

system-control

Can view system-level configuration information and configure it at the **[edit system]** hierarchy level.

Configuration Hierarchy Levels

```
[edit applications]
[edit chassis system-domains]
[edit dynamic-profiles routing-instances instance forwarding-options helpers tftp]
[edit dynamic-profiles routing-instances instance routing-options fate-sharing]
[edit ethernet-switching-options]
[edit forwarding-options helpers bootp]
[edit forwarding-options helpers domain]
[edit forwarding-options helpers port]
[edit forwarding-options helpers tftp]
[edit logical-systems]
[edit logical-systems routing-instances instance forwarding-options helpers bootp]
[edit logical-systems routing-instances instance forwarding-options helpers domain]
[edit logical-systems routing-instances instance forwarding-options helpers port]
[edit logical-systems routing-instances instance forwarding-options helpers tftp]
[edit logical-systems routing-instances instance routing-options fate-sharing]
[edit logical-systems routing-options fate-sharing]
[edit logical-systems system]
[edit poe]
[edit routing-instances instance forwarding-options helpers bootp]
[edit routing-instances instance forwarding-options helpers domain]
[edit routing-instances instance forwarding-options helpers port]
[edit routing-instances instance forwarding-options helpers tftp]
[edit routing-instances instance routing-options fate-sharing]
[edit routing-options fate-sharing]
[edit services]
[edit services ggsn charging charging-log traceoptions]
[edit system]
[edit system archival]
[edit system backup-router]
[edit system compress-configuration-files]
[edit system default-address-selection]
[edit system domain-name]
[edit system domain-search]
[edit system encrypt-configuration-files]
[edit system host-name]
[edit system inet6-backup-router]
[edit system internet-options gre-path-mtu-discovery]
[edit system internet-options ipip-path-mtu-discovery]
[edit system internet-options ipv6-path-mtu-discovery]
[edit system internet-options ipv6-path-mtu-discovery-timeout]
[edit system internet-options ipv6-reject-zero-hop-limit]
[edit system internet-options no-tcp-reset]
[edit system internet-options no-tcp-rfc1323]
[edit system internet-options no-tcp-rfc1323-paws]
[edit system internet-options path-mtu-discovery]
[edit system internet-options source-port upper-limit]
[edit system internet-options source-quench]
[edit system internet-options tcp-drop-synfin-set]
[edit system internet-options tcp-mss]
[edit system license]
[edit system max-configuration-rollback]
[edit system max-configurations-on-flash]
[edit system mirror-flash-on-disk]
[edit system name-server]
[edit system no-multicast-echo]
[edit system no-neighbor-learn]
```

```
[edit system no-redirects]
[edit system ports auxiliary log-out-on-disconnect]
[edit system ports auxiliary port-type]
[edit system ports console log-out-on-disconnect]
[edit system ports console port-type]
[edit system processes]
[edit system saved-core-context]
[edit system saved-core-files]
[edit system services]
[edit system services web-management]
[edit system static-host-mapping]
[edit system syslog]
[edit system time-zone]
[edit virtual-chassis]
[edit vlans]
```

**Related
Documentation**

- [Access Privilege User Permission Flags Overview on page 128](#)
- [Understanding Junos OS Access Privilege Levels on page 123](#)
- [Configuring Access Privilege Levels on page 131](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)
- [system on page 215](#)

trace

Can view trace file settings and configure trace file properties.

Commands

```
clear log
  <clear-log>
monitor
request-monitor-ethernet-delay-measurement
  <request-monitor-ethernet-loss-measurement>
monitor interface
monitor interface traffic
monitor label-switched-path
monitor list
monitor start
monitor static-lsp
monitor stop
show log
  <get-log>
show log user
  <get-syslog-events>
```

**Configuration
Hierarchy Levels**

```
[edit bridge-domains domain multicast-snooping-options traceoptions]
[edit bridge-domains domain protocols igmp-snooping]
[edit bridge-domains domain forwarding-options dhcp-relay traceoptions]
[edit bridge-domains domain protocols igmp-snooping traceoptions]
[edit bridge-domains domain forwarding-options dhcp-relay interface-traceoptions]
```



```

[edit bridge-domains domain multicast-snooping-options traceoptions]
[edit bridge-domains domain protocols igmp-snooping traceoptions]
[edit class-of-service application-traffic-control traceoptions]
[edit demux traceoptions]
[edit dynamic-profiles protocols igmp traceoptions]
[edit dynamic-profiles protocols mld traceoptions]
[edit dynamic-profiles class-of-service application-traffic-control traceoptions]
[edit dynamic-profiles protocols oam ethernet link-fault-management traceoptions]
[dynamic-profiles protocols oam ethernet lmi]
[edit dynamic-profiles protocols router-advertisement traceoptions]
[edit dynamic-profiles protocols oam gre-tunnel traceoptions]
[edit dynamic-profiles routing-instances instance bridge-domains domain
forwarding-options dhcp-relay traceoptions]
[edit dynamic-profiles routing-instances instance bridge-domains domain
multicast-snooping-options traceoptions]
[edit dynamic-profiles routing-instances instance bridge-domains domain protocols
igmp-snooping traceoptions]
[edit dynamic-profiles routing-instances instance forwarding-options dhcp-relay
traceoptions]
[edit dynamic-profiles routing-instances instance multicast-snooping-options
traceoptions]
[edit dynamic-profiles routing-instances instance protocols bgp group neighbor
traceoptions]
[edit dynamic-profiles routing-instances instance protocols bgp group traceoptions]
[edit dynamic-profiles routing-instances instance protocols bgp traceoptions]
[edit dynamic-profiles routing-instances instance protocols esis traceoptions]
[edit dynamic-profiles routing-instances instance protocols igmp-snooping traceoptions]
[edit dynamic-profiles routing-instances instance protocols isis traceoptions]
[edit dynamic-profiles routing-instances instance protocols l2vpn traceoptions]
[edit dynamic-profiles routing-instances instance protocols ldp traceoptions]
[edit dynamic-profiles routing-instances instance protocols msdp group peer
traceoptions]
[edit dynamic-profiles routing-instances instance protocols msdp group traceoptions]
[edit dynamic-profiles routing-instances instance protocols msdp peer traceoptions]
[edit dynamic-profiles routing-instances instance protocols msdp traceoptions]
[edit dynamic-profiles routing-instances instance protocols mvpn traceoptions]
[edit dynamic-profiles routing-instances instance protocols ospf traceoptions]
[edit dynamic-profiles routing-instances instance protocols pim traceoptions]
[edit dynamic-profiles routing-instances instance protocols pim-snooping traceoptions]
[edit dynamic-profiles routing-instances instance protocols rip traceoptions]
[edit dynamic-profiles routing-instances instance protocols ripng traceoptions]
[edit dynamic-profiles routing-instances instance protocols router-discovery traceoptions]
[edit dynamic-profiles routing-instances instance protocols vpls traceoptions]
[edit dynamic-profiles routing-instances instance routing-options multicast traceoptions]
[edit dynamic-profiles routing-instances instance routing-options traceoptions]
[edit dynamic-profiles routing-instances instance services mobile-ip traceoptions]
[edit dynamic-profiles routing-instances instance system services dhcp-local-server
traceoptions]
[edit dynamic-profiles routing-options multicast traceoptions]
[edit fabric protocols bgp group neighbor traceoptions]
[edit fabric protocols bgp group traceoptions]
[edit fabric protocols bgp traceoptions]
[edit fabric routing-instances instance routing-options traceoptions]
[edit fabric routing-options traceoptions]
[edit jnx-example traceoptions]
[edit logical-systems bridge-domains domain forwarding-options dhcp-relay

```

```
traceoptions]
[edit logical-systems bridge-domains domain forwarding-options dhcp-relay
interface-traceoptions]
[edit logical-systems bridge-domains domain multicast-snooping-options traceoptions]
[edit logical-systems bridge-domains domain protocols igmp-snooping traceoptions]
[edit logical-systems forwarding-options dhcp-relay traceoptions]
[edit logical-systems protocols ancp traceoptions]
[edit logical-systems protocols bgp group neighbor traceoptions]
[edit logical-systems protocols bgp group traceoptions]
[edit logical-systems protocols bgp traceoptions]
[edit logical-systems protocols dot1x traceoptions]
[edit logical-systems protocols dvmrp traceoptions]
[edit logical-systems protocols esis traceoptions]
[edit logical-systems protocols igmp traceoptions]
[edit logical-systems protocols igmp-host traceoptions]
[edit logical-systems protocols ilmi traceoptions]
[edit logical-systems protocols igmp-snooping traceoptions]
[edit logical-systems protocols isis traceoptions]
[edit logical-systems protocols l2circuit traceoptions]
[edit logical-systems protocols l2iw traceoptions]
[edit logical-systems protocols lacp traceoptions]
[edit logical-systems protocols layer2-control traceoptions]
[edit logical-systems protocols ldp traceoptions]
[edit logical-systems protocols mld traceoptions]
[edit dynamic-profiles protocols oam ethernet fnp traceoptions]
[edit logical-systems protocols mld-host traceoptions]
[edit logical-systems protocols mpls label-switched-path oam traceoptions]
[edit logical-systems protocols mpls label-switched-path primary oam traceoptions]
[edit logical-systems protocols mpls label-switched-path secondary oam traceoptions]
[edit logical-systems protocols mpls oam traceoptions]
[edit logical-systems protocols msdp group peer traceoptions]
[edit logical-systems protocols msdp group traceoptions]
[edit logical-systems protocols msdp peer traceoptions]
[edit logical-systems protocols msdp traceoptions]
[edit logical-systems protocols neighbor-discovery secure traceoptions]
[edit logical-systems protocols oam ethernet fnp traceoptions]
[edit logical-systems protocols oam ethernet link-fault-management traceoptions]
[edit logical-systems protocols oam ethernet lmi traceoptions]
[edit logical-systems protocols ospf traceoptions]
[edit logical-systems protocols pcep pce traceoptions]
[edit logical-systems protocols pcep pce-group traceoptions]
[edit logical-systems protocols pcep traceoptions]
[edit logical-systems protocols pim traceoptions]
[edit logical-systems protocols pim-snooping traceoptions]
[edit logical-systems protocols ppp monitor-session]
[edit logical-systems protocols ppp traceoptions]
[edit logical-systems protocols ppp-service traceoptions]
[edit logical-systems protocols pppoe traceoptions]
[edit logical-systems protocols rip traceoptions]
[edit logical-systems protocols ripng traceoptions]
[edit logical-systems protocols router-advertisement traceoptions]
[edit logical-systems protocols router-discovery traceoptions]
[edit logical-systems protocols rsvp lsp-set traceoptions]
[edit logical-systems protocols rsvp traceoptions]
[edit logical-systems routing-instances instance bridge-domains domain
multicast-snooping-options traceoptions]
```

```
[edit logical-systems routing-instances instance bridge-domains domain protocols
igmp-snooping traceoptions]
[edit logical-systems routing-instances instance forwarding-options dhcp-relay
traceoptions]
[edit logical-systems routing-instances instance multicast-snooping-options
traceoptions]
[edit logical-systems routing-instances instance protocols bgp group neighbor
traceoptions]
[edit logical-systems routing-instances instance protocols bgp group traceoptions]
[edit logical-systems routing-instances instance protocols bgp traceoptions]
[edit logical-systems routing-instances instance protocols esis traceoptions]
[edit logical-systems routing-instances instance protocols igmp-snooping traceoptions]
[edit logical-systems routing-instances instance protocols isis traceoptions]
[edit logical-systems routing-instances instance protocols l2vpn traceoptions]
[edit logical-systems routing-instances instance protocols ldp traceoptions]
[edit logical-systems routing-instances instance protocols msdp group peer traceoptions]
[edit logical-systems routing-instances instance protocols msdp group traceoptions]
[edit logical-systems routing-instances instance protocols msdp peer traceoptions]
[edit logical-systems routing-instances instance protocols msdp traceoptions]
[edit logical-systems routing-instances instance protocols mvpn traceoptions]
[edit logical-systems routing-instances instance protocols ospf traceoptions]
[edit logical-systems routing-instances instance protocols pim traceoptions]
[edit logical-systems routing-instances instance protocols rip traceoptions]
[edit logical-systems routing-instances instance protocols ripng traceoptions]
[edit logical-systems routing-instances instance protocols router-discovery traceoptions]
[edit logical-systems routing-instances instance protocols vpls traceoptions]
[edit logical-systems routing-instances instance routing-options multicast traceoptions]
[edit logical-systems routing-instances instance routing-options validation group session]
[edit logical-systems routing-instances instance routing-options validation traceoptions]
[edit logical-systems routing-instances instance routing-options traceoptions]
[edit logical-systems routing-instances instance services mobile-ip traceoptions]
[edit logical-systems routing-instances instance system services dhcp-local-server
traceoptions]
[edit logical-systems routing-instances instance system services dhcp-local-server
interface-traceoptions]
[edit logical-systems routing-options multicast traceoptions]
[edit logical-systems routing-options traceoptions]
[edit logical-systems services mobile-ip traceoptions]
[edit logical-systems system services dhcp-local-server traceoptions]
[edit logical-systems system services dhcp-local-server interface-traceoptions]
[edit multicast-snooping-options traceoptions]
[edit protocols ancp traceoptions]
[edit protocols bgp group neighbor traceoptions]
[edit protocols bgp group traceoptions]
[edit protocols bgp traceoptions]
[edit protocols dot1x traceoptions]
[edit protocols dvmrp traceoptions]
[edit protocols esis traceoptions]
[edit protocols igmp traceoptions]
[edit protocols igmp-host traceoptions]
[edit protocols igmp-snooping traceoptions]
[edit protocols ilmi traceoptions]
[edit protocols isis traceoptions]
[edit protocols l2circuit traceoptions]
[edit protocols l2iw traceoptions]
[edit protocols lacp traceoptions]
```

```
[edit protocols layer2-control traceoptions]
[edit protocols ldp traceoptions]
[edit protocols mld traceoptions]
[edit protocols mld-host traceoptions]
[edit protocols mpls label-switched-path oam traceoptions]
[edit protocols mpls label-switched-path primary oam traceoptions]
[edit protocols mpls label-switched-path secondary oam traceoptions]
[edit protocols mpls oam traceoptions]
[edit protocols msdp group peer traceoptions]
[edit protocols msdp group traceoptions]
[edit protocols msdp peer traceoptions]
[edit protocols msdp traceoptions]
[edit protocols neighbor-discovery secure traceoptions]
[edit protocols protocols oam ethernet fnp]
[edit protocols oam ethernet connectivity-fault-management traceoptions]
[edit protocols oam ethernet link-fault-management traceoptions]
[edit protocols oam ethernet lmi traceoptions]
[edit protocols ospf traceoptions]
[edit protocols pim traceoptions]
[edit protocols ppp monitor-session]
[edit protocols ppp traceoptions]
[edit protocols ppp-service traceoptions]
[edit protocols pppoe traceoptions]
[edit protocols rip traceoptions]
[edit protocols ripng traceoptions]
[edit protocols router-advertisement traceoptions]
[edit protocols router-discovery traceoptions]
[edit protocols rsvp lsp-set traceoptions]
[edit protocols rsvp traceoptions]
[edit routing-instances instance bridge-domains domain multicast-snooping-options
traceoptions]
[edit routing-instances instance bridge-domains domain protocols igmp-snooping
traceoptions]
[edit routing-instances instance multicast-snooping-options traceoptions]
[edit routing-instances instance protocols bgp group neighbor traceoptions]
[edit routing-instances instance protocols bgp group traceoptions]
[edit routing-instances instance protocols bgp traceoptions]
[edit routing-instances instance protocols esis traceoptions]
[edit routing-instances instance protocols igmp-snooping traceoptions]
[edit routing-instances instance protocols isis traceoptions]
[edit routing-instances instance protocols l2vpn traceoptions]
[edit routing-instances instance protocols ldp traceoptions]
[edit routing-instances instance protocols msdp group peer traceoptions]
[edit routing-instances instance protocols msdp group traceoptions]
[edit routing-instances instance protocols msdp peer traceoptions]
[edit routing-instances instance protocols msdp traceoptions]
[edit routing-instances instance protocols mvpn traceoptions]
[edit routing-instances instance protocols ospf traceoptions]
[edit routing-instances instance protocols pim traceoptions]
[edit routing-instances instance protocols pim-snooping traceoptions]
[edit routing-instances instance protocols rip traceoptions]
[edit routing-instances instance protocols ripng traceoptions]
[edit routing-instances instance protocols router-discovery traceoptions]
[edit routing-instances instance protocols vpls traceoptions]
[edit routing-instances instance routing-options multicast traceoptions]
[edit routing-instances instance routing-options traceoptions]
```

```

[edit routing-options multicast traceoptions]
[edit routing-options traceoptions]
[edit security idp traceoptions]
[edit security pki traceoptions]
[edit services adaptive-services-pics traceoptions]
[edit services captive-portal-content-delivery]
[edit services l2tp traceoptions]
[edit services server-load-balance traceoptions]
[edit services logging traceoptions]
[edit services mobile-ip traceoptions]
[edit services ssl traceoptions]
[edit system accounting traceoptions]
[edit system auto-configuration traceoptions]
[edit system ddos-protection traceoptions]
[edit system license traceoptions]
[edit system processes app-engine-management-service traceoptions]
[edit system processes app-engine-virtual-machine-management-service traceoptions]
[edit system processes datapath-trace-service traceoptions]
[edit system processes dhcp-service interface-traceoptions]
[edit system processes dhcp-service traceoptions]
[edit system processes diameter-service traceoptions]
[edit system processes general-authentication-service traceoptions]
[edit system processes mac-validation traceoptions]
[edit system processes mag-service traceoptions]
[edit system processes process-monitor traceoptions]
[edit system processes resource-cleanup traceoptions]
[edit system processes sdk-service traceoptions]
[edit system processes static-subscribers traceoptions]
[edit system services database-replication traceoptions]
[edit system services dhcp traceoptions]
[edit system services local-policy-decision-function traceoptions]
[edit system services outbound-ssh traceoptions]
[edit system services service-deployment traceoptions]
[edit system services subscriber-management traceoptions]
[edit system services subscriber-management-helper traceoptions]

```

- Related Documentation**
- [Access Privilege User Permission Flags Overview on page 128](#)
 - [Understanding Junos OS Access Privilege Levels on page 123](#)
 - [Configuring Access Privilege Levels on page 131](#)
 - [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
 - [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)
 - [trace-control on page 223](#)

trace-control

Can modify trace file settings and configure trace file properties.

**Configuration
Hierarchy Levels**

```

[edit bridge-domains domain forwarding-options dhcp-relay interface-traceoptions]
[edit bridge-domains domain forwarding-options dhcp-relay traceoptions]
[edit bridge-domains domain multicast-snooping-options traceoptions]

```

```
[edit bridge-domains domain protocols igmp-snooping traceoptions]
[edit demux traceoptions]
[edit dynamic-profiles protocols igmp traceoptions]
[edit dynamic-profiles protocols mld traceoptions]
[edit dynamic-profiles protocols oam ethernet link-fault-management traceoptions]
[edit dynamic-profiles protocols oam ethernet lmi]
[edit dynamic-profiles protocols router-advertisement traceoptions]
[edit dynamic-profiles protocols oam gre-tunnel traceoptions]
[edit dynamic-profiles routing-instances instance bridge-domains domain
forwarding-options dhcp-relay traceoptions]
[edit dynamic-profiles routing-instances instance bridge-domains domain
multicast-snooping-options traceoptions]
[edit dynamic-profiles routing-instances instance bridge-domains domain protocols
igmp-snooping traceoptions]
[edit dynamic-profiles routing-instances instance forwarding-options dhcp-relay
traceoptions]
[edit dynamic-profiles routing-instances instance multicast-snooping-options
traceoptions]
[edit dynamic-profiles routing-instances instance protocols bgp group neighbor
traceoptions]
[edit dynamic-profiles routing-instances instance protocols bgp group traceoptions]
[edit dynamic-profiles routing-instances instance protocols bgp traceoptions]
[edit dynamic-profiles routing-instances instance protocols esis traceoptions]
[edit dynamic-profiles routing-instances instance protocols igmp-snooping traceoptions]
[edit dynamic-profiles routing-instances instance protocols isis traceoptions]
[edit dynamic-profiles routing-instances instance protocols l2vpn traceoptions]
[edit dynamic-profiles routing-instances instance protocols ldp traceoptions]
[edit dynamic-profiles routing-instances instance protocols msdp group peer
traceoptions]
[edit dynamic-profiles routing-instances instance protocols msdp group traceoptions]
[edit dynamic-profiles routing-instances instance protocols msdp peer traceoptions]
[edit dynamic-profiles routing-instances instance protocols msdp traceoptions]
[edit dynamic-profiles routing-instances instance protocols mvpn traceoptions]
[edit dynamic-profiles routing-instances instance protocols ospf traceoptions]
[edit dynamic-profiles routing-instances instance protocols pim traceoptions]
[edit dynamic-profiles routing-instances instance protocols rip traceoptions]
[edit dynamic-profiles routing-instances instance protocols ripng traceoptions]
[edit dynamic-profiles routing-instances instance protocols router-discovery traceoptions]
[edit dynamic-profiles routing-instances instance protocols vpls traceoptions]
[edit dynamic-profiles routing-instances instance routing-options multicast traceoptions]
[edit dynamic-profiles routing-instances instance routing-options traceoptions]
[edit dynamic-profiles routing-instances instance services mobile-ip traceoptions]
[edit dynamic-profiles routing-instances instance system services dhcp-local-server
traceoptions]
[edit dynamic-profiles routing-options multicast traceoptions]
[edit fabric protocols bgp group neighbor traceoptions]
[edit fabric protocols bgp group traceoptions]
[edit fabric protocols bgp traceoptions]
[edit fabric routing-instances instance routing-options traceoptions]
[edit fabric routing-options traceoptions]
[edit forwarding-options dhcp-relay interface-traceoptions]
[edit forwarding-options dhcp-relay traceoptions]
[edit jnx-example traceoptions]
[edit logical-systems bridge-domains domain forwarding-options dhcp-relay
interface-traceoptions]
[edit logical-systems bridge-domains domain forwarding-options dhcp-relay
```

```

traceoptions]
[edit logical-systems bridge-domains domain multicast-snooping-options traceoptions]
[edit logical-systems bridge-domains domain protocols igmp-snooping traceoptions]
[edit logical-systems forwarding-options dhcp-relay traceoptions]
[edit logical-systems protocols ancp traceoptions]
[edit logical-systems protocols bgp group neighbor traceoptions]
[edit logical-systems protocols bgp group traceoptions]
[edit logical-systems protocols bgp traceoptions]
[edit logical-systems protocols dot1x traceoptions]
[edit logical-systems protocols dvmrp traceoptions]
[edit logical-systems protocols esis traceoptions]
[edit logical-systems protocols igmp traceoptions]
[edit logical-systems protocols igmp-host traceoptions]
[edit logical-systems protocols ilmi traceoptions]
[edit logical-systems protocols isis traceoptions]
[edit logical-systems protocols l2circuit traceoptions]
[edit logical-systems protocols l2iw traceoptions]
[edit logical-systems protocols lacp traceoptions]
[edit logical-systems protocols layer2-control traceoptions]
[edit logical-systems protocols ldp traceoptions]
[edit logical-systems protocols mld traceoptions]
[edit logical-systems protocols mld-host traceoptions]
[edit logical-systems protocols mpls label-switched-path oam traceoptions]
[edit logical-systems protocols mpls label-switched-path primary oam traceoptions]
[edit logical-systems protocols mpls label-switched-path secondary oam traceoptions]
[edit logical-systems protocols mpls oam traceoptions]
[edit logical-systems protocols msdp group peer traceoptions]
[edit logical-systems protocols msdp group traceoptions]
[edit logical-systems protocols msdp peer traceoptions]
[edit logical-systems protocols msdp traceoptions]
[edit logical-systems protocols neighbor-discovery secure traceoptions]
[edit logical-systems protocols oam ethernet link-fault-management traceoptions]
[edit logical-systems protocols oam ethernet lmi traceoptions]
[edit logical-systems protocols ospf traceoptions]
[edit logical-systems protocols pim traceoptions]
[edit logical-systems protocols ppp monitor-session]
[edit logical-systems protocols ppp traceoptions]
[edit logical-systems protocols ppp-service traceoptions]
[edit logical-systems protocols pppoe traceoptions]
[edit logical-systems protocols rip traceoptions]
[edit logical-systems protocols ripng traceoptions]
[edit logical-systems protocols router-advertisement traceoptions]
[edit logical-systems protocols router-discovery traceoptions]
[edit logical-systems protocols rsvp traceoptions]
[edit logical-systems routing-instances instance bridge-domains domain
forwarding-options dhcp-relay interface-traceoptions]
[edit logical-systems routing-instances instance bridge-domains domain
forwarding-options dhcp-relay traceoptions]
[edit logical-systems routing-instances instance bridge-domains domain
multicast-snooping-options traceoptions]
[edit logical-systems routing-instances instance bridge-domains domain protocols
igmp-snooping traceoptions]
[edit logical-systems routing-instances instance forwarding-options dhcp-relay
traceoptions]
[edit logical-systems routing-instances instance multicast-snooping-options
traceoptions]

```

```
[edit logical-systems routing-instances instance protocols bgp group neighbor
traceoptions]
[edit logical-systems routing-instances instance protocols bgp group traceoptions]
[edit logical-systems routing-instances instance protocols bgp traceoptions]
[edit logical-systems routing-instances instance protocols esis traceoptions]
[edit logical-systems routing-instances instance protocols igmp-snooping traceoptions]
[edit logical-systems routing-instances instance protocols isis traceoptions]
[edit logical-systems routing-instances instance protocols l2vpn traceoptions]
[edit logical-systems routing-instances instance protocols ldp traceoptions]
[edit logical-systems routing-instances instance protocols msdp group peer traceoptions]
[edit logical-systems routing-instances instance protocols msdp group traceoptions]
[edit logical-systems routing-instances instance protocols msdp peer traceoptions]
[edit logical-systems routing-instances instance protocols msdp traceoptions]
[edit logical-systems routing-instances instance protocols mvpn traceoptions]
[edit logical-systems routing-instances instance protocols ospf traceoptions]
[edit logical-systems routing-instances instance protocols pim traceoptions]
[edit logical-systems routing-instances instance protocols rip traceoptions]
[edit logical-systems routing-instances instance protocols ripng traceoptions]
[edit logical-systems routing-instances instance protocols router-discovery traceoptions]
[edit logical-systems routing-instances instance protocols vpls traceoptions]
[edit logical-systems routing-instances instance routing-options multicast traceoptions]
[edit logical-systems routing-instances instance routing-options traceoptions]
[edit logical-systems routing-instances instance services mobile-ip traceoptions]
[edit logical-systems routing-instances instance system services dhcp-local-server
interface-traceoptions]
[edit logical-systems routing-instances instance system services dhcp-local-server
traceoptions]
[edit logical-systems routing-options multicast traceoptions]
[edit logical-systems routing-options traceoptions]
[edit logical-systems services mobile-ip traceoptions]
[edit logical-systems system services dhcp-local-server interface-traceoptions]
[edit logical-systems system services dhcp-local-server traceoptions]
[edit multicast-snooping-options traceoptions]
[edit protocols ancp traceoptions]
[edit protocols bgp group neighbor traceoptions]
[edit protocols bgp group traceoptions]
[edit protocols bgp traceoptions]
[edit protocols dot1x traceoptions]
[edit protocols dvmrp traceoptions]
[edit protocols esis traceoptions]
[edit protocols igmp traceoptions]
[edit protocols igmp-host traceoptions]
[edit protocols ilmi traceoptions]
[edit protocols isis traceoptions]
[edit protocols l2circuit traceoptions]
[edit protocols l2iw traceoptions]
[edit protocols lacp traceoptions]
[edit protocols layer2-control traceoptions]
[edit protocols ldp traceoptions]
[edit protocols mld traceoptions]
[edit protocols mld-host traceoptions]
[edit protocols mpls label-switched-path oam traceoptions]
[edit protocols mpls label-switched-path primary oam traceoptions]
[edit protocols mpls label-switched-path secondary oam traceoptions]
[edit protocols mpls oam traceoptions]
[edit protocols msdp group peer traceoptions]
```



```
[edit protocols msdp group traceoptions]
[edit protocols msdp peer traceoptions]
[edit protocols msdp traceoptions]
[edit protocols neighbor-discovery secure traceoptions]
[edit protocols oam ethernet connectivity-fault-management traceoptions]
[edit protocols oam ethernet link-fault-management traceoptions]
[edit protocols oam ethernet lmi traceoptions]
[edit protocols ospf traceoptions]
[edit protocols pim traceoptions]
[edit protocols ppp monitor-session]
[edit protocols ppp traceoptions]
[edit protocols ppp-service traceoptions]
[edit protocols pppoe traceoptions]
[edit protocols rip traceoptions]
[edit protocols ripng traceoptions]
[edit protocols router-advertisement traceoptions]
[edit protocols router-discovery traceoptions]
[edit protocols rsvp traceoptions]
[edit routing-instances instance bridge-domains domain forwarding-options dhcp-relay
interface-traceoptions]
[edit routing-instances instance bridge-domains domain forwarding-options dhcp-relay
traceoptions]
[edit routing-instances instance bridge-domains domain multicast-snooping-options
traceoptions]
[edit routing-instances instance bridge-domains domain protocols igmp-snooping
traceoptions]
[edit routing-instances instance forwarding-options dhcp-relay traceoptions]
[edit routing-instances instance forwarding-options dhcp-relay interface-traceoptions]
[edit routing-instances instance multicast-snooping-options traceoptions]
[edit routing-instances instance protocols bgp group neighbor traceoptions]
[edit routing-instances instance protocols bgp group traceoptions]
[edit routing-instances instance protocols bgp traceoptions]
[edit routing-instances instance protocols esis traceoptions]
[edit routing-instances instance protocols igmp-snooping traceoptions]
[edit routing-instances instance protocols isis traceoptions]
[edit routing-instances instance protocols l2vpn traceoptions]
[edit routing-instances instance protocols ldp traceoptions]
[edit routing-instances instance protocols msdp group peer traceoptions]
[edit routing-instances instance protocols msdp group traceoptions]
[edit routing-instances instance protocols msdp peer traceoptions]
[edit routing-instances instance protocols msdp traceoptions]
[edit routing-instances instance protocols mvpn traceoptions]
[edit routing-instances instance protocols ospf traceoptions]
[edit routing-instances instance protocols pim traceoptions]
[edit routing-instances instance protocols rip traceoptions]
[edit routing-instances instance protocols ripng traceoptions]
[edit routing-instances instance protocols router-discovery traceoptions]
[edit routing-instances instance protocols vpls traceoptions]
[edit routing-instances instance routing-options multicast traceoptions]
[edit routing-instances instance routing-options traceoptions]
[edit routing-instances instance system services dhcp-local-server interface-traceoptions]
[edit routing-instances instance system services dhcp-local-server traceoptions]
[edit routing-options multicast traceoptions]
[edit routing-options traceoptions]
[edit security idp traceoptions]
[edit security pki traceoptions]
```

```
[edit services adaptive-services-pics traceoptions]
[edit services captive-portal-content-delivery]
[edit system ddos-protection traceoptions]
[edit services l2tp traceoptions]
[edit services logging traceoptions]
[edit services mobile-ip traceoptions]
[edit services server-load-balance traceoptions]
[edit services ssl traceoptions]
[edit system accounting traceoptions]
[edit system auto-configuration traceoptions]
[edit system license traceoptions]
[edit system processes datapath-trace-service traceoptions]
[edit system processes diameter-service traceoptions]
[edit system processes general-authentication-service traceoptions]
[edit system processes mac-validation traceoptions]
[edit system processes process-monitor traceoptions]
[edit system processes resource-cleanup traceoptions]
[edit system processes sdk-service traceoptions]
[edit system processes static-subscribers traceoptions]
[edit system services database-replication traceoptions]
[edit system services dhcp traceoptions]
[edit system services dhcp-local-server traceoptions]
[edit system services dhcp-local-server interface-traceoptions]
[edit system services local-policy-decision-function traceoptions]
[edit system services outbound-ssh traceoptions]
[edit system services service-deployment traceoptions]
[edit system services subscriber-management traceoptions]
[edit system services subscriber-management-helper traceoptions]
```

- Related Documentation**
- [Access Privilege User Permission Flags Overview on page 128](#)
 - [Understanding Junos OS Access Privilege Levels on page 123](#)
 - [Configuring Access Privilege Levels on page 131](#)
 - [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
 - [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)
 - [trace on page 218](#)

view

Can view current system-wide, routing table, and protocol-specific values and statistics.

Commands

```
clear ipv6 router-advertisement
<clear-ipv6-router-advertisement-information>

show
show accounting

show accounting profile
<get-accounting-profile-information>

show accounting records
<get-accounting-record-information>
```

```
show amt
show amt statistics
  <get-amt-statistics>
show amt summary
  <get-amt-summary>
show amt tunnel
  <get-amt-tunnel-information>
show amt tunnel gateway-address
  <get-amt-tunnel-gateway-address>
show amt tunnel tunnel-interface
  <get-amt-tunnel-interface>
show ancp
show ancp cos
  <get-ancp-cos-information>

show ancp cos last-update
  <get-ancp-cos-last-update-information>

show ancp cos pending-update
  <get-ancp-cos-pending-information>

show ancp neighbor
  <get-ancp-neighbor-information>

show ancp subscriber
  <get-ancp-subscriber-information>

show ancp subscriber identifier
  <get-ancp-subscriber-identifier-information>
show ancp subscriber neighbor
show app-engine
show app-engine info
show app-engine processes
show app-engine resource-usage
show app-engine status
show app-engine virtual-machine package
show app-engine virtual-machine vm-instance
show aps
  <get-aps-information>

show aps group
  <get-aps-group-information>
show aps interface
  <get-aps-interface-information>
show arp
  <get-arp-table-information>

show as-path
  <get-as-path>
show as-path domain
  <get-as-path-domain>
show auto-configuration
show auto-configuration interfaces
show bfd
show bfd session
```

```
<get-bfd-session-information>

show bfd session address
  <get-bfd-session-address>
show bfd session discriminator
  <get-bfd-session-discriminator>
show bfd session prefix
  <get-bfd-session-prefix>
show bgp
show bgp bmp
  <get-bgp-monitoring-protocol-statistics>
show bgp group
  <get-bgp-group-information>

show bgp group rtf
  <get-bgp-rtf-information>

show bgp group traffic-statistics
  <get-bgp-traffic-statistics-information>

show bgp neighbor
  <get-bgp-neighbor-information>

show bgp neighbor orf
  <get-bgp-orf-information>

show bgp replication
  <get-bgp-replication-information>
show bgp summary
  <get-bgp-summary-information>

show bridge
show bridge domain
  <get-bridge-instance-information>

show bridge domain operational
  <get-operational-bridge-instance-information>
  <get-bridge-flood-information>
show bridge flood
show bridge flood event-queue
  <get-bridge-domain-event-queue-information>

show bridge flood route
show bridge flood route all-ce-flood
  <get-show-bridge-domain-all-ce-flood-route-information>

show bridge flood route all-ve-flood
  <get-show-bridge-domain-ve-flood-route-information>
show bridge flood route alt-root-flood
  <get-bridge-domain-alt-root-flood-route-information>
show bridge flood route bd-flood
  <get-bridge-domain-bd-flood-route-information>
show bridge flood route mlp-flood
  <get-bridge-domain-mlp-flood-route-information>
show bridge flood route re-flood
  <get-bridge-domain-re-flood-route-information>
```

```
show bridge mac-table
  <get-bridge-mac-table>
show bridge mac-table interface
  <get-bridge-interface-mac-table>
show bridge statistics
  <get-bridge-statistics-information>
show chassis
show chassis adc
show chassis alarms
  <get-alarm-information>
show chassis alarms fpc
  <get-fpc-alarm-information>
show chassis beacon
  <get-chassis-beacon-information>
show chassis beacon cb
  <get-chassis-cb-beacon-information>
show chassis environment adc
show chassis environment ccg
  <get-environment-ccg-information>
show chassis cfeb
  <get-cfeb-information>
show chassis cip
show chassis craft-interface
  <get-craft-information>
show chassis environment
  <get-environment-information>
show chassis environment cb
  <get-environment-cb-information>
show chassis environment cip
  <get-environment-cip-information>
show chassis environment feb
  <get-environment-feb-information>
show chassis environment fan
show chassis environment fpc
  <get-environment-fpc-information>
show chassis environment fpm
  <get-environment-fpm-information>
show chassis environment mcs
  <get-environment-mcs-information>
show chassis environment pcg
  <get-environment-pcg-information>
show chassis environment pdu
  <get-environment-pdu-information>
show chassis environment pem
  <get-environment-pem-information>
show chassis environment psm
show chassis environment psu
  <get-environment-psu-information>
show chassis environment routing-engine
  <get-environment-re-information>
show chassis environment scg
  <get-environment-scg-information>
show chassis environment sfb
show chassis environment sfm
  <get-environment-sfm-information>
```

```
show chassis environment sib
  <get-environment-sib-information>

show chassis environment sib f13
show chassis environment sib f2s
show chassis ethernet-switch
show chassis ethernet-switch errors
show chassis ethernet-switch statistics
show chassis fabric
show chassis fabric degraded-fabric-reachability
show chassis fabric device
  <get-chassis-fabric-information-device>
show chassis fabric connectivity
  <get-chassis-fabric-connectivity-information>
show chassis fabric destinations
  <get-fm-fabric-destinations-state>
show chassis fabric errors
show chassis fabric errors autoheal
  <get-fm-plane-autoheal-errors>
show chassis fabric errors fpc
  <get-fm-fpc-errors>

show chassis fabric errors sib
  <get-fm-sib-errors>

show chassis fabric errors sib f13
show chassis fabric errors sib f2s
show chassis fabric feb
show chassis fabric fpcs
  <get-fm-fpc-state-information>

show chassis fabric links
  <get-chassis-fabric-link-information>
show chassis fabric map
show chassis fabric plane
  <get-fm-plane-state-information>

show chassis fabric plane-location
show chassis fabric reachability
  <get-fm-fabric-reachability-information>
show chassis fabric sibs
  <get-fm-sib-state-information>
show chassis fabric spray-weights
  <get-chassis-fabric-spray-weight-information>
show chassis fabric spray-weights from
show chassis fabric spray-weights to
show chassis fabric summary
  <get-fm-state-information>

show chassis fabric topology
  <get-chassis-fabric-topology-information>
show chassis fabric unreachable-destinations
  <get-fm-unreachable-dest-information>
show chassis fan
show chassis feb
  <get-feb-brief-information>
```

```
show chassis feb detail
  <get-feb-information>

show chassis firmware
  <get-firmware-information>

show chassis firmware detail
  <get-firmware-information-detail>
show chassis forwarding
  <get-fwdd-information>

show chassis fpc
  <get-fpc-information>

show chassis fpc pic-status
  <get-pic-information>

show chassis fpc-feb-connectivity
  <get-fpc-feb-connectivity-information>

show chassis hardware
  <get-chassis-inventory>
show chassis hss
show chassis hss link-quality
show chassis in-service-upgrade
show chassis ioc-npc-connectivity
  <get-ioc-npc-connectivity-information>

show chassis lccs
  <get-fru-information>

show chassis location
  <get-chassis-location>

show chassis location fpc
show chassis location interface
show chassis location interface by-name
  <get-interface-location-name-information>

show chassis location interface by-slot
  <get-interface-location-information>
show chassis mac-addresses
show chassis multicast-loadbalance
  <get-chassis-ae-lb-information>

show chassis network-services
  <network-services>

show chassis nonstop-upgrade
show chassis pic
  <get-pic-detail>

show chassis power
  <get-power-usage-information>
```

```
show chassis power detail
<get-power-usage-information-detail>
show chassis power sequence
show chassis power upgrade

show chassis power-ratings
<get-power-management>

show chassis psd
<get-psd-information>

show chassis redundancy
show chassis redundancy feb
<get-feb-redundancy-information>

show chassis redundancy feb errors
<get-feb-redundancy-error-information>

show chassis redundancy feb redundancy-group
<get-feb-redundancy-group-information>

show chassis redundant-power-system
<get-rps-chassis-information>

show chassis routing-engine
<get-route-engine-information>

show chassis routing-engine bios
<get-bios-version-information>
show chassis scb
<get-scb-information>

show chassis sfm
<get-sfm-information>

show chassis sfm detail
show chassis sibs
<get-sib-information>

show chassis spmb
<get-spmb-information>

show chassis spmb sibs
<get-spmb-sib-information>

show chassis ssb
<get-ssb-information>

show chassis synchronization
<get-clock-synchronization-information>

show chassis synchronization backup
show chassis synchronization master
show chassis temperature-thresholds
<get-temperature-threshold-information>
show chassis vcpu
```



```
show chassis zones
  <get-chassis-zones-information>
show class-of-service
  <get-cos-information>

show class-of-service adaptive-shaper
  <get-cos-adaptive-shaper-information>

show class-of-service application-traffic-control
show class-of-service application-traffic-control counter
show class-of-service application-traffic-control statistics
show class-of-service application-traffic-control statistics rate-limiter
show class-of-service application-traffic-control statistics rule
  <get-appqos-rule-statistics>
show class-of-service classifier
  <get-cos-classifier-information>

show class-of-service code-point-aliases
  <get-cos-code-point-map-information>

show class-of-service congestion-notification
  <get-cos-congestion-notification-information>
show class-of-service drop-profile
  <get-cos-drop-profile-information>

show class-of-service fabric
show class-of-service fabric scheduler-map
  <get-cos-fabric-scheduler-map-information>

show class-of-service fabric statistics
  <get-fabric-queue-information>

show class-of-service forwarding-class
  <get-cos-forwarding-class-information>

show class-of-service forwarding-class-set
  <get-cos-forwarding-class-set-information>
show class-of-service forwarding-table
  <get-cos-table-information>

show class-of-service forwarding-table classifier
  <get-cos-classifier-table-information>

show class-of-service forwarding-table classifier mapping
  <get-cos-classifier-table-map-information>

show class-of-service forwarding-table drop-profile
  <get-cos-red-information>

show class-of-service forwarding-table fabric
show class-of-service forwarding-table fabric scheduler-map
  <get-cos-fwtab-fabric-scheduler-map-information>

show class-of-service forwarding-table forwarding-class-map
  <get-cos-forwarding-class-map-table-information>
```

show class-of-service forwarding-table forwarding-class-map mapping
<get-cos-forwarding-class-map-interface-table-information>

show class-of-service forwarding-table loss-priority-map
<get-cos-loss-priority-map-table-information>

show class-of-service forwarding-table loss-priority-map mapping
<get-cos-loss-priority-map-table-binding-information>

show class-of-service forwarding-table loss-priority-rewrite
<get-cos-loss-priority-rewrite-table-information>

show class-of-service forwarding-table loss-priority-rewrite mapping
<get-cos-loss-priority-rewrite-table-binding-information>

show class-of-service forwarding-table policer
<get-cos-policer-table-map-information>

show class-of-service forwarding-table rewrite-rule
<get-cos-rewrite-table-information>

show class-of-service forwarding-table rewrite-rule mapping
<get-cos-rewrite-table-map-information>

show class-of-service forwarding-table scheduler-map
<get-cos-scheduler-map-table-information>

show class-of-service forwarding-table shaper
<get-cos-shaper-table-map-information>

show class-of-service forwarding-table translation-table
<get-cos-translation-table-information>

show class-of-service forwarding-table translation-table mapping
<get-cos-translation-table-mapping-information>

show class-of-service fragmentation-map
<get-cos-fragmentation-map-information>

show class-of-service interface
<get-cos-interface-map-information>

show class-of-service interface-set
<get-cos-interface-set-map-information>

show class-of-service l2tp-session
<get-cos-l2tp-session-map-information>

show class-of-service loss-priority-map
<get-cos-loss-priority-map-information>

show class-of-service loss-priority-rewrite
<get-cos-loss-priority-rewrite-information>

show class-of-service multi-destination
<get-cos-multi-destination-information>

show class-of-service rewrite-rule
<get-cos-rewrite-information>

```
show class-of-service routing-instance
  <get-cos-routing-instance-map-information>

show class-of-service scheduler-map
  <get-cos-scheduler-map-information>

show class-of-service traffic-control-profile
  <get-cos-traffic-control-profile-information>

show class-of-service translation-table
  <get-cos-translation-table-map-information>

show class-of-service virtual-channel
  <get-cos-virtual-channel-information>

show class-of-service virtual-channel-group
  <get-cos-virtual-channel-group-information>

show cli
show cli authorization
  <get-authorization-information>

show cli directory
show cli history
show configuration
show connections
  <get-ccc-information>
show database-replication
show database-replication statistics
  <get-database-replication-statistics-information>

show database-replication summary
  <get-database-replication-summary-information>
show ddos-protection
show ddos-protection protocols
  <get-ddos-protocols-information>
show ddos-protection protocols amtv4
show ddos-protection protocols amtv4 aggregate
show ddos-protection protocols amtv4 aggregate culprit-flows
show ddos-protection protocols amtv4 culprit-flows
show ddos-protection protocols amtv4 flow-detection
show ddos-protection protocols amtv4 parameters
show ddos-protection protocols amtv4 statistics
show ddos-protection protocols amtv4 violations
show ddos-protection protocols amtv6
show ddos-protection protocols amtv6 aggregate
show ddos-protection protocols amtv6 aggregate culprit-flows
show ddos-protection protocols amtv6 culprit-flows
show ddos-protection protocols amtv6 flow-detection
show ddos-protection protocols amtv6 statistics
show ddos-protection protocols amtv6 violations

show ddos-protection protocols ancp
  <get-ddos-ancp-information>
```

```
show ddos-protection protocols ancp aggregate  
  <get-ddos-ancp-aggregate>
```

```
show ddos-protection protocols ancp parameters  
  <get-ddos-ancp-parameters>
```

```
show ddos-protection protocols ancp statistics  
  <get-ddos-ancp-statistics>
```

```
show ddos-protection protocols ancp violations  
  <get-ddos-ancp-violations>
```

```
show ddos-protection protocols ancpv6  
  <get-ddos-ancpv6-information>  
show ddos-protection protocols ancpv6 aggregate  
  get-ddos-ancpv6-aggregate  
show ddos-protection protocols ancpv6 parameters  
  get-ddos-ancpv6-parameters  
show ddos-protection protocols ancpv6 statistics  
  get-ddos-ancpv6-statistics  
show ddos-protection protocols ancpv6 violations  
  get-ddos-ancpv6-violations  
show ddos-protection protocols arp  
  get-ddos-arp-information  
show ddos-protection protocols arp aggregate  
  get-ddos-arp-aggregate  
show ddos-protection protocols arp parameters  
  get-ddos-arp-parameters  
show ddos-protection protocols arp statistics  
  get-ddos-arp-statistics  
show ddos-protection protocols arp violations  
  get-ddos-arp-violations  
show ddos-protection protocols atm  
  get-ddos-atm-information  
show ddos-protection protocols atm aggregate  
  get-ddos-atm-aggregate  
show ddos-protection protocols atm parameters  
  get-ddos-atm-parameters  
show ddos-protection protocols atm statistics  
  get-ddos-atm-statistics  
show ddos-protection protocols atm violations  
  get-ddos-atm-violations  
show ddos-protection protocols bfd  
  get-ddos-bfd-information  
show ddos-protection protocols bfd aggregate  
  get-ddos-bfd-aggregate  
show ddos-protection protocols bfd parameters  
  get-ddos-bfd-parameters  
show ddos-protection protocols bfd statistics  
  get-ddos-bfd-statistics  
show ddos-protection protocols bfd violations  
  get-ddos-bfd-violations
```

```
show ddos-protection protocols bfdv6
  get-ddos-bfdv6-information
show ddos-protection protocols bfdv6 aggregate
  get-ddos-bfdv6-aggregate
show ddos-protection protocols bfdv6 parameters
  get-ddos-bfdv6-parameters
show ddos-protection protocols bfdv6 statistics
  get-ddos-bfdv6-statistics
show ddos-protection protocols bfdv6 violations
  get-ddos-bfdv6-violations
show ddos-protection protocols bgp
  get-ddos-bgp-information
show ddos-protection protocols bgp aggregate
  get-ddos-bgp-aggregate
show ddos-protection protocols bgp parameters
  get-ddos-bgp-parameters
show ddos-protection protocols bgp statistics
  get-ddos-bgp-statistics
show ddos-protection protocols bgp violations
  get-ddos-bgp-violations
show ddos-protection protocols bgpv6
  get-ddos-bgpv6-information
show ddos-protection protocols bgpv6 aggregate
  get-ddos-bgpv6-aggregate
show ddos-protection protocols bgpv6 parameters
  get-ddos-bgpv6-parameters
show ddos-protection protocols bgpv6 statistics
  get-ddos-bgpv6-statistics
show ddos-protection protocols bgpv6 violations
  get-ddos-bgpv6-violations
show ddos-protection protocols demux-autosense
  get-ddos-demuxauto-information
show ddos-protection protocols demux-autosense aggregate
  get-ddos-demuxauto-aggregate
show ddos-protection protocols demux-autosense parameters
  get-ddos-demuxauto-parameters
show ddos-protection protocols demux-autosense statistics
  get-ddos-demuxauto-statistics
show ddos-protection protocols demux-autosense violations
  get-ddos-demuxauto-violations
show ddos-protection protocols dhcpv4
  get-ddos-dhcpv4-information
show ddos-protection protocols dhcpv4 ack
  get-ddos-dhcpv4-ack
show ddos-protection protocols dhcpv4 aggregate
  get-ddos-dhcpv4-aggregate
show ddos-protection protocols dhcpv4 bad-packets
  get-ddos-dhcpv4-bad-pack
show ddos-protection protocols dhcpv4 bootp
  get-ddos-dhcpv4-bootp
show ddos-protection protocols dhcpv4 decline
  get-ddos-dhcpv4-decline
show ddos-protection protocols dhcpv4 discover
  get-ddos-dhcpv4-discover
show ddos-protection protocols dhcpv4 force-renew
  get-ddos-dhcpv4-forcerenew
```

```
show ddos-protection protocols dhcpv4 inform
  get-ddos-dhcpv4-inform
show ddos-protection protocols dhcpv4 lease-active
  get-ddos-dhcpv4-leaseact
show ddos-protection protocols dhcpv4 lease-query
  get-ddos-dhcpv4-leasequery
show ddos-protection protocols dhcpv4 lease-unassigned
  get-ddos-dhcpv4-leaseuna
show ddos-protection protocols dhcpv4 lease-unknown
  get-ddos-dhcpv4-leaseunk
show ddos-protection protocols dhcpv4 nak
  get-ddos-dhcpv4-nak
show ddos-protection protocols dhcpv4 no-message-type
  get-ddos-dhcpv4-no-msgtype
show ddos-protection protocols dhcpv4 offer
  get-ddos-dhcpv4-offer
show ddos-protection protocols dhcpv4 offer culprit-flows
show ddos-protection protocols dhcpv4 parameters
  get-ddos-dhcpv4-parameters
show ddos-protection protocols dhcpv4 release
  get-ddos-dhcpv4-release
show ddos-protection protocols dhcpv4 renew
  get-ddos-dhcpv4-renew
show ddos-protection protocols dhcpv4 request
  get-ddos-dhcpv4-request
show ddos-protection protocols dhcpv4 statistics
  get-ddos-dhcpv4-statistics
show ddos-protection protocols dhcpv4 unclassified
  get-ddos-dhcpv4-unclass
show ddos-protection protocols dhcpv4 violations
  get-ddos-dhcpv4-violations
show ddos-protection protocols dhcpv6
  get-ddos-dhcpv6-information
show ddos-protection protocols dhcpv6 advertise
  get-ddos-dhcpv6-advertise
show ddos-protection protocols dhcpv6 advertise culprit-flows
show ddos-protection protocols dhcpv6 aggregate
  get-ddos-dhcpv6-aggregate
show ddos-protection protocols dhcpv6 confirm
  get-ddos-dhcpv6-confirm
show ddos-protection protocols dhcpv6 decline
  get-ddos-dhcpv6-decline
show ddos-protection protocols dhcpv6 information-request
  get-ddos-dhcpv6-info-req
show ddos-protection protocols dhcpv6 leasequery
  get-ddos-dhcpv6-leasequery
show ddos-protection protocols dhcpv6 leasequery culprit-flows
show ddos-protection protocols dhcpv6 leasequery-data
  get-ddos-dhcpv6-leaseq-da
show ddos-protection protocols dhcpv6 leasequery-done
  get-ddos-dhcpv6-leaseq-do
show ddos-protection protocols dhcpv6 leasequery-reply
  get-ddos-dhcpv6-leaseq-re
show ddos-protection protocols dhcpv6 parameters
  get-ddos-dhcpv6-parameters
show ddos-protection protocols dhcpv6 rebind
```

```
get-ddos-dhcpv6-rebind
show ddos-protection protocols dhcpv6 reconfigure
get-ddos-dhcpv6-reconfig
show ddos-protection protocols dhcpv6 relay-forward
get-ddos-dhcpv6-relay-for
show ddos-protection protocols dhcpv6 relay-reply
get-ddos-dhcpv6-relay-rep
show ddos-protection protocols dhcpv6 release
get-ddos-dhcpv6-release
show ddos-protection protocols dhcpv6 renew
get-ddos-dhcpv6-renew
show ddos-protection protocols dhcpv6 reply
get-ddos-dhcpv6-reply
show ddos-protection protocols dhcpv6 request
get-ddos-dhcpv6-request
show ddos-protection protocols dhcpv6 solicit
get-ddos-dhcpv6-solicit
show ddos-protection protocols dhcpv6 statistics
get-ddos-dhcpv6-statistics
show ddos-protection protocols dhcpv6 unclassified
get-ddos-dhcpv6-unclass
show ddos-protection protocols dhcpv6 unclassified culprit-flows
show ddos-protection protocols dhcpv6 violations
get-ddos-dhcpv6-violations
show ddos-protection protocols diameter
get-ddos-diameter-information
show ddos-protection protocols diameter aggregate
get-ddos-diameter-aggregate
show ddos-protection protocols diameter parameters
get-ddos-diameter-parameters
show ddos-protection protocols diameter statistics
get-ddos-diameter-statistics
show ddos-protection protocols diameter violations
get-ddos-diameter-violations
show ddos-protection protocols dns
get-ddos-dns-information
show ddos-protection protocols dns aggregate
get-ddos-dns-aggregate
show ddos-protection protocols dns parameters
get-ddos-dns-parameters
show ddos-protection protocols dns statistics
get-ddos-dns-statistics
show ddos-protection protocols dns violations
get-ddos-dns-violations
show ddos-protection protocols dtcp
get-ddos-dtcp-information
show ddos-protection protocols dtcp aggregate
get-ddos-dtcp-aggregate
show ddos-protection protocols dtcp aggregate culprit-flows
show ddos-protection protocols dtcp parameters
get-ddos-dtcp-parameters
show ddos-protection protocols dtcp statistics
get-ddos-dtcp-statistics
show ddos-protection protocols dtcp violations
get-ddos-dtcp-violations
show ddos-protection protocols dynamic-vlan
```

```
get-ddos-dynvlan-information
show ddos-protection protocols dynamic-vlan aggregate
  get-ddos-dynvlan-aggregate
show ddos-protection protocols dynamic-vlan parameters
  get-ddos-dynvlan-parameters
show ddos-protection protocols dynamic-vlan statistics
  get-ddos-dynvlan-statistics
show ddos-protection protocols dynamic-vlan violations
  get-ddos-dynvlan-violations
show ddos-protection protocols egpv6
  get-ddos-egpv6-information
show ddos-protection protocols egpv6 aggregate
  get-ddos-egpv6-aggregate
show ddos-protection protocols egpv6 parameters
  get-ddos-egpv6-parameters
show ddos-protection protocols egpv6 statistics
  get-ddos-egpv6-statistics
show ddos-protection protocols egpv6 violations
  get-ddos-egpv6-violations
show ddos-protection protocols eoam
  get-ddos-eoam-information
show ddos-protection protocols eoam aggregate
  get-ddos-eoam-aggregate
show ddos-protection protocols eoam parameters
  get-ddos-eoam-parameters
show ddos-protection protocols eoam statistics
  get-ddos-eoam-statistics
show ddos-protection protocols eoam violations
  get-ddos-eoam-violations
show ddos-protection protocols esmc
  get-ddos-esmc-information
show ddos-protection protocols esmc aggregate
  get-ddos-esmc-aggregate
show ddos-protection protocols esmc parameters
  get-ddos-esmc-parameters
show ddos-protection protocols esmc statistics
  get-ddos-esmc-statistics
show ddos-protection protocols esmc violations
  get-ddos-esmc-violations
show ddos-protection protocols fab-probe
  <get-ddos-fab-probe-information>
show ddos-protection protocols fab-probe aggregate
  <get-ddos-fab-probe-aggregate>
show ddos-protection protocols fab-probe parameters
  <get-ddos-fab-probe-parameters>
show ddos-protection protocols fab-probe statistics
  <get-ddos-fab-probe-statistics>
show ddos-protection protocols fab-probe violations
  <get-ddos-fab-probe-violations>
show ddos-protection protocols firewall-host
  get-ddos-fw-host-information
show ddos-protection protocols firewall-host aggregate
  get-ddos-fw-host-aggregate
show ddos-protection protocols firewall-host parameters
  get-ddos-fw-host-parameters
show ddos-protection protocols firewall-host statistics
```



```
get-ddos-fw-host-statistics
show ddos-protection protocols firewall-host violations
get-ddos-fw-host-violations

show ddos-protection protocols ftp
  get-ddos-ftp-information
show ddos-protection protocols ftp aggregate
  get-ddos-ftp-aggregate
show ddos-protection protocols ftp parameters
  get-ddos-ftp-parameters
show ddos-protection protocols ftp statistics
  get-ddos-ftp-statistics
show ddos-protection protocols ftp violations
  get-ddos-ftp-violations
show ddos-protection protocols ftpv6
  get-ddos-ftp6-information
show ddos-protection protocols ftpv6 aggregate
  get-ddos-ftp6-aggregate
show ddos-protection protocols ftpv6 parameters
  get-ddos-ftp6-parameters
show ddos-protection protocols ftpv6 statistics
  get-ddos-ftp6-statistics
show ddos-protection protocols ftpv6 violations
  get-ddos-ftp6-violations
show ddos-protection protocols gre
  get-ddos-gre-information
show ddos-protection protocols gre aggregate
  get-ddos-gre-aggregate
show ddos-protection protocols gre parameters
  get-ddos-gre-parameters
show ddos-protection protocols gre statistics
  get-ddos-gre-statistics
show ddos-protection protocols gre violations
  get-ddos-gre-violations
show ddos-protection protocols icmp
  get-ddos-icmp-information
show ddos-protection protocols icmp aggregate
  get-ddos-icmp-aggregate
show ddos-protection protocols icmp parameters
  get-ddos-icmp-parameters
show ddos-protection protocols icmp statistics
  get-ddos-icmp-statistics
show ddos-protection protocols icmp violations
  get-ddos-icmp-violations
show ddos-protection protocols icmpv6
  <get-ddos-icmpv6-information>
show ddos-protection protocols icmpv6 aggregate
  <get-ddos-icmpv6-aggregate>
show ddos-protection protocols icmpv6 aggregate culprit-flows
show ddos-protection protocols icmpv6 parameters
  <get-ddos-icmpv6-parameters>
show ddos-protection protocols icmpv6 statistics
  <get-ddos-icmpv6-statistics>
show ddos-protection protocols icmpv6 violations
  <get-ddos-icmpv6-violations>
```

```
show ddos-protection protocols igmp
  get-ddos-igmp-information
show ddos-protection protocols igmp aggregate
  get-ddos-igmp-aggregate
show ddos-protection protocols igmp aggregate culprit-flows
show ddos-protection protocols igmp parameters
  get-ddos-igmp-parameters
show ddos-protection protocols igmp statistics
  get-ddos-igmp-statistics
show ddos-protection protocols igmp violations
  get-ddos-igmp-violations
show ddos-protection protocols igmp-snoop
  get-ddos-igmp-snoop-information
show ddos-protection protocols igmp-snoop aggregate
  get-ddos-igmp-snoop-aggregate
show ddos-protection protocols igmp-snoop parameters
  get-ddos-igmp-snoop-parameters
show ddos-protection protocols igmp-snoop statistics
  get-ddos-igmp-snoop-statistics
show ddos-protection protocols igmp-snoop violations
  get-ddos-igmp-snoop-violations
show ddos-protection protocols igmpv4v6
  get-ddos-igmpv4v6-information
show ddos-protection protocols igmpv4v6 aggregate
  get-ddos-igmpv4v6-aggregate
show ddos-protection protocols igmpv4v6 aggregate culprit-flows
show ddos-protection protocols igmpv4v6 parameters
  get-ddos-igmpv4v6-parameters
show ddos-protection protocols igmpv4v6 statistics
  get-ddos-igmpv4v6-statistics
show ddos-protection protocols igmpv4v6 violations
  get-ddos-igmpv4v6-violations
show ddos-protection protocols igmpv6
  get-ddos-igmpv6-information
show ddos-protection protocols igmpv6 aggregate
  get-ddos-igmpv6-aggregate
show ddos-protection protocols igmpv6 parameters
  get-ddos-igmpv6-parameters
show ddos-protection protocols igmpv6 statistics
  get-ddos-igmpv6-statistics
show ddos-protection protocols igmpv6 violations
  get-ddos-igmpv6-violations
show ddos-protection protocols ip-fragments
  get-ddos-ip-frag-information
show ddos-protection protocols ip-fragments aggregate
  get-ddos-ip-frag-aggregate
show ddos-protection protocols ip-fragments first-fragment
  get-ddos-ip-frag-first-frag
show ddos-protection protocols ip-fragments parameters
  get-ddos-ip-frag-parameters
show ddos-protection protocols ip-fragments statistics
  get-ddos-ip-frag-statistics
show ddos-protection protocols ip-fragments trail-fragment
  get-ddos-ip-frag-trail-frag
show ddos-protection protocols ip-fragments violations
  get-ddos-ip-frag-violations
```

```
show ddos-protection protocols ip-options
  get-ddos-ip-opt-information
show ddos-protection protocols ip-options aggregate
  get-ddos-ip-opt-aggregate
show ddos-protection protocols ip-options non-v4v6
<get-ddos-ip-opt-non-v4v6>
show ddos-protection protocols ip-options parameters
  get-ddos-ip-opt-parameters
show ddos-protection protocols ip-options router-alert
  get-ddos-ip-opt-rt-alert
show ddos-protection protocols ip-options statistics
  get-ddos-ip-opt-statistics
show ddos-protection protocols ip-options unclassified
  get-ddos-ip-opt-unclass
show ddos-protection protocols ip-options violations
  get-ddos-ip-opt-violations
show ddos-protection protocols ipv4-unclassified
  get-ddos-ipv4-uncls-information
show ddos-protection protocols ipv4-unclassified aggregate
  get-ddos-ipv4-uncls-aggregate
show ddos-protection protocols ipv4-unclassified parameters
  get-ddos-ipv4-uncls-parameters
show ddos-protection protocols ipv4-unclassified statistics
  get-ddos-ipv4-uncls-statistics
show ddos-protection protocols ipv4-unclassified violations
  get-ddos-ipv4-uncls-violations
show ddos-protection protocols ipv6-unclassified
  get-ddos-ipv6-uncls-information
show ddos-protection protocols ipv6-unclassified aggregate
  get-ddos-ipv6-uncls-aggregate
show ddos-protection protocols ipv6-unclassified parameters
  get-ddos-ipv6-uncls-parameters
show ddos-protection protocols ipv6-unclassified statistics
  get-ddos-ipv6-uncls-statistics
show ddos-protection protocols ipv6-unclassified violations
  get-ddos-ipv6-uncls-violations
show ddos-protection protocols isis
  get-ddos-isis-information
show ddos-protection protocols isis aggregate
  get-ddos-isis-aggregate
show ddos-protection protocols isis parameters
  get-ddos-isis-parameters
show ddos-protection protocols isis statistics
  get-ddos-isis-statistics
show ddos-protection protocols isis violations
  get-ddos-isis-violations
show ddos-protection protocols jfm
  get-ddos-jfm-information
show ddos-protection protocols jfm aggregate
  get-ddos-jfm-aggregate
show ddos-protection protocols jfm parameters
  get-ddos-jfm-parameters
show ddos-protection protocols jfm statistics
  get-ddos-jfm-statistics
show ddos-protection protocols jfm violations
  get-ddos-jfm-violations
```

```
show ddos-protection protocols l2tp
  get-ddos-l2tp-information
show ddos-protection protocols l2tp aggregate
  get-ddos-l2tp-aggregate
show ddos-protection protocols l2tp parameters
  get-ddos-l2tp-parameters
show ddos-protection protocols l2tp statistics
  get-ddos-l2tp-statistics
show ddos-protection protocols l2tp violations
  get-ddos-l2tp-violations
show ddos-protection protocols lacp
  get-ddos-lacp-information
show ddos-protection protocols lacp aggregate
  get-ddos-lacp-aggregate
show ddos-protection protocols lacp parameters
  get-ddos-lacp-parameters
show ddos-protection protocols lacp statistics
  get-ddos-lacp-statistics
show ddos-protection protocols lacp violations
  get-ddos-lacp-violations
show ddos-protection protocols ldp
  get-ddos-ldp-information
show ddos-protection protocols ldp aggregate
  get-ddos-ldp-aggregate
show ddos-protection protocols ldp parameters
  get-ddos-ldp-parameters
show ddos-protection protocols ldp statistics
  get-ddos-ldp-statistics
show ddos-protection protocols ldp violations
  get-ddos-ldp-violations
show ddos-protection protocols ldpv6
  get-ddos-ldpv6-information
show ddos-protection protocols ldpv6 aggregate
  get-ddos-ldpv6-aggregate
show ddos-protection protocols ldpv6 parameters
  get-ddos-ldpv6-parameters
show ddos-protection protocols ldpv6 statistics
  get-ddos-ldpv6-statistics
show ddos-protection protocols ldpv6 violations
  get-ddos-ldpv6-violations
show ddos-protection protocols lldp
  get-ddos-lldp-information
show ddos-protection protocols lldp aggregate
  get-ddos-lldp-aggregate
show ddos-protection protocols lldp parameters
  get-ddos-lldp-parameters
show ddos-protection protocols lldp statistics
  get-ddos-lldp-statistics
show ddos-protection protocols lldp violations
  get-ddos-lldp-violations
show ddos-protection protocols lmp
  get-ddos-lmp-information
show ddos-protection protocols lmp aggregate
  get-ddos-lmp-aggregate
show ddos-protection protocols lmp parameters
  get-ddos-lmp-parameters
```

```
show ddos-protection protocols lmp statistics
  get-ddos-lmp-statistics
show ddos-protection protocols lmp violations
  get-ddos-lmp-violations
show ddos-protection protocols lmpv6
  get-ddos-lmpv6-information
show ddos-protection protocols lmpv6 aggregate
  get-ddos-lmpv6-aggregate
show ddos-protection protocols lmpv6 parameters
  get-ddos-lmpv6-parameters
show ddos-protection protocols lmpv6 statistics
  get-ddos-lmpv6-statistics
show ddos-protection protocols lmpv6 violations
  get-ddos-lmpv6-violations
show ddos-protection protocols mac-host
  get-ddos-mac-host-information
show ddos-protection protocols mac-host aggregate
  get-ddos-mac-host-aggregate
show ddos-protection protocols mac-host parameters
  get-ddos-mac-host-parameters
show ddos-protection protocols mac-host statistics
  get-ddos-mac-host-statistics
show ddos-protection protocols mac-host violations
  get-ddos-mac-host-violations
show ddos-protection protocols mlp
  get-ddos-mlp-information
show ddos-protection protocols mlp aggregate
  get-ddos-mlp-aggregate
show ddos-protection protocols mlp aging-exception
  get-ddos-mlp-aging-exc
show ddos-protection protocols mlp packets
  get-ddos-mlp-packets
show ddos-protection protocols mlp parameters
  get-ddos-mlp-parameters
show ddos-protection protocols mlp statistics
  get-ddos-mlp-statistics
show ddos-protection protocols mlp unclassified
  get-ddos-mlp-unclass
show ddos-protection protocols mlp violations
  get-ddos-mlp-violations
show ddos-protection protocols msdp
  get-ddos-msdp-information
show ddos-protection protocols msdp aggregate
  get-ddos-msdp-aggregate
show ddos-protection protocols msdp parameters
  get-ddos-msdp-parameters
show ddos-protection protocols msdp statistics
  get-ddos-msdp-statistics
show ddos-protection protocols msdp violations
  get-ddos-msdp-violations
show ddos-protection protocols msdpv6
  get-ddos-msdpv6-information
show ddos-protection protocols msdpv6 aggregate
  get-ddos-msdpv6-aggregate
show ddos-protection protocols msdpv6 parameters
  get-ddos-msdpv6-parameters
```

```
show ddos-protection protocols msdpv6 statistics
  get-ddos-msdpv6-statistics
show ddos-protection protocols msdpv6 violations
  get-ddos-msdpv6-violations
show ddos-protection protocols multicast-copy
  get-ddos-mcast-copy-information
show ddos-protection protocols multicast-copy aggregate
  get-ddos-mcast-copy-aggregate
show ddos-protection protocols multicast-copy parameters
  get-ddos-mcast-copy-parameters
show ddos-protection protocols multicast-copy statistics
  get-ddos-mcast-copy-statistics
show ddos-protection protocols multicast-copy violations
  get-ddos-mcast-copy-violations
show ddos-protection protocols mvrp
  get-ddos-mvrp-information
show ddos-protection protocols mvrp aggregate
  get-ddos-mvrp-aggregate
show ddos-protection protocols mvrp parameters
  get-ddos-mvrp-parameters
show ddos-protection protocols mvrp statistics
  get-ddos-mvrp-statistics
show ddos-protection protocols mvrp violations
  get-ddos-mvrp-violations
show ddos-protection protocols ntp
  get-ddos-ntp-information
show ddos-protection protocols ntp aggregate
  get-ddos-ntp-aggregate
show ddos-protection protocols ntp parameters
  get-ddos-ntp-parameters
show ddos-protection protocols ntp statistics
  get-ddos-ntp-statistics
show ddos-protection protocols ntp violations
  get-ddos-ntp-violations
show ddos-protection protocols oam-lfm
  get-ddos-oam-lfm-information
show ddos-protection protocols oam-lfm aggregate
  get-ddos-oam-lfm-aggregate
show ddos-protection protocols oam-lfm parameters
  get-ddos-oam-lfm-parameters
show ddos-protection protocols oam-lfm statistics
  get-ddos-oam-lfm-statistics
show ddos-protection protocols oam-lfm violations
  get-ddos-oam-lfm-violations
show ddos-protection protocols ospf
  get-ddos-ospf-information
show ddos-protection protocols ospf aggregate
  get-ddos-ospf-aggregate
show ddos-protection protocols ospf parameters
  get-ddos-ospf-parameters
show ddos-protection protocols ospf statistics
  get-ddos-ospf-statistics
show ddos-protection protocols ospf violations
  get-ddos-ospf-violations
show ddos-protection protocols ospfv3v6
  get-ddos-ospfv3v6-information
```

```
show ddos-protection protocols ospfv3v6 aggregate
  get-ddos-ospfv3v6-aggregate
show ddos-protection protocols ospfv3v6 parameters
  get-ddos-ospfv3v6-parameters
show ddos-protection protocols ospfv3v6 statistics
  get-ddos-ospfv3v6-statistics
show ddos-protection protocols ospfv3v6 violations
  get-ddos-ospfv3v6-violations
show ddos-protection protocols parameters
  get-ddos-protocols-parameters
show ddos-protection protocols pfe-alive
  get-ddos-pfe-alive-information
show ddos-protection protocols pfe-alive aggregate
  get-ddos-pfe-alive-aggregate
show ddos-protection protocols pfe-alive parameters
  get-ddos-pfe-alive-parameters
show ddos-protection protocols pfe-alive statistics
  get-ddos-pfe-alive-statistics
show ddos-protection protocols pfe-alive violations
  get-ddos-pfe-alive-violations
show ddos-protection protocols pim
  get-ddos-pim-information
show ddos-protection protocols pim aggregate
  get-ddos-pim-aggregate
show ddos-protection protocols pim aggregate culprit-flows
show ddos-protection protocols pim parameters
  get-ddos-pim-parameters
show ddos-protection protocols pim statistics
  get-ddos-pim-statistics
show ddos-protection protocols pim violations
  get-ddos-pim-violations

show ddos-protection protocols pimv6
  <get-ddos-pimv6-information>
show ddos-protection protocols pimv6 aggregate
  <get-ddos-pimv6-aggregate>
show ddos-protection protocols pimv6 aggregate culprit-flows
show ddos-protection protocols pimv6 parameters
  <get-ddos-pimv6-parameters>
show ddos-protection protocols pimv6 statistics
  <get-ddos-pimv6-statistics>
show ddos-protection protocols pimv6 violations
  <get-ddos-pimv6-violations>

show ddos-protection protocols pmvrp
  get-ddos-pmvrp-information
show ddos-protection protocols pmvrp aggregate
  get-ddos-pmvrp-aggregate
show ddos-protection protocols pmvrp parameters
  get-ddos-pmvrp-parameters
show ddos-protection protocols pmvrp statistics
  get-ddos-pmvrp-statistics
show ddos-protection protocols pmvrp violations
  get-ddos-pmvrp-violations
```

```
show ddos-protection protocols pos
  get-ddos-pos-information
show ddos-protection protocols pos aggregate
  get-ddos-pos-aggregate
show ddos-protection protocols pos aggregate culprit-flows
show ddos-protection protocols pos parameters
  get-ddos-pos-parameters
show ddos-protection protocols pos statistics
  get-ddos-pos-statistics
show ddos-protection protocols pos violations
  get-ddos-pos-violations
show ddos-protection protocols ppp
  get-ddos-ppp-information
show ddos-protection protocols ppp aggregate
  get-ddos-ppp-aggregate
show ddos-protection protocols ppp authentication
  get-ddos-ppp-auth
show ddos-protection protocols ppp authentication culprit-flows
show ddos-protection protocols ppp ipcp
  get-ddos-ppp-ipcp
show ddos-protection protocols ppp ipv6cp
  get-ddos-ppp-ipv6cp
show ddos-protection protocols ppp isis
  get-ddos-ppp-isis
show ddos-protection protocols ppp isis culprit-flows
show ddos-protection protocols ppp lcp
  get-ddos-ppp-lcp
show ddos-protection protocols ppp lcp culprit-flows
show ddos-protection protocols ppp mplscp
  get-ddos-ppp-mplscp
show ddos-protection protocols ppp mplscp culprit-flows
show ddos-protection protocols ppp parameters
  get-ddos-ppp-parameters
show ddos-protection protocols ppp statistics
  get-ddos-ppp-statistics
show ddos-protection protocols ppp unclassified
<get-ddos-ppp-unclass>
show ddos-protection protocols ppp violations
  get-ddos-ppp-violations
show ddos-protection protocols pppoe
  get-ddos-pppoe-information
show ddos-protection protocols pppoe aggregate
  get-ddos-pppoe-aggregate
show ddos-protection protocols pppoe padi
  get-ddos-pppoe-padi
show ddos-protection protocols pppoe padm
  get-ddos-pppoe-padm
show ddos-protection protocols pppoe padn
  get-ddos-pppoe-padn
show ddos-protection protocols pppoe pado
  get-ddos-pppoe-pado
show ddos-protection protocols pppoe padr
  get-ddos-pppoe-padr
show ddos-protection protocols pppoe pads
  get-ddos-pppoe-pads
show ddos-protection protocols pppoe padt
```



```
get-ddos-pppoe-padt
show ddos-protection protocols pppoe parameters
get-ddos-pppoe-parameters
show ddos-protection protocols pppoe statistics
get-ddos-pppoe-statistics
show ddos-protection protocols pppoe violations
get-ddos-pppoe-violations
show ddos-protection protocols ptp
get-ddos-ntp-information
show ddos-protection protocols ntp aggregate
get-ddos-ntp-aggregate
show ddos-protection protocols ntp aggregate culprit-flows
show ddos-protection protocols ntp parameters
get-ddos-ntp-parameters
show ddos-protection protocols ntp statistics
get-ddos-ntp-statistics
show ddos-protection protocols ntp violations
get-ddos-ntp-violations
show ddos-protection protocols pvstp
get-ddos-pvstp-information
show ddos-protection protocols pvstp aggregate
get-ddos-pvstp-aggregate
show ddos-protection protocols pvstp parameters
get-ddos-pvstp-parameters
show ddos-protection protocols pvstp statistics
get-ddos-pvstp-statistics
show ddos-protection protocols pvstp violations
get-ddos-pvstp-violations
show ddos-protection protocols radius
get-ddos-radius-information
show ddos-protection protocols radius accounting
get-ddos-radius-account
show ddos-protection protocols radius aggregate
get-ddos-radius-aggregate
show ddos-protection protocols radius accounting culprit-flows
show ddos-protection protocols radius authorization
get-ddos-radius-auth
show ddos-protection protocols radius parameters
get-ddos-radius-parameters
show ddos-protection protocols radius server
get-ddos-radius-server
show ddos-protection protocols radius statistics
get-ddos-radius-statistics
show ddos-protection protocols radius violations
get-ddos-radius-violations
show ddos-protection protocols redirect
get-ddos-redirect-information
show ddos-protection protocols redirect aggregate
get-ddos-redirect-aggregate
show ddos-protection protocols redirect parameters
get-ddos-redirect-parameters
show ddos-protection protocols redirect statistics
get-ddos-redirect-statistics
show ddos-protection protocols redirect violations
get-ddos-redirect-violations
```

```
show ddos-protection protocols reject
  <get-ddos-reject-information>
show ddos-protection protocols reject aggregate
  <get-ddos-reject-aggregate>
show ddos-protection protocols reject parameters
  <get-ddos-reject-parameters>
show ddos-protection protocols reject statistics
  <get-ddos-reject-statistics>
show ddos-protection protocols reject violations
  <get-ddos-reject-violations>
show ddos-protection protocols rejectv6show ddos-protection protocols rejectv6
aggregate
show ddos-protection protocols rejectv6 aggregate culprit-flows
show ddos-protection protocols rejectv6 flow-detection
show ddos-protection protocols rejectv6 parameters
show ddos-protection protocols rejectv6 statistics
show ddos-protection protocols rejectv6 violations
show ddos-protection protocols rip
  get-ddos-rip-information
show ddos-protection protocols rip aggregate
  get-ddos-rip-aggregate
show ddos-protection protocols rip aggregate culprit-flows
show ddos-protection protocols rip culprit-flows
show ddos-protection protocols rip parameters
  get-ddos-rip-parameters
show ddos-protection protocols rip statistics
  get-ddos-rip-statistics
show ddos-protection protocols rip violations
  get-ddos-rip-violations
show ddos-protection protocols ripv6
  get-ddos-ripv6-information
show ddos-protection protocols ripv6 aggregate
  get-ddos-ripv6-aggregate
show ddos-protection protocols ripv6 aggregate culprit-flows
show ddos-protection protocols ripv6 parameters
  get-ddos-ripv6-parameters
show ddos-protection protocols ripv6 statistics
  get-ddos-ripv6-statistics
show ddos-protection protocols ripv6 violations
  get-ddos-ripv6-violations
show ddos-protection protocols rsvp
  get-ddos-rsvp-information
show ddos-protection protocols rsvp aggregate
  get-ddos-rsvp-aggregate
show ddos-protection protocols rsvp aggregate culprit-flows
show ddos-protection protocols rsvp parameters
  get-ddos-rsvp-parameters
show ddos-protection protocols rsvp statistics
  get-ddos-rsvp-statistics
show ddos-protection protocols rsvp violations
  get-ddos-rsvp-violations
show ddos-protection protocols rsvpv6
  get-ddos-rsvpv6-information
show ddos-protection protocols rsvpv6 aggregate
  get-ddos-rsvpv6-aggregate
```

```
show ddos-protection protocols rsvpv6 aggregate culprit-flows
show ddos-protection protocols rsvpv6 parameters
  get-ddos-rsvpv6-parameters
show ddos-protection protocols rsvpv6 statistics
  get-ddos-rsvpv6-statistics
show ddos-protection protocols rsvpv6 violations
  get-ddos-rsvpv6-violations
show ddos-protection protocols sample
<get-ddos-sample-information>
show ddos-protection protocols sample aggregate
<get-ddos-sample-aggregate>
show ddos-protection protocols sample aggregate culprit-flows
show ddos-protection protocols sample host
<get-ddos-sample-host>
show ddos-protection protocols sample parameters
<get-ddos-sample-parameters>
show ddos-protection protocols sample pfe
<get-ddos-sample-pfe>
show ddos-protection protocols sample pfe culprit-flows
show ddos-protection protocols sample statistics
<get-ddos-sample-statistics>
show ddos-protection protocols sample syslog
show ddos-protection protocols sample tap
<get-ddos-sample-tap>
show ddos-protection protocols sample tap culprit-flows
show ddos-protection protocols sample violations
<get-ddos-sample-violations>
show ddos-protection protocols services
  get-ddos-services-information
show ddos-protection protocols services aggregate
  get-ddos-services-aggregate
show ddos-protection protocols services parameters
  get-ddos-services-parameters
show ddos-protection protocols services statistics
  get-ddos-services-statistics
show ddos-protection protocols services violations
  get-ddos-services-violations
show ddos-protection protocols snmp
  get-ddos-snmp-information
show ddos-protection protocols snmp aggregate
  get-ddos-snmp-aggregate
show ddos-protection protocols snmp aggregate culprit-flows
show ddos-protection protocols snmp parameters
  get-ddos-snmp-parameters
show ddos-protection protocols snmp statistics
  get-ddos-snmp-statistics
show ddos-protection protocols snmp violations
  get-ddos-snmp-violations
show ddos-protection protocols snmpv6
  get-ddos-snmpv6-information
show ddos-protection protocols snmpv6 aggregate
  get-ddos-snmpv6-aggregate
show ddos-protection protocols snmpv6 aggregate culprit-flows
show ddos-protection protocols snmpv6 parameters
  get-ddos-snmpv6-parameters
show ddos-protection protocols snmpv6 statistics
```

```
get-ddos-snmpv6-statistics
show ddos-protection protocols snmpv6 violations
get-ddos-snmpv6-violations
show ddos-protection protocols ssh
get-ddos-ssh-information
show ddos-protection protocols ssh aggregate
get-ddos-ssh-aggregate
show ddos-protection protocols ssh parameters
get-ddos-ssh-parameters
show ddos-protection protocols ssh statistics
get-ddos-ssh-statistics
show ddos-protection protocols ssh violations
get-ddos-ssh-violations
show ddos-protection protocols sshv6
get-ddos-sshv6-information
show ddos-protection protocols sshv6 aggregate
get-ddos-sshv6-aggregate
show ddos-protection protocols sshv6 parameters
get-ddos-sshv6-parameters
show ddos-protection protocols sshv6 statistics
get-ddos-sshv6-statistics
show ddos-protection protocols sshv6 violations
get-ddos-sshv6-violations
show ddos-protection protocols statistics
get-ddos-protocols-statistics
show ddos-protection protocols stp
get-ddos-stp-information
show ddos-protection protocols stp aggregate
get-ddos-stp-aggregate
show ddos-protection protocols stp parameters
get-ddos-stp-parameters
show ddos-protection protocols stp statistics
get-ddos-stp-statistics
show ddos-protection protocols stp violations
get-ddos-stp-violations
show ddos-protection protocols tacacs
get-ddos-tacacs-information
show ddos-protection protocols tacacs aggregate
get-ddos-tacacs-aggregate
show ddos-protection protocols tacacs parameters
get-ddos-tacacs-parameters
show ddos-protection protocols tacacs statistics
get-ddos-tacacs-statistics
show ddos-protection protocols tacacs violations
get-ddos-tacacs-violations
show ddos-protection protocols tcp-flags
get-ddos-tcp-flags-information
show ddos-protection protocols tcp-flags aggregate
get-ddos-tcp-flags-aggregate
show ddos-protection protocols tcp-flags established
get-ddos-tcp-flags-establish
show ddos-protection protocols tcp-flags initial
get-ddos-tcp-flags-initial
show ddos-protection protocols tcp-flags parameters
get-ddos-tcp-flags-parameters
show ddos-protection protocols tcp-flags statistics
```

```
get-ddos-tcp-flags-statistics
show ddos-protection protocols tcp-flags unclassified
get-ddos-tcp-flags-unclass
show ddos-protection protocols tcp-flags violations
get-ddos-tcp-flags-violations
show ddos-protection protocols telnet
get-ddos-telnet-information
show ddos-protection protocols telnet aggregate
get-ddos-telnet-aggregate
show ddos-protection protocols telnet aggregate culprit-flows
show ddos-protection protocols telnet parameters
get-ddos-telnet-parameters
show ddos-protection protocols telnet statistics
get-ddos-telnet-statistics
show ddos-protection protocols telnet violations
get-ddos-telnet-violations
show ddos-protection protocols telnetv6
get-ddos-telnetv6-information
show ddos-protection protocols telnetv6 aggregate
get-ddos-telnetv6-aggregate
show ddos-protection protocols telnetv6 aggregate culprit-flows
show ddos-protection protocols telnetv6 parameters
get-ddos-telnetv6-parameters
show ddos-protection protocols telnetv6 statistics
get-ddos-telnetv6-statistics
show ddos-protection protocols telnetv6 violations
get-ddos-telnetv6-violations
show ddos-protection protocols ttl
get-ddos-ttl-information
show ddos-protection protocols ttl aggregate
get-ddos-ttl-aggregate
show ddos-protection protocols ttl parameters
get-ddos-ttl-parameters
show ddos-protection protocols ttl statistics
get-ddos-ttl-statistics
show ddos-protection protocols ttl violations
get-ddos-ttl-violations
show ddos-protection protocols tunnel-fragment
get-ddos-tun-frag-information
show ddos-protection protocols tunnel-fragment aggregate
get-ddos-tun-frag-aggregate
show ddos-protection protocols tunnel-fragment aggregate culprit-flows
show ddos-protection protocols tunnel-fragment parameters
get-ddos-tun-frag-parameters
show ddos-protection protocols tunnel-fragment statistics
get-ddos-tun-frag-statistics
show ddos-protection protocols tunnel-fragment violations
get-ddos-tun-frag-violations
show ddos-protection protocols unclassified
<get-ddos-uncls-information>
show ddos-protection protocols unclassified aggregate
<get-ddos-uncls-aggregate>
show ddos-protection protocols unclassified parameters
<get-ddos-uncls-parameters>
show ddos-protection protocols unclassified statistics
<get-ddos-uncls-statistics>
```

```
show ddos-protection protocols unclassified violations
<get-ddos-uncls-violations>
show ddos-protection protocols violations
  get-ddos-protocols-violations
show ddos-protection protocols virtual-chassis
  get-ddos-vchassis-information
show ddos-protection protocols virtual-chassis aggregate
  get-ddos-vchassis-aggregate
show ddos-protection protocols virtual-chassis aggregate culprit-flows
show ddos-protection protocols virtual-chassis control-high
  get-ddos-vchassis-control-hi
show ddos-protection protocols virtual-chassis control-low
  get-ddos-vchassis-control-lo
show ddos-protection protocols virtual-chassis parameters
  get-ddos-vchassis-parameters
show ddos-protection protocols virtual-chassis statistics
  get-ddos-vchassis-statistics
show ddos-protection protocols virtual-chassis unclassified
  get-ddos-vchassis-unclass
show ddos-protection protocols virtual-chassis vc-packets
  get-ddos-vchassis-vc-packets
show ddos-protection protocols virtual-chassis vc-ttl-errors
  get-ddos-vchassis-vc-ttl-err
show ddos-protection protocols virtual-chassis violations
  get-ddos-vchassis-violations
show ddos-protection protocols vrrp
  get-ddos-vrrp-information
show ddos-protection protocols vrrp aggregate
  get-ddos-vrrp-aggregate
show ddos-protection protocols vrrp aggregate culprit-flows
show ddos-protection protocols vrrp parameters
  get-ddos-vrrp-parameters
show ddos-protection protocols vrrp statistics
  get-ddos-vrrp-statistics
show ddos-protection protocols vrrp violations
  get-ddos-vrrp-violations
show ddos-protection protocols vrrpv6
  get-ddos-vrrpv6-information
show ddos-protection protocols vrrpv6 aggregate
  get-ddos-vrrpv6-aggregate
show ddos-protection protocols vrrpv6 aggregate culprit-flows
show ddos-protection protocols vrrpv6 parameters
  get-ddos-vrrpv6-parameters
show ddos-protection protocols vrrpv6 statistics
  get-ddos-vrrpv6-statistics
show ddos-protection protocols vrrpv6 violations
  get-ddos-vrrpv6-violations
show ddos-protection statistics
  get-ddos-statistics-information
show ddos-protection version
  get-ddos-version
show dhcp
show dhcp relay
show dhcp relay binding
  <get-dhcp-relay-binding-information>
```

```
show dhcp relay binding interface
<get-dhcp-relay-interface-bindings>
show dhcp relay statistics
<get-dhcp-relay-statistics-information>

show dhcp server
show dhcp server binding
<get-dhcp-server-binding-information>

show dhcp server binding interface
<get-dhcp-relay-binding-interface>
show dhcp server statistics
<get-dhcp-server-statistics-information>
show dhcp statistics
<get-dhcp-service-statistics-information>
show dhcpv6
show dhcpv6 relay
show dhcpv6 relay binding
<get-dhcpv6-relay-binding-information>
show dhcpv6 relay binding interface
<get-dhcpv6-relay-binding-interface>
show dhcpv6 relay statistics
<get-dhcpv6-relay-statistics-information>
show dhcpv6 server
show dhcpv6 server binding
<get-dhcpv6-server-binding-information>

show dhcpv6 server binding interface
<get-dhcpv6-server-binding-interface>
show dhcpv6 server statistics
<get-dhcpv6-server-statistics-information>
show dhcpv6 statistics
<get-dhcpv6-service-statistics-information>
show diameter
<get-diameter-information>

show diameter function
<get-diameter-function-information>

show diameter function statistics
<get-diameter-function-statistics>

show diameter instance
<get-diameter-instance-information>

show diameter network-element
<get-diameter-network-element-information>

show diameter network-element map
<get-diameter-network-element-map-information>

show diameter peer
<get-diameter-peer-information>

show diameter peer map
<get-diameter-peer-map-information>
```

```
show diameter peer statistics
  <get-diameter-peer-statistics>

show diameter route
  <get-diameter-route-information>

show dot1x
show dot1x authentication-failed-users
  <get-dot1x-authentication-failed-users>

show dot1x interface
  <get-dot1x-interface-information>

show dot1x static-mac-address
  <get-dot1x-static-mac-addresses>

show dot1x static-mac-address interface
  <get-dot1x-interface-mac-addresses>

show dvmrp
show dvmrp interfaces
  <get-dvmrp-interfaces-information>

show dvmrp neighbors
  <get-dvmrp-neighbors-information>

show dvmrp prefix
  <get-dvmrp-prefix-information>

show dvmrp prunes
  <get-dvmrp-prunes-information>

show dynamic-tunnels
show dynamic-tunnels database
  <get-dynamic-tunnels-database>
show esis
show esis adjacency
  <get-esis-adjacency-information>

show esis interface
  <get-esis-interface-information>

show esis statistics
  <get-esis-statistics-information>

show event-options
show event-options event-scripts
show event-options event-scripts policies
  <get-event-scripts-policies>

show extension-provider
show extension-provider system
show extension-provider system connections
  <get-mspinfo-connections>
```



```
show extension-provider system packages
  <get-mspinfo-packages>

show extension-provider system processes
  <get-mspinfo-processes>

show extension-provider system processes brief
  <get-mspinfo-processes-brief>

show extension-provider system processes extensive
  <get-mspinfo-processes-extensive>

show extension-provider system uptime
  <get-mspinfo-uptime>

show extension-provider system virtual-memory
  <get-core-key-list>
  <get-fabric-summary-information>
  <get-key-vg-binding>
  <get-mac-ip-binding-information>
  <get-mc-ccpc-src-mod-filters>
  <get-mc-edge-map-to-key-binding>
  <get-mc-edge-key-to-map-binding>
  <get-mc-edge-vg-portmap>
  <get-mc-nsf>
  <get-mc-root-key-to-map-binding>
  <get-mc-root-map-to-key-binding>
  <get-mc-root-vg-pfemap>
  <get-mc-vccpdf-adjacency-database>
  <get-mspinfo-virtual-memory>
get-fabric-statistics
get-fabric-summary-information
  <get-vlan-domain-map-information>
show forwarding-options
show forwarding-options next-hop-group
  <get-forwarding-options-next-hop-group>
show forwarding-options port-mirroring
  <get-forwarding-options-port-mirroring>
show helper
show helper statistics
  <get-helper-statistics-information>
show hfrf
show hfrf profiles
show iccp
  <get-inter-chassis-control-protocol-information>
show igmp
show igmp group
  <get-igmp-group-information>

show igmp interface
  <get-igmp-interface-information>

show igmp output-group
  <get-igmp-output-group-information>

show igmp snooping
```

```
show igmp snooping interface
  <get-igmp-snooping-interface-information>

show igmp snooping interface bridge-domain
<get-igmp-snooping-bridge-domain-interface>
show igmp snooping membership
  <get-igmp-snooping-membership-information>

show igmp snooping membership bridge-domain
show igmp snooping statistics
  <get-igmp-snooping-statistics-information>

show igmp snooping statistics bridge-domain
<get-igmp-snooping-bridge-domain-membership>
show igmp statistics
  <get-igmp-statistics-information>

show ike
show ike security-associations
  <get-ike-security-associations-information>

show ilmi
<get-ilmi-information>
show ilmi interface
<get-ilmi-interface-information>
show ilmi statistics
<get-ilmi-statistics>
show ingress-replication
  <get-ingress-replication-information>
show interfaces
  <get-interface-information>

show interfaces controller
<get-interface-controller-information>
show interfaces destination-class
  <get-destination-class-statistics>

show interfaces destination-class all
<get-all-destination-class-statistics>
show interfaces diagnostics
show interfaces diagnostics optics
  <get-interface-optics-diagnostics-information>

show interfaces far-end-interval
  <show-interfaces-far-end-interval>
show interfaces filters
  <get-interface-filter-information>

show interfaces interface-set
<get-interface-set-information>
show interfaces interface-set queue
  <get-interface-set-queue-information>

show interfaces interval
  <show-interfaces-interval>
show interfaces load-balancing
```

```
<interface-load-balancing>
show interfaces mac-database
  <get-mac-database>

show interfaces mc-ae
  <get-mc-ae-interface-information>
show interfaces policers
  <get-interface-policer-information>

show interfaces queue
  <get-interface-queue-information>

show interfaces redundancy
  <get-redundancy-status>
show interfaces redundancy detail
  <get-redundancy-status-details>
show interfaces routing
show interfaces source-class
  <get-source-class-statistics>

show interfaces source-class all
  <get-all-source-class-statistics>
show interfaces targeting
  <get-targeting-information>
show ipsec
show ipsec redundancy
show ipsec redundancy interface
  <get-ipsec-pic-redundancy-information>

show ipsec redundancy security-associations
  <get-ipsec-tunnel-redundancy-information>

show ipsec security-associations
  <get-security-associations-information>

show ipv6
show ipv6 neighbors
  <get-ipv6-nd-information>

show ipv6 router-advertisement
  <get-ipv6-ra-information>

show isis
show isis adjacency
  <get-isis-adjacency-information>

show isis authentication
  <get-isis-authentication-information>

show isis backup
show isis backup coverage
  <get-isis-backup-coverage-information>

show isis backup label-switched-path
  <get-isis-backup-lsp-information>
```

```
show isis backup spf

show isis backup spf results
  <get-isis-backup-spf-results-information>

show isis context-identifier
  <get-isis-context-identifier-information>

show isis context-identifier identifier
  <get-isis-context-identifier-origin-information>
show isis database
  <get-isis-database-information>

show isis hostname
  <get-isis-hostname-information>

show isis interface
  <get-isis-interface-information>

show isis overview
  <get-isis-overview-information>

show isis route
  <get-isis-route-information>

show isis spf
show isis spf brief
  <get-isis-spf-results-brief-information>

show isis spf log
  <get-isis-spf-log-information>

show isis spf results
  <get-isis-spf-results-information>

show isis statistics
  <get-isis-statistics-information>

show l2-learning
show l2-learning backbone-instance
  <get-l2-learning-backbone-instance>
show l2-learning global-information
  <get-l2-learning-global-information>
show l2-learning global-mac-count
  <get-l2-learning-global-mac-count>
show l2-learning instance
  <get-l2-learning-routing-instances>
show l2-learning interface
  <get-l2-learning-interface-information>
show l2-learning mac-move-buffer
  <get-l2-learning-mac-move-buffer-information>
show l2-learning provider-instance
  <get-l2-learning-provider-instance>
show l2-learning redundancy-groups
  <get-l2-learning-redundancy-groups>
show l2-learning remote-backbone-edge-bridges
```

```
<get-l2-learning-remote-backbone-edge-bridges>
show l2circuit
show l2circuit connections
  <get-l2ckt-connection-information>

show l2cpd
show l2cpd task
  <get-l2cpd-task-information>
show l2cpd task io
  <get-l2cpd-tasks-io-statistics>
show l2cpd task memory
  <get-l2cpd-task-memory>
show l2cpd task replication
  <get-l2cpd-replication-information>
show l2vpn
show l2vpn connections
  <get-l2vpn-connection-information>

show lacp
show lacp interfaces
  <get-lacp-interface-information>
show lacp statistics
show lacp statistics interfaces
  <get-lacp-interface-statistics>
show lacp timeouts
show ldp
show ldp database
  <get-ldp-database-information>

show ldp fec-filters
  <get-ldp-fec-filters-information>

show ldp interface
  <get-ldp-interface-information>

show ldp neighbor
  <get-ldp-neighbor-information>

show ldp oam
  <get-ldp-oam-information>
show ldp overview
  <get-ldp-overview-information>
show ldp path
  <get-ldp-path-information>

show ldp route
  <get-ldp-route-information>

show ldp session
  <get-ldp-session-information>

show ldp statistics
  <get-ldp-statistics-information>

show ldp traffic-statistics
  <get-ldp-traffic-statistics-information>
```

```
show link-management
  <get-lm-information>

show link-management peer
  <get-lm-peer-information>

show link-management routing
  <get-lm-routing-information>

show link-management routing peer
  <get-lm-routing-peer-information>

show link-management routing resource
  <get-lm-routing-resource-information>

show link-management routing te-link
  <get-lm-routing-te-link-information>

show lldp
  <get-lldp-information>

show lldp detail
  <get-lldp-information-detail>

show lldp local-information
  <get-lldp-local-info>

show lldp neighbors
  <get-lldp-neighbors-information>

show lldp neighbors interface
  <get-lldp-interface-neighbors>
show lldp remote-global-statistics
  <get-lldp-remote-global-statistics>

show lldp statistics
  <get-lldp-statistics-information>

show lldp statistics interface
  <get-lldp-interface-statistics>
show link-management statistics
  <get-lm-statistics-information>

show link-management statistics peer
  <get-lm-peer-statistics>

show link-management te-link
  <get-lm-te-link-information>

show mac-rewrite
show mac-rewrite interface
  <get-mac-rewrite-interface-information>
show mld
show mld group
  <get-mld-group-information>
```

```
show mld interface
  <get-mld-interface-information>

show mld output-group
  <get-mld-output-group-information>

show mld statistics
  <get-mld-statistics-information>

show mobile-ip
show mobile-ip home-agent
show mobile-ip home-agent binding
  <get-mip-binding-information>

show mobile-ip home-agent binding ip-address
  <get-ip-mip-binding-information>

show mobile-ip home-agent binding nai
  <get-nai-mip-binding-information>

show mobile-ip home-agent binding summary
  <get-summary-mip-binding-information>

show mobile-ip home-agent interface
  <get-mip-ha-interface-information>

show mobile-ip home-agent overview
  <get-mip-ha-overview-information>

show mobile-ip home-agent traffic
  <get-mip-ha-traffic-information>

show mobile-ip home-agent virtual-network
  <get-mip-ha-virtual-network-information>

show mobile-ip tunnel
  <get-mip-tunnel-information>
show mobile-ip wimax
show mobile-ip wimax release
  <get-mip-wimax-release-information>

show mpls
show mpls admin-groups
  <get-mpls-admin-group-information>

show mpls admin-groups-extended
  <get-mpls-admin-group-extended-information>
show mpls call-admission-control
  <get-mpls-call-admission-control-information>

show mpls context-identifier
  <get-mpls-context-identifier-information>

show network-access domain-map
show network-access domain-map statistics
```

```
<get-domain-map-statistics>
show mpls cspf
  <get-mpls-cspf-information>

show mpls diffserv-te
  <get-mpls-diffserv-te-information>
show mpls egress-protection
show mpls interface
  <get-mpls-interface-information>

show mpls lsp
  <get-mpls-lsp-information>

show mpls lsp autobandwidth
  <get-mpls-lsp-autobandwidth>
show mpls srlg
  <get-mpls-srlg-information>
show oam ethernet fnp
show oam ethernet fnp interface
show oam ethernet fnp messages
show oam ethernet fnp status
  <get-fnp-status>
show mpls lsp defaults
  <get-mpls-lsp-defaults-information>

show mpls path
  <get-mpls-path-information>

show mpls static-lsp
  <get-mpls-static-lsp-information>
show mpls traceroute
show mpls traceroute database
show mpls traceroute database ldp
  <get-mpls-traceroute-database-ldp>
show msdp
  <get-msdp-information>
show msdp source
  <get-msdp-source-information>

show msdp source-active
  <get-msdp-source-active-information>

show msdp statistics
  <get-msdp-statistics-information>

show multicast
show multicast backup-pe-groups
  <get-multicast-backup-pe-groups-information>

show multicast backup-pe-groups address
  <get-multicast-backup-pe-address-information>

show multicast backup-pe-groups group
  <get-multicast-backup-pe-group-information>
show multicast flow-map
  <get-multicast-flow-maps-information>
```



```
show multicast interface
  <get-multicast-interface-information>

show multicast next-hops
  <get-multicast-next-hops-information>

show multicast pim-to-igmp-proxy
  <get-multicast-pim-to-igmp-proxy-information>

show multicast pim-to-mld-proxy
  <get-multicast-pim-to-mld-proxy-information>

show multicast route
  <get-multicast-route-information>

show multicast rpf
  <get-multicast-rpf-information>

show multicast scope
  <get-multicast-scope-information>

show multicast sessions
  <get-multicast-sessions-information>

show multicast snooping
show multicast snooping next-hops
  <get-multicast-snooping-next-hops-information>

show multicast snooping route
  <get-multicast-snooping-route-information>

show multicast statistics
  <get-multicast-statistics-information>

show multicast usage
  <get-multicast-usage-information>

show mvpn
show mvpn c-multicast
  <get-mvpn-c-multicast-route>
show mvpn instance
  <get-mvpn-instance-information>

show mvpn neighbor
  <get-mvpn-neighbor-information>
show mvrp
  <get-mvrp-information>

show mvrp applicant-state
  <get-mvrp-applicant-information>

show mvrp dynamic-vlan-memberships
  <get-mvrp-dynamic-vlan-memberships>

show mvrp interface
```

```
<get-mvrp-interface-information>

show mvrp registration-state
  <get-mvrp-registration-state>

show mvrp statistics
  <get-mvrp-interface-statistics>

show network-access
show network-access aaa
show network-access aaa radius-servers
  <get-radius-servers-table>
show network-access aaa statistics
  <get-aaa-module-statistics>

show network-access aaa statistics address-assignment
show network-access aaa statistics address-assignment client
  <get-address-assignment-client-statistics>
show network-access aaa statistics address-assignment pool
  <get-address-assignment-pool-statistics>
show network-access aaa subscribers
  <get-aaa-subscriber-table>

show network-access aaa subscribers session-id

show network-access aaa subscribers statistics
  <get-aaa-subscriber-statistics>

show network-access aaa terminate-code
  <get-aaa-terminate-code>
show network-access aaa terminate-code aaa
  <get-aaa-terminate-code-aaa>
show network-access aaa terminate-code dhcp
  <get-aaa-terminate-code-dhcp>
show network-access aaa terminate-code l2tp
  <get-aaa-terminate-code-l2tp>
show network-access aaa terminate-code ppp
  <get-aaa-terminate-code-ppp>
show network-access aaa terminate-code reverse
  <get-aaa-terminate-code-reverse>
show network-access aaa terminate-code reverse aaa
  <get-aaa-terminate-code-reverse-aaa>
show network-access aaa terminate-code reverse dhcp
  <get-aaa-terminate-code-reverse-dhcp>
show network-access aaa terminate-code reverse l2tp
  <get-aaa-terminate-code-reverse-l2tp>
show network-access aaa terminate-code reverse ppp
  <get-aaa-terminate-code-reverse-ppp>
show network-access address-assignment
show network-access address-assignment pool
  <get-address-assignment-pool-table>

show network-access requests
show network-access requests pending
  <get-authentication-pending-table>
```

```
show network-access requests statistics
  <get-authentication-statistics>

show network-access securid-node-secret-file
  <get-node-secret-file-table>

show ntp
show ntp associations
show ntp status
show oam
show oam ethernet
show oam ethernet connectivity-fault-management
show oam ethernet connectivity-fault-management delay-statistics
  <get-cfm-delay-statistics>

show oam ethernet connectivity-fault-management forwarding-state
show oam ethernet connectivity-fault-management forwarding-state instance
  <get-cfm-forwarding-state-instance-information>

show oam ethernet connectivity-fault-management forwarding-state interface
  <get-cfm-forwarding-state-interface-information>

show oam ethernet connectivity-fault-management interfaces
  <get-cfm-interfaces-information>
show oam ethernet connectivity-fault-management loss-statistics
  <get-cfm-loss-statistics>
show oam ethernet connectivity-fault-management mep-database
  <get-cfm-mep-database>

show oam ethernet connectivity-fault-management mep-statistics
  <get-cfm-mep-statistics>

show oam ethernet connectivity-fault-management mip
  <get-cfm-mip-information>

show oam ethernet connectivity-fault-management path-database
  <get-cfm-linktrace-path-database>

show oam ethernet connectivity-fault-management policer
  <get-evc-information>

show oam ethernet connectivity-fault-management sla-iterator-statistics
  <get-cfm-iterator-statistics>
show oam ethernet evc
  <get-evc-information>
show oam ethernet link-fault-management
  <get-lfmd-information>

show oam ethernet lmi
  <get-elmi-information>

show oam ethernet lmi statistics
  <get-elmi-statistics>

show ospf
show ospf backup
```

```
show ospf backup coverage
  <get-ospf-backup-coverage-information>

show ospf backup lsp
  <get-ospf-backup-lsp-information>

show ospf backup neighbor
  <get-ospf-backup-neighbor-information>

show ospf backup spf
  <get-ospf-backup-spf-information>

show ospf context-identifier
  <get-ospf-context-id-information>

show ospf database
  <get-ospf-database-information>

show ospf interface
  <get-ospf-interface-information>

show ospf io-statistics
  <get-ospf-io-statistics-information>

show ospf log
  <get-ospf-log-information>

show ospf neighbor
  <get-ospf-neighbor-information>

show ospf overview
  <get-ospf-overview-information>

show ospf route
  <get-ospf-route-information>

show ospf statistics
  <get-ospf-statistics-information>

show ospf3
show ospf3 backup
show ospf3 backup coverage
  <get-ospf3-backup-coverage-information>

show ospf3 backup lsp
  <get-ospf3-backup-lsp-information>

show ospf3 backup neighbor
  <get-ospf3-backup-neighbor-information>

show ospf3 backup spf
  <get-ospf3-backup-spf-information>

show ospf3 database
  <get-ospf3-database-information>
```

```
show ospf3 interface
  <get-ospf3-interface-information>

show ospf3 io-statistics
  <get-ospf3-io-statistics-information>

show ospf3 log
  <get-ospf3-log-information>

show ospf3 neighbor
  <get-ospf3-neighbor-information>

show ospf3 overview
  <get-ospf3-overview-information>

show ospf3 route
  <get-ospf3-route-information>

show ospf3 statistics
  <get-ospf3-statistics-information>

show passive-monitoring
  <get-passive-monitoring-information>

show passive-monitoring error
  <get-passive-monitoring-error-information>

show passive-monitoring flow
  <get-passive-monitoring-flow-information>

show passive-monitoring memory
  <get-passive-monitoring-memory-information>

show passive-monitoring status
  <get-passive-monitoring-status-information>

show passive-monitoring usage
  <get-passive-monitoring-usage-information>
show path-computation-client
show path-computation-client active-pce
show path-computation-client statistics
show pfe
show pfe cfeb
show pfe feb
show pfe fpc
show pfe fwdd
show pfe lcc
show pfe next-hop
show pfe pfem
show pfe pfem detail
show pfe pfem extensive
show pfe route
show pfe route clnp
show pfe route clnp table
show pfe route inet6
show pfe route inet6 table
```

```
show pfe route ip
show pfe route ip table
show pfe route iso
show pfe route iso table
show pfe scb
show pfe sfm
show pfe ssb
show pfe statistics
show pfe statistics fabric
show pfe statistics ip
show pfe statistics ip6
show pfe statistics traffic
    <get-pfe-statistics>

show pfe statistics traffic cpu
show pfe statistics traffic cpu fpe
show pfe statistics traffic egress-queues
show pfe statistics traffic egress-queues fpc
show pfe statistics traffic multicast
show pfe statistics traffic multicast fpcshow pfe statistics traffic protocol
show pfe terse
    <get-pfe-information>

show pfe version brief
show pfe version detail
show pgm
show pgm negative-acknowledgments
    <get-pgm-nak>

show pgm source-path-messages
    <get-pgm-source-path-messages>

show pgm statistics
    <get-pgm-statistics>

show pim
show pim bidirectional
show pim bidirectional df-election
    <get-pim-bidir-df-election-information>
show pim bidirectional df-election interface
    <get-pim-bidir-df-election-interface-information>
show pim bootstrap
    <get-pim-bootstrap-information>

show pim interfaces
    <get-pim-interfaces-information>

show pim join
    <get-pim-join-information>

show pim mdt
    <get-pim-mdt-information>

show pim mdt data-mdt-joins
    <get-pim-data-mdt-join-information>
show pim mvpn
```

```
<get-pim-mvpn-information>

show pim neighbors
  <get-pim-neighbors-information>

show pim rps
  <get-pim-rps-information>
show pim snooping
show pim snooping interfaces
show pim snooping join
show pim snooping neighbors
show pim snooping statistics
show pim source
  <get-pim-source-information>

show pim statistics
  <get-pim-statistics-information>

show policy
show policy conditions
show policy damping
show ppp
show ppp address-pool
  <get-ppp-address-pool-information>

show ppp interface
  <get-ppp-interface-information>

show ppp statistics
  <get-ppp-statistics-information>

show ppp summary
  <get-ppp-summary-information>

show pppoe
show pppoe interfaces
  <get-pppoe-interface-information>
show pppoe lockout
  <get-pppoe-lockout-information>

show pppoe service-name-tables
  <get-pppoe-service-name-table-information>

show pppoe statistics
  <get-pppoe-statistics-information>

show pppoe underlying-interfaces
  <get-pppoe-underlying-interface-information>

show pppoe version
  <get-pppoe-version>

show protection-group
show protection-group ethernet-aps
  <show-protection-group-ethernet-aps>
show protection-group ethernet-ring
```

```
show protection-group ethernet-ring aps
  <get-raps-pdu-information>
show protection-group ethernet-ring data-channel
  <get-ring-data-channel-information>
show protection-group ethernet-ring interface
  <get-ring-interface-information>
show protection-group ethernet-ring node-state
  <get-raps-state-machine-information>
show protection-group ethernet-ring node-state
show protection-group ethernet-ring statistics
  <get-ring-tatistics>
show protection-group ethernet-ring vlan
  <get-ring-vlan-information>
show ptp
show ptp clock
  get-ptp-clock>
show ptp global-information
  get-ptp-global-information>
show ptp hybrid
show ptp hybrid config
  <get-ptp-hybrid-mapping>
show ptp hybrid status
  <get-ptp-hybrid-status>
show ptp last-tod-update
  <get-last-tod-update>
show ptp lock-status
  get-ptp-lock-status>
show ptp master
  <get-ptp-master>
show ptp port
  <get-ptp-port>
show ptp quality-level-mapping
  <get-ptp-quality-level-mapping>
show ptp slave
  <get-ptp-slave>
show ptp statistics
  <get-ptp-statistics>
show r2cp
show r2cp interfaces
  <get-r2cp-interface-information>
show r2cp radio
  <get-r2cp-radio-information>
show r2cp sessions
  <get-r2cp-session-information>
show r2cp statistics
  <get-r2cp-statistics>
show redundant-power-system
show redundant-power-system led
show redundant-power-system multi-backup
  <get-rps-scale-information>
show redundant-power-system network
  <get-rps-network-information>
show redundant-power-system power-supply
show redundant-power-system status
show redundant-power-system upgrade
  <get-rps-upgrade-information>
```



```
show redundant-power-system version
show rip
show rip general-statistics
  <get-rip-general-statistics-information>

show rip neighbor
  <get-rip-neighbor-information>

show rip statistics
  <get-rip-statistics-information>
show rip statistics peer
  <get-rip-peer-information>
show ripng
show ripng general-statistics
  <get-ripng-general-statistics-information>

show ripng neighbor
  <get-ripng-neighbor-information>
show ripng statistics
  <get-ripng-statistics-information>
show route
  <get-route-information>

show route export
  <get-rlexport-table-information>

show route export instance
  <get-rlexport-instance-information>

show route localization
  <get-fib-localization-information>
show route export vrf-target
  <get-rlexport-target-information>

show route flow
show route flow validation
  <get-rtflow-dep-information>

show route forwarding-table
  <get-forwarding-table-information>

show route instance
  <get-instance-information>

show route instance operational
  <get-operational-routing-instance-information>

show route martians
  <get-route-martians>
show route resolution
  <get-route-resolution-information>
show route resolution summary
  <get-route-resolution-summary>
show route resolution unresolved
show route rib-groups
  <get-route-rib-groups>
```

```
show route snooping
<get-route-snooping-information>
show route snooping summary
<get-route-snooping-summary>
show route summary
<get-route-summary-information>

show rsvp
show rsvp interface
<get-rsvp-interface-information>

show rsvp neighbor
<get-rsvp-neighbor-information>

show rsvp session
<get-rsvp-session-information>

show rsvp statistics
<get-rsvp-statistics-information>

show rsvp version
<get-rsvp-version-information>

show sap
show sap listen
<get-sap-listen-information>

show services
show services accounting
<get-service-accounting-information>

show services accounting aggregation
<get-service-accounting-aggregation-information>

show services accounting aggregation as
<get-service-accounting-aggregation-as-information>

show services accounting aggregation destination-prefix
<get-service-accounting-aggregation-destination-prefix-information>

show services accounting aggregation protocol-port
<get-service-accounting-aggregation-protocol-port-information>

show services accounting aggregation source-destination-prefix
<get-service-accounting-aggregation-source-destination-prefix-information>

show services accounting aggregation source-prefix
<get-service-accounting-aggregation-source-prefix-information>

show services accounting aggregation template
<get-service-accounting-aggregation-template-information>

show services accounting errors
<get-service-accounting-errors-information>

show services accounting flow
```

```
<get-service-accounting-flow-information>

show services accounting flow-detail
  <get-service-accounting-flow-detail>

show services accounting memory
  <get-service-accounting-memory-information>

show services accounting packet-size-distribution
  <get-packet-distribution-information>

show services accounting status
  <get-service-accounting-status-information>

show services accounting usage
  <get-service-accounting-usage-information>

show services alg
show services alg conversations
  <get-service-msp-alg-conversation-information>
show services alg sip-globals
  <get-service-msp-alg-sip-globals-information>
show services alg statistics
show services application-aware-access-list
show services application-aware-access-list flows
show services application-aware-access-list flows interface
  <get-application-aware-access-list-flows-interface>
show services application-aware-access-list flows subscriber
  <get-application-aware-access-list-flows-subscriber>
show services application-aware-access-list statistics
show services application-aware-access-list statistics interface
  <get-application-aware-access-list-statistics-interface>
show services application-aware-access-list statistics subscriber
  <get-application-aware-access-list-statistics-subscriber>
show services application-identification
show services application-identification application
show services application-identification application detail
  <get-appid-application-signature-detail>
show services application-identification application summary
  <get-appid-application-signature-summary>
show services application-identification application-system-cache
  <get-appid-application-system-cache>

show services application-identification counter
  <get-appid-counter>
show services application-identification counter ssl-encrypted-sessions
  <get-appid-counter-encrypted>
show services application-identification group
show services application-identification group detail
  <get-appid-application-group-detail>
show services application-identification group summary
  <get-appid-application-group-summary>
show services application-identification statistics
show services application-identification statistics application-groups
  <get-appid-application-group-statistics>
```

```
show services application-identification statistics applications
  <get-appid-application-statistics>
show services application-identification version
  <get-appid-package-version>

show services border-signaling-gateway
show services border-signaling-gateway accounting
show services border-signaling-gateway accounting statistics
  <get-service-border-signaling-gateway-charging-statistics>
show services border-signaling-gateway accounting status
  <get-service-border-signaling-gateway-charging-status>
show services border-signaling-gateway admission-control
  <get-service-border-signaling-gateway-statistics-admission-control>

show services border-signaling-gateway by-call-context-id
  <get-service-bsg-information-by-call-context-id>

show services border-signaling-gateway by-contact
  <get-service-border-signaling-gateway-information-by-contact>

show services border-signaling-gateway by-request-uri
  <get-service-border-signaling-gateway-information-by-request-uri>

show services border-signaling-gateway calls
  <get-service-border-signaling-gateway-statistics-calls>

show services border-signaling-gateway calls-duration
  <get-service-border-signaling-gateway-calls-duration>

show services border-signaling-gateway calls-failed

how services border-signaling-gateway charging
show services border-signaling-gateway charging statistics
  <get-service-border-signaling-gateway-charging-statistics>
show services border-signaling-gateway charging status
  <get-service-border-signaling-gateway-charging-status>
show services border-signaling-gateway denied-messages
  <get-service-bsg-denied-messages>

show services border-signaling-gateway embedded-spdf
  <get-service-border-signaling-gateway-embedded-spdf>

show services border-signaling-gateway embedded-spdf status
  <get-service-border-signaling-gateway-embedded-spdf-status>

show services border-signaling-gateway name-resolution-cache

show services border-signaling-gateway name-resolution-cache all
  <get-service-border-signaling-gateway-name-resolution-cache-all>

show services border-signaling-gateway name-resolution-cache by-fqdn
  <get-border-signaling-gateway-name-resolution-cache-by-fqdn>
show services border-signaling-gateway status
  <get-service-bsg-status-information>
show services captive-portal-content-delivery
```

```
show services captive-portal-content-delivery pic
  <get-cpcd-pic-information>
show services captive-portal-content-delivery profile
  <get-cpcd-profile>
show services captive-portal-content-delivery rule
  <get-cpcd-rule>
show services captive-portal-content-delivery ruleset
  <get-cpcd-rule-set>
show services captive-portal-content-delivery sset
  <get-cpcd-service-set>
show services captive-portal-content-delivery statistics
  <get-cpcd-pic-statistics>
show services captive-portal-content-delivery statistics interface
show services cos
show services cos statistics
  <get-service-cos-statistics-information>

show services cos statistics diffserv
  <get-service-cos-diffserv-statistics>

show services cos statistics forwarding-class
  <get-service-cos-forwarding-class-statistics>

show services crtp
  <get-service-crtp-params-information>

show services crtp extensive
  <get-service-crtp-extensive-information>

show services crtp flows
  <get-service-crtp-flow-table-information>

show services dynamic-flow-capture
show services dynamic-flow-capture content-destination
  <get-services-dynamic-flow-capture-content-destination-information>

show services dynamic-flow-capture control-source
  <get-services-dynamic-flow-capture-control-source-information>

show services dynamic-flow-capture statistics
  <get-services-dfc-statistics-information>
show services fips
show services fips pic
show services fips pic status
  <get-fips-pic-status-information>

show services flow-collector
  <get-services-flow-collector-information>

show services flow-collector file
  <get-services-flow-collector-file-information>

show services flow-collector input
  <get-services-flow-collector-input-information>

show services flow-table
```

```
show services flow-table statistics
  <get-flow-table-statistics-information>

show services flows
  <get-service-msp-flow-table-information>

show services ggsn
show services ggsn diagnostics
show services ggsn diagnostics pdp
  <get-pdp-diagnostics-per-apn>

show services ggsn statistics
  <get-ggsn-statistics>

show services ggsn statistics apn
  <get-ggsn-apn-statistics-information>

show services ggsn statistics charging
  <get-ggsn-charging-statistics-information>

show services ggsn statistics gtp
  <get-ggsn-gtp-statistics-information>

show services ggsn statistics gtp-prime
  <get-ggsn-gtp-prime-statistics-information>

show services ggsn statistics imsi
  <get-ggsn-imsi-user-information>

show services ggsn statistics l2tp-tunnel
  <get-ggsn-l2tp-tunnel-statistics-information>

show services ggsn statistics msisdn
show services ggsn statistics radius
  <get-ggsn-radius-statistics-information>

show services ggsn statistics sgsn
  <get-ggsn-sgsn-statistics-information>

show services ggsn status
  <get-ggsn-interface-information>

show services ggsn trace
show services ggsn trace all
  <get-ggsn-trace>

show services ggsn trace imsi
  <get-ggsn-imsi-trace>

show services ggsn trace msisdn
  <get-ggsn-msisdn-trace>

show services ids
show services ids destination-table
  <get-service-ids-destination-table-information>
```

```
show services ids pair-table
  <get-service-ids-pair-table-information>

show services ids source-table
  <get-service-ids-source-table-information>

show services inline
show services inline nat
show services inline nat pool
  <get-inline-nat-pool-information>
show services inline nat statistics
  <get-inline-nat-statistics-information>
show services inline softwire
show services inline softwire statistics
  <get-inline-service-sw-statistics-information>
show services ipsec-vpn
show services ipsec-vpn ike
show services ipsec-vpn ike security-associations
  <get-ike-services-security-associations-information>

show services ipsec-vpn ipsec
show services ipsec-vpn ipsec security-associations
  <get-services-security-associations-information>

show services ipsec-vpn ipsec statistics
  <get-services-ipsec-statistics-information>

show services l2tp
show services l2tp destination
  <get-l2tp-destination-information>
show services l2tp disconnect-cause-summary<
  <get-l2tp-disconnect-cause-summary>
show services l2tp multilink
  <get-l2tp-multilink-information>

show services l2tp radius
show services l2tp radius accounting
show services l2tp radius accounting servers
  <get-services-l2tp-radius-accounting-servers-information>

show services l2tp radius accounting statistics
  <get-services-l2tp-radius-accounting-statistics-information>

show services l2tp radius authentication
show services l2tp radius authentication servers
  <get-services-l2tp-radius-authentication-servers-information>

show services l2tp radius authentication statistics
  <get-services-l2tp-radius-authentication-statistics-information>

show services l2tp radius servers
  <get-services-l2tp-radius-authentication-accounting-servers-information>

show services l2tp radius statistics
  <get-services-l2tp-radius-authentication-accounting-statistics-information>
```

```
show services l2tp session
  <get-l2tp-session-information>

show services l2tp summary
  <get-l2tp-summary-information>

show services l2tp tunnel
  <get-l2tp-tunnel-information>

show services l2tp user
  <get-l2tp-user-information>
show services link-services
show services link-services cpu-usage
  <get-link-services-cpu-usage>

show services local-policy-decision-function
show services local-policy-decision-function flows
show services local-policy-decision-function flows interface
  <get-local-policy-decision-function-flows-interface>
show services local-policy-decision-function flows subscriber
  <get-local-policy-decision-function-flows-subscriber>
show services local-policy-decision-function statistics
show services local-policy-decision-function statistics interface
  <get-local-policy-decision-function-statistics-interface>
show services local-policy-decision-function statistics subscriber
  <get-local-policy-decision-function-statistics-subscriber>
show services logging
show services logging history
show services logging history client
show services logging logfiles
show services nat
show services nat ipv6-multicast-interfaces
  <get-service-nat-ipv6-multicast-information>

show services nat deterministic-nat
show services nat deterministic-nat internal-host
show services nat deterministic-nat nat-port-block
show services nat mappings
  <get-service-nat-mapping-address-pooling-paired>
show services nat mappings brief
  <get-service-nat-mapping-brief>
show services nat mappings detail
show services nat mappings endpoint-independent
  <get-service-nat-mapping-endpoint-independent>
show services nat mappings brief
  <get-service-nat-mapping-brief>
show services nat mappings detail
  <get-service-nat-mapping-detail>
show services nat mappings pcp
show services nat mappings summary
  <get-service-nat-mapping-summary>
show services nat pool
  <get-service-nat-pool-information>
show services pcp
show services pgcp
show services pgcp active-configuration
```



```
<get-pgcpd-active-configuration>

show services pgcp active-configuration gateway
  <get-service-pgcp-active-configuration-gateway>

show services pgcp conversations
  <get-service-pgcp-conversation-information>

show services pgcp conversations gateway
  <get-service-pgcp-conversation-information-gateway>

show services pgcp flows
  <get-service-pgcp-flow-table-information>

show services pgcp flows gateway
  <get-service-pgcp-flow-table-information-gateway>

show services pgcp gate
  <get-service-pgcp-gate>

show services pgcp gate gateway
  <get-service-pgcp-gate-gateway>

show services pgcp gates
  <get-service-pgcp-gates>

show services pgcp gates gateway
  <get-service-pgcp-gates-gateway>

show services pgcp root-termination
  <get-services-pgcpd-root-termination>

show services pgcp root-termination gateway
  <get-services-pgcpd-root-termination-gateway>

show services pgcp statistics
  <get-service-pgcp-statistics>

show services pgcp statistics gateway
  <get-service-pgcp-statistics-gateway>

show services pgcp terminations
  <get-service-pgcp-terminations>

show services pgcp terminations gateway
  <get-service-pgcp-terminations-gateway>

show services rpm
show services rpm active-servers
  <get-active-servers>

show services rpm history-results
  <get-history-results>

show services rpm probe-results
  <get-probe-results>
```

```
show services rpm twamp
  <twamp-information>
show services rpm twamp server
  <twamp-server-information>
show services rpm twamp server connection
  <twamp-server-connection-information>
show services rpm twamp server session
  <twamp-server-session-information>
show services server-load-balance
show services server-load-balance external-manager
show services server-load-balance external-manager information
show services server-load-balance external-manager statistics
  <get-external-manager-statistics-information>
show services server-load-balance hash-table
  <get-hash-table-information>
show services server-load-balance health-monitor
show services server-load-balance health-monitor information
  <get-real-server-health-monitor-information>
show services server-load-balance health-monitor statistics
  <get-real-server-health-monitor-statistics-information>
show services server-load-balance real-server
show services server-load-balance real-server statistics
  <get-real-server-statistics-information>
show services server-load-balance real-server-group
show services server-load-balance real-server-group information
  <get-real-server-group-information>
show services server-load-balance real-server-group statistics
  <get-real-server-group-statistics-information>
show services server-load-balance sticky
  <get-sticky-table-information>
show services server-load-balance virtual-server
show services server-load-balance virtual-server information
  <get-virtual-server-information>
show services server-load-balance virtual-server statistics
  <get-virtual-server-statistics-information>
show services service-identification
show services service-identification header-redirect
show services service-identification header-redirect statistics
  <get-header-redirect-set-statistics-information>

show services service-identification statistics
  <get-service-identification-statistics-information>

show services service-identification uri-redirect
show services service-identification uri-redirect statistics
  <get-uri-redirect-set-statistics-information>

show services service-sets
show services service-sets cpu-usage
  <get-service-set-cpu-statistics>

show services service-sets memory-usage
  <get-service-set-memory-statistics>

show services service-sets memory-usage zone
```

```
show services service-sets plug-ins
  <get-service-set-plugin-summary>

show services service-sets statistics
show services service-sets statistics packet-drops
  <get-service-set-packet-drop-statistics>

show services service-sets statistics syslog
  <get-service-set-syslog-statistics>
show services service-sets statistics tcp-mss
  <get-service-set-tcp-mss-statistics>

show services service-sets summary
  <get-service-set-summary-information>

show services sessions
  <get-msp-session-table>

show services softwire
  <get-service-softwire-table-information>

show services softwire flows
  <get-service-fwnat-flow-table-information>

show services softwire statistics
  <get-service-softwire-statistics-information>

show services stateful-firewall
show services stateful-firewall flow-analysis
  <get-service-flow-analysis-information>
show services stateful-firewall conversations
  <get-service-sfw-conversation-information>

show services stateful-firewall flows
  <get-service-sfw-flow-table-information>
show services stateful-firewall redundancy-statistics
  <get-service-sfw-redundancy-statistics>

show services stateful-firewall sip-call
  <get-service-sfw-sip-call-information>

show services stateful-firewall sip-register
  <get-service-sfw-sip-register-information>

show services stateful-firewall statistics
  <get-service-sfw-statistics-information>

show services stateful-firewall statistics application-protocol
  <et-sfw-application-protocol-statistics>
show services stateful-firewall subscriber-analysis
  <get-service-subs-analysis-information>
show services subscriber
show services subscriber bandwidth
show services subscriber bandwidth client-id
  <get-services-subscriber-bandwidth-by-session-id>
```

```
show services subscriber bandwidth interface
  <get-services-subscriber-bandwidth-by-interface>
show services subscriber bandwidth ip-address
  <get-services-subscriber-bandwidth-by-ip-address>
show services subscriber bandwidth service-interface
  <get-services-subscriber-bandwidth-by-service-interface>
show services subscriber dynamic-policies
  <get-services-subscriber-dynamic-policies>
show services subscriber flows
  <get-services-subscriber-flows>
show services subscriber sessions
  <get-services-subscriber-session>
show services subscriber statistics
  <get-services-subscriber-statistics>
show snmp
show snmp health-monitor
  <get-health-monitor-information>

show snmp health-monitor alarms
  <get-health-monitor-alarm-information>

show snmp health-monitor logs
  <get-health-monitor-log-information>

show snmp inform-statistics
  <get-snmp-inform-statistics>

show snmp mib
show snmp mib get
  <get-snmp-object>

show snmp mib get-next
  <get-next-snmp-object>

show snmp mib walk
  <get-walk-snmp-object>

show snmp proxy
show snmp rmon
  <get-rmon-information>

show snmp rmon alarms
  <get-rmon-alarm-information>

show snmp rmon events
  <get-rmon-event-information>

show snmp rmon history
  <get-rmon-history-information>

show snmp rmon logs
  <get-rmon-log-information>

show snmp statistics
  <get-snmp-information>
```

```
show snmp v3
  <get-snmp-v3-information>

show snmp v3 access
  <get-snmp-v3-access-information>

show snmp v3 community
  <get-snmp-v3-community-information>

show snmp v3 general
  <get-snmp-v3-general-information>

show snmp v3 groups
  <get-snmp-v3-group-information>

show snmp v3 notify
  <get-snmp-v3-notify-information>

show snmp v3 notify filter
  <get-snmp-v3-notify-filter-information>

show snmp v3 target
  <get-snmp-v3-target-information>

show snmp v3 target address
  <get-snmp-v3-target-address-information>

show snmp v3 target parameters
  <get-snmp-v3-target-parameters-information>

show snmp v3 users
  <get-snmp-v3-usm-user-information>

show spanning-tree
show spanning-tree bridge
  <get-stp-bridge-information>
show spanning-tree interface
  <get-stp-interface-information>
show spanning-tree mstp
show spanning-tree mstp configuration
  <get-mstp-configuration-information>
show spanning-tree statistics
  <get-stp-interface-statistics>
show spanning-tree statistics bridge
show spanning-tree statistics interface
show spanning-tree statistics routing-instance
  <get-stp-routing-instance-statistics>
show spanning-tree stp-buffer
show static-subscribers
show static-subscribers sessions
<show subscribers
  <get-subscribers>
show subscribers summary
  <get-subscribers-summary>
<get-syslog-filenames>
```

```
show synchronous-ethernet
show synchronous-ethernet esmc
show synchronous-ethernet esmc statistics
show synchronous-ethernet esmc transmit
show synchronous-ethernet global-information
show system
show system alarms
    <get-system-alarm-information>

show system auto-snapshot
show system boot-messages
show system buffers
show system certificate
show system commit
    <get-commit-information>

show system commit server
    <get-commit-server-information>
show system commit server queue
    <get-commit-server-queue-information>
show system configuration
show system configuration archival
    <get-system-archival>

show system configuration rescue
    <get-rescue-information>

show system connections
show system core-dumps
    <get-system-core-dumps>
show system core-dumps core-file-info
    <get-core-file-information>

show system core-dumps kernel-crashinfo
show system core-dumps transfer-status
show system diagnostics
show system diagnostics inventory
show system diagnostics usage
show system directory-usage
    <get-directory-usage-information>

show system firmware
    <get-system-firmware-information>

show system license
    <get-license-summary-information>

show system license installed
    <get-license-information>
show system license key-content
show system license keys
    <get-license-key-information>

show system license usage
    <get-license-usage-summary>
show system login
```

```
show system login logout
  <get-system-login-logout-information>
show system memory
<show system processes
show system processes brief
show system processes extensive
show system processes health
  <get-process-health-information>

show system processes providers
show system processes resource-limits
  <get-system-process-resource-limits>
show system processes summary
show system queues
show system reboot
show system resource-cleanup
show system resource-cleanup processes
  <get-system-resource-cleanup-processes-information>

show system rollback
  <get-rollback-information>

show system services
show system services dhcp
show system services dhcp binding
  <get-dhcp-binding-information>

show system services dhcp conflict
  <get-dhcp-conflict-information>

show system services dhcp global
  <get-dhcp-global-information>

show system services dhcp pool
  <get-dhcp-pool-information>

show system services dhcp statistics
  <get-dhcp-statistics-information>

show system services reverse
  <get-system-services-reverse-information>

show system services service-deployment
  <get-service-deployment-service-information>

show system snapshot
  <get-snapshot-information>

show system software
show system software backup
  <get-package-backup-information>
  <get-software-installation-status>
show system software recovery-package

show system statistics
  <get-statistics-information>
```

```
show system statistics bridge
<get-system-bridge-statistics>
show system statistics extended
show system statistics vpls
show system storage
  <get-system-storage>
show system storage partitions
  <get-system-storage-partitions>
show system subscriber-management
show system subscriber-management summary
show system switchover
  <get-switchover-information>

show system uptime
  <get-system-uptime-information>

show system users
  <get-system-users-information>

show system virtual-memory
show task
show task io
show task memory
show task replication
  <get-routing-task-replication-state>
show task snooping
show task snooping io
show task snooping memory
  <get-snooping-task-memory-information>
show ted
show ted database
  <get-ted-database-information>

show ted link
  <get-ted-link-information>

show ted protocol
  <get-ted-protocol-information>

show version
  <get-software-information>

show vpls
show vpls connections
  <get-vpls-connection-information>

show vpls flood
show vpls flood event-queue
  <get-vpls-event-queue-information>

show vpls flood route
show vpls flood route all-ce-flood
  <get-vpls-all-ce-flood-route-information>

show vpls flood route all-flood
```



```
<get-vpls-all-flood-route-information>

show vpls flood route alt-root-flood
  <get-vpls-alt-root-flood-route-information>

show vpls flood route ce-flood
  <get-vpls-ce-flood-route-information>

show vpls flood route mlp-flood
  <get-vpls-mlp-flood-route-information>

show vpls flood route re-flood
  <get-vpls-re-flood-route-information>

show vpls mac-table
  <get-vpls-mac-table>

show vpls mac-table interface
  <get-vpls-interface-mac-table>

show vpls statistics
  <get-vpls-statistics-information>

show vrrp
show vrrp interface
show vrrp track
test interface
test interface fdl-line-loop
test interface fdl-line-loop ansi
test interface fdl-line-loop ansi initiate
test interface fdl-line-loop ansi terminate
test interface fdl-line-loop bellcore
test interface fdl-line-loop bellcore initiate
test interface fdl-line-loop bellcore terminate
test interface fdl-payload-loop
test interface fdl-payload-loop ansi
test interface fdl-payload-loop ansi initiate
test interface fdl-payload-loop ansi terminate
test interface fdl-payload-loop bellcore
test interface fdl-payload-loop bellcore initiate
test interface fdl-payload-loop bellcore terminate
test interface inband-line-loop
test interface inband-line-loop ansi
test interface inband-line-loop ansi initiate
test interface inband-line-loop ansi terminate
test interface inband-line-loop bellcore
test interface inband-line-loop bellcore initiate
test interface inband-line-loop bellcore terminate
test interface inband-line-loop initiate
test interface inband-line-loop terminate
test interface inband-payload-loop
test interface inband-payload-loop ansi
test interface inband-payload-loop ansi initiate
test interface inband-payload-loop ansi terminate
test interface inband-payload-loop bellcore
test interface inband-payload-loop bellcore initiate
```

```
test interface inband-payload-loop bellcore terminate
test msdp
test msdp dependent-peers
test msdp rpf-peer
test policy
<
```

**Configuration
Hierarchy Levels**

```
[edit dynamic-profiles routing-instances instance services mobile-ip home-agent
enable-service]
[edit logical-systems routing-instances instance services mobile-ip home-agent
enable-service]
[edit logical-systems services mobile-ip home-agent enable-service]
[edit routing-instances instance services mobile-ip home-agent enable-service]
[edit services mobile-ip home-agent enable-service]
```

**Related
Documentation**

- [Access Privilege User Permission Flags Overview on page 128](#)
- [Understanding Junos OS Access Privilege Levels on page 123](#)
- [Configuring Access Privilege Levels on page 131](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)

view-configuration

Can view all of the configuration excluding secrets, system scripts, and event options.



NOTE: Only users with the maintenance permission can view commit script, op script, or event script configuration.

Commands

No associated CLI commands.

**Configuration
Hierarchy Levels**

No associated CLI configuration hierarchy levels and statements.

**Related
Documentation**

- [Access Privilege User Permission Flags Overview on page 128](#)
- [Understanding Junos OS Access Privilege Levels on page 123](#)
- [Configuring Access Privilege Levels on page 131](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 132](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 130](#)

Administration

- [Operational Mode Commands on page 292](#)

Operational Mode Commands

show cli authorization

Syntax show cli authorization

Release Information Command introduced before Junos OS Release 7.4.

Description Display the permissions for the current user.

```
user@host> show cli authorization
Current user: 'root' login: 'boojum' class '(root)'
Permissions:
Permissions:
  admin          -- Can view user accounts
  admin-control-- Can modify user accounts
  clear          -- Can clear learned network info
  configure      -- Can enter configuration mode
  control        -- Can modify any config
  edit          -- Can edit full files
  field          -- Can use field debug commands
  floppy         -- Can read and write the floppy
  interface      -- Can view interface configuration
  interface-control-- Can modify interface configuration
  network        -- Can access the network
  reset          -- Can reset/restart interfaces and daemons
  routing        -- Can view routing configuration
  routing-control-- Can modify routing configuration
  shell          -- Can start a local shell
  snmp           -- Can view SNMP configuration
  snmp-control-- Can modify SNMP configuration
  system         -- Can view system configuration
  system-control-- Can modify system configuration
  trace          -- Can view trace file settings
  trace-control-- Can modify trace file settings
  view           -- Can view current values and statistics
  maintenance    -- Can become the super-user
  firewall       -- Can view firewall configuration
  firewall-control-- Can modify firewall configuration
  secret         -- Can view secret statements
  secret-control-- Can modify secret statements
  rollback       -- Can rollback to previous configurations
  security       -- Can view security configuration
  security-control-- Can modify security configuration
  access         -- Can view access configuration
  access-control-- Can modify access configuration
  view-configuration-- Can view all configuration (not including secrets)
  flow-tap       -- Can view flow-tap configuration
  flow-tap-control-- Can modify flow-tap configuration
  idp-profiler-operation-- Can Profiler data
  pgcp-session-mirroring-- Can view pgcp session mirroring configuration
  pgcp-session-mirroring-control-- Can modify pgcp session mirroring configuration
  storage        -- Can view fibre channel storage protocol configuration
  storage-control-- Can modify fibre channel storage protocol configuration
  all-control    -- Can modify any configuration
```

Required Privilege Level view

- Related** • *show cli*
Documentation • *show cli directory*

CHAPTER 4

System Basics

- [Overview on page 295](#)
- [Configuration on page 301](#)
- [Administration on page 446](#)

Overview

- [Installation Overview on page 295](#)
- [Licenses Overview on page 296](#)

Installation Overview

- [Junos OS Package Names on page 295](#)

Junos OS Package Names

You upgrade the Juniper Networks Junos operating system (Junos OS) on a Juniper Networks EX Series Ethernet Switch by copying a software package to your switch or another system on your local network, then install the new software package on the switch.

A software package name is in the following format:

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

where:

- ***package-name*** is the name of the package—for example, ***jinstall-ex-4200***.
- ***m.n*** is the software release, with ***m*** representing the major release number and ***n*** representing the minor release number—for example, ***9.5***.
- ***Z*** indicates the type of software release, where ***R*** indicates released software and ***B*** indicates beta-level software.
- ***x.y*** represents the version of the major software release (***x***) and an internal tracking number (***y***)—for example, ***1.6***.
- ***domestic-signed*** is appended to all EX Series package names. For most Junos packages, ***domestic*** is used for the United States and Canada and ***export*** for worldwide distribution. However, for EX Series software, ***domestic*** is used for worldwide distribution as well.

A sample EX Series software package name is:

`jinstall-ex-4200-9.5R1.6-domestic-signed.tgz`

**Related
Documentation**

- *Installing Software on EX Series Switches (J-Web Procedure)*
- *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*
- *Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)*
- [Downloading Software Packages from Juniper Networks on page 301](#)
- *Understanding Software Installation on EX Series Switches*

Licenses Overview

- [Understanding Software Licenses for EX Series Switches on page 296](#)

Understanding Software Licenses for EX Series Switches

To enable and use some of the Juniper Networks operating system (Junos OS) features, you must purchase, install, and manage separate software licenses. If the switch has the appropriate software license, you can configure and use these features.

The Junos OS feature license (that is, the purchased authorization code) is universal. However, to conform to Junos OS feature licensing requirements, you must install a unique license key (a combination of the authorization code and the switch's serial number) on each switch.

For a Virtual Chassis deployment, two license keys are recommended for redundancy—one for the device in the master role and the other for the device in the backup role:

- In an EX8200 Virtual Chassis, the devices in the master and backup roles are always XRE200 External Routing Engines.
- In all other Virtual Chassis, the devices in the master and backup roles are switches.

You do not need additional license keys for Virtual Chassis member switches that are in the linecard role or for the redundant Routing Engine (RE) modules or the redundant Switch Fabric and Routing Engine (SRE) modules in an EX8200 member switch.

This topic describes:

- [Purchasing a Software Feature License on page 297](#)
- [Features Requiring a License on EX2200 Switches on page 297](#)
- [Features Requiring a License on EX3300 Switches on page 298](#)
- [Features Requiring a License on EX3200, EX4200, EX4500, EX4550, EX6200, EX8200, and EX9200 Switches on page 299](#)
- [License Warning Messages on page 300](#)

Purchasing a Software Feature License

The following sections list features that require separate licenses. To purchase a software license, contact your Juniper Networks sales representative (<http://www.juniper.net/us/en/contact-us/sales-offices>). You will be asked to supply the chassis serial number of your switch; you can obtain the serial number by running the `show chassis hardware` command.



NOTE: You are required to provide a 12-digit serial number when purchasing a license for an XRE200 External Routing Engine in an EX8200 Virtual Chassis.

The serial number listed on the XRE200 External Routing Engine serial ID label is 16 digits long. Use the last 12 digits of the 16-digit serial number to purchase the license.

You can use the `show chassis hardware` command output to display the 12-digit serial number of the XRE200 External Routing Engine to use when you purchase the license.

Features Requiring a License on EX2200 Switches

For Juniper Networks EX2200 Ethernet Switches, the following features can be added to basic Junos OS by installing an enhanced feature license (EFL):

- Bidirectional Forwarding Detection (BFD)
- Connectivity fault management (IEEE 802.1ag)
- IGMP (Internet Group Management Protocol) version 1 (IGMPv1), IGMPv2, and IGMPv3
- OSPFv1/v2 (with four active interfaces)
- Protocol Independent Multicast (PIM) dense mode, PIM source-specific mode, PIM sparse mode
- Q-in-Q tunneling (IEEE 802.1ad)
- Real-time performance monitoring (RPM)
- Virtual Router
- Virtual Router Redundancy Protocol (VRRP)

Table 42 on page 297 lists the EFLs that you can purchase for EX2200 switch models. If you have the license, you can run all of the above-mentioned enhanced software features on your EX2200 switch.

Table 42: Junos OS EFL Part Number on EX2200 Switches

Switch Model	EFL Part Number
EX2200-C-12P-2G EX2200-C-12T-2G	EX-12-EFL

Table 42: Junos OS EFL Part Number on EX2200 Switches (*continued*)

Switch Model	EFL Part Number
EX2200-24T-4G EX2200-24P-4G EX2200-24T-DC-4G	EX-24-EFL
EX2200-48T-4G EX2200-48P-4G	EX-48-EFL

Features Requiring a License on EX3300 Switches

Two types of licenses are available on Juniper Networks EX3300 Ethernet Switches: enhanced feature licenses (EFLs) and advanced feature licenses (AFLs).

To use the following features on the EX3300 switches, you must install an EFL:

- Bidirectional Forwarding Detection (BFD)
- IGMP (Internet Group Management Protocol) version 1 (IGMPv1), IGMPv2, and IGMPv3
- IPv6 routing protocols: Multicast Listener Discovery version 1 and 2 (MLD v1/v2), OSPFv3, PIM multicast, VRRPv6, virtual router support for unicast and filter-based forwarding (FBF)
- OSPFv1/v2
- Protocol Independent Multicast (PIM) dense mode, PIM source-specific mode, PIM sparse mode
- Q-in-Q tunneling (IEEE 802.1ad)
- Virtual Router
- Virtual Router Redundancy Protocol (VRRP)

[Table 43 on page 298](#) lists the EFLs that you can purchase for EX3300 switch models. If you have the license, you can run all of the above-mentioned enhanced software features on your EX3300 switch.

Table 43: Junos OS EFL Part Number on EX3300 Switches

Switch Model	EFL Part Number
EX3300-24T EX3300-24P EX3300-24T-DC	EX-24-EFL
EX3300-48T EX3300-48T-BF EX3300-48P	EX-48-EFL

To use the following feature on EX3300 switches, you must install an AFL:

- BGP/MBGP

- IPv6 routing protocols: IPv6 BGP and IPv6 for MBGP
- Virtual routing and forwarding (VRF) BGP

[Table 44 on page 299](#) lists the AFLs that you can purchase for EX3300 switch models. For EX3300 switches, you must purchase and install a corresponding EFL along with the AFL in order to enable the advanced license features. If you have both these licenses, you can run all of the above-mentioned advanced software features on your EX3300 switch.

Table 44: Junos OS AFL Part Number on EX3300 Switches

Switch Model	AFL Part Number
EX3300-24T EX3300-24P EX3300-24T-DC	EX-24-AFL
EX3300-48T EX3300-48T-BF EX3300-48P	EX-48-AFL

Features Requiring a License on EX3200, EX4200, EX4500, EX4550, EX6200, EX8200, and EX9200 Switches

To use the following features on Juniper Networks EX3200, EX4200, EX4500, EX4550, EX8200, and EX9200 Ethernet Switches, you must install an advanced feature license (AFL):

- Border Gateway Protocol (BGP) and multiprotocol BGP (MBGP)
- Intermediate System-to-Intermediate System (IS-IS)
- IPv6 routing protocols: IS-IS for IPv6, IPv6 BGP, IPv6 for MBGP
- Logical systems (available only on EX9200 switches)
- MPLS with RSVP-based label-switched paths (LSPs) and MPLS-based circuit cross-connects (CCCs) (Not supported on EX9200 switches)

To use the following features on Juniper Networks EX6200 Ethernet Switches, you must install an advanced feature license (AFL):

- Border Gateway Protocol (BGP)
- Intermediate System-to-Intermediate System (IS-IS)
- IPv6 routing protocols: IS-IS for IPv6, IPv6 BGP

[Table 45 on page 300](#) lists the AFLs that you can purchase for EX3200, EX4200, EX4500, EX4550, EX6200, EX8200, and EX9200 switches. If you have the license, you can run all of the above-mentioned advanced software features on your EX3200, EX4200, EX4500, EX4550, EX6200, EX8200, or EX9200 switch.

Table 45: Junos OS AFL Part Number on EX3200, EX4200, EX4500, EX4550, EX6200, EX8200, and EX9200 Switches

Switch Model	AFL Part Number
EX3200-24P EX3200-24T EX4200-24F EX4200-24P EX4200-24PX EX4200-24T	EX-24-AFL
EX3200-48P EX3200-48T EX4200-48F EX4200-48P EX4200-48PX EX4200-48T	EX-48-AFL
EX4500-40F-BF EX4500-40F-BF-C EX4500-40F-FB EX4500-40F-FB-C	EX-48-AFL
EX4550	EX4550-AFL
EX6210	EX6210-AFL
EX8208	EX8208-AFL
EX8216	EX8216-AFL
EX-XRE200	EX-XRE200-AFL
EX9204	EX9204-AFL
EX9208	EX9208-AFL
EX9214	EX9214-AFL

License Warning Messages

For using features that require a license, you must install and configure a license key. To obtain a license key, use the contact information provided in your certificate.

If you have not purchased the AFL or EFL and installed the license key, you receive warnings when you try to commit the configuration:

```
[edit protocols]
  'bgp'
    warning: requires 'bgp' license
error: commit failed: (statements constraint check failed)
```

The system generates system log (**syslog**) alarm messages notifying you that the feature requires a license—for example:

```
Sep 3 05:59:11 craftd[806]: Minor alarm set, BGP Routing Protocol usage
requires a license
Sep 3 05:59:11 alarmd[805]: Alarm set: License color=YELLOW, class=CHASSIS,
reason=BGP Routing Protocol usage requires a license
Sep 3 05:59:11 alarmd[805]: LICENSE_EXPIRED: License for feature bgp(47) expired
```

Output of the **show system alarms** command displays the active alarms:

```
user@switch> show system alarms
1 alarm currently active
Alarm time           Class  Description
2009-09-03 06:00:11 UTC  Minor  BGP Routing Protocol usage requires a license
```

Related Documentation

- *Managing Licenses for the EX Series Switch (CLI Procedure)*
- *Managing Licenses for the EX Series Switch (J-Web Procedure)*
- *Monitoring Licenses for the EX Series Switch*
- *License Key Components for the EX Series Switch*
- [EX Series Switch Software Features Overview on page 63](#)

Configuration

- [Junos OS Packages on page 301](#)
- [Statement Hierarchies on page 302](#)

Junos OS Packages

- [Downloading Software Packages from Juniper Networks on page 301](#)

Downloading Software Packages from Juniper Networks

You can download Junos OS packages from the Juniper Networks website to upgrade software on your EX Series switch.

Before you begin to download software upgrades, ensure that you have a Juniper Networks Web account and a valid support contract. To obtain an account, complete the registration form at the Juniper Networks website: <https://www.juniper.net/registration/Register.jsp>.

To download software upgrades from Juniper Networks:

1. Using a Web browser, follow the links to the download URL on the Juniper Networks webpage. For EX Series, there are not separate software packages for Canada the U.S. and other locations. Therefore, select **Canada and U.S. Version** regardless of your location:
 - <https://www.juniper.net/support/downloads/junos.html>
2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.

3. Using the J-Web interface or the CLI, select the appropriate software package for your application. See [“Junos OS Package Names” on page 295](#).
4. Download the software to a local host or to an internal software distribution site.

**Related
Documentation**

- *Installing Software on EX Series Switches (J-Web Procedure)*
- *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*
- *Understanding Software Installation on EX Series Switches*

Statement Hierarchies

- [\[edit accounting-options\] Hierarchy Level on page 303](#)
- [\[edit chassis\] Hierarchy Level on page 305](#)
- [\[edit class-of-service\] Hierarchy Level on page 313](#)
- [\[edit dynamic-profiles\] Hierarchy Level on page 317](#)
- [\[edit event-options\] Hierarchy Level on page 318](#)
- [\[edit firewall\] Hierarchy Level on page 320](#)
- [\[edit forwarding-options\] Hierarchy Level on page 333](#)
- [\[edit forwarding-options accounting\] Hierarchy Level on page 334](#)
- [\[edit forwarding-options dhcp-relay\] Hierarchy Level on page 334](#)
- [\[edit forwarding-options enhanced-hash-key\] Hierarchy Level on page 338](#)
- [\[edit forwarding-options family\] Hierarchy Level on page 339](#)
- [\[edit forwarding-options fast-reroute-priority\] Hierarchy Level on page 340](#)
- [\[edit forwarding-options hash-key\] Hierarchy Level on page 340](#)
- [\[edit forwarding-options helpers\] Hierarchy Level on page 341](#)
- [\[edit forwarding-options load-balance\] Hierarchy Level on page 343](#)
- [\[edit forwarding-options next-hop-group\] Hierarchy Level on page 344](#)
- [\[edit forwarding-options port-mirroring\] Hierarchy Level on page 344](#)
- [\[edit forwarding-options rpf-loose-mode-discard\] Hierarchy Level on page 345](#)
- [\[edit forwarding-options sampling\] Hierarchy Level on page 346](#)
- [\[edit groups\] Hierarchy Level on page 348](#)
- [\[edit interfaces\] Hierarchy Level on page 348](#)
- [\[edit logical-systems\] Hierarchy Level on page 359](#)
- [\[edit multi-chassis\] Hierarchy Level on page 360](#)
- [\[edit policy-options\] Hierarchy Level on page 360](#)
- [\[edit protocols\] Hierarchy Level on page 365](#)
- [\[edit protocols bgp\] Hierarchy Level on page 368](#)
- [\[edit protocols isis\] Hierarchy Level on page 378](#)
- [\[edit protocols l2-learning\] Hierarchy Level on page 380](#)

- [\[edit protocols lacp\] Hierarchy Level on page 380](#)
- [\[edit protocols layer2-control\] Hierarchy Level on page 381](#)
- [\[edit protocols ldp\] Hierarchy Level on page 381](#)
- [\[edit protocols mpls\] Hierarchy Level on page 386](#)
- [\[edit protocols mstp\] Hierarchy Level on page 386](#)
- [\[edit protocols ospf\] Hierarchy Level on page 391](#)
- [\[edit protocols ospf3\] Hierarchy Level on page 395](#)
- [\[edit protocols pim\] Hierarchy Level on page 398](#)
- [\[edit protocols rip\] Hierarchy Level on page 402](#)
- [\[edit protocols ripng\] Hierarchy Level on page 404](#)
- [\[edit protocols router-advertisement\] Hierarchy Level on page 405](#)
- [\[edit protocols router-discovery\] Hierarchy Level on page 406](#)
- [\[edit protocols rstp\] Hierarchy Level on page 406](#)
- [\[edit protocols sap\] Hierarchy Level on page 407](#)
- [\[edit protocols vrrp\] Hierarchy Level on page 407](#)
- [\[edit protocols vstp\] Hierarchy Level on page 408](#)
- [Layer 2 Routing Instances Configuration Hierarchy on page 409](#)
- [\[edit routing-options\] Hierarchy Level on page 411](#)
- [\[edit security\] Hierarchy Level on page 421](#)
- [\[edit security alarms\] Hierarchy Level on page 421](#)
- [\[edit security authentication-key-chains\] Hierarchy Level on page 421](#)
- [\[edit security certificates\] Hierarchy Level on page 422](#)
- [\[edit security ike\] Hierarchy Level on page 422](#)
- [\[edit security ipsec\] Hierarchy Level on page 423](#)
- [\[edit security log\] Hierarchy Level on page 424](#)
- [\[edit security pki\] Hierarchy Level on page 424](#)
- [\[edit security ssh-known-hosts\] Hierarchy Level on page 425](#)
- [\[edit security traceoptions\] Hierarchy Level on page 425](#)
- [\[edit snmp\] Hierarchy Level on page 426](#)
- [\[edit switch-options\] Hierarchy Level on page 430](#)
- [\[edit system\] Hierarchy Level on page 431](#)
- [\[edit vlans\] Hierarchy Level on page 445](#)

[\[edit accounting-options\] Hierarchy Level](#)

```
accounting-options {  
  class-usage-profile profile-name {  
    destination-classes {  
      destination-class-name;  
    }  
  }  
}
```

```
    file filename;
    interval minutes;
    source-classes {
        source-class-name;
    }
}
file filename {
    archive-sites {
        site-name;
    }
    files number;
    nonpersistent;
    size bytes;
    start-time time;
    transfer-interval minutes;
}
filter-profile profile-name {
    counters {
        counter-name;
    }
    file filename;
    interval minutes;
}
interface-profile profile-name {
    fields {
        input-bytes;
        input-errors;
        input-multicast;
        input-packets;
        input-unicast;
        output-bytes;
        output-errors;
        output-multicast;
        output-packets;
        output-unicast;
        rpf-check-bytes;
        rpf-check-packets;
        rpf-check6-bytes;
        rpf-check6-packets;
        unsupported-protocol;
    }
    file filename;
    interval minutes;
}
mib-profile profile-name {
    file filename;
    interval minutes;
    object-names {
        mib-object-name;
    }
    operation (get | get-next | walk);
}
policy-decision-statistics-profile profile-name {
    application-aware-access-list-fields {
        address;
        application;
    }
}
```

```

        application-group;
        input-bytes;
        input-interface;
        input-packets;
        mask;
        output-bytes;
        output-packets;
        subscriber-name;
        timestamp;
        vrf-name;
    }
    file filename;
}
routing-engine-profile profile-name {
    fields {
        field-name;
    }
    file filename;
    interval minutes;
}
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

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

```

chassis {
    aggregated-devices {
        ethernet {
            device-count number;
            lacp {
                link-protection {
                    non-revertive;
                }
                system-priority;
            }
        }
        sonet {
            device-count number;
        }
        maximum-links maximum-links-limit;
    }
    alarm {
        ds1 {
            ais (ignore | red | yellow);
            ylw (ignore | red | yellow);
        }
        ethernet {
            link-down (ignore | red | yellow);
        }
        integrated-services {
            failure (ignore | red | yellow);
        }
        management-ethernet {
            link-down (ignore | red | yellow);
        }
    }
}

```

```

}
relay
  input {
    port port-number {
      mode (close | open);
      trigger (ignore | red | yellow;
    }
  }
  output {
    port port-number {
      input-relay input-relay;
      mode (close | open);
      temperature;
    }
  }
serial {
  cts-absent (ignore | red | yellow);
  dcd-absent (ignore | red | yellow);
  dsr-absent (ignore | red | yellow);
  loss-of-rx-clock (ignore | red | yellow);
  loss-of-tx-clock (ignore | red | yellow);
  tm-absent (ignore | red | yellow);
}
services {
  hw-down (ignore | red | yellow);
  linkdown (ignore | red | yellow);
  pic-hold-reset (ignore | red | yellow);
  pic-reset (ignore | red | yellow);
  rx-errors (ignore | red | yellow);
  sw-down (ignore | red | yellow);
  tx-errors (ignore | red | yellow);
}
sonet {
  (ais-l | ais-p | ber-sd | ber-sf | locd | lol | lop-p | los | pll | plm-p | rfi-l | rfl-p | uneq-p)
  (ignore | red | yellow);
}
t3 {
  (ais | exz | ferf | idle | lcv | lof | los | pll | ylw) (ignore | red | yellow);
}
}
cluster {
  control-link-recovery;
  control-ports {
    fpc slot-number port port-number;
  }
  heartbeat-interval milliseconds;
  heartbeat-threshold number;
  redundancy-group {
    ... the redundancy-group subhierarchy appears at the end of the [edit chassis cluster]
    hierarchy ...
  }
}
reth-count number;
traceoptions {
  file <filename> <files number> <match regular-expression> <size maximum-file-size>
  <world-readable | no-world-readable>;
  flag flag;

```



```

    level severity;
    no-remote-trace;
}
redundancy-group group-number {
    gratuitous-arp-count number;
    hold-down-interval seconds;
    interface-monitor {
        interface-name weight number;
    }
    ip-monitoring {
        family {
            inet {
                ipv4-address {
                    interface rethindex.logical-unit-number secondary-ip-address ipv4-address;
                    weight number;
                }
            }
        }
        global-threshold number;
        global-weight number;
        retry-count count;
        retry-interval interval;
    }
    node node-number priority priority-number;
    preempt;
}
config-button {
    no-clear;
    no-rescue;
}
container-devices {
    device-count number;
}
craft-lockout;
disable-power-management;
disk-partition partition-name (/config | /var) {
    level (full | high) {
        free-space threshold-value (mb | percent);
    }
}
enhanced-policer;
extended-statistics;
fabric {
    degraded {
        action-fpc-restart-disable;
        degraded-fabric-detection-enable
        degraded-fpc-bad-plane-threshold number-bad-planes;
    }
    redundancy-mode (increased-bandwidth | redundant);
}
filter;
fpc slot-number {
    ... the fpc subhierarchy appears after the main [edit chassis] hierarchy ...
}
fpc-feb-connectivity {
    fpc slot-number feb (slot-number | none);
}

```

```
}
fpc-resync;
fru-poweron-sequence sequence;
lcc index {
    ... the lcc subhierarchy appears after the main [edit chassis] hierarchy ...
}
maximum-ecmp value;
memory-enhanced {
    filter;
    route;
    vpn-label;
}
network-services (ethernet | enhanced-ethernet | ip | enhanced-ip);
(packet-scheduling | no-packet-scheduling);
pem {
    minimum number;
}
policer-drop-probability-low;
ppp-subscriber-services (disable | enable);
redundancy {
    cfeb slot (always | preferred);
    failover {
        on-disk-failure;
        on-loss-of-keepalives;
    }
    feb {
        redundancy-group group-name {
            description description;
            feb slot-number <backup | primary>;
            no-auto-failover;
        }
    }
    graceful-switchover;
    keepalive-time seconds;
    routing-engine slot-number (backup | disabled | master);
    sfm slot-number (always | preferred);
    ssb slot-number (always | preferred);
}
route-memory-enhanced;
route-localization {
    inet (chassis);
    inet6;
}
routing-engine {
    bios {
        no-auto-upgrade;
    }
    on-disk-failure disk-failure-action (halt | reboot);
}
sfm slot-number {
    power off;
}
sib {
    minimum number;
}
(source-route | no-source-route);
```

```

state [
    cb-upgrade [on | off];
}
synchronization { # for M Series and T Series routers
    primary (external-a | external-b);
    secondary (external-a | external-b);
    signal-type (e1 | t1);
    switching-mode (non-revertive | revertive);
    transmitter-enable;
    validation-interval seconds;
    y-cable-line-termination;
}
synchronization { # for MX80 and MX240 routers
    clock-mode (auto-select | free-run);
    esmc-transmit {
        interfaces (all | interface-name);
    }
    hold-interval {
        configuration-change seconds;
        restart seconds;
        switchover seconds;
    }
    network-option (option-1 | option-2);
    quality-mode-enable;
    selection-mode (configured-quality|received-quality);
    source {
        (external-a | external-b) {
            priority number;
            quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
            request (force-switch | lockout);
        }
        interfaces interface-name {
            priority number;
            quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
            request (force-switch | lockout);
            wait-to-restore minutes;
        }
    }
    switchover-mode (revertive | non-revertive);
}
synchronization { # for ACX Series routers
    clock-mode (auto-select | free-run);
    esmc-transmit {
        interfaces (all | interface-name);
    }
    hold-interval {
        configuration-change seconds;
        restart seconds;
        switchover seconds;
    }
    network-option (option-1 | option-2);
    quality-mode-enable;
    selection-mode (configured-quality | received-quality);
    source {
        (bits | gps) {
            priority number;

```

```

        quality-level (prc | prs | sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
        request (force-switch | lockout);
    }
    interfaces interface-name {
        priority number;
        quality-level (prc | prs | sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
        request (force-switch | lockout);
        wait-to-restore minutes;
    }
}
switchover-mode (non-revertive | revertive);
}
system-domains {
    protected-system-domains psdnumerical-index {
        control-plane-bandwidth-percent percent;
        control-slot-numbers [ slot-numbers ];
        control-system-id control-system-id;
        description description;
        fpcs [ slot-numbers ];
    }
    root-domain-id root-domain-id;
}
vrf-mtu-check;
}

chassis {
    fpc slot-number {
        number-of-ports active-ports;
        offline;
        pic slot-number {
            ... the pic subhierarchy appears after the main [edit chassis fpc slot-number] hierarchy
            ...
        }
        port-mirror-instance port-mirror-instance-name;
        power (off | on);
        sampling-instance instance-name;
    }

    fpc slot-number {
        pic slot-number {
            adaptive-services {
                service-package (layer-2 | layer-3 | ...the following extension-provider subhierarchy
                ...);
            }
            extension-provider {
                control-cores number;
                data-cores number;
                data-flow-affinity {
                    hash-key (layer-3 | layer-4);
                }
            }
            channelization;
            forwarding-db-size megabytes;
            object-cache-size megabytes;
            package package-name;
            policy-db-size megabytes;
            syslog {
                facility {

```

```

        severity;
        destination (pic-console | routing-engine);
    }
}
wired-process-mem-size megabytes;
}
}
aggregated-devices {
    ima {
        device-count number;
    }
}
aggregate-ports;
atm-cell-relay-accumulation;
atm-l2circuit-mode (aal5 | cell | trunk trunk);
cel {
    e1 port-number {
        channel-group group-number timeslots slot-number;
    }
}
ct3 {
    port port-number {
        t1 link-number {
            channel-group group-number timeslots slot-number;
        }
    }
}
ethernet {
    pic-mode (enhanced-switching | routing | switching);
}
fibre-channel {
    port port-number;
    port-range port-range-low port-range-high
}
egress-policer-overhead bytes;
forwarding-mode {
    sa-multicast;
    vlan-steering {
        vlan-rule (high-low | odd-even);
    }
}
framing (e1 | e3 | sdh | sonet | t1 | t3);
idle-cell-format {
    itu-t;
    payload-pattern payload-pattern-byte;
}
ingress-policer-overhead bytes;
inline-services {
    bandwidth (1g | 10g);
}
linerate-mode;
max-queues-per-interface (4 | 8);
mlfr-uni-nni-bundles number;
no-concatenate;
no-multi-rate;
port port-number {

```

```

    framing (e1 | e3 | sdh | sonet | t1 | t3);
    forwarding-mode {
        sa-multicast;
    }
    speed ( oc3-stm1 | oc12-stm4 | oc48-stm16);
}
port-mirror-instance port-mirror-instance-name;
q-pic-large-buffer {
    (large-scale | small-scale);
}
red-buffer-occupancy {
    weighted-averaged <instant-usage-weight-exponent weight-value>;
}
shdsl {
    pic-mode (1-port-atm | 2-port-atm);
}
sparse-dlcis;
traffic-manager {
    egress-shaping-overhead number;
    ingress-shaping-overhead number;
    mode {
        egress-only;
        ingress-and-egress;
        session-shaping;
    }
}
tunnel-queuing;
tunnel-services {
    bandwidth (1g | 10g | 20g | 40g);
    tunnel-only;
}
vtmapping (itu-t | klm);
}
}

chassis {
    lcc index {
        fpc slot-number {
            ... the fpc subhierarchy appears after the main [edit chassis lcc index] hierarchy ...
        }
        offline;
        online-expected;
    }
}

lcc index {
    fpc slot-number {
        pic slot-number {
            ... the pic subhierarchy appears after the main [edit chassis lcc index fpc slot-number] hierarchy ...
        }
        power (off | on);
        sampling-instance instance-name;
    }
}

```

```

fpc slot-number {
  pic slot-number {
    aggregate-ports;
    atm-cell-relay-accumulation;
    atm-l2circuit-mode (aal5 | cell | trunk trunk);
    framing (e1 | e3 | sdh | sonet | t1 | t3);
    idle-cell-format {
      itu-t;
      payload-pattern payload-pattern-byte;
    }
    linerate-mode;
    max-queues-per-interface (4 | 8);
    no-concatenate;
    no-mcast-replication;
    no-pre-classifier;
    port port-number {
      framing (e1 | e3 | sdh | sonet | t1 | t3);
    }
    q-pic-large-buffer {
      (large-scale | small-scale);
    }
    red-buffer-occupancy {
      weighted-averaged <instant-usage-weight-exponent weight-value>;
    }
    shdsl {
      pic-mode (1-port-atm | 2-port-atm);
    }
    traffic-manager {
      egress-shaping-overhead bytes;
      ingress-shaping-overhead bytes;
      mode {
        egress-only;
        ingress-and-egress;
      }
    }
  }
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

[\[edit class-of-service\] Hierarchy Level](#)

```

class-of-service {
  classifiers {
    type classifier-name {
      forwarding-class class-name {
        loss-priority (high | low | medium-high | medium-low) code-points [ aliases bits ];
      }
      import (classifier-name | default);
    }
  }
  code-point-aliases {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) {
      alias-name bits;
    }
  }
}

```

```

    }
  }
  drop-profiles {
    profile-name {
      fill-level percentage drop-probability percentage;
      interpolate {
        drop-probability value;
        fill-level value;
      }
    }
  }
  fabric {
    scheduler-map {
      priority (high | low) scheduler scheduler-name;
    }
  }
  forwarding-class-map {
    map-name {
      class class-name queue-num queue-number <restricted-queue queue-number>;
    }
  }
  forwarding-classes {
    class class-name policing-priority (normal | premium) queue-num queue-number
      priority (high | low);
    queue queue-number class-name policing-priority (normal | premium) priority (high |
      low);
  }
  forwarding-policy {
    class class-name {
      classification-override {
        forwarding-class class-name;
      }
    }
    next-hop-map map-name {
      forwarding-class class-name {
        discard;
        lsp-next-hop [ lsp-regular-expressions ];
        next-hop [ next-hop-names ];
        non-lsp-next-hop;
      }
    }
  }
  fragmentation-maps {
    map-name {
      forwarding-class class-name {
        drop-timeout milliseconds;
        fragment-threshold bytes;
        multilink-class number;
        no-fragmentation;
      }
    }
  }
  host-outbound-traffic {
    dscp-code-point value;
    forwarding-class class-name;
    ieee-802.1 {

```



```

    default value;
    rewrite-rules;
}
tcp {
    raise-internet-control-priority;
}
}
interfaces {
    ... the interfaces subhierarchy appears after the main [edit class-of-service] hierarchy
    ...
}
}
restricted-queues {
    forwarding-class class-name queue-number;
}
rewrite-rules {
    (dscp | dscp-ipv6 | exp | frame-relay-de | ieee-802.1 | ieee-802.1ad | inet-precedence)
    rewrite-rule {
        forwarding-class class-name {
            loss-priority level code-point (alias | bits);
        }
        import (rewrite-rule | default);
    }
}
}
routing-instances routing-instance-name {
    classifiers {
        dscp (classifier-name | default);
        dscp-ipv6 (classifier-name | default);
        exp (classifier-name | default);
        ieee-208.1 (classifier-name | default | encapsulated | vlan-tag (inner | outer));
    }
}
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
}
schedulers {
    scheduler-name {
        adjust-minimum value;
        adjust-percent value;
        buffer-size (exact | percent percentage | remainder);
        drop-profile-map loss-priority (any | high | low | medium-high | medium-low)
            protocol any;
        excess-priority (high | low | medium-high | medium-low);
        excess-rate (percent percentage | proportion proportion);
        priority (high | low | medium-high | medium-low | strict-high);
        shaping-rate (bps | percent percentage | burst-size size);
        transmit-rate (bps | percent percentage | remainder) <exact | rate-limit>;
    }
}
}
traceoptions {
    file <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}

```

```

}
traffic-control-profiles {
  profile-name {
    adjust-minimum rate;
    delay-buffer-rate (bps | cps cps | percent percentage);
    excess-rate (percent percentage | proportion value);
    guaranteed-rate (bps | percent percentage) <burst-size bytes>;
    overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
    scheduler-map map-name;
    shaping-rate (bps | percent percentage) <burst-size bytes>;
  }
}
tri-color;
}

class-of-service {
  interfaces {
    interface-name {
      excess-bandwidth-share (equal | proportional value);
      input-excess-bandwidth-share (equal | proportional value);
      input-scheduler-map map-name;
      input-shaping-rate bps;
      input-traffic-control-profile profile-name;
      output-forwarding-class-map map-name;
      output-traffic-control-profile profile-name;
      scheduler-map map-name;
      scheduler-map-chassis (map-name | derived);
      shaping-rate bps;
      unit (logical-unit-number | *) {
        classifiers {
          dscp (classifier-name | default) {
            family [ inet mpls ];
          }
          dscp-ipv6 (classifier-name | default) {
            family [ inet mpls ];
          }
          exp (classifier-name | default);
          ieee-208.1 (classifier-name | default) <vlan-tag (inner | outer)>;
          ieee-208.1ad (classifier-name | default);
          inet-precedence (classifier-name | default);
        }
        forwarding-class class-name;
        input-scheduler-map map-name;
        input-shaping-rate bps;
        input-traffic-control-profile profile-name shared-instance instance-name;
        loss-priority-maps {
          (map-name | default);
        }
        loss-priority-rewrites {
          (map-name | default);
        }
        output-forwarding-class-map map-name;
        output-traffic-control-profile profile-name shared-instance instance-name;
        rewrite-rules {
          dscp (rule-name | default) <protocol mpls>;
          dscp-ipv6 (rule-name | default);
        }
      }
    }
  }
}

```

```

exp (rule-name | default) <protocol [ mpls-any | mpls-inet-both |
    mpls-inet-both-non-vpn ]>;
exp-push-push-push default;
exp-swap-push-push default;
ieee-802.1 (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
ieee-802.1ad (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
inet-precedence (rewrite-name | default) <protocol mpls>;
}
scheduler-map map-name;
shaping-rate bps;
translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 |
    to-exp-from-exp | to-inet-precedence-from-inet-precedence) table-name;
}
}
interface-set interface-set-name {
    excess-bandwidth-share (equal | proportional value);
    input-excess-bandwidth-share (equal | proportional value);
    input-traffic-control-profile profile-name;
    input-traffic-control-profile-remaining profile-name;
    internal-node;
    output-traffic-control-profile profile-name;
    output-traffic-control-profile-remaining profile-name;
}
}
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

[\[edit dynamic-profiles\] Hierarchy Level](#)

```

dynamic-profiles {
    profile-name {
        class-of-service
        ... statements from those in [edit class-of-service] Hierarchy Level.
        firewall
        ... statements from those in [edit firewall] Hierarchy Level.
        interfaces
        ... statements from those in [edit interfaces] Hierarchy Level.
        policy-options
        ... statements from those in [edit policy-options] Hierarchy Level.
        predefined-variable-defaults variable-name default-value
        profile-variable-set variable-set-name dynamic-variable-name substitute-variable-name
        protocols
        ... statements from those in [edit protocols] Hierarchy Level.
        routing-instances
        ... statements from those in [edit routing-instances] Hierarchy Level.
        routing-options
        ... statements from those in [edit routing-options] Hierarchy Level.
        services
        ... statements from those in [edit services] Hierarchy Level.
        variables {
            variable-name {
                default-value default-value;
                equals expression;
                mandatory;
            }
        }
    }
}

```

```
    }
    uid;
    uid-reference;
  }
}
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit dynamic-profiles routing-instances] Hierarchy Level*
- *[edit dynamic-profiles routing-options] Hierarchy Level*
- *[edit dynamic-profiles variables] Hierarchy Level*

[edit event-options] Hierarchy Level

```
event-options {
  destinations {
    destination-name {
      archive-sites {
        url <password password>;
      }
      transfer-delay seconds;
    }
  }
  event-script {
    file filename {
      checksum (md5 | sha-256 | sha1) hash;
      refresh;
      refresh-from url;
      remote-execution {
        remote-hostname {
          passphrase user-password;
          username user-login;
        }
      }
      source url;
    }
    max-datasize
    refresh;
    refresh-from url;
    traceoptions {
      file <filename> <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
      flag flag;
      no-remote-trace;
    }
  }
  generate-event event-name {
    time-interval seconds;
    time-of-day hh:mm:ss;
  }
  max-policies policies;
  policy policy-name {
    ... the policy subhierarchy appears after the main [edit event-options] hierarchy ...
  }
}
```

```

}
traceoptions {
  file <filename> <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
}

event-options {
  policy policy-name {
    attributes-match {
      event1.attribute-name equals event2.attribute-name;
      event.attribute-name matches regular-expression;
      event1.attribute-name starts-with event2.attribute-name;
    }
    events [ events ];
    then {
      change-configuration {
        commands {
          "command";
        }
        commit-options {
          check <synchronize>;
          force;
          log "comment-string";
          synchronize;
        }
        retry count number interval seconds;
        user-name username;
      }
      event-script filename {
        arguments {
          argument-name argument-value;
        }
        destination destination-name {
          retry-count number retry-interval seconds;
          transfer-delay seconds;
        }
        output-filename filename;
        output-format (text | xml);
        user-name username;
      }
      execute-commands {
        commands {
          "command";
        }
        destination destination-name {
          retry-count number retry-interval seconds;
          transfer-delay seconds;
        }
        output-filename filename;
        output-format (text | xml);
        user-name username;
      }
    }
    ignore;
  }
}

```

```
priority-override {
    facility facility-type;
    severity severity-level;
}
raise-trap;
upload filename (filename | committed) destination destination-name {
    retry-count number retry-interval seconds;
    transfer-delay seconds;
    user-name username;
}
}
within seconds {
    events [ events ];
    not events [ events ];
    trigger (after number | on number | until number);
}
}
```

Related Documentation

- [Notational Conventions Used in Junos OS Configuration Hierarchies](#)

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

Several statements in the **[edit firewall]** hierarchy are valid at numerous locations within the hierarchy. To make the complete hierarchy easier to read, the repeated statements are listed in the following sections, which are referenced at the appropriate locations in “[Complete \[edit firewall\] Hierarchy](#)” on page 325.

- [Common Firewall Actions on page 320](#)
- [Common IP Firewall Actions on page 321](#)
- [Common IPv4 Firewall Actions on page 321](#)
- [Common IP Firewall Match Conditions on page 322](#)
- [Common IPv4 Firewall Match Conditions on page 323](#)
- [Common Layer 2 Firewall Match Conditions on page 323](#)
- [Complete \[edit firewall\] Hierarchy on page 325](#)

Common Firewall Actions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit firewall\] Hierarchy](#)” on page 325 instead of the statements being repeated.

- **[edit firewall family (*any* | *ccc* | *ethernet-switching* | *inet* | *inet6* | *mpls* | *vpls*) filter *filter-name* term *term-name* then]**
- **[edit firewall filter *filter-name* term *term-name* then]**

The common firewall actions are as follows:

```
count counter-name;
forwarding-class class-name;
loss-priority (high | low | medium-high | medium-low);
```

```

next term;
policer policer-name;
three-color-policer policer-name {
    (single-rate single-rate-policer-name | two-rate two-rate-policer-name);
}

```

Common IP Firewall Actions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- **[edit firewall family inet filter *filter-name* term *term-name* then]**
- **[edit firewall family inet6 filter *filter-name* term *term-name* then]**
- **[edit firewall filter *filter-name* term *term-name* then]**

The common IP firewall actions are as follows:

```

log;
logical-system logical-system-name <routing-instance routing-instance-name>
    <topology topology-name>;
port-mirror;
port-mirror-instance instance-name;
routing-instance routing-instance-name <topology topology-name>;
sample;
service-filter-hit;
syslog;
topology topology-name;

```

Common IPv4 Firewall Actions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- **[edit firewall family inet filter *filter-name* term *term-name* then]**
- **[edit firewall filter *filter-name* term *term-name* then]**

The common IP version 4 (IPv4) firewall actions are as follows:

```

(accept | discard <accounting collector-name> | reject <administratively-prohibited |
    bad-host-tos | bad-network-tos | fragmentation-needed | host-prohibited |
    host-unknown | host-unreachable | network-prohibited | network-unknown |
    network-unreachable | port-unreachable | precedence-cutoff | precedence-violation |
    protocol-unreachable | source-host-isolated | source-route-failed | tcp-reset>);
ipsec-sa sa-name;
load-balance sa-name;
next-hop-group group-name;
prefix-action action-name;

```

Common IP Firewall Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit firewall\] Hierarchy](#)” on page 325 instead of the statements being repeated.

- [edit firewall family inet dialer-filter *filter-name* term *term-name* from] (with the exceptions noted at this level in “[Complete \[edit firewall\] Hierarchy](#)” on page 325)
- [edit firewall family inet filter *filter-name* term *term-name* from]
- [edit firewall family inet6 dialer-filter *filter-name* term *term-name* from] (with the exceptions noted at this level in “[Complete \[edit firewall\] Hierarchy](#)” on page 325)
- [edit firewall family inet6 filter *filter-name* term *term-name* from]
- [edit firewall filter *filter-name* term *term-name* from]

The common IP firewall match conditions are as follows:

```
address {
  ip-prefix</prefix-length> <except>;
}
destination-address {
  ip-prefix</prefix-length> <except>;
}
destination-class [ class-names ] | destination-class-except [ class-names ];
(destination-port [ port-names ] | destination-port-except [ port-names ]);
destination-prefix-list {
  list-name <except>;
}
(forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
icmp-code [ codes ] | icmp-code-except [ codes ];
icmp-type [ types ] | icmp-type-except [ types ];
interface interface-name;
(interface-group [ group-names ] | interface-group-except [ group-names ]);
interface-set set-name;
(loss-priority [ priorities ] | loss-priority-except [ priorities ]);
(packet-length [ values ] | packet-length-except [ values ]);
(port [ port-names ] | port-except [ port-names ]);
prefix-list {
  list-name <except>;
}
service-filter-hit;
source-address {
  ip-prefix</prefix-length> <except>;
}
(source-class [ class-names ] | source-class-except [ class-names ]);
(source-port [ port-names ] | source-port-except [ port-names ]);
source-prefix-list {
  list-name <except>;
}
tcp-established;
tcp-flags flag;
tcp-initial;
```


Common IPv4 Firewall Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit firewall\] Hierarchy](#)” on page 325 instead of the statements being repeated.

- **[edit firewall family inet dialer-filter *filter-name* term *term-name* from]** (with the exceptions noted at this level in “[Complete \[edit firewall\] Hierarchy](#)” on page 325)
- **[edit firewall family inet filter *filter-name* term *term-name* from]**
- **[edit firewall filter *filter-name* term *term-name* from]**

The common IPv4 firewall match conditions are as follows:

```
(ah-spi [ values ] | ah-spi-except [ values ]);
(dscp [ code-point-values ] | dscp-except [ code-point-values ]);
(esp-spi [ values ] | esp-spi-except [ values ]);
first-fragment;
fragment-flags flag;
(fragment-offset [ offsets ] | fragment-offset-except [ offsets ]);
(ip-options [ option-names ] | ip-options-except [ option-names ]);
is-fragment;
precedence-names ] | precedence-except [ precedence-names ]);
(protocol [ protocol-names ] | protocol-except [ protocol-names ]);
(ttl [ tll-values ] | ttl-except [ tll-values ]);
```

Common Layer 2 Firewall Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit firewall\] Hierarchy](#)” on page 325 instead of the statements being repeated.

- **[edit firewall family ethernet-switching filter *filter-name* term *term-name* from]**
- **[edit firewall family vpls filter *filter-name* term *term-name* from]**

The common Layer 2 firewall match conditions are as follows:

```
destination-mac-address {
    mac-address <except>;
}
(destination-port [ port-names ] | destination-port-except [ port-names ]);
(dscp [ code-point-values ] | dscp-except [ code-point-values ]);
(ether-type [ protocol-types ] | ether-type-except [ protocol-types ]);
(forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
(icmp-code [ codes ] | icmp-code-except [ codes ]);
(icmp-type [ types ] | icmp-type-except [ types ]);
(interface-group [ group-names ] | interface-group-except [ group-names ]);
ip-address {
    ip-prefix</prefix-length> <except>;
}
ip-destination-address {
    ip-prefix</prefix-length> <except>;
}
(ip-precedence [ precedence-names ] | ip-precedence-except [ precedence-names ]);
```

```
(ip-protocol [ protocol-names ] | ip-protocol-except [ protocol-names ]);
ip-source-address ip-prefix </prefix-length>;
(learn-vlan-1p-priority [ priorities ] | learn-vlan-1p-priority [ priorities ]);
(learn-vlan-id [ vlan-ids ] | learn-vlan-id-except [ vlan-ids ]);
(loss-priority [ priorities ] | loss-priority-except [ priorities ]);
(port [ port-names ] | port-except [ port-names ]);
source-mac-address {
    mac-address <except>;
}
(source-port [ port-names ] | source-port-except [ port-names ]);
tcp-flags flag;
(traffic-type [ broadcast known-unicast multicast unknown-unicast ] |
    traffic-type-except [ broadcast known-unicast multicast unknown-unicast ]);
(user-vlan-1p-priority [ priorities ] | user-vlan-1p-priority [ priorities ]);
(user-vlan-id [ vlan-ids ] | user-vlan-id-except [ vlan-ids ]);
(vlan-ether-type [ protocol-types ] | vlan-ether-type-except [ protocol-types ]);
```

Complete [edit firewall] Hierarchy

```

firewall {
  family (any | ccc | ethernet-switching | inet | inet6 | mpls | vpls) {
    ... the family subhierarchies appear after the main [edit firewall] hierarchy ...
  }
  filter filter-name {
    accounting-profile [ profile-names ];
    enhanced-mode;
    interface-shared-with;
    interface-specific;
    physical-interface-policer;
    term term-name {
      filter filter-name;
      from {
        ... statements in Common IP Firewall Match Conditions on page 322 AND
        statements in Common IPv4 Firewall Match Conditions on page 323 ...
      }
      then {
        ... statements in Common Firewall Actions on page 320 AND
        statements in Common IP Firewall Actions on page 321 AND
        statements in Common IPv4 Firewall Actions on page 321 ...
      }
    }
  }
}
hierarchical-policer policer-name {
  aggregate {
    if-exceeding {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    then {
      discard;
      forwarding-class class-name;
      loss-priority (high | low | medium-high | medium-low);
    }
  }
  logical-interface-policer;
  physical-interface-policer;
  premium {
    if-exceeding {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    then {
      discard;
    }
  }
}
}
shared-bandwidth-policer;
interface-set interface-set-name {
  interface-name;
}
load-balance-group group-name {
  next-hop-group [ group-names ];
}

```

```

}
policer policer-name {
  filter-specific;
  if-exceeding {
    (bandwidth-limit bps | bandwidth-percent percentage);
    burst-size-limit bytes;
  }
  logical-bandwidth-policer;
  logical-interface-policer;
  physical-interface-policer;
  then {
    discard;
    forwarding-class class-name;
    loss-priority (high | low | medium-high | medium-low);
  }
}
three-color-policer policer-name {
  action {
    loss-priority high then discard;
  }
  filter-specific;
  logical-interface-policer;
  physical-interface-policer;
  shared-bandwidth-policer;
  single-rate {
    (color-aware | color-blind);
    committed-burst-size bytes;
    committed-information-rate bps;
    excess-burst-size bytes;
  }
  two-rate {
    (color-aware | color-blind);
    committed-burst-size bytes;
    committed-information-rate bps;
    peak-burst-size bytes;
    peak-information-rate bps;
  }
}
}

firewall {
  family any {
    filter filter-name {
      interface-shared;
      term term-name {
        from {
          (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
          interface interface-name;
          interface-set set-name;
          (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
          (packet-length [ values ] | packet-length-except [ values ]);
        }
        then {
          ... statements in Common Firewall Actions on page 320 PLUS ...
          (accept | discard);
        }
      }
    }
  }
}

```

```

    }
  }
}

firewall {
  family ccc {
    filter filter-name {
      accounting-profile [ profile-names ];
      physical-interface-filter;
      interface-specific;
      term term-name {
        filter filter-name;
        from {
          (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
          (interface-group [ group-names ] | interface-group-except [ group-names ]);
          (learn-vlan-1p-priority [ priorities ] | learn-vlan-1p-priority [ priorities ]);
          (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
          (user-vlan-1p-priority [ priorities ] | user-vlan-1p-priority [ priorities ]);
        }
        then {
          ... statements in Common Firewall Actions on page 320 PLUS ...
          (accept | discard);
          port-mirror-instance instance-name;
        }
      }
    }
  }
}

firewall {
  family ethernet-switching {
    filter filter-name {
      interface-specific;
      term term-name {
        from {
          destination-address {
            ip-prefix</prefix-length>;
          }
          destination-mac-address {
            mac-address;
          }
          destination-port [ port-names ];
          destination-prefix-list {
            list-name;
          }
          dot1q-tag [ tag-values ];
          dot1q-user-priority [ priority-values ];
          dscp [ code-point-values ];
          ether-type [ protocol-names ];
          fragment-flags flag;
          icmp-code [ codes ];
          icmp-type [ types ];
          interface interface-name;
          is-fragment;

```

```

precedence [ precedence-names ];
protocol [ protocol-names ];
source-address {
    ip-prefix < / prefix-length >;
}
source-mac-address {
    mac-address;
}
source-port [ port-names ];
source-prefix-list {
    list-name;
}
tcp-established;
tcp-flags flag;
tcp-initial;
vlan [ vlan-names ];
}
then {
    (accept | discard);
    analyzer analyzer-name;
    count counter-name;
    forwarding-class class-name;
    interface interface-name;
    log;
    loss-priority (high | low);
    policer policer-name;
    syslog;
    vlan vlan-name;
}
}
}
}
}
}

firewall {
    family inet {
        dialer-filter filter-name {
            accounting-profile [ profile-names ];
            term term-name {
                from {
                    ... statements in Common IP Firewall Match Conditions on page 322 AND
                    statements in Common IPv4 Firewall Match Conditions on page 323 EXCEPT
                    FOR ...
                    (ah-spi [ values ] | ah-spi-except [ values ]); # NOT valid at this level
                    (destination-class [ class-names ] |
                     destination-class-except [ class-names ]); # NOT valid at this level
                    interface interface-name; # NOT valid at this level
                    (loss-priority [ priorities ] | loss-priority-except [ priorities ]); # NOT valid at
                     this level
                    service-filter-hit; # NOT valid at this level
                    (source-class [ class-names ] | source-class-except [ class-names ]); # NOT
                     valid at this level
                }
            }
            then {
                (ignore | note);
                log;
            }
        }
    }
}

```

```

        sample;
        syslog;
    }
}
filter filter-name {
    accounting-profile [ profile-names ];
    interface-specific;
    term term-name {
        filter filter-name;
        from {
            ... statements in Common IP Firewall Match Conditions on page 322 AND
               statements in Common IPv4 Firewall Match Conditions on page 323 ...
        }
        then {
            ... statements in Common Firewall Actions on page 320 AND
               statements in Common IP Firewall Actions on page 321 AND
               statements in Common IPv4 Firewall Actions on page 321 ...
        }
    }
}
prefix-action name {
    count;
    destination-prefix-length prefix-length;
    filter-specific;
    policer policer-name;
    source-prefix-length prefix-length;
    subnet-prefix-length prefix-length;
}
service-filter filter-name {
    term term-name {
        from {
            address {
                ip-prefix</prefix-length>;
            }
            (ah-spi [ values ] | ah-spi-except [ values ]);
            destination-address {
                ip-prefix</prefix-length>;
            }
            (destination-port [ port-names ] | destination-port-except [ port-names ]);
            destination-prefix-list {
                list-name;
            }
            (esp-spi [ values ] | esp-spi-except [ values ]);
            first-fragment;
            fragment-flags flag;
            (fragment-offset [ offsets ] | fragment-offset-except [ offsets ]);
            (interface-group [ group-names ] | interface-group-except [ group-names ]);
            (ip-options [ option-names ] | ip-options-except [ option-names ]);
            is-fragment;
            (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
            (port [ port-names ] | port-except [ port-names ]);
            prefix-list {
                list-name;
            }
            (protocol [ protocol-names ] | protocol-except [ protocol-names ]);
        }
    }
}

```

```

        source-address {
            ip-prefix </prefix-length>;
        }
        (source-port [ port-names ] | source-port-except [ port-names ]);
        source-prefix-list {
            list-name;
        }
        tcp-flags flag-name;
    }
    then {
        count counter-name;
        log;
        port-mirror;
        sample;
        (service | skip);
    }
}
}
simple-filter filter-name {
    term term-name {
        from {
            destination-address ip-prefix </prefix-length>;
            destination-port port-name;
            forwarding-class [ class-names ];
            protocol protocol-name;
            source-address ip-prefix </prefix-length>;
            source-port port-name;
        }
        then {
            forwarding-class class-name;
            loss-priority (high | low | medium-high | medium-low);
            policer policer-name;
        }
    }
}
}
}
}
firewall {
    family inet6 {
        dialer-filter filter-name {
            accounting-profile [ profile-names ];
            term term-name {
                from {
                    ... statements in Common IP Firewall Match Conditions on page 322 PLUS ...
                    (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
                    ... BUT NOT ...
                    (destination-class [ class-names ] |
                     destination-class-except [ class-names ]); # NOT valid at this level
                    (forwarding-class [ class-names ] |
                     forwarding-class-except [ class-names ]); # NOT valid at this level
                    interface interface-name; # NOT valid at this level
                    (interface-group [ group-names ] | interface-group-except [ group-names ]); #
                     NOT valid at this level
                    (loss-priority [ priorities ] | loss-priority-except [ priorities ]); # NOT valid at
                     this level
                }
            }
        }
    }
}

```



```

        service-filter-hit; # NOT valid at this level
        (source-class [ class-names ] | source-class-except [ class-names ]); # NOT
            valid at this level
        tcp-established; # NOT valid at this level
        tcp-flags flag; # NOT valid at this level
        tcp-initial; # NOT valid at this level
    }
    then {
        (ignore | note);
        log;
        sample;
        syslog;
    }
}
}
filter filter-name {
    accounting-profile [ profile-names ];
    interface-specific;
    term term-name {
        filter filter-name;
        from {
            ... statements in Common IP Firewall Match Conditions on page 322 PLUS ...
            (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
            (traffic-class [ code-point-values ] | traffic-class-except [ code-point-values ]);
        }
        then {
            ... statements in Common Firewall Actions on page 320 AND
            statements in Common IP Firewall Actions on page 321 PLUS ...
            (accept | discard | reject <address-unreachable | administratively-prohibited |
                beyond-scope | fragmentation-needed | no-route | port-unreachable |
                tcp-reset>);
        }
    }
}
}
service-filter filter-name {
    term term-name {
        from {
            address {
                ip-prefix</prefix-length>;
            }
            (ah-spi [ values ] | ah-spi-except [ values ]);
            destination-address {
                ip-prefix</prefix-length>;
            }
            (destination-port [ port-names ] | destination-port-except [ port-names ]);
            destination-prefix-list {
                list-name;
            }
            (esp-spi [ values ] | esp-spi-except [ values ]);
            (interface-group [ group-names ] | interface-group-except [ group-names ]);
            (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
            (port [ port-names ] | port-except [ port-names ]);
            prefix-list {
                list-name;
            }
            source-address {

```

```

        ip-prefix </prefix-length>;
    }
    (source-port [ port-names ] | source-port-except [ port-names ]);
    source-prefix-list {
        list-name;
    }
    tcp-flags flag-name;
}
then {
    count counter-name;
    log;
    port-mirror;
    sample;
    (service | skip);
}
}
}
}
}

firewall {
    family mpls {
        filter filter-name {
            accounting-profile [ profile-names ];
            interface-specific;
            physical-interface-filter;
            term term-name {
                from {
                    (exp [ exp-bits ] | exp-except [ exp-bits ]);
                }
                then {
                    (ignore | note);
                    log;
                    sample;
                    syslog;
                }
            }
        }
    }
    filter filter-name {
        accounting-profile [ profile-names ];
        interface-specific;
        physical-interface-filter;
        term term-name {
            filter filter-name;
            from {
                (exp [ exp-bits ] | exp-except [ exp-bits ]);
                (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
                interface interface-name;
                interface-set set-name;
                (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
            }
            then {
                ... statements in Common Firewall Actions on page 320 PLUS ...
                (accept | discard);
                sample;
            }
        }
    }
}

```

```

    }
  }
}

firewall {
  family vpls {
    filter filter-name {
      accounting-profile [ profile-names ];
      interface-specific;
      term term-name {
        filter filter-name;
        from {
          ... statements in Common Layer 2 Firewall Match Conditions on page 323 ...
        }
        then {
          ... statements in Common Firewall Actions on page 320 PLUS ...
          (accept | discard);
          port-mirror;
          port-mirror-instance instance-name;
        }
      }
    }
  }
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

[\[edit forwarding-options\] Hierarchy Level](#)

Each of the following topics lists the statements at a subhierarchy of the **[edit forwarding-options]** hierarchy.

- [\[edit forwarding-options accounting\] Hierarchy Level on page 334](#)
- [\[edit forwarding-options dhcp-relay\] Hierarchy Level on page 334](#)
- [\[edit forwarding-options enhanced-hash-key\] Hierarchy Level on page 338](#)
- [\[edit forwarding-options family\] Hierarchy Level on page 339](#)
- [\[edit forwarding-options fast-reroute-priority\] Hierarchy Level on page 340](#)
- [\[edit forwarding-options hash-key\] Hierarchy Level on page 340](#)
- [\[edit forwarding-options helpers\] Hierarchy Level on page 341](#)
- [\[edit forwarding-options load-balance\] Hierarchy Level on page 343](#)
- [\[edit forwarding-options next-hop-group\] Hierarchy Level on page 344](#)
- [\[edit forwarding-options port-mirroring\] Hierarchy Level](#)
- [\[edit forwarding-options rpf-loose-mode-discard\] Hierarchy Level on page 345](#)
- [\[edit forwarding-options sampling\] Hierarchy Level on page 346](#)

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

[edit forwarding-options accounting] Hierarchy Level

```
forwarding-options {
  accounting group-name {
    output {
      aggregate-export-interval seconds;
      cflowd hostname {
        aggregation {
          autonomous-system;
          destination-prefix;
          protocol-port;
          source-destination-prefix {
            caida-compliant;
          }
          source-prefix;
        }
        autonomous-system-type (origin | peer);
        port port-number;
        version format;
      }
      flow-active-timeout seconds;
      flow-inactive-timeout seconds;
      interface interface-name {
        engine-id number;
        engine-type number;
        source-address address;
      }
    }
  }
}
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- [\[edit forwarding-options\] Hierarchy Level on page 333](#)

[edit forwarding-options dhcp-relay] Hierarchy Level

```
forwarding-options {
  dhcp-relay {
    active-server-group server-group-name;
  }
  arp-inspection;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-name;
      logical-system-name;
      mac-address;
      option-60;
    }
  }
}
```

```

    option-82 <circuit-id> <remote-id>;
    routing-instance-name;
    user-prefix user-prefix-string;
  }
}
dhcpv6 {
  active-server-group group-name;
}
authentication {
  password password-string;
}
username-include {
  circuit-type;
  client-id;
  delimiter delimiter-character;
  domain-name domain-name-string;
  interface-name;
  logical-system-name;
  relay-agent-interface-id;
  relay-agent-remote-id;
  relay-agent-subscriber-id;
  routing-instance-name;
  user-prefix user-prefix-string;
}
dynamic-profile profile-name {
  aggregate-clients (merge |replace);
  use-primary profile-name;
}
group group-name {
  ... the group subhierarchy appears after the main [edit forwarding-options
    dhcp-relay] hierarchy ...
}
liveness-detecton {
  ... the liveness-detection subhierarchy appears after the main [edit
    forwarding-options dhcp-relay] hierarchy ...
}
overrides {
  ... the overrides subhierarchy appears after the main [edit forwarding-options
    dhcp-relay] hierarchy ...
}
relay-agent-interface-id {
  prefix;
  user-interface-description;
}
relay-option {
  default-action;
  equals;
  option-number;
  starts-with;
}
server-group;
service-profile;
)
duplicate-clients-on-interface;
dynamic-profile profile-name {
  aggregate-clients (merge |replace);
  use-primary profile-name;
}

```

```

}
forward-snooped-clients (all-interfaces | configured-interfaces |
non-configured-interfaces);
group group-name {
  active-server-group server-group-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      logical-system-name;
      mac-address;
      option-60;
      option-82 <circuit-id> <remote-id>;
      routing-instance-name;
      user-prefix user-prefix-string;
    }
    dynamic-profile profile-name {
      aggregate-clients (merge | replace);
      use-primary profile-name;
    }
  }
  interface;
  liveness-detection;
  overrides;
  relay-option;
  relay-option-82 ;
  service-profile;
}
}
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
}
method {
  bfd {
    detection-time {
      threshold milliseconds;
    }
    holddown-interval;
    minimum--interval;
    minimum-receive-interval;
    multiplier;
    no-adaptation;
    session-mode;
  }
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
  version;
}
overrides {
  (allow-snooped-clients | no-allow-snooped-clients);
  always-write-giaddr;
  always-write-option-82;
  client-discover-match <option60-and-option82>;
}

```

```

disable-relay;
interface-client-limit number;
layer2-unicast-replies;
no-allow-snooped-clients;
no-arp;
no-bind-on-request;
no-unicast-replies;
proxy-mode;
replace-ip-source-with giaddr;
send-release-on-delete;
trust-option-82;
}
relay-option {
  default-action {
    drop;
    forward-only;
    local-server-group group-name;
    relay-server-group group-name;
  }
  equals {
    ascii string;
    hexadecimal string;
  }
}
option-number (60 | 77);
}
starts-with {
  ascii string;
  hexadecimal string;
}
relay-option-82 {
  circuit-id (value | ... the following prefix statement ...) {
    prefix {
      host-name;
      logical-system-name;
      routing-instance-name;
    }
    use-interface-description (device | logical);
  }
}
server-group {
  server-group-name {
    ip-address;
  }
}
service-profile name:
}
}

```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- [\[edit forwarding-options\] Hierarchy Level on page 333](#)

[\[edit forwarding-options enhanced-hash-key\] Hierarchy Level](#)

```
enhanced-hash-key {
  family [
    inet
      gtp-tunnel-endpoint-identifier;
      incoming-interface-index;
      no-destination-port;
      no-source-port;
      type-of-service;
    }
    inet6 {
      gtp-tunnel-endpoint-identifier;
      incoming-interface-index;
      no-destination-port;
      no-source-port;
      traffic-class;
    }
    mpls {
      incoming-interface-index;
      label-1-exp;
      no-payload;
    }
    multiservice {
      incoming-interface-index;
      no-mac-addresses;
      no-payload;
    }
  ]
  services-loadbalancing {
    family {
      inet layer-3-services {
        destination-address;
        incoming-interface-index;
        source-address;
      }
    }
    inet6 {
      destination-address;
      incoming-interface-index;
      source-address;
    }
  }
}
symmetric {
  family {
    inet {
      gtp-tunnel-endpoint-identifier;
      incoming-interface-index;
      no-destination-port;
      no-source-port;
      type-of-service;
    }
    inet6 {
      gtp-tunnel-endpoint-identifier;
```



```

        incoming-interface-index;
        no-destination-port;
        no-source-port;
        traffic-class;
    }
}
mpls {
    incoming-interface-index;
    label-1-exp;
    no-payload;
}
multiservice {
    incoming-interface-index;
    no-mac-addresses;
    no-payload;
}
services-loadbalancing {
    family {
        inet layer-3-services {
            destination-address;
            incoming-interface-index;
            source-address;
        }
    }
    inet6 {
        destination-address;
        incoming-interface-index;
        source-address;
    }
}
}

```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - [\[edit forwarding-options\] Hierarchy Level on page 333](#)

[\[edit forwarding-options family\] Hierarchy Level](#)

```

forwarding-options {
    family inet {
        filter {
            input filter-name;
            output filter-name;
        }
    }
    family inet6 {
        filter {
            input filter-name;
            output filter-name;
            route-accounting;
            source-checking;
        }
    }
    family mpls {
        filter {
            input filter-name;
            output filter-name;
        }
    }
}

```

```
    }  
  }  
  family vpls {  
    filter {  
      input filter-name;  
    }  
    flood {  
      input filter-name;  
    }  
  }  
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - [\[edit forwarding-options\] Hierarchy Level on page 333](#)

[\[edit forwarding-options fast-reroute-priority\] Hierarchy Level](#)

```
forwarding-options {  
  fast-reroute-priority (low | medium | high);  
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - [\[edit forwarding-options\] Hierarchy Level on page 333](#)

[\[edit forwarding-options hash-key\] Hierarchy Level](#)

```
forwarding-options {  
  hash-key {  
    family inet {  
      layer-3;  
      layer-4;  
      symmetric-hash {  
        complement;  
      }  
    }  
    family mpls {  
      label-1;  
      label-2;  
      label-3;  
      no-labels;  
      payload {  
        ether-pseudowire;  
        ip {  
          layer-3-only;  
          port-data {  
            destination-lsb;  
            destination-msb;  
            source-lsb;  
            source-msb;  
          }  
        }  
      }  
    }  
  }  
}
```

```

    }
    family multiservice {
        destination-mac;
        payload {
            ip {
                layer-3 {
                    (destination-ip-only | source-ip-only);
                }
                layer-4;
            }
        }
        source-mac;
        symmetric-hash {
            complement;
        }
    }
}

```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- [\[edit forwarding-options\] Hierarchy Level on page 333](#)

[\[edit forwarding-options helpers\] Hierarchy Level](#)

```

forwarding-options {
    helpers {
        bootp {
            client-response-ttl number;
            description text-description;
            dhcp-option82 {
                disable;
                circuit-id {
                    prefix hostname;
                    use-interface-description;
                    use-vlan-id;
                }
                remote-id {
                    prefix (hostname | mac | none);
                    use-interface-description;
                    use-string text-string;
                }
                vendor-id {
                    text-string;
                }
            }
        }
        interface interface-name-or-wildcard {
            broadcast;
            client-response-ttl number;
            description text-description;
            dhcp-option82 {
                ... same statements as at the [edit forwarding-options helpers bootp] hierarchy level ...
            }
            maximum-hop-count number;

```

```
        minimum-wait-time seconds;
        no-listen;
        server address {
            logical-system logical-system-name <routing-instance [ <default>
                routing-instance-names ]>;
            routing-instance [ <default> routing-instance-names ];
        }
    }
    maximum-hop-count number;
    minimum-wait-time seconds;
    relay-agent-option;
    server address {
        logical-system logical-system-name <routing-instance [ <default>
            routing-instance-names ]>;
        routing-instance [ <default> routing-instance-names ];
    }
}

helpers {
    domain {
        description text-description;
        interface {
            interface-name {
                broadcast;
                description text-description;
                no-listen;
                server <address> <logical-system logical-system-name>
                    <routing-instance (default | routing-instance-name)>;
            }
        }
        server <address> <logical-system logical-system-name>
            <routing-instance (default | routing-instance-name)>;
    }
}

helpers {
    port port-number {
        description text-description;
        interface {
            interface-name {
                broadcast;
                description text-description;
                no-listen;
                server <address> <logical-system logical-system-name>
                    <routing-instance (default | routing-instance-name)>;
            }
        }
        server <address> <logical-system logical-system-name>
            <routing-instance (default | routing-instance-name)>;
    }
}

helpers {
    rtsdb-client-traceoptions {
        if-rtsdb
```

```

        flag (all |init |map |routing-socket);
    }
}
helpers {
    tftp {
        description text-description;
        interface {
            interface-name {
                broadcast;
                description text-description;
                no-listen;
                server <address> <logical-system logical-system-name>
                    <routing-instance (default | routing-instance-name)>;
            }
        }
        server <address> <logical-system logical-system-name>
            <routing-instance (default | routing-instance-name)>;
    }
}

helpers {
    traceoptions {
        file <filename> <files number> <match regular-expression> <size maximum-file-size>
            <world-readable | no-world-readable>;
        flag flag;
        level severity;
        no-remote-trace;
    }
}

```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - [\[edit forwarding-options\] Hierarchy Level on page 333](#)

[\[edit forwarding-options load-balance\] Hierarchy Level](#)

You can configure load balancing under the [\[edit forwarding options\]](#) hierarchy.

```

forwarding-options {
    load-balance {
        indexed-next-hop;
        per-flow {
            hash-seed;
        }
        per-prefix {
            hash-seed number;
        }
    }
}

```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - [\[edit forwarding-options\] Hierarchy Level on page 333](#)

[\[edit forwarding-options next-hop-group\] Hierarchy Level](#)

```
forwarding-options {
  next-hop-group group-name {
    group-type {
      layer-2 {
        interface interface-name {
          next-hop address;
        }
      }
      next-hop-subgroup subgroup-name ;
    }
    interface interface-name {
      next-hop address;
    }
    next-hop-subgroup subgroup-name ;
  }
}
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- [\[edit forwarding-options\] Hierarchy Level on page 333](#)

[\[edit forwarding-options port-mirroring\] Hierarchy Level](#)

```
forwarding-options {
  port-mirroring {
    disable;
    disable-all-instances;
    family (ccc | ethernet-switching | inet | inet6 | vpls) {
      output {
        (interface interface-name | next-hop-group group-name);
        no-filter-check;
        routing-instance instance-name;
        vlan (vlan-name | vlan-id) <no-tag>;
      }
    }
    family ccc {
      output {
        interface interface-name
        next-hop-group group-name;
        no-filter-check;
      }
    }
    family ethernet-switching {
      output {
        interface interface-name
        next-hop-group group-name;
        no-filter-check;
      }
    }
    family inet {
      output {
        interface interface-name {
          next-hop ipv4-address;
        }
      }
      next-hop-group group-name;
    }
  }
}
```

```

        no-filter-check;
    }
}
family inet6 {
    output {
        interface interface-name {
            next-hop ipv6-address;
        }
        no-filter-check;
    }
}
family vpls {
    output {
        interface interface-name
        next-hop-group group-name;
        no-filter-check;
    }
}
input {
    maximum-packet-length bytes;
    rate rate;
}
instance instance-name {
    disable;
    family family-name {
        ... same statements as at the [edit forwarding-options port-mirroring family (ccc |
            inet | inet6 | vpls)] hierarchy levels ...
    }
    input {
        ... same statements as at the [edit forwarding-options port-mirroring input] hierarchy
            level ...
    }
}
mirror-once;
traceoptions {
    file <filename> <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
    no-remote-trace;
}
}
}

```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit forwarding-options] Hierarchy Level*

[edit forwarding-options rpf-loose-mode-discard] Hierarchy Level

```

rpf-loose-mode-discard {
    family {
        inet;
        inet6;
    }
}

```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - [\[edit forwarding-options\] Hierarchy Level on page 333](#)

[\[edit forwarding-options sampling\] Hierarchy Level](#)

```
forwarding-options {
  sampling {
    disable;
  }
  family {
    inet {
      disable;
    }
    output {
      aggregate-export-interval seconds;
      extension-service service-name;
      file {
        disable;
        filename filename;
        files number;
        size bytes;
        (stamp | no-stamp);
        (world-readable | no-world-readable);
      }
      flow-active-timeout seconds;
      flow-inactive-timeout seconds;
    }
    flow-server hostname {
      aggregation {
        autonomous-system-type (origin | peer);
        (local-dump | no-local-dump);
        port port-number;
        source-address address;
        version format;
        version9 {
          template template-name;
        }
      }
    }
  }
  interface interface-name {
    engine-id number;
    engine-type number;
    source-address address;
  }
}
inet6 {
  disable;
}
output {
  aggregate-export-interval seconds;
  extension-service service-name;
  flow-active-timeout seconds;
  flow-inactive-timeout seconds;
```



```

    }
    flow-server hostname {
        aggregation {
            autonomous-system-type (origin | peer);
            (local-dump | no-local-dump);
            port port-number;
            source-address address;
            version9 {
                template template-name;
            }
        }
    }
    interface {
        ... same statements as at the [edit forwarding-options sampling family inet]
        hierarchy level ...
    }
}
mpls {
    disable;
}
output {
    aggregate-export-interval seconds;
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
}
flow-server hostname {
    aggregation {
        autonomous-system-type (origin | peer);
        (local-dump | no-local-dump);
        port port-number;
        source-address address;
        version9 {
            template template-name;
        }
    }
    interface {
        ... same statements as at the [edit forwarding-options sampling family
        inet] hierarchy level ...
    }
}
}
input {
    rate number;
    run-length number;
    max-packets-per-second number;
    maximum-packet-length bytes;
}
instance instance-name {
    disable;
    input {
        rate number;
        run-length number;
        max-packets-per-second number;
        maximum-packet-length bytes;
    }
}

```

```
    }
    sample-once;
    traceoptions {
        no-remote-trace;
        file filename <files number> <size bytes> <match expression> <world-readable |
        no-world-readable>;
    }
}
}
```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- [\[edit forwarding-options\] Hierarchy Level on page 333](#)

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

```
groups {
    group-name {
        ... statements from any subhierarchy at the [edit] hierarchy level ...
    }
}
```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

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

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```
interfaces {
    interface-name {
        ... the "interface-name" subhierarchy appears after the main [edit interfaces] hierarchy level ...
    }
    interface-set interface-set-name {
        interface interface-name {
            (unit unit-number | vlan-tags-outer vlan-tag);
        }
    }
    irb (Interfaces) {
        accounting-profile name;
        description text;
        disable;

        (gratuitous-arp-reply | no-gratuitous-arp-reply);
        hold-time up milliseconds down milliseconds;
        mtu bytes;
        no-gratuitous-arp-request;

        traceoptions {
            flag flag;
        }
        (traps | no-traps);
        unit logical-unit-number {
```

```

accounting-profile name;
bandwidth rate;
description text;
disable;
encapsulation type;
family inet {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address ipv4-address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    preferred;
    primary;
    vrrp-group group-id {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        (preempt | no-preempt) {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bits-per-second priority-cost priority;
                priority-cost priority;
            }
            priority-hold-time seconds;
            route prefix/prefix-length routing-instance instance-name priority-cost priority;
        }
        virtual-address [ addresses ];
        vrrp-inherit-from vrrp-group;
    }
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
targeted-broadcast {

```

```
        forward-and-send-to-re;
        forward-only;
    }
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bandwidth priority-cost number;
                priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-inet6-address [addresses];
        virtual-link-local-address ipv6-address;
        vrrp-inherit-from {
            active-group group-number;
            active-interface interface-name;
        }
    }
}
(dad-disable | no-dad-disable);
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
no-redirects;
policer {
    input policer-name;
    output policer-name;
```

```

    }
    rpf-check {
        fail-filter filter-name;
        mode {
            loose;
        }
    }
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    filter {
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    policer {
        input policer-name;
        output policer-name;
    }
}
native-inner-vlan-id vlan-id;
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id-list [vlan-id's];
vlan-id-range [vlan-id-range];
}
}
traceoptions {
    file <filename> <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
    flag flag <disable>;
    no-remote-trace;
}
}
interfaces {
    interface-name {
        disable;
        accounting-profile name;
        aggregated-ether-options {
            ethernet-switch-profile {
                tag-protocol-id [ hexadecimal-identifiers ];
            }
        }
        (flow-control | no-flow-control);
        lacp {
            (active | passive);
            admin-key key;
            fast-failover;
            link-protection {
                disable;
                (revertive | non-revertive);
            }
        }
        periodic (fast | slow);
    }
}

```

```

    system-id mac-address;
    system-priority priority;
}
(link-protection | no-link-protection);
link-speed (100m | 1g | 8g | 10g | 40g | 50g | 80g | 100g | oc192);
logical-interface-fpc-redundancy;
(loopback | no-loopback);
mc-ae {
    chassis-id chassis-id;
    events {
        iccp-peer-down {
            force-icl-down;
            prefer-status-control-active;
        }
    }
    mc-ae-id mc-ae-id;
    mode (active-active | active-standby);
    redundancy-group group-id;
    status-control (active | standby);
}
minimum-links number;
rebalance-periodic {
    start-time time;
    interval number;
}
source-address-filter {
    mac-address;
}
(source-filtering | no-source-filtering);
}
auto-configure {
    remove-when-no-subscribers;
    stacked-vlan-ranges {
        access-profile profile-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                mac-address;
                option-82 ( circuit-id | remote-id);
                radius-realm radius-realm-string;
                user-prefix user-prefix-string;
            }
        }
    }
    dynamic-profile profile-name {
        accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
        ranges (any | low-tag-high-tag),(any | low-tag-high-tag);
    }
}
vlan-ranges {
    access-profile profile-name;
    authentication {
        password password-string;

```

```

username-include {
    circuit-type;
    delimiter delimiter-character;
    domain-name domain-name-string;
    interface-name;
    mac-address;
    option-82;
    radius-realm radius-realm-string;
    user-prefix user-prefix-string;
}
}
dynamic-profile profile-name {
    accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
    ranges (any | low-tag)–(any | high-tag);
}
}
override tag vlan-tag dynamic-profile profile name;
}
encapsulation (ethernet-bridge | ethernet-vpls | extended-vlan-bridge |
    extended-vlan-vpls | flexible-ethernet-services | vlan-vpls);
ether-options {
    802.3ad {
        aex;
        (backup | primary);
        lacp {
            force-up;
            port-priority
        }
    }
}
asynchronous-notification;
(auto-negotiation | no-auto-negotiation);
ethernet-switch-profile {
    ethernet-policer-profile {
        input-priority-map {
            ieee802.1p premium [ values ];
        }
        output-priority-map {
            classifier {
                premium {
                    forwarding-class class-name {
                        loss-priority (high | low);
                    }
                }
            }
        }
    }
}
policer cos-policer-name {
    aggregate {
        bandwidth-limit bps;
        burst-size-limit bytes;
    }
    premium {
        bandwidth-limit bps;
        burst-size-limit bytes;
    }
}
tag-protocol-id;

```

```

    }
    (mac-learn-enable | no-mac-learn-enable);
  }
  (flow-control | no-flow-control);
  ignore-l3-incompletes;
  link-mode (automatic | full-duplex | half-duplex);
  (loopback | no-loopback);
  keepalives <interval seconds> <down-count number> <up-count number>;
  speed (1g | 10m | 100m | 10m-100m | auto-negotiation);
  source-address-filter {
    mac-address;
  }
  source-filtering | no-source-filtering;
}
flexible-vlan-tagging;
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time (up milliseconds | down milliseconds);
interface-transmit-statistics;
(keepalives <down-count number> <interval seconds> <up-count number> |
 no-keepalives);
layer2-policer {
  apply-groups [ group-names ];
  apply-groups-except [ group-names ];
}
link-mode (automatic | full-duplex);
mac mac-address;
mtu bytes;
multi-chassis-protection peer-ip-address {
  interface interface-name;
}
native-vlan-id number;
no-gratuitous-arp-request;
optics-options {
  alarm low-light-alarm {
    (link-down | syslog);
  }
  warning low-light-warning {
    (link-down | syslog);
  }
  wavelength nm;
}
passive-monitor-mode;
per-unit-scheduler;
speed (10m | 100m | 1g | auto | oc3 | oc12 | oc48);
stacked-vlan-tagging;
traceoptions {
  flag flag;
}
transmit-bucket {
  overflow discard;
  rate percentage;
  threshold bytes;
}
(traps | no-traps);
unidirectional;
vlan-tagging;

```



```

}

interface-name {
  unit logical-unit-number {
    disable;
    accept-source-mac {
      mac-address mac-address {
        policer {
          input policer-name;
          output policer-name;
        }
      }
    }
  }
  account-layer2-overhead (Interface Level) {
    value;
    egress bytes;
    ingress bytes;
  }
  accounting-profile name;
  advisory-options {
    downstream-rate rate;
    upstream-rate rate;
  }
  arp-resp (restricted|unrestricted);
  bandwidth rate;
  clear-dont-fragment-bit;
  copy-tos-to-outer-ip-header;
  demux-destination family;
  encapsulation (vlan-bridge | vlan-vpls);
  epd-threshold cells plp1 cells;
  filter filter-name;
  inner-vlan-id-range start start-id end end-id;
  input-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
  }
  interface-shared-with psdnumerical-index;
  layer2-policer {
    input-hierarchical-policer policer-name;
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
  }
  multi-chassis-protection peer-ip-address {
    interface interface-name;
  }
  native-inner-vlan-id number;
  output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
  }
}

```

```

    tag-protocol-id tpid;
    vlan-id number;
}
peer-interface interface-name;
peer-unit unit-number;
plp-to-clp;
proxy-arp <restricted | unrestricted>;
rpm {
    (client | server);
    twamp-server;
}
swap-by-poppush;
vlan-id number;
vlan-id-list [ vlan-id vlan-id-vlan-id ];
vlan-id-range number-number;
vlan-tags (inner <tpid.>vlan-id | inner-list [ vlan-id vlan-id-vlan-id ] |
    inner-range <tpid.>vlan-id-vlan-id) outer <tpid.>vlan-id;
}

unit logical-unit-number {
    family ethernet-switching {
        filter {
            group filter-group-number;
            (input filter-name | input-list [ filter-names ]);
            (output filter-name | output-list [ filter-names ]);
            (inner-vlan-id-list [ vlan-ids ] | vlan-id number | vlan-id-list [ number
                number-number ]);
            interface-mode (access | trunk);
        }
        policer {
            input policer-name;
            output policer-name;
        }
        vlan-rewrite {
            translate old-vlan-id new-vlan-id;
        }
        vlan {
            members [ all vlan-identifiers ];
        }
    }
}

family inet {
    filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
    input-hierarchical-policer policer-name;
    mac-validate (loose | strict);
    mtu bytes;
    no-neighbor-learn;
    no-redirects;
    policer {
        arp policer-template-name;
        input policer-name;
        output policer-name;
    }
    primary;
}

```

```

receive-options-packets;
receive-ttl-exceeded;
rpf-check {
    fail-filter filter-name;
    mode loose;
}
sampling {
    (input | output | input output);
}
simple-filter {
    input filter-name;
}
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
unnumbered-address interface-name <destination address>
    <destination-profile profile-name> <preferred-source-address address>;
}

family inet6 {
    address ipv6-address {
        destination destination-address;
        eui-64;
        ndp ipv6-address <l2-interface interface-name> <(mac mac-address |
            multicast-mac multicast-mac-address) <publish>>;
        preferred;
        primary;
        vrrp-inet6-group group-number {
            (accept-data | no-accept-data);
            fast-interval milliseconds;
            inet6-advertise-interval seconds;
            (no-preempt; | ... the following preempt statement ...)
            preempt {
                hold-time seconds;
            }
            priority number;
            track {
                interface interface-name {
                    bandwidth-threshold bits-per-second priority-cost priority;
                    priority-cost priority;
                }
                priority-hold-time seconds;
                route ip-address-prefix/prefix-length routing-instance instance-name
                    priority-cost priority;
            }
            virtual-inet6-address [ addresses ];
            virtual-link-local-address ipv6-address;
            vrrp-inherit-from {
                active-group group-number;
                active-interface interface-name;
            }
        }
    }
    (dad-disable | no-dad-disable);
}

```

```
filter {
  group filter-group-number;
  (input filter-name | input-list [ filter-names ]);
  (output filter-name | output-list [ filter-names ]);
}
input-hierarchical-policer policer-name;
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
policer {
  input policer-name;
  output policer-name;
}
rpf-check {
  fail-filter filter-name;
  mode loose;
}
sampling {
  (input | output | input output);
}
unnumbered-address interface-name preferred-source-address address;
}

family iso {
  address iso-address;
  mtu bytes;
}

family mlfrr-end-to-end {
  bundle logical-interface-name;
}

family mpls {
  filter {
    group filter-group-number;
    (input filter-name | input-list [ filter-names ]);
    (output filter-name | output-list [ filter-names ]);
  }
  input-hierarchical-policer policer-name;
  maximum-labels maximum-labels;
  mtu bytes;
  policer {
    input policer-name;
    output policer-name;
  }
}

family vpls {
  core-facing;
  filter {
    group filter-group-number;
```

```

        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
    policer {
        input policer-name;
        output policer-name;
    }
}
}
}

```

Related Documentation

- [Notational Conventions Used in Junos OS Configuration Hierarchies](#)

[\[edit logical-systems\] Hierarchy Level](#)

As indicated in the following hierarchy, you can include at this hierarchy level several of the hierarchies that can be included at the **[edit]** hierarchy level. However, some statements in a subhierarchy are not valid for logical systems. To learn which statements can be included under **[edit logical-systems *logical-system-name*]** on your device, issue the **set ?** command at the hierarchy level of interest.

```

logical-systems {
  logical-system-name {
    access {
      address-assignment {
        ... same statements as in the address-assignment subhierarchy in [edit access]
        Hierarchy Level ...
      }
    }
    access-profile profile-name;
    bridge-domains {
      ... (MX Series only) same statements as in [edit bridge-domains] Hierarchy Level ...
    }
    bridge-domains {
      ... (MX Series only) same statements as in [edit bridge-domains] Hierarchy Level ...
    }
    firewall {
      ... same statements as in several subhierarchies in [edit firewall] Hierarchy Level ...
    }
    forwarding-options {
      ... same statements as in [edit forwarding-options dhcp-relay] Hierarchy Level ...
    }
    interfaces {
      interface-name {
        unit logical-unit-number {
          ... some of the statements in the unit subhierarchy in [edit interfaces] Hierarchy
          Level ...
        }
      }
    }
    policy-options {
      ... same statements as in \[edit policy-options\] Hierarchy Level on page 360 ...
    }
  }
}

```

```
}
protocols {
  ... same statements as in [edit protocols] Hierarchy Level ...
}
routing-instances {
  ... most statements in [edit routing-instances] Hierarchy Level ...
}
routing-options {
  ... most statements in [edit routing-options] Hierarchy Level ...
}
services {
  mobile-ip {
    ... same statements as in [edit services mobile-ip] Hierarchy Level ...
  }
}
switch-options {
  ... (MX Series only) same statements as in [edit switch-options] Hierarchy Level ...
}
system {
  services {
    dhcp-local-server {
      ... same statements as in the services dhcp-local-server subhierarchy in [edit
        system] Hierarchy Level ...
    }
  }
  syslog {
    ... most statements in syslog subhierarchy in [edit system] Hierarchy Level...
  }
}
}
```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

[\[edit multi-chassis\] Hierarchy Level](#)

```
multi-chassis {
  multi-chassis-protection ipv4-address {
    interface interface-name;
  }
}
```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

[\[edit policy-options\] Hierarchy Level](#)

Several statements in the **[edit policy-options]** hierarchy are valid at numerous locations within the hierarchy. To make the complete hierarchy easier to read, the repeated

statements are listed in the following sections, which are referenced at the appropriate locations in [“Complete \[edit policy-options\] Hierarchy” on page 363](#).

- [Common Policy Terms on page 361](#)
- [Common Policy Match Conditions on page 362](#)
- [Common Ingress Policy Match Conditions on page 363](#)
- [Complete \[edit policy-options\] Hierarchy on page 363](#)

Common Policy Terms

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Common Ingress Policy Match Conditions” on page 363](#) and [“Complete \[edit policy-options\] Hierarchy” on page 363](#) instead of the statements being repeated.

- [edit policy-options policy-statement *policy-name* from prefix-list-filter *prefix-list-name* (exact | longer | orlonger)]
- [edit policy-options policy-statement *policy-name* from route-filter *ip-prefix*</*prefix-length*> (exact | longer | orlonger | through *ip-prefix*</*prefix-length*> | upto /*prefix-length*)]
- [edit policy-options policy-statement *policy-name* from source-address-filter *ip-prefix*</*prefix-length*> (exact | longer | orlonger | through *ip-prefix*</*prefix-length*> | upto /*prefix-length*)]
- [edit policy-options policy-statement *policy-name* term *term-name* from prefix-list-filter *prefix-list-name* (exact | longer | orlonger)]
- [edit policy-options policy-statement *policy-name* term *term-name* from route-filter *ip-prefix*</*prefix-length*> (exact | longer | orlonger | through *ip-prefix*</*prefix-length*> | upto /*prefix-length*)]
- [edit policy-options policy-statement *policy-name* term *term-name* from source-address-filter *ip-prefix*</*prefix-length*> (exact | longer | orlonger | through *ip-prefix*</*prefix-length*> | upto /*prefix-length*)]
- [edit policy-options policy-statement *policy-name* then]
- [edit policy-options policy-statement *policy-name* term *term-name* then]

The common policy terms are as follows:

```
(accept | reject);
aigp-originate distance;
as-path-expand (as-number | last-as) <count number>;
as-path-prepend as-number;
class class-name;
color (preference | add number | subtract number);
color2 (preference | add number | subtract number);
community (add | delete | set | + | - | =) community-name;
cos-next-hop-map map-name;
damping list-name;
default-action (accept | reject);
destination-class class-name;
```

```

external {
    type (1 | 2);
}
forwarding-class class-name;
install-nexthop <strict> (lsp [ lsp-names ] | lsp-regex [ regular-expressions ] |
    static-lsp [ lsp-names ] | static-lsp-regex [ regular-expressions ])
    <except (lsp [ lsp-names ] | lsp-regex [ regular-expressions ] | static-lsp [ lsp-names ] |
    static-lsp-regex [ regular-expressions ])>;
load-balance per-packet;
local-preference (preference | add number | subtract number);
metric (metric-value | add number | igp <metric-offset> | minimum-igp <metric-offset> |
    subtract number | ... the following complex expression ...);
expression {
    metric (multiplier number | offset number | multiplier number offset number);
    metric2 (multiplier number | offset number | multiplier number offset number);
}
metric2 (metric-value | add number | subtract number);
metric3 (metric-value | add number | subtract number);
metric4 (metric-value | add number | subtract number);
next (policy | term);
next-hop (ip-address | discard | next-table routing-table-name | peer-address | reject |
    self);
origin (egp | igp | incomplete);
preference (preference | add number | subtract number);
preference2 (preference | add number | subtract number);
priority (high | low | medium);
source-class class-name;
tag (tag-number | add number | subtract number);
tag2 (tag-number | add number | subtract number);
trace;

```

Common Policy Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit policy-options\] Hierarchy” on page 363](#) instead of the statements being repeated.

- [edit policy-options policy-statement *policy-name* from]
- [edit policy-options policy-statement *policy-name* term *term-name* from]
- [edit policy-options policy-statement *policy-name* term *term-name* to]
- [edit policy-options policy-statement *policy-name* to]

The common policy match conditions are as follows:

```

area area-id;
as-path [ regular-expression-names ];
as-path-group [ as-path-group-names ];
color preference;
color2 preference;
community [ community-names ];
external {
    type (1 | 2);
}
family family-name;

```



```

instance instance-name;
interface [ interface-names ];
level isis-level;
local-preference value;
metric metric-value;
metric2 metric-value;
metric3 metric-value;
metric4 metric-value;
neighbor [ ip-addresses ];
next-hop [ ip-addresses ];
nlri-route-type route-type-number;
origin (egp | igp | incomplete);
policy [ policy-names ];
preference preference;
preference2 preference;
protocol [ protocol-names ];
rib routing-table-name;
tag [ tag-numbers ];
tag2 tag-number;

```

Common Ingress Policy Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit policy-options\] Hierarchy” on page 363](#) instead of the statements being repeated at each level.

- [edit policy-options policy-statement *policy-name* from]
- [edit policy-options policy-statement *policy-name* term *term-name* from]

The common ingress policy match conditions are as follows:

```

aggregate-contributor;
condition [ conditions ];
multicast-scope (scope-value | global | link-local | node-local | organization-local |
  site-local) <orhigher | orlower>;
next-hop-type merged;
prefix-list prefix-list-name;
prefix-list-filter prefix-list-name (exact | longer | orlonger) {
  ... statements in Common Policy Terms on page 361 ...;
}
route-filter ip-prefix</prefix-length> (exact | longer | orlonger |
  through ip-prefix</prefix-length> | upto /prefix-length) {
  ... statements in Common Policy Terms on page 361 ...;
}
route-type (external | internal);
rtf-prefix-list name route-targets
source-address-filter ip-prefix</prefix-length> (exact | longer | orlonger |
  through ip-prefix</prefix-length> | upto /prefix-length) {
  ... statements in Common Policy Terms on page 361 ...;
}
state (active | inactive);

```

Complete [edit policy-options] Hierarchy

The statement hierarchy in this section can also be included at the [edit logical-systems *logical-system-name*] hierarchy level.

```
policy-options {
  as-path name regular-expression;
  as-path-group group-name {
    as-path name regular-expression;
  }
  community name {
    invert-match;
    members [ community-ids ];
  }
  condition condition-name {
    if-route-exists address table table-name;
    route-active-on (node0 | node1);
  }
  damping name {
    disable;
    half-life minutes;
    max-suppress minutes;
    reuse number;
    suppress number;
  }
  policy-statement policy-name {
    from {
      ... statements in Common Policy Match Conditions on page 362 AND
      ... statements in Common Ingress Policy Match Conditions on page 363 ...
    }
    term term-name {
      from {
        ... statements in Common Policy Match Conditions on page 362 AND
        ... statements in Common Ingress Policy Match Conditions on page 363 ...
      }
      to {
        ... statements in Common Policy Match Conditions on page 362 ...
      }
      then {
        ... statements in Common Policy Terms on page 361 ...
      }
    }
    to {
      ... statements in Common Policy Match Conditions on page 362 ...
    }
    then {
      ... statements in Common Policy Terms on page 361 ...
    }
  }
  prefix-list list-name {
    ip-prefix </prefix-length>;
    apply-path path;
  }
}
```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

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

Each of the following topics lists the statements at a subhierarchy of the [\[edit protocols\]](#) hierarchy.

- [\[edit protocols bfd\] Hierarchy Level](#)
- [\[edit protocols bgp\] Hierarchy Level on page 368](#)
- [\[edit protocols dot1x\] Hierarchy Level](#)
- [\[edit protocols dvmrp\] Hierarchy Level](#)
- [\[edit protocols igmp\] Hierarchy Level](#)
- [igmp-snooping](#)
- [\[edit protocols isis\] Hierarchy Level on page 378](#)
- [\[edit protocols l2-learning\] Hierarchy Level on page 380](#)
- [\[edit protocols lacp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols layer2-control\] Hierarchy Level on page 381](#)
- [\[edit protocols ldp\] Hierarchy Level on page 381](#)
- [\[edit protocols lldp\] Hierarchy Level](#)
- [\[edit protocols mld\] Hierarchy Level](#)
- [\[edit protocols mpls\] Hierarchy Level on page 386](#)
- [\[edit protocols msdp\] Hierarchy Level](#)
- [\[edit protocols mstp\] Hierarchy Level on page 386](#)
- [\[edit protocols mvrp\] Hierarchy Level](#)
- [\[edit protocols neighbor-discovery\] Hierarchy Level](#)
- [\[edit protocols oam\] Hierarchy Level](#)
- [\[edit protocols ospf\] Hierarchy Level on page 391](#)
- [\[edit protocols ospf3\] Hierarchy Level on page 395](#)
- [\[edit protocols pim\] Hierarchy Level on page 398](#)
- [\[edit protocols protection-group\] Hierarchy Level](#)
- [\[edit protocols rip\] Hierarchy Level on page 402](#)
- [\[edit protocols ripng\] Hierarchy Level on page 404](#)
- [\[edit protocols router-advertisement\] Hierarchy Level on page 405](#)
- [\[edit protocols router-discovery\] Hierarchy Level on page 406](#)
- [\[edit protocols rstp\] Hierarchy Level on page 406](#)
- [\[edit protocols sap\] Hierarchy Level on page 407](#)

- [\[edit protocols vrrp\] Hierarchy Level on page 407](#)
- [\[edit protocols vstp\] Hierarchy Level on page 408](#)

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

bfd

Syntax	<pre> bfd { traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],</p> <p>[edit protocols],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure trace options for Bidirectional Forwarding Protocol (BFD) traffic.
Default	If you do not include this statement, no BFD tracing operations are performed.
Options	<p>disable—(Optional) Disable the BFD tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the /var/log directory. We recommend that you place global routing protocol tracing output in the routing-log file.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. These are the BFD protocol tracing options:</p> <ul style="list-style-type: none"> • adjacency—Trace adjacency messages. • all—Trace all options for BFD. • error—Trace all errors. • event—Trace all events. • issu—Trace in-service software upgrade (ISSU) packet activity.

- **nsr-packet**—Trace non-stop-routing (NSR) packet activity.
- **nsr-synchronization**—Trace NSR synchronization events.
- **packet**—Trace all packets.
- **pipe**—Trace pipe messages.
- **pipe-detail**—Trace pipe messages in detail.
- **ppm-packet**—Trace packet activity by periodic packet management (PPM).
- **state**—Trace state transitions.
- **timer**—Trace timer processing.

match *regular-expression*—(Optional) Regular expression for lines to be logged.

no-world-readable—(Optional) Prevent any user from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the trace file again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level	routing and trace—To view this statement in the configuration.
	routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for Static Routes on page 3197

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

Several statements in the **[edit protocols mpls]** hierarchy are valid at numerous locations within it. To make the complete hierarchy easier to read, the repeated statements are listed in “[Common BGP Family Options](#)” on page 369 and that section is referenced at the appropriate locations in “[Complete \[edit protocols bgp\] Hierarchy](#)” on page 369.

- [Common BGP Family Options on page 369](#)
- [Complete \[edit protocols bgp\] Hierarchy on page 369](#)

Common BGP Family Options

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit protocols bgp\] Hierarchy](#)” on page 369 instead of the statements being repeated.

- [edit protocols bgp family inet (any | flow | labeled-unicast | multicast | unicast)]
- [edit protocols bgp family inet6 (any | labeled-unicast | multicast | unicast)]
- [edit protocols bgp family (inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) signaling]
- [edit protocols bgp family inet-vpn (any | flow | multicast | unicast)]
- [edit protocols bgp family inet6-vpn (any | multicast | unicast)]
- [edit protocols bgp family iso-vpn unicast]

The common BGP family options are as follows:

```
accepted-prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
}
damping;
loops number;
prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
}
rib-group group-name;
topology name {
    community {
        target identifier;
    }
}
```

Complete [edit protocols bgp] Hierarchy

The statement hierarchy listed in this section can also be included at the [edit logical-systems *logical-system-name*] hierarchy level.

```
protocols {
    bgp {
        disable;
        accept-remote-nexthop;
        advertise-external <conditional>;
        advertise-from-main-vpn-tables;
        advertise-inactive;
        (advertise-peer-as | no-advertise-peer-as);
        authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
        authentication-key key;
        authentication-key-chain key-chain;
        bfd-liveness-detection {
            authentication {
                algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
```

```

    key-chain key-chain-name;
    loose-check;
}
detection-time {
    threshold milliseconds;
}
holddown-interval milliseconds;
minimum-interval milliseconds;
minimum-receive-interval milliseconds;
multiplier number;
no-adaptation;
session-mode (automatic | multihop | single-hop);
transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
}
version (1 | automatic);
}
cluster cluster-identifier;
damping;
description text-description;
export [ policy-names ];
family family-name {
    ... the family subhierarchies appear after the main [edit protocols bgp] hierarchy ...
}
graceful-restart {
    disable;
    restart-time seconds;
    stale-routes-time seconds;
}
group group-name {
    ... the group subhierarchy appears after the main [edit protocols bgp] hierarchy ...
}
hold-time seconds;
idle-after-switch-over (seconds | forever);
import [ policy-names ];
include-mp-next-hop;
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <loops number> <alias> <private>;
local-interface interface-name;
local-preference local-preference;
log-updown;
metric-out (metric | igp (delay-med-update | offset) | minimum-igp offset);
mtu-discovery;
multihop {
    no-nexthop-change;
    ttl tvl-value;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
outbound-route-filter {
    bgp-orf-cisco-mode;
    prefix-based {

```



```

        accept {
            inet;
            inet6;
        }
    }
}
passive;
path-selection {
    always-compare-med;
    as-path-ignore;
    cisco-non-deterministic;
    external-router-id;
    med-plus-igp {
        igp-multiplier number;
        med-multiplier number;
    }
}
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
vpn-apply-export;
}

bgp {
    family inet {
        (any | multicast) {
            ... statements in Common BGP Family Options on page 369 ...
        }
        flow {
            ... statements in Common BGP Family Options on page 369 PLUS ...
            no-validate [ validation-procedure-names ];
        }
        labeled-unicast {
            ... statements in Common BGP Family Options on page 369 PLUS ...
            add-path {
                receive;
                send {
                    path-count number;
                    prefix-policy [ policy-names ];
                }
            }
            aggregate-label {
                community community-name;
            }
            aigp [disable];
            explicit-null connected-only;
            per-group-label;
            per-prefix-label;
            resolve-vpn;
            rib (inet.3 | inet6.3);

```

```

        traffic-statistics {
            file filename <files number> <size maximum-file-size> <world-readable |
                no-world-readable>;
            interval seconds;
        }
    }
    unicast {
        ... statements in Common BGP Family Options on page 369 PLUS ...
        add-path {
            receive;
            send {
                path-count number;
                prefix-policy [ policy-names ];
            }
        }
        topology name {
            community target identifier;
        }
    }
}

bgp {
    family inet6 {
        (any | multicast) {
            ... statements in Common BGP Family Options on page 369 ...
        }
        labeled-unicast {
            ... statements in Common BGP Family Options on page 369 PLUS ...
            add-path {
                receive;
                send {
                    path-count number;
                    prefix-policy [ policy-names ];
                }
            }
            aggregate-label {
                community community-name;
            }
            aigp [disable];
            explicit-null;
            per-group-label;
            traffic-statistics {
                file filename <files number> <size maximum-file-size> <world-readable |
                    no-world-readable>;
                interval seconds;
            }
        }
        unicast {
            ... statements in Common BGP Family Options on page 369 PLUS ...
            topology name {
                community target identifier;
            }
        }
    }
}

```

```

bgp {
  family (inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) {
    signaling {
      ... statements in Common BGP Family Options on page 369 ...
    }
  }
}

bgp {
  family inet-vpn {
    (any | multicast | unicast) {
      ... statements in Common BGP Family Options on page 369 PLUS ...
      aggregate-label <community community-name>;
    }
    flow {
      ... statements in Common BGP Family Options on page 369 ...
    }
  }
}

bgp {
  family inet6-vpn {
    (any | multicast | unicast) {
      ... statements in Common BGP Family Options on page 369 PLUS ...
      aggregate-label <community community-name>;
    }
  }
}

bgp {
  family iso-vpn {
    unicast {
      ... statements in Common BGP Family Options on page 369 PLUS ...
      aggregate-label <community community-name>;
    }
  }
}

bgp {
  family route-target {
    accepted-prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    proxy-generate <route-target-policy route-target-policy-name>;
  }
}

bgp {

```

```
group group-name {
  ... same statements as at the [edit protocols bgp] hierarchy level PLUS ...
  allow [ all ip-prefix</prefix-length> ];
  as-override;
  multipath <multiple-as>;
  neighbor address {
    ... the neighbor subhierarchy appears after the main [edit protocols bgp group
      group-name] hierarchy ...
  }
  type (external | internal);
  ... BUT NOT ...
  disable; # NOT valid at this level
  group group-name { ... } # NOT valid at this level
  path-selection { ... } # NOT valid at this level
}

group group-name {
  neighbor address {
    ... same statements as at the [edit protocols bgp] hierarchy level PLUS ...
    as-override;
    multipath <multiple-as>;
    ... BUT NOT ...
    disable; # NOT valid at this level
    group group-name { ... } # NOT valid at this level
    neighbor address { ... } # NOT valid at this level
    path-selection { ... } # NOT valid at this level
  }
}
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

dvmrp

Syntax	<pre> dvmrp { disable; export [<i>policy-names</i>]; import [<i>policy-names</i>]; interface <i>interface-name</i> { disable; hold-time <i>seconds</i>; metric <i>metric</i>; mode (forwarding unicast-routing); } rib-group <i>group-name</i>; traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Enable DVMRP on the router or switch.
Default	DVMRP is disabled on the router or switch.
Options	The statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring DVMRP</i>

igmp

Syntax `igmp {`
 `accounting;`
 `interface interface-name {`
 `disable;`
 `(accounting | no-accounting);`
 `group-limit limit;`
 `group-policy [policy-names];`
 `group-threshold`
 `immediate-leave;`
 `log-interval`
 `oif-map map-name;`
 `passive;`
 `promiscuous-mode;`
 `ssm-map ssm-map-name;`
 `ssm-map-policy ssm-map-policy-name;`
 `static {`
 `group multicast-group-address {`
 `exclude;`
 `group-count number;`
 `group-increment increment;`
 `source ip-address {`
 `source-count number;`
 `source-increment increment;`
 `}`
 `}`
 `}`
 `version version;`
 `}`
 `query-interval seconds;`
 `query-last-member-interval seconds;`
 `query-response-interval seconds;`
 `robust-count number;`
 `traceoptions {`
 `file filename <files number> <size size> <world-readable | no-world-readable>;`
 `flag flag <flag-modifier> <disable>;`
 `}`
 `}`

Hierarchy Level `[edit logical-systems logical-system-name protocols],`
 `[edit protocols]`

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 12.1 for the QFX Series.
 Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Enable IGMP on the router or switch. IGMP must be enabled for the router or switch to receive multicast packets.

The remaining statements are explained separately.

Default	IGMP is disabled on the router or switch. IGMP is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP).
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Enabling IGMP on page 3661

igmp-snooping

Syntax	<pre> igmp-snooping { vlan <i>vlan-id</i> { immediate-leave; interface <i>interface-name</i> { group-limit <i>limit</i>; host-only-interface; immediate-leave; multicast-router-interface; static { group <i>ip-address</i> { source <i>ip-address</i>; } } } } proxy { source-address <i>ip-address</i>; } query-interval <i>seconds</i>; query-last-member-interval <i>seconds</i>; query-response-interval <i>seconds</i>; robust-count <i>number</i>; } </pre>
Hierarchy Level	[edit protocols]
Release Information	Statement introduced in Junos OS Release 8.5.
Description	Enable IGMP snooping on the router or switch.
Default	IGMP snooping is disabled on the router or switch.
Options	The statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding IGMP Snooping</i> • <i>IGMP Snooping in MC-LAG Active-Active on MX Series Routers Overview</i>

[edit protocols isis] Hierarchy Level

The following statement hierarchy can also be included at the [edit protocols isis] hierarchy level.

```
protocols {
  isis {
    disable;
    clns-routing;
    context-identifier ip-address</prefix> {
      level (1 | 2) <disable>;
    }
    export [ policy-names ];
    graceful-restart {
      disable;
      helper-disable;
      restart-duration seconds;
    }
    ignore-attached-bit;
    interface interface-name {
      ... the interface subhierarchy appears after the main [edit protocols isis] hierarchy ...
    }
    label-switched-path name level level metric metric;
    level (1 | 2) {
      disable;
      authentication-key key;
      authentication-type authentication;
      external-preference preference;
      no-csnp-authentication;
      no-hello-authentication;
      no-psnp-authentication;
      preference preference;
      prefix-export-limit number;
      wide-metrics-only;
    }
    loose-authentication-check;
    lsp-lifetime seconds;
    max-areas number;
    no-adjacency-holddown;
    no-authentication-check;
    no-ipv4-routing;
    no-ipv6-routing;
    overload {
      advertise-high-metrics;
      timeout seconds;
    }
    reference-bandwidth reference-bandwidth;
    rib-group {
      inet group-name;
      inet6 group-name;
    }
    spf-options {
      delay milliseconds;
      holddown milliseconds;
      rapid-runs number;
    }
  }
}
```



```

}
topologies {
  ipv4-multicast;
  ipv6-multicast;
  ipv6-unicast;
}
traceoptions {
  file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
traffic-engineering {
  disable;
  family inet {
    shortcuts {
      multicast-rpf-routes;
    }
  }
  family inet6 {
    shortcuts;
  }
}
ignore-lsp-metrics;
}

isis {
  interface interface-name {
    disable;
    bfd-liveness-detection {
      authentication {
        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
          meticulous-keyed-sha-1 | simple-password);
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      version (1 | automatic);
    }
  }
  checksum;
  csnp-interval (seconds | disable);
  hello-padding (adaptive | loose | strict);
  ldp-synchronization {
    disable;
    hold-time seconds;
  }
  level (1 | 2) {

```

```
    disable;
    hello-authentication-key key;
    hello-authentication-type authentication;
    hello-interval seconds;
    hold-time seconds;
    ipv4-multicast-metric number;
    ipv6-multicast-metric number;
    ipv6-unicast-metric number;
    metric metric;
    passive;
    priority number;
    te-metric metric;
  }
  link-protection;
  lsp-interval milliseconds;
  mesh-group (value | blocked);
  no-adjacency-down-notification;
  no-eligible-backup;
  no-ipv4-multicast;
  no-ipv6-multicast;
  no-ipv6-unicast;
  no-unicast-topology;
  node-link-protection;
  passive;
  point-to-point;
}
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

[\[edit protocols l2-learning\] Hierarchy Level](#)

```
protocols {
  l2-learning {
    global-mac-limit {
      limit;
      packet-action drop;
    }
    global-mac-statistics;
    global-mac-table-aging-time seconds;
    global-no-mac-learning;
  }
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

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

The following statement hierarchy can also be included at the [\[edit logical-systems logical-system-name\]](#) hierarchy level.

```

protocols {
  lacp {
    traceoptions {
      file <filename> <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
      flag flag;
      no-remote-trace;
    }
    ppm (centralized | distributed)
  }
}

```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

[\[edit protocols layer2-control\] Hierarchy Level](#)

The following statement hierarchy can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```

protocols {
  layer2-control {
    bpdu-block {
      disable-timeout seconds;
      interface [ interface-names ];
    }
    mac-rewrite {
      interface interface-name {
        protocol {
          cdp;
          stp;
          vtp;
        }
      }
    }
    nonstop-bridging;
    traceoptions {
      file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
      flag flag <disable>;
    }
  }
}

```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

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

The following statement hierarchy can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```

protocols {

```

```
ldp {
  (deaggregate | no-deaggregate);
  dod-request-policy [ policy-names ];
  egress-policy [ policy-names ];
  explicit-null;
  export [ policy-names ];
  graceful-restart {
    disable;
    helper-disable;
    maximum-neighbor-reconnect-time seconds;
    maximum-neighbor-recovery-time seconds;
    reconnect-time seconds;
    recovery-time seconds;
  }
  igp-synchronization holddown-interval seconds;
  import [ policy-names ];
  interface interface-name {
    (allow-subnet-mismatch | no-allow-subnet-mismatch);
    disable;
    hello-interval seconds;
    hold-time seconds;
    transport-address (interface | router-id);
  }
  keepalive-interval seconds;
  keepalive-timeout seconds;
  l2-smart-policy;
  log-updown {
    trap disable;
  }
  next-hop {
    merged {
      policy [ policy-names ];
    }
  }
  no-forwarding;
  oam {
    ... the oam subhierarchy appears after the main [edit protocols ldp] hierarchy ...
  }
  policing {
    fec class-address {
      ingress-traffic filter-name;
      transit-traffic filter-name;
    }
  }
  preference preference;
  session destination-address {
    authentication-algorithm algorithm;
    authentication-key key;
    authentication-key-chain key-chain;
    downstream-on-demand;
  }
  session-protection <timeout seconds>;
  strict-targeted-hellos;
  targeted-hello {
    hello-interval seconds;
    hold-time seconds;
  }
}
```

```

}
traceoptions {
  file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
track-igp-metric;
traffic-statistics {
  file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
  interval seconds;
  no-penultimate-hop;
}
transport-address (address | interface | router-id);
}

ldp {
  oam {
    bfd-liveness-detection {
      detection-time {
        threshold milliseconds;
      }
    }
    ecmp;
    failure-action (remove-nexthop | remove-route);
    holddown-interval milliseconds;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    version (1 | automatic);
  }
  fec class-address {
    bfd-liveness-detection {
      ... same statements as at the [edit protocols ldp oam bfd-liveness-detection]
        hierarchy level ...
    }
    no-bfd-liveness-detection;
    periodic-traceroute {
      ... same statements as at the [edit protocols ldp oam periodic-traceroute]
        hierarchy level PLUS ...
      disable;
    }
  }
}
ingress-policy [ policy-names ];
periodic-traceroute {
  exp cos-value;
  fanout next-hops;
  frequency minutes;
  paths number;
  retries number;
  source address;
  ttl number;
}

```

```
        wait seconds;
      }
    }
  }
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

lldp

Syntax

```
lldp {
  advertisement-interval seconds;
  disable;
  hold-multiplier number;
  interface (all | interface-name) {
    disable;
  }
  lldp-configuration-notification-interval seconds;
  port-id-subtype {
    interface-name;
    locally-assigned;
  }
  ptopo-configuration-maximum-hold-time seconds;
  ptopo-configuration-trap-interval seconds;
  traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag <disable>;
  }
}
```

Hierarchy Level [edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information Statement introduced in Junos OS Release 9.6.

Description (MX Series and T Series routers and EX Series switches only) Specify LLDP configuration parameters.

Options The statements are explained separately.

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

- Related Documentation**
- [Configuring LLDP on page 3246](#)

mld

```

Syntax  mld {
        accounting;
        interface interface-name {
            (accounting | no-accounting);
            disable;
            group-limit limit;
            group-policy [ policy-names ];
            immediate-leave;
            oif-map [ map-names ];
            passive;
            ssm-map ssm-map-name;
            ssm-map-policy ssm-map-policy-name;
            static {
                group multicast-group-address {
                    exclude;
                    group-count number;
                    group-increment increment;
                    source ip-address {
                        source-count number;
                        source-increment increment;
                    }
                }
            }
            version version;
        }
        maximum-transmit-rate packets-per-second;
        query-interval seconds;
        query-last-member-interval seconds;
        query-response-interval seconds;
        robust-count number;
        traceoptions {
            file filename <files number> <size size> <world-readable | no-world-readable>;
            flag flag <flag-modifier> <disable>;
        }
    }

```

Hierarchy Level [edit logical-systems *logical-system-name* protocols],
[edit protocols]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable MLD on the routing device. MLD must be enabled for the routing device to receive multicast packets.

Default MLD is disabled on the routing device. MLD is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP).

Options The statements are explained separately.

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

Related Documentation

- *Enabling MLD*

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

- [Complete \[edit protocols mpls\] Hierarchy on page 386](#)

Complete [edit protocols mpls] Hierarchy

The statement hierarchy listed in this section can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```
protocols {
  mpls {
    disable;
    interface (interface-name | all) {
      always-mark-connection-protection-tlv
      disable;
      admin-group [ group-names ];
      srlg srlg-name
      static {
        protection-revert-time seconds
      }
      switch-away-lsps;
    }
    ipv6-tunneling;
    priority setup-priority hold-priority;
    traceoptions {
      file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
      flag flag;
    }
  }
}
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

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

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```
protocols {
  mstp {
    disable;
    backup-bridge-priority priority;
    bpdu-block-on-edge;
    bpdu-destination-mac-address provider-bridge-group;
    bridge-priority priority;
    configuration-name configuration-name;
```



```

forward-delay seconds;
hello-time seconds;
interface interface-name {
    bpdu-timeout-action {
        alarm;
        block;
    }
    cost cost;
    edge;
    mode (point-to-point | shared);
    no-root-port;
    priority interface-priority;
}
max-age seconds;
max-hops hops;
msti identifier {
    backup-bridge-priority priority;
    bridge-priority priority;
    interface interface-name {
        cost cost;
        priority interface-priority;
    }
    vlan [ vlan-ids ];
}
priority-hold-time seconds;
revision-level revision-level;
system-id mac-address {
    ip-address ip-address</prefix-length>;
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
    flag flag <disable>;
}
vpls-flush-on-topology-change;
}
}

```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

neighbor-discovery

Syntax neighbor-discovery {
 no-dad-on-state-change ;
 onlink-subnet-only;
 secure {
 security-level {
 (default | secure-messages-only);
 }
 cryptographic-address {
 key-length *number*;
 key-pair *pathname*;
 }
 timestamp {
 clock-drift *number*;
 known-peer-window *number*;
 new-peer-window *number*;
 }
 traceoptions {
 file *filename* <files *number*> <match *regular-expression*> <size *size*> <world-readable |
 no-world-readable>;
 flag *flag*;
 no-remote-trace;
 }
 }
 }

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 9.3.
 Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Enable Secure Neighbor Discovery.

 The remaining statements are explained separately.

Default Disabled

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

Related Documentation • *Example: Configuring Secure IPv6 Neighbor Discovery*

oam

```

Syntax  oam {
        ethernet{
            connectivity-fault-management {
                action-profile profile-name {
                    action {
                        interface-down;
                    }
                    default-actions {
                        interface-down;
                    }
                    event {
                        adjacency-loss;
                    }
                }
            }
            linktrace {
                age (30m | 10m | 1m | 30s | 10s);
                path-database-size path-database-size;
            }
            maintenance-domain domain-name {
                level number;
                mip-half-function (none | default | explicit);
                name-format (character-string | none | dns | mac+2oct);
                maintenance-association ma-name {
                    continuity-check {
                        hold-interval minutes;
                        interface-status-tlv;
                        interval (10m | 10s | 1m | 1s | 100ms);
                        loss-threshold number;
                        port-status-tlv;
                    }
                    mep mep-id {
                        auto-discovery;
                        direction down;
                        interface interface-name;
                        remote-mep mep-id {
                            action-profile profile-name;
                        }
                    }
                }
            }
        }
        performance-monitoring {
            sla-iterator-profiles {
                profile-name {
                    calculation-weight {
                        delay delay-value;
                        delay-variation delay-variation-value;
                    }
                    cycle-time cycle-time-value;
                    iteration-period iteration-period-value;
                    measurement-type two-way-delay;
                    passive;
                }
            }
        }
    }

```

```

    }
  }
  traceoptions {
    file filename <files number> <match regex> <size size> <world-readable |
      no-world-readable>;
    flag flag ;
    no-remote-trace;
  }
}
link-fault-management {
  action-profile profile-name;
  action {
    syslog;
    link-down;
    send-critical-event
  }
  event {
    link-adjacency-loss;
    link-event-rate {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
  }
}
interface interface-name {
  link-discovery (active | passive);
  pdu-interval interval;
  pdu-threshold threshold-value;
  remote-loopback;
  event-thresholds {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
  negotiation-options {
    allow-remote-loopback;
    no-allow-link-events;
  }
}
traceoptions {
  file filename <files number> <match regex> <size size> <world-readable |
    no-world-readable>;
  flag flag ;
  no-remote-trace;
}
}
}

```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 9.4 for EX Series switches.
connectivity-fault-management introduced in Junos OS Release 10.2 for EX Series switches.

Description	Provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support for Ethernet interfaces on EX Series switches or configure connectivity fault management (CFM) for IEEE 802.1ag Operation, Administration, and Management (OAM) support on the switches.
	The remaining statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches</i> • <i>Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches</i> • <i>Configuring Ethernet OAM Link Fault Management (CLI Procedure)</i> • <i>Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)</i>

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

The following statement hierarchy can also be included at the [\[edit logical-systems logical-system-name\]](#) hierarchy level.

```

protocols {
  ospf {
    disable;
    area area-id {
      ... the area subhierarchy appears after the main [edit protocols ospf] hierarchy ...
    }
    backup-spf-options {
      disable;
      downstream-paths-only;
      no-install;
    }
    database-protection {
      ignore-count number;
      ignore-time seconds;
      maximum-lsa number;
      reset-time seconds;
      warning-only;
      warning-threshold percent;
    }
    export [ policy-names ];
    external-preference preference;
    graceful-restart {
      disable;
      helper-disable <both | restart-signaling | standard>;
      no-strict-lsa-checking;
      notify-duration seconds;
      restart-duration seconds;
    }
    import [ policy-names ];
    lsa-refresh-interval;
    no-nssa-abr;
    no-rfc-1583;
  }
}

```

```

overload <timeout seconds>;
preference preference;
prefix-export-limit number;
reference-bandwidth reference-bandwidth;
rib-group group-name;
spf-options {
    delay milliseconds;
    holddown milliseconds;
    rapid-runs number;
}
topology (default | ipv4-multicast | name) {
    backup-spf-options {
        disable;
        downstream-paths-only;
        no-install;
    }
    overload;
    prefix-export-limit number;
    spf-options {
        delay milliseconds;
        holddown milliseconds;
        rapid-runs number;
    }
    topology-id number;
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
traffic-engineering {
    advertise-unnumbered-interfaces;
    credibility-protocol-preference;
    ignore-lsp-metrics;
    multicast-rpf-routes;
    no-topology;
    shortcuts <lsp-metric-into-summary>;
}
}

ospf {
    area area-id {
        area-range ip-prefix</prefix-length> <exact> <override-metric metric> <restrict>;
        context-identifier identifier
        interface interface-name {
            ... the interface subhierarchy appears after the main [edit ospf area area-id] hierarchy
            level ...
        }
        label-switched-path name {
            disable;
            metric metric;
            topology (name | default | ipv4-multicast) {
                disable;
                metric metric;
            }
        }
    }
}

```

```

network-summary-export [ policy-names ];
network-summary-import [ policy-names ];
nssa {
  area-range ip-prefix </prefix-length> <exact> <override-metric metric> <restrict>;
  default-lsa {
    default-metric metric;
    metric-type type;
    type-7;
  }
  (summaries | no-summaries);
}
peer-interface interface-name {
  disable;
  authentication {
    md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
    simple-password key-string;
  }
  dead-interval seconds;
  demand-circuit;
  flood-reduction;
  hello-interval seconds;
  no-neighbor-down-notification;
  retransmit-interval seconds;
  transit-delay seconds;
}
stub <default-metric metric> <summaries | no-summaries>;
virtual-link neighbor-id router-id transit-area area-id {
  disable;
  authentication {
    md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
    simple-password key-string;
  }
  dead-interval seconds;
  demand-circuit;
  flood-reduction;
  hello-interval seconds;
  ipsec-sa sa-name;
  no-neighbor-down-notification;
  retransmit-interval seconds;
  topology (name | default | ipv4-multicast) {
    disable;
    metric metric;
  }
  transit-delay seconds;
}
}

area area-id {
  interface interface-name {
    disable;
    authentication {
      md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
      simple-password key-string;
    }
    bandwidth-based-metrics {
      bandwidth value metric number;
    }
  }
}

```

```
}
bfd-liveness-detection {
  authentication {
    algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
      meticulous-keyed-sha-1 | simple-password);
    key-chain key-chain-name;
    loose-check;
  }
  detection-time {
    threshold milliseconds;
  }
  full-neighbors-only;
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
  version (1 | automatic);
}
dead-interval seconds;
demand-circuit;
dynamic-neighbors;
flood-reduction;
hello-interval seconds;
interface-type (nbma | p2mp | p2p);
ipsec-sa sa-name;
ldp-synchronization {
  disable;
  hold-time seconds;
}
(link-protection | node-link-protection);
metric metric;
neighbor address <eligible>;
no-eligible-backup;
no-interface-state-traps;
no-neighbor-down-notification;
passive {
  traffic-engineering {
    remote-node-id address;
  }
}
poll-interval seconds;
priority number;
retransmit-interval seconds;
secondary;
te-metric metric;
topology (name | default | ipv4-multicast) {
  disable;
  bandwidth-based-metrics {
    bandwidth value;
    metric number;
  }
  metric metric;
}
```



```

    }
    transit-delay seconds;
  }
}
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

[edit protocols ospf3] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```

protocols {
  ospf3 {
    disable;
    area area-id {
      ... the area subhierarchy appears after the main [edit protocols ospf3] hierarchy ...
    }
    backup-spf-options {
      disable;
      downstream-paths-only;
      no-install;
    }
    database-protection {
      ignore-count number;
      ignore-time seconds;
      maximum-lsa number;
      reset-time seconds;
      warning-only;
      warning-threshold percent;
    }
    export [ policy-names ];
    external-preference preference;
    graceful-restart {
      disable;
      helper-disable;
      no-strict-lsa-checking;
      notify-duration seconds;
      restart-duration seconds;
    }
    import [ policy-names ];
    lsa-refresh-interval;
    no-nssa-abr;
    no-rfc-1583;
    overload <timeout seconds>;
    preference preference;
    prefix-export-limit number;
    realm (ipv4-multicast | ipv4-unicast | ipv6-multicast | ipv6-unicast) {
      ... the realm subhierarchies appear after the main [edit protocols ospf3] hierarchy ...
    }
    reference-bandwidth reference-bandwidth;
  }
}

```

```

rib-group group-name;
spf-options {
    delay milliseconds;
    holddown milliseconds;
    no-ignore-our-externals;
    rapid-runs number;
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
traffic-engineering {
    ignore-lsp-metrics;
    shortcuts <lsp-metric-into-summary>;
}
}

ospf3 {
    area area-id {
        area-range ip-prefix </prefix-length> <exact> <override-metric metric> <restrict>;
        inter-area-prefix-export [ policy-names ];
        inter-area-prefix-import [ policy-names ];
        interface interface-name {
            ... the interface subhierarchy appears after the main [edit ospf3 area area-id]
            hierarchy level ...
        }
        nssa {
            area-range ip-prefix </prefix-length> <exact> <override-metric metric> <restrict>;
            default-lsa {
                default-metric metric;
                metric-type type;
                type-7;
            }
            (summaries | no-summaries);
        }
        stub <default-metric metric> <summaries | no-summaries>;
        virtual-link neighbor-id router-id transit-area area-id {
            disable;
            dead-interval seconds;
            demand-circuit;
            flood-reduction;
            hello-interval seconds;
            ipsec-sa sa-name;
            retransmit-interval seconds;
            transit-delay seconds;
        }
    }
}

area area-id {
    interface interface-name {
        disable;
        bandwidth-based-metrics {
            bandwidth value metric number;
        }
        bfd-liveness-detection {

```

```

    authentication {
        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
            meticulous-keyed-sha-1 | simple-password);
        key-chain key-chain-name;
        loose-check;
    }
    detection-time {
        threshold milliseconds;
    }
    full-neighbors-only;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
dead-interval seconds;
demand-circuit;
flood-reduction;
hello-interval seconds;
interface-type (p2mp-over-lan | p2p);
ipsec-sa sa-name;
(link-protection | node-link-protection);
metric metric;
no-eligible-backup;
own-router-lsa;
passive {
    traffic-engineering {
        remote-node-id address;
    }
}
priority number;
retransmit-interval seconds;
transit-delay seconds;
}
}
}

ospf3 {
    realm (ipv4-multicast| ipv6-multicast) {
        ... same statements as at the [edit protocols ospf3] hierarchy level, EXCEPT FOR ...
        area area-id {
            interface interface-name {
                no-eligible-backup;  # NOT valid at this level
            }
            virtual-link { ... }  # NOT valid at this level
        }
        backup-spf-options { ... }  # NOT valid at this level
        realm realm-identifier { ... }  # NOT valid at this level
        traffic-engineering { ... }  # NOT valid at this level
    }
}
}

```

```
ospf3 {
  realm ipv4-unicast {
    ... same statements as at the [edit protocols ospf3] hierarchy level, PLUS ...
    area area-id {
      interface interface-name {
        ldp-synchronization {
          disable;
          hold-time seconds;
        }
      }
    }
  }

  ... BUT NOT ...
  area area-id {
    virtual-link { ... } # NOT valid at this level
  }
  realm realm-identifier { ... } # NOT valid at this level
  traffic-engineering { ... } # NOT valid at this level
}
}
```

```
ospf3 {
  realm ipv6-unicast {
    disable;
    backup-spf-options {
      disable;
      downstream-paths-only;
      no-install;
    }
  }
}
}
```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

[edit protocols pim] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```
protocols {
  pim {
    disable;
    assert-timeout seconds;
    default-vpn-source {
      interface-name interface-name;
    }
    dense-groups {
      address <announce | reject>;
    }
    dr-election-on-p2p;
```

```

export [ policy-names ];
family (inet | inet6) {
    disable;
}
graceful-restart {
    disable;
    no-bidirectional-mode;
    restart-duration seconds;
}
import [ policy-names ];
interface interface-name {
    ... the interface subhierarchy appears after the main [edit protocols pim] hierarchy ...
    family (inet | inet6) {
        disable;
    }
}
join-load-balance;
join-prune-timeout seconds;
nonstop-routing {
    disable;
}
override-interval milliseconds;
propagation-delay milliseconds;
reset-tracking-bit;
rib-group {
    inet group-name;
    inet6 group-name;
}
rp {
    ... the rp subhierarchy appears after the main [edit protocols pim] hierarchy ...
}
sglimit {
    family (inet | inet6) {
        log-interval seconds;
        maximum limit;
        threshold value;
    }
    log-interval seconds;
    maximum limit;
    threshold value;
}
spt-threshold {
    infinity [ policy-names ];
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
    flag (route | state) <flag-modifier> <disable> <filter <match-on prefix>
        <policy [ policy-names ]>>;
}
}

pim {

```

```

interface interface-name {
  accept-remote-source;
  disable;
  bfd-liveness-detection {
    authentication {
      algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
        meticulous-keyed-sha-1 | simple-password);
      key-chain key-chain-name;
      loose-check;
    }
    detection-time {
      threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    version (1 | automatic);
  }
  bidirectional {
    df-election {
      backoff-period milliseconds;
      offer-period milliseconds;
      robustness-count number;
    }
  }
  family (inet | inet6) {
    disable;
  }
  hello-interval seconds;
  bidirectional-sparse | bidirectional-sparse-dense mode (bidirectional-sparse |
    bidirectional-sparse-dense | dense | sparse | sparse-dense);
  neighbor-policy [ policy-names ];
  override-interval milliseconds;
  priority number;
  propagation-delay milliseconds;
  reset-tracking-bit;
  version (1 | 2);
}
}

pim {
  rp {
    auto-rp {
      (announce | discovery | mapping);
      (mapping-agent-election | no-mapping-agent-election);
    }
    bidirectional {
      address address {
        group-ranges {
          destination-ip-prefix </prefix-length>;
        }
      }
    }
  }
}

```

```

        hold-time seconds;
        priority number;
    }
}
bootstrap {
    family (inet | inet6) {
        export [ policy-names ];
        import [ policy-names ];
        priority number;
    }
}
bootstrap-export [ policy-names ];
bootstrap-import [ policy-names ];
bootstrap-priority number;
dr-register-policy [ policy-names ];
embedded-rp {
    group-ranges {
        ip-prefix</prefix-length>;
    }
    maximum-rps limit;
}
group-rp-mapping {
    family (inet | inet6) {
        log-interval seconds;
        maximum limit;
        threshold value;
    }
}
log-interval seconds;
maximum limit;
threshold value;
}
}
local {
    ... the local subhierarchy appears after the main [edit protocols pim rp] hierarchy ...
}
register-limit {
    family (inet | inet6) {
        log-interval seconds;
        maximum limit;
        threshold value;
    }
}
log-interval seconds;
maximum limit;
threshold value;
}
rp-register-policy [ policy-names ];
static {
    address address {
        group-ranges {
            ip-prefix</prefix-length>;
        }
        override;
        version (1 | 2);
    }
}

```

```
    }
  }
}

rp {
  local {
    disable;
    address address;
    family (inet | inet6) {
      disable;
      address address;
      anycast-pim {
        local-address address;
        rp-set {
          address address <forward-msdp-sa>;
        }
      }
    }
    group-ranges {
      ip-prefix </prefix-length>;
    }
    hold-time seconds;
    override;
    priority number;
  }
  group-ranges {
    ip-prefix </prefix-length>;
  }
  hold-time seconds;
  override;
  priority number;
}
}
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

[edit protocols rip] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```
protocols {
  rip {
    authentication-key password;
    authentication-type type;
    (check-zero | no-check-zero);
    graceful-restart {
      disable;
      restart-time seconds;
    }
  }
  group group-name {
    ... the group subhierarchy appears after the main [edit protocols rip] hierarchy ...
  }
}
```



```

}
holddown seconds;
import [ policy-names ];
message-size number;
metric-in metric;
receive (both | none | version-1 | version-2);
rib-group group-name;
route-timeout seconds;
send (broadcast | multicast | none | version-1);
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
update-interval seconds;
}

rip {
    group group-name {
        bfd-liveness-detection {
            authentication {
                algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (1 | automatic);
        }
        demand-circuit;
        export [ policy-names ];
        import [ policy-names ];
        max-retrans-time seconds;
        metric-out metric;
        neighbor interface-name {
            ... the neighbor subhierarchy appears after the main [edit protocols rip group
                group-name] hierarchy level ...
        }
        preference preference;
        route-timeout seconds;
        update-interval seconds;
    }

    group group-name {
        neighbor neighbor-name {
            any-sender;

```

```
authentication-key password;
authentication-type type;
bfd-liveness-detection {
    ... same statements as at the [edit protocols rip group group-name
        bfd-liveness-detection] hierarchy level ...
}
(check-zero | no-check-zero);
demand-circuit;
import [ policy-names ];
max-retrans-time seconds;
message-size number;
metric-in metric;
receive (both | none | version-1 | version-2);
route-timeout seconds;
send (broadcast | multicast | none | version-1);
update-interval seconds;
}
}
}
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

[edit protocols ripng] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```
protocols {
  ripng {
    graceful-restart {
      disable;
      restart-time seconds;
    }
    group group-name {
      export [ policy-names ];
      import [ policy-names ];
      metric-out metric;
      neighbor neighbor-name {
        import [ policy-names ];
        metric-in metric;
        receive <none>;
        route-timeout seconds;
        send <none>;
        update-interval seconds;
      }
      preference number;
      route-timeout seconds;
      update-interval seconds;
    }
    holddown seconds;
    import [ policy-names ];
    metric-in metric;
```

```

receive <none>;
route-timeout seconds;
send <none>;
update-interval seconds;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
}

```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

[\[edit protocols router-advertisement\] Hierarchy Level](#)

The following statement hierarchy can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```

protocols {
    router-advertisement {
        interface interface-name {
            current-hop-limit number;
            default-lifetime seconds;
            (link-mtu | no-link-mtu);
            (managed-configuration | no-managed-configuration);
            max-advertisement-interval seconds;
            min-advertisement-interval seconds;
            (other-stateful-configuration | no-other-stateful-configuration);
            prefix prefix {
                (autonomous | no-autonomous);
                (on-link | no-on-link);
                preferred-lifetime seconds;
                valid-lifetime seconds;
            }
            reachable-time milliseconds;
            retransmit-timer milliseconds;
            virtual-router-only;
        }
        traceoptions {
            file filename <files number> <size maximum-file-size> <world-readable |
            no-world-readable>;
            flag flag;
        }
    }
}

```

- Related Documentation**
- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*
 - *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

[\[edit protocols router-discovery\] Hierarchy Level](#)

The following statement hierarchy can also be included at the [\[edit logical-systems logical-system-name\]](#) hierarchy level.

```
protocols {
  router-discovery {
    disable;
    address address {
      (advertise | ignore);
      (broadcast | multicast);
      (ineligible | priority number);
    }
    interface interface-name {
      lifetime seconds;
      max-advertisement-interval seconds;
      min-advertisement-interval seconds;
    }
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier> <disable>;
    }
  }
}
```

Related Documentation

- [Notational Conventions Used in Junos OS Configuration Hierarchies](#)
- [\[edit protocols\] Hierarchy Level](#)

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

The following statement hierarchy can also be included at the [\[edit logical-systems logical-system-name\]](#) hierarchy level.

```
protocols {
  rstp {
    disable;
    backup-bridge-priority priority;
    bpdu-block-on-edge;
    bpdu-destination-mac-address provider-bridge-group;
    bridge-priority priority;
    extended-system-id id;
    force-version stp;
    forward-delay seconds;
    hello-time seconds;
    interface interface-name {
      bpdu-timeout-action {
        alarm;
        block;
      }
    }
    cost cost;
    edge;
    mode (point-to-point | shared);
    no-root-port;
    priority interface-priority;
```

```

    }
    max-age seconds;
    priority-hold-time seconds;
    system-id mac-address {
        ip-address ip-address </prefix-length>;
    }
    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
        flag flag <disable>;
    }
    vpls-flush-on-topology-change;
}

```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

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

The following statement hierarchy can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```

protocols {
    sap {
        disable;
        listen address <port port>;
    }
}

```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

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

The following statement hierarchy can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```

protocols {
    vrrp {
        asymmetric-hold-time;
        delegate-processing;
        failover-delay milliseconds;
        global-advertisements-threshold advertisement-value;
        skew-timer-disable;
        startup-silent-period seconds;
        traceoptions {
            file <filename> <files number> <match regular-expression> <microsecond-stamp>
            <size maximum-file-size> <world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
        version-3;
    }
}

```

```
}  
}
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*
- *Junos OS Hierarchy and RFC Reference*
- *Junos® OS Ethernet Interfaces*
- *Junos® OS Network Interfaces*

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

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```
protocols {  
  vstp {  
    disable;  
    bpdu-block-on-edge;  
    force-version stp;  
    interface interface-name {  
      access-trunk  
      bpdu-timeout-action {  
        alarm;  
        block;  
      }  
      cost cost;  
      edge;  
      mode (point-to-point | shared);  
      no-root-port;  
      priority interface-priority;  
    }  
    priority-hold-time seconds;  
    system-id mac-address {  
      ip-address ip-address </prefix-length>;  
    }  
    vlan vlan-id {  
      ... the vlan subhierarchy appears after the main [edit protocols vstp] hierarchy level ...  
    }  
    vpls-flush-on-topology-change;  
  }  
  
  vstp {  
    vlan vlan-id {  
      backup-bridge-priority priority;  
      bridge-priority priority;  
      forward-delay seconds;  
      hello-time seconds;  
      interface interface-name {  
        ... same statements as at the [edit protocols vstp interface interface-name] hierarchy level ...  
      }  
      max-age seconds;  
    }  
  }  
}
```

```

    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
        flag flag <disable>;
    }
}
}
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

Layer 2 Routing Instances Configuration Hierarchy

Use the **vpls** routing instance type for point-to-multipoint LAN implementations between a set of sites in a VPN.

To configure routing instances for Layer 2 networks, include the following statements:

```

routing-instances {
    routing-instance-name {
        access {
            address-assignment {
                ... same statements as in the address-assignment subhierarchy in [edit access]
                Hierarchy Level ...
            }
            address-protection;
            description text;
            egress-protection {
                context-identifier context-id;
            }
            forwarding-options {
                ...forwarding-options...
            }
            instance-role role;
            instance-type type;
            interface interface-name;
            l2-domain-id-for-l3 id;
            l2vpn-id community;
            layer3-domain-identifier identifier;
            multicast-snooping-options {
                ... same statements as in [edit multicast-snooping-options] Hierarchy Level EXCEPT
                FOR ...
                traceoptions {...} # NOT valid at this level
            }
            no-irb-layer-2-copy;
            no-local-switching;
            no-vrf-advertise;
            no-vrf-propagate-ttl;
            pbb-options {
                default-bvlan bvlan;
                peer-instance instance;
                vlan-id vlan-id isid-list [ isid-numbers ]
            }
        }
    }
}

```

```
protocols {
  ... the protocols subhierarchy appears after the main [edit routing-instances
    routing-instance-name] hierarchy ...
}
provider-tunnel {
  ... the provider-tunnel subhierarchy appears after the main [edit routing-instances
    routing-instance-name] hierarchy ...
}
route-distinguisher (as-number:number | ip-address:number);
routing-interface interface;
routing-options {
  ... the routing-options subhierarchy appears after the main [edit routing-instances
    routing-instance-name] hierarchy ...
}
service-groups {
  service-group-name {
    pbb-service-options {
      default-isid isid-number;
      isid isid-number vlan-id-list [ vlan-ids ];
      mac-address mac-address;
    }
    service-type type;
  }
}
services {
  mobile-ip {
    ... same statements as in [edit services mobile-ip] Hierarchy Level ...
  }
}
switch-options {
  ... same statements as in [edit switch-options] Hierarchy Level ...
}
vlan-id (id | all | none);
vlan-model one-to-one;
vlan-tags outer <tpid.>vlan-id inner <tpid.>vlan-id;
[edit vlans] Hierarchy Level on page 445 {
  ... same statements as in [edit vlans] Hierarchy Level ...
}
vrf-advertise-selective {
  family {
    inet-mvpn;
    inet6-mvpn;
  }
}
vrf-export [ policy-names ];
vrf-import [ policy-names ];
vrf-propagate-ttl;
vrf-table-label;
vrf-target {
  export community-name;
  import community-name;
}
protocols {
  ... protocols-configuration ...
}
routing-options {
```



```

...routing-options-configuration ...
}
bridge-domains {
  bridge-domain-name {
    domain-type bridge;
    interface interface-name;
    routing-interface routing-interface-name;
    vlan-id (Bridge Domain or VLAN) (none | all | number);
    vlan-tags outer number inner number;
    bridge-options {
      interface-mac-limit limit {
        packet-action drop;
      }
      interface interface-name {
        interface-mac-limit limit {
          packet-action drop;
        }
      }
      mac-statistics;
      mac-table-size limit {
        packet-action drop;
      }
      no-mac-learning;
      static-mac mac-address;
    }
  }
}
}
}

```

With the exception of the **instance-type virtual-switch** statement (which configures a virtual-switch routing instance), you can include the statements at the following hierarchy levels:

- **[edit]**
- **[edit logical-systems *logical-system-name*]**

The **instance-type virtual-switch** statement is not supported at the **[edit logical-systems *logical-system-name*]** hierarchy level.

Related Documentation

- [Routing Instances Overview](#)
- [Layer 2 Routing Instance Types](#)
- [Configuring a Layer 2 Virtual Switch on page 1847](#)
- [Configuring a Layer 2 Control Protocol Routing Instance](#)

[edit routing-options] Hierarchy Level

Several statements in the **[edit routing-options]** hierarchy are valid at numerous locations within the hierarchy. To make the complete hierarchy easier to read, the repeated statements are listed in “[Common Routing Options](#)” on page 412 and that section is

referenced at the appropriate locations in “[Complete \[edit routing-options\] Hierarchy](#)” on page 413.

- [Common Routing Options](#) on page 412
- [Complete \[edit routing-options\] Hierarchy](#) on page 413

Common Routing Options

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit routing-options\] Hierarchy](#)” on page 413 instead of the statements being repeated.

- [\[edit routing-options aggregate defaults\]](#)
- [\[edit routing-options aggregate route *ip-prefix* </prefix-length>\]](#)
- [\[edit routing-options generate defaults\]](#)
- [\[edit routing-options generate route *ip-prefix* </prefix-length>\]](#)
- [\[edit routing-options static defaults\]](#)
- [\[edit routing-options static route *ip-prefix* </prefix-length>\]](#)

The common routing options are as follows:

```
(active | passive);
as-path {
    aggregator as-number address;
    atomic-aggregate;
    origin (egp | igp | incomplete);
    path path-identifier;
}
color metric <type metric-type>;
color2 metric <type metric-type>;
community [ community-id no-advertise no-export no-export-subconfed ];
metric metric <type metric-type>;
metric2 metric <type metric-type>;
metric3 metric <type metric-type>;
metric4 metric <type metric-type>;
passive;
preference preference-value <type metric-type>;
preference2 preference-value <type metric-type>;
tag metric <type metric-type>;
tag2 metric <type metric-type>;
```

Complete [edit routing-options] Hierarchy

The statement hierarchy in this section can also be included at the [edit logical-systems *logical-system-name*] hierarchy level.

```

routing-options {
  access {
    route ip-prefix</prefix-length> {
      metric metric;
      next-hop [ addresses ];
      preference preference-value;
      qualified-next-hop address;
      tag route-tag;
    }
  }
  access-internal {
    route ip-prefix</prefix-length> {
      next-hop [ addresses ];
      qualified-next-hop address;
      tag route-tag;
    }
  }
  admin-groups-extended group-name {
    group-value group-identifier;
  }
  admin-groups-extended-range {
    maximum maximum-number;
    minimum minimum-number;
  }
  aggregate {
    defaults {
      ... statements in Common Routing Options on page 412 PLUS ...
      (brief | full);
      discard;
    }
    route ip-prefix</prefix-length> {
      ... statements in Common Routing Options on page 412 PLUS ...
      (brief | full);
      discard;
      policy [ policy-names ];
    }
  }
  auto-export {
    disable;
    family inet {
      disable;
      flow {
        disable;
        rib-group rib-group;
      }
      multicast {
        disable;
        rib-group rib-group;
      }
      unicast {
        disable;
      }
    }
  }
}

```

```
        rib-group rib-group;
    }
}
family inet6 {
    disable;
    multicast {
        disable;
        rib-group rib-group;
    }
    unicast {
        disable;
        rib-group rib-group;
    }
}
family iso {
    disable;
    unicast {
        disable;
        rib-group rib-group;
    }
}
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
autonomous-system autonomous-system <asdot-notation> <loops number>;
no-bfd-triggered-local-repair;
bgp-orf-cisco-mode;
bmp {
    memory-limit bytes;
    station-address (ip-address | name);
    station-port-number port-number;
    statistics-timeout seconds;
}
confederation as-number members [ as-numbers ];
dynamic-tunnels tunnel-name {
    destination-networks prefix;
    gre;
    rsvp-te entry-name {
        destination-networks network-prefix;
        label-switched-path-template {
            default-template;
            template-name;
        }
    }
}
source-address address;
}
fate-sharing {
    group group-name {
        cost value;
        from {
            address <to address>;
        }
    }
}
```

```

}
flow {
  firewall-install-disable;
  route name {
    match {
      destination address;
      destination-port [ afs bgp biff bootpc bootps cmd cvspserver dhcp domain eklogin
        ekshell exec finger ftp ftp-data http https ident imap kerberos-sec klogin kpasswd
        krb-prop krbupdate kshell ldap ldp login mobileip-agent mobilip-mn msdp
        netbios-dgm netbios-ns netbios-ssn nfsd nntp ntalk ntp pop3 pptp printer radacct
        radius rip rkinit smtp snmp snmptrap snpp socks ssh sunrpc syslog tacacs
        tacacs-ds talk telnet tftp timed who xdmcp ];
      dscp [ code-points ];
      fragment [ don't-fragment first-fragment is-fragment last-fragment
        not-a-fragment ];
      icmp-code [ communication-prohibited-by-filtering destination-host-prohibited
        destination-host-unknown fragmentation-needed host-precedence-violation
        host-unreachable host-unreachable-for-tos ip-header-bad network-unreachable
        network-unreachable-for-tos port-unreachable precedence-cutoff-in-effect
        protocol-unreachable redirect-for-host redirect-for-network
        redirect-for-tos-and-host redirect-for-tos-and-net required-option-missing
        source-host-isolated source-route-failed ttl-eq-zero-during-reassembly
        ttl-eq-zero-during-transit ];
      icmp-type [ echo-reply echo-request info-reply info-request mask-reply
        mask-request parameter-problem redirect router-advertisement router-solicit
        source-quench time-exceeded timestamp timestamp-reply unreachable ];
      packet-length [ values ];
      port [ ... same values as for the preceding destination-port statement ... ];
      protocol [ ah esp gre icmp igmp ipip ospf pim rsvp sctp tcp udp ];
      source address;
      source-port [ ... same values as for the preceding destination-port statement ... ];
      tcp-flags [ ack fin push rst syn urgent ];
    }
    then {
      (accept | discard);
      community community-name;
      next-term;
      rate-limit value;
      routing-instance routing-instance-name;
      sample;
    }
  }
}
term order (legacy | standard);
validation {
  traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
}
forwarding-table {
  chained-composite-next-hop {
    ingress {
      l3vpn {
        extended-space;

```

```

    }
  }
}
export [ policy-name ];
indexed-next-hop;
(indirect-next-hop | no-indirect-next-hop);
(indirect-next-hop-change-acknowledgements |
  no-indirect-next-hop-change-acknowledgements;)
krt-nexthop-ack-timeout interval;
unicast-reverse-path (active-paths | feasible-paths);
}
generate {
  defaults {
    ... statements in Common Routing Options on page 412 PLUS ...
    (brief | full);
    discard;
  }
  route ip-prefix</prefix-length> {
    ... statements in Common Routing Options on page 412 PLUS ...
    (brief | full);
    discard;
    policy [ policy-names ];
  }
}
graceful-restart {
  disable;
  restart-duration seconds;
}
host-fast-reroute {
  global-arp-prefix-limit number;
  global-supplementary-blackout-timer minutes;
}
instance-export [ policy-names ];
instance-import [ policy-names ];
interface interface-name { # In the routing-instance only
  arp-prefix-limit number;
  link-protection;
  supplementary-blackout-timer minutes;
}
interface-routes {
  family (inet | inet6) {
    export {
      lan;
      point-to-point;
    }
    import [ policy-names ];
  }
  rib-group {
    inet group-name;
    inet6 group-name;
  }
}
martians {
  ip-prefix</prefix-length> (exact | longer | orlonger |
    prefix-length-range /minimum-prefix-length–/maximum-prefix-length |
    through ip-prefix</prefix-length> | upto /prefix-length> <allow>;

```

```

}
maximum-paths path-limit <log-only | threshold value> <log-interval seconds>;
maximum-prefixes prefix-limit <log-only | threshold value> <log-interval seconds>;
med-igp-update-interval minutes;
multicast {
  ... the multicast subhierarchy appears after the main [edit routing-options] hierarchy ...
}
nonstop-routing;
options {
  mark seconds;
  syslog {
    level level;
    upto level;
  }
}
ppm {
  no-delegate-processing;
}
resolution {
  rib routing-table-name {
    import [ policy-names ];
    resolution-ribs [ routing-table-names ];
  }
  tracefilter [ filter-policy-names ];
  traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
rib routing-table-name {
  access {
    ... same statements as at the [edit routing-options access] hierarchy level ...
  }
  access-internal {
    ... same statements as at the [edit routing-options access-internal] hierarchy level ...
  }
  aggregate {
    ... same statements as at the [edit routing-options aggregate] hierarchy level ...
  }
  generate {
    ... same statements as at the [edit routing-options generate] hierarchy level ...
  }
  martians {
    ip-prefix</prefix-length> (exact | longer | orlonger |
      prefix-length-range /minimum-prefix-length–/maximum-prefix-length |
      through ip-prefix</prefix-length> | upto /prefix-length) <allow>;
  }
  maximum-paths path-limit <log-only | threshold value> <log-interval seconds>;
  maximum-prefixes prefix-limit <log-only | threshold value> <log-interval seconds>;
  static {
    ... same statements as at the [edit routing-options static] hierarchy level ...
  }
}
rib-groups {
  group-name {

```

```
        export-rib table-name;
        import-policy [ policy-names ];
        import-rib [ table-names ];
    }
}
route-distinguisher-id address;
route-record;
router-id address;
source-routing {
    ip;
    ipv6;
}
srlg {
    srlg-name {
        srlg-cost srlg-cost;
        srlg-value srlg-value;
    }
}
static {
    ... the static subhierarchy appears after the main [edit routing-options] hierarchy ...
}
topologies {
    family (inet | inet6) {
        topology topology-name;
    }
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <disable>;
}
validation {
    group group-name {
        max-sessions number;
        session address {
            hold-time seconds;
            local-address local-ip-address;
            port port-number;
            preference number;
            record-lifetime seconds;
            refresh-time seconds;
        }
    }
}
notification-rib value;
static {
    record destination {
        maximum-length prefix-length {
            origin-autonomous-system as-number {
                validation-state (invalid | valid);
            }
        }
    }
}
}
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
```



```

    }
  }
}

routing-options {
  multicast {
    asm-override-ssm;
    backup-pe-group group-name {
      backups [ addresses ];
      local-address address;
    }
    flow-map flow-map-name {
      bandwidth <bps> <adaptive>;
      forwarding-cache {
        timeout (never <non-discard-entry-only> | minutes);
      }
      policy [ policy-names ];
      redundant-sources [ addresses ];
    }
    forwarding-cache {
      family (inet | inet6) {
        threshold {
          log-warning value;
          suppress value <reuse value>;
        }
        threshold {
          log-warning value;
          suppress value <reuse value>;
        }
        timeout minutes;
      }
    }
    interface interface-name {
      maximum-bandwidth bps;
      no-qos-adjust;
      reverse-oif-mapping {
        no-qos-adjust;
      }
      subscriber-leave-timer seconds;
    }
  }
  pim-to-igmp-proxy {
    upstream-interface [ interface-names ];
  }
  pim-to-mld-proxy {
    upstream-interface [ interface-names ];
  }
  rpf-check-policy [ policy-names ];
  scope scope-name {
    interface [ interface-names ];
    prefix ip-prefix </prefix-length>;
  }
  scope-policy [ policy-names ];
  ssm-groups [ ip-prefix </prefix-length> ];
  ssm-map ssm-map-name {
    policy [ policy-names ];
    source [ addresses ];
  }
}

```

```

    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable |
            no-world-readable>;
        flag flag <disable>;
    }
}

routing-options {
    static {
        defaults {
            ... statements in Common Routing Options on page 412 PLUS ...
            (install | no-install);
            (readvertise | no-readvertise);
            (resolve | no-resolve);
            (retain | no-retain);
        }
        rib-group group-name;
        route destination-prefix {
            ... statements in Common Routing Options on page 412 PLUS ...
            backup-pe-group group-name;
            bfd-liveness-detection {
                detection-time {
                    threshold milliseconds;
                }
                holddown-interval milliseconds;
                local-address ip-address;
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                minimum-receive-ttl milliseconds;
                multiplier number;
                neighbor address;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                version (1 | automatic);
            }
            (discard | next-hop [ addresses ] | next-table address | receive | reject);
            (install | no-install);
            lsp-next-hop {
                metric metric;
                preference preference;
            }
            p2mp-lsp-next-hop lsp-name {
                metric metric;
                preference preference;
            }
            (readvertise | no-readvertise);
            (resolve | no-resolve);
            (retain | no-retain);
            static-lsp-next-hop lsp-name {
                metric metric;
                preference preference-value;
            }
        }
    }
}

```

```

    }
  }
}

```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

[edit security] Hierarchy Level

Each of the following topics lists the statements at a subhierarchy of the **[edit security]** hierarchy.

- [\[edit security alarms\] Hierarchy Level on page 421](#)
- [\[edit security authentication-key-chains\] Hierarchy Level on page 421](#)
- [\[edit security certificates\] Hierarchy Level on page 422](#)
- [\[edit security ike\] Hierarchy Level on page 422](#)
- [\[edit security ipsec\] Hierarchy Level on page 423](#)
- [\[edit security log\] Hierarchy Level on page 424](#)
- [\[edit security pki\] Hierarchy Level on page 424](#)
- [\[edit security ssh-known-hosts\] Hierarchy Level on page 425](#)
- [\[edit security traceoptions\] Hierarchy Level on page 425](#)

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

[edit security alarms] Hierarchy Level

```

security {
  alarms {
    audible;
    potential-violation {
      policy;
      replay-attacks;
    }
  }
}

```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- [\[edit security\] Hierarchy Level on page 421](#)

[edit security authentication-key-chains] Hierarchy Level

```

security {
  authentication-key-chains {
    key-chain key-chain-name {
      key key {
        secret secret-data;
        start-time yyyy-mm-dd.hh:mm:ss;
      }
    }
  }
}

```

```
}
```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- [\[edit security\] Hierarchy Level on page 421](#)

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

```
security {  
  certificates {  
    cache-size bytes;  
    cache-timeout-negative seconds;  
    certification-authority ca-profile-name {  
      ca-name ca-identity;  
      crl file-name;  
      encoding (binary | pem);  
      enrollment-url url-name;  
      file certificate-filename;  
      ldap-url url-name;  
    }  
    enrollment-retry attempts;  
    local certificate-filename {  
      certificate-key-string;  
      load-key-file URL key-filename;  
    }  
    maximum-certificates number;  
    path-length certificate-path-length;  
  }  
}
```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- [\[edit security\] Hierarchy Level on page 421](#)

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

```
security {  
  ike {  
    policy ike-peer-address {  
      description description;  
      encoding (binary | pem);  
      identity identity-name;  
      local-certificate certificate-filename;  
      local-key-pair private-public-key-file;  
      mode (aggressive | main);  
      pre-shared-key (ascii-text key | hexadecimal key);  
      proposals [ proposal-names ];  
    }  
    proposal ike-proposal-name {  
      authentication-algorithm (md5 | sha-256 | sha1);  
      authentication-method (dsa-signatures | pre-shared-keys | rsa-signatures);  
      description description;  
      dh-group (group1 | group2 | group5);  
      encryption-algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);  
      lifetime-seconds seconds;  
    }  
  }  
}
```

```

    }
  }
}

```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- [\[edit security\] Hierarchy Level on page 421](#)

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

```

security {
  ipsec {
    policy ipsec-policy-name {
      description text-description;
      perfect-forward-secrecy {
        keys (group1 | group14 | group2 | group5);
      }
      proposals [ proposal-names ];
    }
    proposal ipsec-proposal-name {
      authentication-algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
      description text-description;
      encryption-algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
      lifetime-seconds seconds;
      protocol (ah | bundle | esp);
    }
    security-association sa-name {
      description text-description;
      dynamic {
        ipsec-policy policy-name;
        replay-window-size (32 | 64);
      }
      manual {
        direction (bidirectional | inbound | outbound) {
          authentication {
            algorithm (hmac-md5-96 | hmac-sha1-96);
            key (ascii-text key | hexadecimal key);
          }
          auxiliary-spi spi-index;
          encryption {
            algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
            key (ascii-text key | hexadecimal key);
          }
          protocol (ah | bundle | esp);
          spi spi-index;
          encryption {
            algorithm 3des-cbc;
            key ascii-text ascii-text-string;
          }
        }
      }
      mode (transport | tunnel);
    }
  }
}

```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - [\[edit security\] Hierarchy Level on page 421](#)

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

```
security {  
  log {  
    cache {  
      exclude exclude-name{  
        destination-address;  
        destination-port;  
        event-id;  
        failure;  
        interface-name;  
        policy-name;  
        process;  
        protocol;  
        source-address;  
        source-port;  
        success;  
        username;  
      }  
      limit limit;  
    }  
  }  
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - [\[edit security\] Hierarchy Level on page 421](#)

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

```
security {  
  pki {  
    auto-re-enrollment {  
      certificate-id certificate-id {  
        ca-profile ca-profile-name;  
        challenge-password password;  
        re-enroll-trigger-time-percentage percentage;  
        re-generate-keypair;  
        validity-period days;  
      }  
    }  
    ca-profile ca-profile-name {  
      administrator {  
        email-address email-address;  
      }  
      ca-identity ca-identifier;  
      enrollment {  
        retry attempts;  
        retry-interval seconds;  
        url url;  
      }  
      revocation-check {  
        disable;  
      }  
    }  
  }  
}
```

```

    crt {
        disable on-download-failure;
        refresh-interval hours;
        url url {
            password password;
        }
    }
}
routing-instance;
}
traceoptions {
    file <filename> <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
    flag flag;
}
}
}

```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - [\[edit security\] Hierarchy Level on page 421](#)

[\[edit security ssh-known-hosts\] Hierarchy Level](#)

```

security {
    ssh-known-hosts {
        fetch-from-server server-name;
        host hostname {
            dsa-key dsa-key;
            ecdsa-sha2-nistp256-key ecdsa-sha2-nistp256-key;
            ecdsa-sha2-nistp384-key ecdsa-sha2-nistp384-key;
            ecdsa-sha2-nistp521-key ecdsa-sha2-nistp521-key;
            rsa-key rsa-key;
            rsa1-key rsa1-key;
        }
        load-key-file key-file;
    }
}

```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - [\[edit security\] Hierarchy Level on page 421](#)

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

```

security {
    traceoptions {
        file filename <files number> <size size>;
        flag all;
        flag database;
        flag general;
        flag ike;
        flag parse;
        flag policy-manager;
        flag routing-socket;
    }
}

```

```
        flag timer;
    }
}
```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- [\[edit security\] Hierarchy Level on page 421](#)

[edit snmp] Hierarchy Level

```
snmp {
  client-list list-name {
    address {
      restrict;
    }
  }
  community community-name {
    authorization (read-only | read-write);
    client-list-name list-name;
    clients {
      address <restrict>;
    }
    logical-system logical-system-name {
      routing-instance instance-name;
    }
    routing-instance instance-name {
      client-list-name list-name;
      clients {
        address <restrict>;
      }
    }
  }
  view view-name;
}
contact contact-information;
description description;
engine-id {
  (local engine-id | use-default-ip-address | use-mac-address);
}
filter-duplicates;
filter-interfaces {
  interfaces
  all-internal-interfaces;
  interface 1;
  interface 2;
}
health-monitor {
  falling-threshold percentage;
  idp {
    falling-threshold;
    interval seconds;
    rising-threshold;
  }
  interval seconds;
  rising-threshold percentage;
}
```



```

interface [ interface-names ];
location location;
logical-system-trap-filter;
name system-name;
nonvolatile {
    commit-delay seconds;
}
rmon {
    alarm index {
        description description;
        falling-event-index index;
        falling-threshold integer;
        falling-threshold-interval seconds;
        interval seconds;
        request-type (get-next-request | get-request | walk-request);
        rising-event-index index;
        rising-threshold integer;
        sample-type (absolute-value | delta-value);
        startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
        syslog-subtag text-string;
        variable oid-variable;
    }
    event index {
        community community-name;
        description description;
        type (log | log-and-trap | none | snmptrap);
    }
    history index {
        bucket-size number;
        interface interface-name;
        interval seconds;
        owner owner-name;
    }
}
routing-instance-access {
    access-list {
        routing-instance-name <restrict>;
    }
}
traceoptions {
    file <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
trap-group group-name {
    categories {
        authentication;
        chassis;
        chassis-cluster;
        configuration;
        link;
        otn-alarms {
            alarm-name;
        }
    }
    remote-operations;
}

```

```
    rmon-alarm;
    routing;
    services;
    sonet-alarms {
        alarm-name;
    }
    startup;
    vrrp-events;
}
destination-port port-number;
logical-system logical-system-name {
    routing-instance instance-name;
}
routing-instance instance-name;
targets {
    address;
}
version (all | v1 | v2);
}
trap-options {
    agent-address outgoing-interface;
    enterprise-oid;
    logical-system logical-system-name {
        routing-instance instance-name;
    }
    routing-instance instance-name {
        source-address (address | lo0);
    }
    source-address address;
}
v3 {
    ... the v3 subhierarchy appears after the main [edit snmp] hierarchy level ...
}
view view-name {
    oid object-identifier <exclude | include>;
}
}

snmp {
    v3 {
        notify name {
            tag tag-name;
            type (inform | trap);
        }
        notify-filter profile-name {
            oid oid <exclude | include>;
        }
        snmp-community community-index {
            community-name community-name;
            context context-name;
            security-name security-name;
            tag tag-name;
        }
        target-address target-address-name {
            address address;
            address-mask address-mask;
        }
    }
}
```

```

logical-system logical-system-name {
    routing-instance routing-instance-name;
}
port port-number;
retry-count number;
routing-instance routing-instance-name;
tag-list tag-list;
target-parameters parameter-name;
timeout seconds;
}
target-parameters parameter-name {
    notify-filter profile-name;
    parameters {
        message-processing-model (v1 | v2c | v3);
        security-level (authentication | none | privacy);
        security-model (usm | v1 | v2c);
        security-name security-name;
    }
}
usm {
    local-engine {
        user username {
            authentication-md5 {
                authentication-password password;
            }
            authentication-none;
            authentication-sha {
                authentication-password password;
            }
            privacy-3des {
                privacy-password password;
            }
            privacy-aes128 {
                privacy-password password;
            }
            privacy-des {
                privacy-password password;
            }
            privacy-none;
        }
    }
    remote-engine engine-id {
        user username {
            authentication-md5 {
                authentication-password password;
            }
            authentication-none;
            authentication-sha {
                authentication-password password;
            }
            privacy-3des {
                privacy-password password;
            }
            privacy-aes128 {
                privacy-password password;
            }
        }
    }
}

```

```
        privacy-des {
            privacy-password password;
        }
        privacy-none;
    }
}
vacm {
    access {
        group group-name {
            context-prefix prefix {
                security-model (any | usm | v1 | v2c) {
                    security-level (authentication | none | privacy) {
                        context-match (exact | prefix);
                        notify-view view-name;
                        read-view view-name;
                        write-view view-name;
                    }
                }
            }
        }
    }
    default-context-prefix prefix {
        security-model (any | usm | v1 | v2c) {
            security-level (authentication | none | privacy) {
                context-match (exact | prefix);
                notify-view view-name;
                read-view view-name;
                write-view view-name;
            }
        }
    }
}
security-to-group {
    security-model (usm | v1 | v2c) {
        security-name security-name {
            group group-name;
        }
    }
}
```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *Understanding the Integrated Local Management Interface*

[edit switch-options] Hierarchy Level

```
switch-options {
    interface interface-name {
        interface-mac-limit {
            number-of-addresses;
            packet-action drop;
        }
    }
}
```

```

    no-mac-learning;
  }
  interface-mac-limit {
    number-of-addresses;
    packet-action drop;
  }
  mac-statistics;
  mac-table-size {
    number-of-addresses;
    packet-action drop;
  }
  no-mac-learning;
  service-id number;
}

```

[edit system] Hierarchy Level

```

system {
  accounting {
    destination {
      radius {
        server {
          server-address {
            accounting-port port-number;
            max-outstanding-requests
            port port-number;
            retry number;
            secret password;
            source-address address;
            timeout seconds;
          }
        }
      }
    }
  }
  tacplus {
    server {
      server-address {
        port port-number;
        secret password;
        single-connection;
        source-address address;
        timeout seconds;
      }
    }
  }
  events [ change-log interactive-commands login ];
}
allow-6pe-traceroute;
allow-v4mapped-packets;
archival {
  configuration {
    archive-sites {
      ftp://<username>:<password>@<host>:<port>/<url-path>;
      scp://<username>:<password>@<host>:<port>/<url-path>;
    }
  }
  transfer-interval interval;
}

```

```
        transfer-on-commit;
    }
}
arp {
    aging-timer minutes;
    gratuitous-arp-delay;
    gratuitous-arp-on-ifup;
    interfaces {
        logical-interface-name {
            aging-timer minutes;
        }
    }
    passive-learning;
    purging;
}
authentication-order [ authentication-methods ];
auto-configuration {
    traceoptions {
        file <filename> <files number> <match regular-expression> <size size>
            <world-readable | no-world-readable>;
        flag <all | auth | configuration | ;interfaces | io | rtsock | ui>
        level level;
        no-remote-trace;
    }
}
backup-router address <destination [ destination-addresses ]>;
commit {
    fast-synchronize;
    synchronize;
    server {
        commit-interval number;
        days-to-keep-error-logs number;
        maximum-aggregate-pool number;
        maximum-entries number;
        traceoptions {
            file <filename> <files number> <match regular-expression> <size size>
                <world-readable | no-world-readable>;
            flag <all | auth | configuration | ;interfaces | io | rtsock | ui>
            level level;
            no-remote-trace;
        }
    }
}
(compress-configuration-files | no-compress-configuration-files);
ddos-protection {
    global {
        disable-fpc;
        disable-logging;
        disable-routing-engine;
        flow-detection;
        flow-report-rate;
        violation-report-rate;
    }
    protocols protocol-group (aggregate | packet-type) {
        bandwidth packets-per-second;
        burst size;
```

```

disable-fpc;
disable-logging;
disable-routing-engine;
fpc {
    bandwidth-scale percentage;
    burst-scale percentage;
    disable-fpc;
}
priority level;
recover-time seconds;
flow-detection {
    flow-detect-time detect-period;
    no-flow-logging;
    timeout-active-flows enable-period;
    flow-level-bandwidth;
    flow-level-control (all | keep-all | police);
    flow-detection-mode (always-on | automatic | disabled);
    physical-interface;
    flow-recover-time recover-period;
    flow-timeout-time timeout-period;
    subscriber;
}
}
traceoptions{
    file filename <files number> <match regular-expression > <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
}
default-address-selection;
diag-port-authentication (encrypted-password "password" | plain-text-password);
dynamic-profile-options {
    versioning;
}
domain-name domain-name;
domain-search [ domain-list ];
donot-disable-ip6op-ondad;
extensions {
    providers {
        provider-id {
            license-type license deployment-scope [ deployments ];
        }
    }
}
resource-limits {
    package package-name {
        resources {
            cpu {
                priority number;
                time seconds;
            }
            file {
                core-size bytes;
                open number;
                size bytes;
            }
        }
    }
}

```

```
    }
    memory {
      data-size bytes;
      locked-in bytes;
      resident-set-size bytes;
      socket-buffers bytes;
      stack-size bytes;
    }
  }
}
process process-ui-name {
  resources {
    cpu {
      priority number;
      time seconds;
    }
    file {
      core-size bytes;
      open number;
      size bytes;
    }
    memory {
      data-size bytes;
      locked-in bytes;
      resident-set-size bytes;
      socket-buffers bytes;
      stack-size bytes;
    }
  }
}
}
}
fips {
  level level;
}
host-name hostname;
inet6-backup-router ipv6-address <destination address>;
internet-options {
  (gre-path-mtu-discovery | no-gre-path-mtu-discovery);
  icmpv4-rate-limit bucket-size number packet-rate rate;
  icmpv6-rate-limit bucket-size number packet-rate rate;
  (ipip-path-mtu-discovery | no-ipip-path-mtu-discovery);
  (ipv6-path-mtu-discovery | noipv6-path-mtu-discovery);
  ipv6-path-mtu-discovery-timeout;
  no-tcp-rfc1323-paws;
  no-tcp-rfc1323;
  (path-mtu-discovery | no-path-mtu-discovery);
  source-port upper-limit port-number;
  (source-quench | no-source-quench);
  tcp-drop-synfin-set;
}
kernel-replication;
license {
  autoupdate {
    url URL;
    password password;
  }
}
```



```

    }
    renew before-expiration number;
    interval number
    traceoptions {
        file <filename> <files number> <size maximum-file-size> <world-readable |
            no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
location {
    altitude feet;
    building name;
    country-code code;
    floor number;
    hcoord horizontal-coordinate;
    lata service-area;
    latitude degrees;
    longitude degrees;
    npa-nxx number;
    postal-code postal-code;
    rack number;
    vcoord vertical-coordinate;
}
login {
    announcement "text";
    class class-name {
        access-end "hh<:mm:<ss>>";
        access-start "hh<:mm:<ss>>";
        allow-commands "regular-expression";
        ( allow-configuration | allow-configuration-regexps ) "regular expression 1" "regular
            expression 2";
        allowed-days [ sunday monday tuesday wednesday thursday friday saturday ];
        configuration-breadcrumbs;
        deny-commands "regular-expression";
        ( deny-configuration | deny-configuration-regexps ) "regular expression 1" "regular
            expression 2";
        idle-timeout minutes;
        logical-system logical-system-name;
        login-alarms;
        login-script filename;
        login-tip;
        permissions [ permissions ];
        security-role [ security-role ];
    }
    deny-sources (address address | apply-groups | apply-groups-except) ;
    message "text";
    password {
        change-type (character-sets | set-transitions);
        format (des | md5 | sha1);
        maximum-length length;
        minimum-changes number;
        minimum-length length;
        minimum-lower-cases number;
        minimum-numeric number;
        minimum-punctuations number;
    }
}

```

```
    minimum-upper-cases number;
  }
  retry-options {
    backoff-factor number;
    backoff-threshold number;
    maximum-time number;
    minimum-time number;
    tries-before-disconnect number;
  }
  user username {
    authentication {
      (encrypted-password "password" | plain-text-password);
      load-key-file filename;
      ssh-dsa "public-key" <from hostname>;
      ssh-ecdsa "public-key" <from hostname>;
      ssh-rsa "public-key" <from hostname>;
    }
    class class-name;
    full-name "complete-name";
    uid uid-value;
  }
}
max-configurations-on-flash number;
mirror-flash-on-disk;
name-server {
  address;
}
nd-maxmcast-solicit
nd-retransmit-timer
no-multicast-echo;
no-neighbor-learn;;
no-ping-record-route;
no-ping-time-stamp;
no-redirects;
no-redirects-ipv6;
ntp {
  authentication-key key-number type md5 value password;
  boot-server address;
  broadcast <address> <key key-number> <tvl value> <version value>;
  broadcast-client;
  multicast-client <address>;
  peer address <key key-number> <prefer> <version value>;
  server address <key key-number> <prefer> <version value>;
  source-address source-address;
  trusted-key [ key-numbers ];
}
pic-console-authentication {
  (encrypted-password "encrypted-password" | plain-text-password);
}
ports {
  auxiliary {
    disable;
    insecure;
    type (ansi | small-xterm | vt100 | xterm);
    port-type (mini-usb | rj45);
  }
}
```

```

    }
    console {
        disable;
        insecure;
        log-out-on-disconnect;
        type (ansi | small-xterm | vt100 | xterm);
    }
}
processes {
    process-name (enable | disable) failover (alternate-media | other-routing-engine);
    command path;
    timeout seconds;
}
proxy {
    password password;
    port port-number;
    server (hostname | ip-address);
    username username;
}
radius-options {
    attributes {
        nas-ip-address address;
    }
    password-protocol mschap-v2;
}
radius-server {
    server-address {
        accounting-port port-number;
        max-outstanding-requests number;
        port port-number;
        retry number;
        secret password;
        source-address source-address;
        timeout seconds;
    }
}
root-authentication {
    (encrypted-password "password" | plain-text-password);
    load-key-file filename;
    ssh-dsa "public-key" <from hostname>;
    ssh-ecdsa "public-key" <from hostname>;
    ssh-rsa "public-key" <from hostname>;
}
(saved-core-context | no-saved-core-context);
saved-core-files number;
scripts {
    load-scripts-from-flash;
    commit {
        allow-transients;
        direct-access;
        file filename.xml {
            checksum (md5 | sha-256 | sha1) hash;
            optional;
            refresh;
            refresh-from url;
            source url;
        }
    }
}

```

```
    }
    max-datasize
    refresh;
    refresh-from url;
    traceoptions {
        file <filename> <files number> <size maximum-file-size> <world-readable |
            no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
op {
    file filename.xml {
        arguments {
            argument-name {
                description descriptive-text;
            }
        }
        checksum (md5 | sha-256 | sha1) hash;
        command filename-alias;
        description descriptive-text;
        refresh;
        refresh-from url;
        source url;
    }
    max-datasize
    no-allow-url
    refresh;
    refresh-from url;
    traceoptions {
        file <filename> <files number> <size maximum-file-size> <world-readable |
            no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
}
static-host-mapping {
    hostname {
        alias [ aliases ];
        inet [ addresses ];
        inet6 [ addresses ];
        sysid system-identifier;
    }
}
syslog {
    allow-duplicates;
    archive <binary-data | no-binary-data> <files number> <size size> <world-readable |
        no-world-readable>;
    console {
        any | authorization | change-log | conflict-log | daemon | dfc | external | firewall | ftp
            | interactive-commands | kernel | ntp | pfe | security | user) (alert | any | critical |
            emergency | error | info | none | notice | warning);
    }
    file filename {
        facility severity;
```

```

allow-duplicates;
any (alert | any | critical | emergency | error | info | none | notice | warning);
archive <archive-sites {ftp-url <password password>}> <files number> <size size>
  <start-time "YYYY-MM-DD.hh:mm"> <transfer-interval minutes> <world-readable |
  no-world-readable>;
authorization (alert | any | critical | emergency | error | info | none | notice | warning);
change-log (alert | any | critical | emergency | error | info | none | notice | warning);
conflict-log (alert | any | critical | emergency | error | info | none | notice | warning);
daemon (alert | any | critical | emergency | error | info | none | notice | warning);
dfc (alert | any | critical | emergency | error | info | none | notice | warning);
explicit-priority;
external (alert | any | critical | emergency | error | info | none | notice | warning);
firewall (alert | any | critical | emergency | error | info | none | notice | warning);
ftp (alert | any | critical | emergency | error | info | none | notice | warning);
interactive-commands (alert | any | critical | emergency | error | info | none | notice
  | warning);
kernel (alert | any | critical | emergency | error | info | none | notice | warning);
match "regular-expression";
ntp (alert | any | critical | emergency | error | info | none | notice | warning);
pfe (alert | any | critical | emergency | error | info | none | notice | warning);
security (alert | any | critical | emergency | error | info | none | notice | warning);
structured-data {
  brief
}
host (hostname | other-routing-engine | scc-master) {
  facility severity;
  authorization (alert | any | critical | emergency | error | info | none | notice | warning);
  change-log (alert | any | critical | emergency | error | info | none | notice | warning);
  conflict-log (alert | any | critical | emergency | error | info | none | notice | warning);
  daemon (alert | any | critical | emergency | error | info | none | notice | warning);
  dfc (alert | any | critical | emergency | error | info | none | notice | warning);
  explicit-priority;
  external (alert | any | critical | emergency | error | info | none | notice | warning);
  facility-override facility;
  firewall (alert | any | critical | emergency | error | info | none | notice | warning);
  ftp (alert | any | critical | emergency | error | info | none | notice | warning);
  interactive-commands (alert | any | critical | emergency | error | info | none | notice
    | warning);
  kernel (alert | any | critical | emergency | error | info | none | notice | warning);
  log-prefix string;
  match "regular-expression";
  ntp (alert | any | critical | emergency | error | info | none | notice | warning);
  pfe (alert | any | critical | emergency | error | info | none | notice | warning);
  security (alert | any | critical | emergency | error | info | none | notice | warning);
  source-address source-address;
  structured-data {
    brief
    user (username | *) {
  }
  log-rotate-frequency minutes;
  server;
  source-address address;
  time-format (year | millisecond | year millisecond);
  user (username | *) {
    facility severity;
    match "regular-expression";

```

```

    }
  }
  tacplus-options {
    (exclude-cmd-attribute | no-cmd-attribute-value);
    service-name service-name;
  }
  tacplus-server {
    server-address {
      port port-number;
      secret password;
      single-connection;
      source-address source-address;
      timeout seconds;
    }
  }
  time-zone (GMT | GMT+hour-offset | GMT-hour-offset | zone-name);
  tracing destination-override syslog host address;
  use-imported-time-zones;
}
}
system {
  services {
    database-replication {
      traceoptions {
        file <filename> <files number> <match regular-expression>
          <size maximum-file-size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
      }
    }
  }
  dhcp-local-server {
    authentication {
      password password;
      username-include {
        circuit-type;
        delimiter delimiter-character;
        domain-name domain-name;
        logical-system-name;
        mac-address;
        option-60;
        option-82 <circuit-id> <remote-id>;
        routing-instance-name;
        user-prefix user-prefix;
      }
    }
    duplicate-clients-on-interface;
    dynamic-profile (profile-name | junos-default-profile) <aggregate-clients <merge |
      replace> | use-primary primary-profile-name>;
    forward-snooped-clients (all-interfaces | configured-interfaces |
      non-configured-interfaces);
    group group-name {
      dynamic-profile (profile-name | junos-default-profile) <aggregate-clients <merge |
        replace> | use-primary primary-profile-name>;
      interface interface-name {
        exclude;
        overrides {

```

```

...same statements as at the [edit system services dhcp-local-server overrides]
hierarchy level ...
}
trace;
upto upto-interface-name;
}
}
overrides {
  client-discover-match <option60-and-option82>;
  interface-client-limit number;
  no-arp;
  process-inform {
    pool pool-name;
  }
}
pool-match-order {
  external-authority;
  ip-address-first;
  option-82;
}
reconfigure {
  attempts attempt-count;
  clear-on-abort;
  strict;
  timeout timeout-value;
  token token-value;
  trigger {
    radius-disconnect;
  }
}
traceoptions {
  file <filename> <files number> <match regular-expression>
    <size maximum-file-size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
}
dhcpv4-profiles profile-name {
  bind-interface interface-name;
  dead-server-retry-interval interval-in-seconds;
  dead-server-successive-retry-attempt number-of-attempts;
  dhcp-server-selection-algorithm (highest-priority-server | round-robin);
  lease-time time-in-seconds;
  pool-name pool-name;
  retransmission-attempt number-of-attempts;
  retransmission-interval interval-in-seconds;
  servers ip-address {
    priority value;
  }
}
}
dhcpv6-profiles profile-name {
  bind-interface interface-name;
  lease-time time-in-seconds;
  pool-name pool-name;
  retransmission-attempt number-of-attempts;
  retransmission-interval interval-in-seconds;
}

```

```
}
traceoptions {
  file <filename> <files number> <match regular-expression>
    <size maximum-file-size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
}
finger {
  connection-limit limit;
  rate-limit limit;
}
flow-tap-dtcp {
  ssh {
    connection-limit limit;
    rate-limit limit;
  }
}
ftp {
  connection-limit limit;
  rate-limit limit;
}
local-policy-decision-function {
  statistics {
    aacl-statistics-profile profile-name {
      aacl-fields {
        address;
        all-fields;
        application;
        application-group;
        input-bytes;
        input-interface;
        input-packets;
        ipv6-address
        ipv6-prefix-length
        mask;
        output-bytes;
        output-packets;
        subscriber-name;
        timestamp;
        vrf-name;
      }
      file filename;
      record-type (delta | interim);
    }
  }
  file filename {
    archive-sites {
      url;
    }
  }
  files number;
  size bytes;
  transfer-interval minutes;
}
record-type (data | interim);
}
traceoptions {
```



```

    file <filename> <files number> <match regular-expression>
      <size maximum-file-size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
netconf {
  ssh {
    connection-limit limit;
    port port;
    rate-limit limit;
  }
  traceoptions {
    file <filename> <files number> <match regular-expression> <size size>
      <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
    on-demand;
  }
}
outbound-ssh {
  client client-id {
    address {
      port port-number;
      retry number;
      timeout seconds;
    }
    device-id device-id;
    keep-alive {
      retry number;
      timeout seconds;
    }
    reconnect-strategy (in-order | sticky);
    secret secret;
    services netconf;
  }
  traceoptions {
    file <filename> <files number> <match regular-expression>
      <size maximum-file-size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
resource-monitor {
  resource-category jtree {
    resource-type free-dwords {
      low-watermark number;
      high-watermark number;
    }
    resource-type free-pages {
      low-watermark number;
      high-watermark number;
    }
  }
}
no-throttle;
no-logging;

```

```
high-threshold number;
traceoptions {
  file filename <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
}
service-deployment {
  local-certificate certificate-name;
  servers {
    server-address {
      port port-number;
      security-options {
        (ssl3 | tls);
      }
      user username;
    }
  }
  source-address source-address;
  traceoptions {
    flag flag;
  }
}
ssh {
  ciphers [ cipher-1 cipher-2 cipher-3 ...]
  client-alive-count-max seconds;
  client-alive-interval seconds;
  connection-limit limit;
  hostkey-algorithm limit;
  key-exchange limit;
  macs limit;
  max-sessions-per-connection number;
  no-tcp-forwarding;
  protocol-version [v1 v2];
  rate-limit limit;
  root-login (allow | deny | deny-password);
}
subscriber-management {
  enforce-strict-scale-limit-license;
  gres-route-flush-delay;
  maintain-subscriber {
    interface-delete;
  }
  traceoptions {
    file filename <files number> <match regular-expression> <size maximum-file-size>
      <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
traceoptions {
  file filename <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
```

```

    }
telnet {
    connection-limit limit;
    rate-limit limit;
}
tftp-server {
    connection-limit limit;
    rate-limit limit;
}
xnm-clear-text {
    connection-limit limit;
    rate-limit limit;
}
xnm-ssl {
    connection-limit limit;
    local-certificate certificate-name;
    rate-limit limit;
}
}

```

Related Documentation

- [Notational Conventions Used in Junos OS Configuration Hierarchies](#)

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

```

vlans {
  vlan-name {
    description text-description;
    domain-type bridge;
    forwarding-options {
      filter {
        input filter-name;
      }
      flood {
        input filter-name;
      }
    }
    interface interface-name;
    l3-interface interface-name;
    multicast-snooping-options {
      ... same statements as in multicast-snooping-options ...
    }
    no-irb-layer-2-copy;
    service-id number;
    switch-options {
      ... the switch-options subhierarchy appears after the main [edit vlans vlan-name]
        hierarchy ...
    }
    vlan-id (all | none | number);
    vlan-id-list [ vlan-id-numbers ];
    vlan-tags outer <tpid.>vlan-id <inner <tpid.>vlan-id>;
  }

  vlan-name {
    switch-options {
      interface interface-name {

```

```
    interface-mac-limit {  
        limit;  
        packet-action drop;  
    }  
    no-mac-learning;  
    static-mac mac-address {  
        vlan-id number;  
    }  
}  
interface-mac-limit {  
    limit;  
    packet-action drop;  
}  
mac-statistics;  
mac-table-size {  
    number-of-addresses;  
    packet-action drop;  
}  
no-mac-learning;  
}  
}
```

Administration

- [Operational Commands: CLI Interface on page 446](#)
- [Operational Commands: Software Installation on page 465](#)
- [Operational Commands: System Monitoring on page 516](#)

Operational Commands: CLI Interface

set cli complete-on-space

Syntax	set cli complete-on-space (off on)
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Set the command-line interface (CLI) to complete a partial command entry when you type a space or a tab. This is the default behavior of the CLI.
Options	off —Turn off command completion. on —Allow either a space or a tab to be used for command completion.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>CLI User Interface Overview</i> • show cli on page 456
List of Sample Output	set cli complete-on-space on page 447
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

set cli complete-on-space

In the following example, pressing the Spacebar changes the partial command entry from **com** to **complete-on-space**. The example shows how adding the keyword **off** at the end of the command disables command completion.

```
user@host> set cli com<Space>
user@host>set cli complete-on-space off
Disabling complete-on-space
```

set cli directory

Syntax	set cli directory <i>directory</i>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Set the current working directory.
Options	<i>directory</i> —Pathname of the working directory.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• <i>CLI User Interface Overview</i>• show cli directory on page 462
List of Sample Output	set cli directory on page 448
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

set cli directory

```
user@host> set cli directory /var/home/regress
Current directory: /var/home/regress
```

set cli idle-timeout

Syntax	set cli idle-timeout < <i>minutes</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Set the maximum time that an individual session can be idle before the user is logged off the router or switch.
Options	<i>minutes</i> —(Optional) Maximum idle time. The range of values, in minutes, is 0 through 100,000. If you do not issue this command, and the user's login class does not specify this value, the user is never forced off the system after extended idle times. Setting the value to 0 disables the timeout.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>CLI User Interface Overview</i> • show cli on page 456
List of Sample Output	set cli idle-timeout on page 449
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

set cli idle-timeout

```
user@host> set cli idle-timeout 60
Idle timeout set to 60 minutes
```

set cli prompt

Syntax	set cli prompt <i>string</i>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Set the prompt so that it is displayed within the CLI.
Options	<i>string</i> —CLI prompt string. To include spaces in the prompt, enclose the string in quotation marks. By default, the string is <i>username@hostname</i> .
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• CLI User Interface Overview• show cli on page 456
List of Sample Output	set cli prompt on page 450
Output Fields	When you enter this command, the new CLI prompt is displayed.

Sample Output

set cli prompt

```
user@host> set cli prompt lab1-router>  
lab1-router>
```


set cli restart-on-upgrade

Syntax	set cli restart-on-upgrade string (off on)
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	For an individual session, set the CLI to prompt you to restart the router or switch after upgrading the software.
Options	off —Disables the prompt. on —Enables the prompt.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• <i>CLI User Interface Overview</i>• show cli on page 456
List of Sample Output	set cli restart-on-upgrade on page 451
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

set cli restart-on-upgrade

```
user@host> set cli restart-on-upgrade on
Enabling restart-on-upgrade
```

set cli screen-length

Syntax	<code>set cli screen-length <i>length</i></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Set terminal screen length.
Options	<i>length</i> —Number of lines of text that the terminal screen displays (0 through 10,000). The default is 24.
Additional Information	The point at which the ---(more)--- prompt appears on the screen is a function of this setting and the settings for the <code>set cli screen-width</code> and <code>set cli terminal</code> commands.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• <i>CLI User Interface Overview</i>• set cli screen-width on page 453• set cli terminal on page 454• show cli on page 456
List of Sample Output	set cli screen-length on page 452
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

set cli screen-length

```
user@host> set cli screen-length 75
Screen length set to 75
```

set cli screen-width

Syntax	<code>set cli screen-width <i>width</i></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Set the terminal screen width.
Options	<i>width</i> —Number of characters (0 through 1024) in a line. The default is 80.
Additional Information	The point at which the ---(more)--- prompt appears on the screen is a function of this setting and the settings for the <code>set cli screen-length</code> and <code>set cli terminal</code> commands.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>CLI User Interface Overview</i> • set cli screen-length on page 452 • set cli terminal on page 454 • show cli on page 456
List of Sample Output	set cli screen-width on page 453
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

set cli screen-width

```
user@host> set cli screen-width
Screen width set to 132
```

set cli terminal

Syntax	set cli terminal <i>terminal-type</i>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Set the terminal type.
Options	<i>terminal-type</i> —Type of terminal that is connected to the Ethernet management port: <ul style="list-style-type: none">• ansi—ANSI-compatible terminal (80 characters by 24 lines)• small-xterm—Small xterm window (80 characters by 24 lines)• vt100—VT100-compatible terminal (80 characters by 24 lines)• xterm—Large xterm window (80 characters by 65 lines)
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• <i>CLI User Interface Overview</i>• set cli screen-length on page 452• set cli screen-width on page 453• show cli on page 456
List of Sample Output	set cli terminal on page 454
Output Fields	This command provides no output.

Sample Output

set cli terminal

```
user@host> set cli terminal xterm
```

set cli timestamp

Syntax	set cli timestamp (format <i>timestamp-format</i> disable)
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Set a timestamp for CLI output.
Options	<p>format <i>timestamp-format</i>—Set the date and time format for the timestamp. The timestamp format you specify can include the following placeholders in any order:</p> <ul style="list-style-type: none"> • %m—Two-digit month • %d—Two-digit date • %T—Six-digit hour, minute, and seconds <p>disable—Remove the timestamp from the CLI.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>CLI User Interface Overview</i> • show cli on page 456
List of Sample Output	set cli timestamp on page 455
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

set cli timestamp

```
user@host> set cli timestamp format '%m-%d-%T'
'04-21-17:39:13'
CLI timestamp set to: '%m-%d-%T'
```

show cli

Syntax	show cli
Syntax (QFX Series)	show cli <authorization> <directory> <history <i>count</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display configured CLI settings.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show cli on page 457
Output Fields	Table 46 on page 456 lists the output fields for the show cli command. Output fields are listed in the approximate order in which they appear.

Table 46: show cli Output Fields

Field Name	Field Description
CLI complete-on-space	Capability to complete a partial command entry when you type a space or a tab: on or off .
CLI idle-timeout	Maximum time that an individual session can be idle before the user is logged out from the router or switch. When this feature is enabled, the number of minutes is displayed. Otherwise, the state is disabled .
CLI restart-on-upgrade	CLI is set to prompt you to restart the router or switch after upgrading the software: on or off .
CLI screen-length	Number of lines of text that the terminal screen displays.
CLI screen-width	Number of characters in a line on the terminal screen.
CLI terminal	Terminal type.
CLI is operating in	Mode: enhanced .
CLI timestamp	Date and time format for the timestamp. If the timestamp is not set, the state is disabled .
CLI working directory	Pathname of the working directory.

Sample Output

show cli

```
user@host> show cli
CLI complete-on-space set to on
CLI idle-timeout disabled
CLI restart-on-upgrade set to on
CLI screen-length set to 47
CLI screen-width set to 132
CLI terminal is 'vt100'
CLI is operating in enhanced mode
CLI timestamp disabled
CLI working directory is '/var/home/regress'
```

show cli authorization

Syntax	show cli authorization
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the permissions for the current user.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show cli authorization on page 460
Output Fields	Table 47 on page 458 lists the output fields for the show cli authorization command. In the table, all possible permissions are displayed and output fields are listed in alphabetical order.

Table 47: show cli authorization Output Fields

Field Name	Field Description
access	Can view access configuration information.
access-control	Can modify access configuration.
admin	Can view user account information.
admin-control	Can modify user account information.
clear	Can clear learned network information.
configure	Can enter configuration mode.
control	Can modify any configuration.
edit	Can edit configuration files.
field	Reserved for field (debugging) support.
firewall	Can view firewall configuration information.
firewall-control	Can modify firewall configuration information.
floppy	Can read from and write to removable media.
flow-tap	Can view flow-tap configuration information.

Table 47: show cli authorization Output Fields (*continued*)

Field Name	Field Description
flow-tap-control	Can configure flow-tap configuration information.
idp-profiler-operation	Can configure Profiler data.
interface	Can view interface configuration information.
interface-control	Can modify interface configuration information.
maintenance	Can perform system maintenance.
network	Can access the network by entering the ping , ssh , telnet , and traceroute commands.
pgcp-session-mirroring	Can view Packet Gateway Control Protocol session mirroring configuration.
pgcp-session-mirroring-control	Can modify Packet Gateway Control Protocol session mirroring configuration all-control.
reset	Can reset or restart interfaces and system processes.
rollback	Can roll back to previous configurations.
routing	Can view routing configuration information.
routing-control	Can modify routing configuration information.
secret	Can view passwords and authentication keys in the configuration.
secret-control	Can modify passwords and authentication keys in the configuration.
security	Can view security configuration information.
security-control	Can modify security configuration information.
shell	Can start a local shell.
snmp	Can view SNMP configuration information.
snmp-control	Can modify SNMP configuration information.
system	Can view system configuration information.
system-control	Can modify system configuration information.
trace	Can view trace file settings information.

Table 47: show cli authorization Output Fields (*continued*)

Field Name	Field Description
trace-control	Can modify trace file settings information.
view	Can view current values and statistics.
view-configuration	Can view all configuration information (not including secrets).

Sample Output

show cli authorization

```

user@host> show cli authorization
Current user: 'remote' login: 'user' class ''
Permissions:
  admin      -- Can view user accounts
  admin-control-- Can modify user accounts
  clear      -- Can clear learned network information
  configure  -- Can enter configuration mode
  control    -- Can modify any configuration
  edit       -- Can edit full files
  field      -- Special for field (debug) support
  floppy     -- Can read and write from the floppy
  interface  -- Can view interface configuration
  interface-control-- Can modify interface configuration
  network    -- Can access the network
  reset      -- Can reset/restart interfaces and daemons
  routing    -- Can view routing configuration
  routing-control-- Can modify routing configuration
  shell      -- Can start a local shell
  snmp       -- Can view SNMP configuration
  snmp-control-- Can modify SNMP configuration
  system     -- Can view system configuration
  system-control-- Can modify system configuration
  trace      -- Can view trace file settings
  trace-control-- Can modify trace file settings
  view       -- Can view current values and statistics
  maintenance -- Can become the super-user
  firewall   -- Can view firewall configuration
  firewall-control-- Can modify firewall configuration
  secret     -- Can view secret configuration
  secret-control-- Can modify secret configuration
  rollback   -- Can rollback to previous configurations
  security   -- Can view security configuration
  security-control-- Can modify security configuration
  access     -- Can view access configuration
  access-control-- Can modify access configuration
  view-configuration-- Can view all configuration (not including secrets)
  flow-tap   -- Can view flow-tap configuration
  flow-tap-control-- Can configure flow-tap service
Individual command authorization:
  Allow regular expression: none
  Deny regular expression: none
  Allow configuration regular expression: none
  Deny configuration regular expression: none

```


show cli directory

Syntax	show cli directory
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the current working directory.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show cli directory on page 462
Output Fields	Table 48 on page 462 lists the output fields for the show cli directory command. Output fields are listed in the approximate order in which they appear.

Table 48: show cli directory Output Fields

Field Name	Field Description
Current directory	Pathname of the current working directory.

Sample Output

show cli directory

```
user@host> show cli directory
Current directory: /var/home/regress
```

show cli history

Syntax	<code>show cli history</code> <code><count></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display a list of previous CLI commands.
Options	none —Display all previous CLI commands. count —(Optional) Maximum number of commands to display.
Required Privilege Level	view
List of Sample Output	show cli history on page 463
Output Fields	Table 49 on page 463 lists the output fields for the show cli history command. Output fields are listed in the approximate order in which they appear.

Table 49: show cli history Output Fields


Field Name	Field Description
<i>timestamp</i>	Time at which the command was entered.
<i>command-syntax</i>	Command that was entered.

Sample Output

show cli history

```
user@host> show cli history
11:14:14 -- show arp
11:22:10 -- show cli authorization
11:27:12 -- show cli history
```

start shell

Syntax	<code>start shell (csh sh)</code> <code><user username></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Exit from the CLI environment and create a UNIX-level shell. To return to the CLI, type exit from the shell.
<div> NOTE:<ul style="list-style-type: none">To issue this command, the user must have the required login access privileges configured by including the permissions statement at the [edit system login class class-name] hierarchy level.UNIX wheel group membership or permissions are no longer required to issue this command.</div>	
Options	csh —Create a UNIX C shell. sh —Create a UNIX Bourne shell. user username —(Optional) Start the shell as another user.
Additional Information	When you are in the shell, the shell prompt has the following format: <code>username@hostname%</code> An example of the prompt is: <code>root@host%</code>
Required Privilege Level	shell and maintenance
List of Sample Output	start shell csh on page 464
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

start shell csh

```
user@host> start shell csh
%
exit
```

```
%  
username@hostname% start shell sh  
%  
exit  
user@host>
```

Operational Commands: Software Installation

request system license add

Syntax	<code>request system license add (<i>filename</i> terminal)</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Add a license key.
Options	<i>filename</i> —License key from a file or URL. Specify the filename or the URL where the key is located. <i>terminal</i> —License key from the terminal.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• <i>Adding New Licenses (CLI Procedure)</i>
List of Sample Output	request system license add on page 466
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system license add

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


request system license delete

Syntax	<code>request system license delete <i>license-id</i></code>
Syntax (QFX Series)	<code>request system license delete <i>license-identifier</i></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Delete a license key. You can delete only one license at a time.
Options	<i>license-id</i> —License ID that uniquely identifies a license key. <i>license-identification</i> —(QFX Series) License ID that uniquely identifies a license key.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • Deleting a License (CLI Procedure)
List of Sample Output	request system license delete on page 467
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system license delete

```
user@host> request system license delete G03000002223
```

request system license save

Syntax	<code>request system license save (<i>filename</i> terminal)</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Save installed license keys to a file or URL.
Options	<i>filename</i> —License key from a file or URL. Specify the filename or the URL where the key is located. <i>terminal</i> —License key from the terminal.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• <i>Saving License Keys</i>
List of Sample Output	request system license save on page 468
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system license save

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

request system reboot

Syntax	request system reboot <at <i>time</i> > <both-routing-engines> <in <i>minutes</i> > <media (compact-flash disk removable-compact-flash usb)> <message " <i>text</i> "> <other-routing-engine>
Syntax (EX Series Switches)	request system reboot <all-members> <at <i>time</i> > <both-routing-engines> <in <i>minutes</i> > <local> <media (external internal)> <member <i>member-id</i> > <message " <i>text</i> "> <other-routing-engine> <slice <i>slice</i> >
Syntax (TX Matrix Router)	request system reboot <all-chassis all-lcc lcc <i>number</i> scc> <at <i>time</i> > <both-routing-engines> <in <i>minutes</i> > <media (compact-flash disk)> <message " <i>text</i> "> <other-routing-engine>
Syntax (TX Matrix Plus Router)	request system reboot <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> > <at <i>time</i> > <both-routing-engines> <in <i>minutes</i> > <media (compact-flash disk)> <message " <i>text</i> "> <other-routing-engine> <partition (1 2 alternate)>
Syntax (MX Series Router)	request system reboot <all-members> <at <i>time</i> > <both-routing-engines> <in <i>minutes</i> > <local> <media (external internal)> <member <i>member-id</i> > <message " <i>text</i> "> <other-routing-engine>
Release Information	Command introduced before Junos OS Release 7.4.

Option **other-routing-engine** introduced in Junos OS Release 8.0.

Command introduced in Junos OS Release 9.0 for EX Series switches.

Option **sfc** introduced for the TX Matrix Plus router in Junos OS Release 9.6.

Option **both-routing-engines** introduced in Junos OS Release 12.1.

Description Reboot the software.

Options **none**—Reboot the software immediately.

all-chassis—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router or TX Matrix Plus router, reboot all routers connected to the TX Matrix or TX Matrix Plus router, respectively.

all-lcc—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router or TX Matrix Plus router, reboot all line card chassis connected to the TX Matrix or TX Matrix Plus router, respectively.

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

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

- **now**—Stop or reboot the software immediately. This is the default.
- **+minutes**—Number of minutes from now to reboot the software.
- **yymmddhhmm**—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.
- **hh:mm**—Absolute time on the current day at which to stop the software, specified in 24-hour time.

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

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

lcc number—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router or TX Matrix Plus router, the number of a specified line card chassis connected to the TX Matrix or TX Matrix Plus router, respectively. Replace **number** with a value from 0 through 3.

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

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

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

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

- **internal**—Reboot the internal flash device.

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

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

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

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

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

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

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

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

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

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

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



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



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

Required Privilege Level maintenance

Related Documentation

- *clear system reboot*
- *request system halt*
- *request system reboot*
- *Rebooting and Halting a QFX Series Product*

List of Sample Output

- [request system reboot on page 472](#)
- [request system reboot \(at 2300\) on page 472](#)
- [request system reboot \(in 2 Hours\) on page 473](#)
- [request system reboot \(Immediately\) on page 473](#)
- [request system reboot \(at 1:20 AM\) on page 473](#)

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

Sample Output

`request system reboot`

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

`request system reboot (at 2300)`

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

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

request system reboot (in 2 Hours)

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

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

request system reboot (Immediately)

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

request system reboot (at 1:20 AM)

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

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

request system snapshot

Syntax	request system snapshot <partition>
Syntax (ACX Series Routers)	request system snapshot <media type> <partition>
Syntax (EX Series Switches)	request system snapshot <all-members local member <i>member-id</i> > <media type> <partition> <re0 re1 routing-engine <i>routing-engine-id</i> > <slice alternate>
Syntax (J Series Routers)	request system snapshot <as-primary> <config-size <i>size</i> > <data-size <i>size</i> > <factory> <media type> <partition> <root-size <i>size</i> > <swap-size <i>size</i> >
Syntax (MX Series Routers)	request system snapshot <all-members> <local> <member <i>member-id</i> > <partition>
Syntax (TX Matrix Routers)	request system snapshot <all-chassis all-lcc lcc <i>number</i> scc> <partition>
Syntax (TX Matrix Plus Routers)	request system snapshot <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> > <partition>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 10.0 for EX Series switches.
Description	<ul style="list-style-type: none">On the router, back up the currently running and active file system partitions to standby partitions that are not running. Specifically, the root file system (/) is backed up to /altroot, and /config is backed up to /altconfig. The root and /config file systems are on the router's flash drive, and the /altroot and /altconfig file systems are on the router's hard drive.On the switch, take a snapshot of the files currently used to run the switch—the complete contents of the root (/) , /altroot, /config, /var, and /var-tmp directories, which include the running Junos OS, the active configuration, and log files.



CAUTION: After you run the `request system snapshot` command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

Options The specific options available depend upon the router or switch:

none—Back up the currently running software as follows:

- On the router, back up the currently running and active file system partitions to standby partitions that are not running. Specifically, the root file system (/) is backed up to `/altroot`, and `/config` is backed up to `/altconfig`. The root and `/config` file systems are on the router's flash drive, and the `/altroot` and `/altconfig` file systems are on the router's hard drive.
- On the switch, take a snapshot of the files currently used to run the switch and copy them to the media that the switch did not boot from. If the switch is booted from internal media, the snapshot is copied to external (USB) media. If the switch is booted from external (USB) media, the snapshot is copied to internal media.
- If the snapshot destination is external media but a USB flash drive is not connected, an error message is displayed.
- If the automatic snapshot procedure is already in progress, the command returns the following error: **Snapshot already in progress. Cannot start manual snapshot.** For additional information about the automatic snapshot feature, see *Understanding Resilient Dual-Root Partitions on Switches*.

all-chassis | all-lcc | lcc *number* —(TX Matrix and TX Matrix Plus router only) (Optional)

- **all-chassis**—On a TX Matrix router, archive data and executable areas for all Routing Engines in the chassis. On a TX Matrix Plus router, archive data and executable areas for all Routing Engines in the chassis.
- **all-lcc**—On a TX Matrix router, archive data and executable areas for all T640 routers (or line-card chassis) connected to a TX Matrix router. On a TX Matrix Plus router, archive data and executable areas for all T1600 routers (or line-card chassis) connected to a TX Matrix Plus router.
- **lcc *number***—On a TX Matrix router, archive data and executable areas for a specific T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, archive data and executable areas for a specific T1600 router (or line-card chassis) that is connected to a TX Matrix Plus router. Replace *number* with a value from 0 through 3.

all-members | local | member *member-id*—(EX Series Virtual Chassis and MX Series routers only) (Optional) Specify where to place the snapshot (archive data and executable areas) in a Virtual Chassis:

- **all-members**—Create a snapshot (archive data and executable areas) for all members of the Virtual Chassis.

- **local**—Create a snapshot (archive data and executable areas) on the member of the Virtual Chassis that you are currently logged into.
- **member *member-id***—Create a snapshot (archive data and executable areas) for the specified member of the Virtual Chassis.

as-primary—(J Series routers only) (Optional) Create a snapshot that can be used to replace the medium in the primary compact flash drive. This option can be used on the removable compact flash only. The option copies the default files that were loaded on the primary compact flash drive when it was shipped from the factory, plus the rescue configuration if one has been set. This option is useful if you have multiple routers and want to use the same software and configuration on each router. After a boot device is created as a primary compact flash drive, it can operate in only a primary compact flash drive slot. This option causes the boot medium to be partitioned.

config-size *size*—(J Series routers only) (Optional) Specify the size of the config partition, in megabytes. The default value is **10** percent of physical memory on the boot partition. The config partition is mounted on **/config**, and the configuration files are stored in this partition. This option causes the boot medium to be partitioned.

data-size *size*—(J Series routers only) (Optional) Specify the size of the data partition, in megabytes. The default value is **0** MB. The data partition is mounted on **/data**. This space is not used by the router, and can be used for extra storage. This option causes the boot medium to be partitioned.

factory—(J Series routers only) (Optional) Copy only default files that were loaded on the primary compact flash drive when it was shipped from the factory, plus the rescue configuration if one has been set. After the boot medium is created with the factory option, it can operate in only the primary compact flash drive.

media *type*—(J Series routers and EX Series switches only)(Optional) Specify the boot device the software is copied to:

- **compact-flash**—Copy software to the primary compact flash drive.
- **external**—(EX Series switches only) Copy software to an external mass storage device, such as a USB flash drive. If a USB drive is not connected, the switch displays an error message.
- **internal**—(EX Series switches only) Copy software to an internal flash drive.
- **removable-compact-flash**—Copy software to the removable compact flash drive.
- **usb**—(ACX Series, M320, T640, MX960, and J Series routers only) Copy software to the device connected to the USB port.

partition—(Optional) Repartition the flash drive before a snapshot occurs. If the partition table on the flash drive is corrupted, the **request system snapshot** command fails and reports errors. The partition option is supported only for restoring the software image from the hard drive to the flash drive.

(Routers only) You cannot issue the **request system snapshot** command when you enable flash disk mirroring. We recommend that you disable flash disk mirroring when you upgrade or downgrade the software. For more information, see the *Junos OS System Basics Configuration Guide*.

(EX Series switches only) If the snapshot destination is the media that the switch did not boot from, you must use the **partition** option.

re0 | re1 | routing-engine routing-engine-id—(EX6200 and EX8200 switches only) Specify where to place the snapshot in a redundant Routing Engine configuration.

- **re0**—Create a snapshot on Routing Engine 0.
- **re1**—Create a snapshot on Routing Engine 1.
- **routing-engine routing-engine-id**—Create a snapshot on the specified Routing Engine.

root-size size—(J Series routers only) (Optional) Specify the size of the root partition, in megabytes. The default value is one-third of the physical memory minus the config, data, and swap partitions. The root partition is mounted on / and does not include configuration files. This option causes the boot medium to be partitioned.

slice alternate—(EX Series switches only) (Optional) Take a snapshot of the active root partition and copy it onto the alternate slice on the boot media.

scc—(TX Matrix routers only) (Optional) Archive data and executable areas for a TX Matrix router (or switch-card chassis).

sfc number—(TX Matrix Plus routers only) (Optional) Archive data and executable areas for a TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

swap-size size—(J Series router only) (Optional) Specify the size of the swap partition, in megabytes. The default value is one-third of the physical memory on a boot medium larger than 128 MB, or 0 MB on a smaller boot device. The swap partition is used for swap files and software failure memory snapshots. Software failure memory snapshots are saved to the boot medium only if it is specified as the dump device in the system dump-device configuration hierarchy. This option causes the boot medium to be partitioned.

- Additional Information**
- (Routers only) Before upgrading the software on the router, when you have a known stable system, issue the **request system snapshot** command to back up the software, including the configuration, to the **/altroot** and **/altconfig** file systems. After you have upgraded the software on the router and are satisfied that the new packages are successfully installed and running, issue the **request system snapshot** command again to back up the new software to the **/altroot** and **/altconfig** file systems.
 - (Routers only) You cannot issue the **request system snapshot** command when you enable flash disk mirroring. We recommend that you disable flash disk mirroring when you upgrade or downgrade the software. For more information, see the *Junos OS System Basics Configuration Guide*

	<ul style="list-style-type: none"> (TX Matrix and TX Matrix Plus routers only) On a routing matrix, if you issue the request system snapshot command on the master Routing Engine, all the master Routing Engines connected to the routing matrix are backed up. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are backed up.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> show system snapshot on page 514 <i>show system auto-snapshot</i>
List of Sample Output	request system snapshot (Routers) on page 478 request system snapshot (EX Series Switches) on page 478 request system snapshot (When the Partition Flag Is On) on page 479 request system snapshot (When Mirroring Is Enabled) on page 479 request system snapshot all-lcc (Routing Matrix) on page 479 request system snapshot all-members (Virtual Chassis) on page 479
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system snapshot (Routers)

```
user@host> request system snapshot
umount: /altroot: not currently mounted
Copying / to /altroot.. (this may take a few minutes)
umount: /altconfig: not currently mounted
Copying /config to /altconfig.. (this may take a few minutes)

The following filesystems were archived: / /config
```

request system snapshot (EX Series Switches)

```
user@switch> request system snapshot partition
Clearing current label...
Partitioning external media (/dev/da1) ...
Partitions on snapshot:

Partition Mountpoint Size Snapshot argument
s1a /altroot 179M none
s2a / 180M none
s3d /var/tmp 361M none
s3e /var 121M none
s4d /config 60M none
Copying '/dev/da0s1a' to '/dev/da1s1a' .. (this may take a few minutes)
Copying '/dev/da0s2a' to '/dev/da1s2a' .. (this may take a few minutes)
Copying '/dev/da0s3d' to '/dev/da1s3d' .. (this may take a few minutes)
Copying '/dev/da0s3e' to '/dev/da1s3e' .. (this may take a few minutes)
Copying '/dev/da0s4d' to '/dev/da1s4d' .. (this may take a few minutes)
The following filesystems were archived: /altroot / /var/tmp /var /config
```

request system snapshot (When the Partition Flag Is On)

```

user@host> request system snapshot partition
Performing preliminary partition checks ...
Partitioning ad0 ...
umount: /altroot: not currently mounted
Copying / to /altroot.. (this may take a few minutes)

The following filesystems were archived: / /config

```

request system snapshot (When Mirroring Is Enabled)

```

user@host> request system snapshot
Snapshot is not possible since mirror-flash-on-disk is configured.

```

request system snapshot all-lcc (Routing Matrix)

```

user@host> request system snapshot all-lcc
lcc0-re0:
-----
Copying '/' to '/altroot' .. (this may take a few minutes)
Copying '/config' to '/altconfig' .. (this may take a few minutes)
The following filesystems were archived: / /config

lcc2-re0:
-----
Copying '/' to '/altroot' .. (this may take a few minutes)
Copying '/config' to '/altconfig' .. (this may take a few minutes)
The following filesystems were archived: / /config

```

request system snapshot all-members (Virtual Chassis)

```

user@switch> request system snapshot all-members media internal
fpc0:
-----
Copying '/dev/da0s2a' to '/dev/da0s1a' .. (this may take a few minutes)
The following filesystems were archived: /

fpc1:
-----
Copying '/dev/da0s2a' to '/dev/da0s1a' .. (this may take a few minutes)
The following filesystems were archived: /

fpc2:
-----
Copying '/dev/da0s2a' to '/dev/da0s1a' .. (this may take a few minutes)
The following filesystems were archived: /

fpc3:
-----
Copying '/dev/da0s2a' to '/dev/da0s1a' .. (this may take a few minutes)
The following filesystems were archived: /

fpc4:
-----
Copying '/dev/da0s2a' to '/dev/da0s1a' .. (this may take a few minutes)
The following filesystems were archived: /

fpc5:
-----

```

```
Copying '/dev/da0s2a' to '/dev/da0s1a' .. (this may take a few minutes)
The following filesystems were archived: /
```

request system software add

Syntax request system software add *package-name*
 <best-effort-load>
 <delay-restart>
 <force>
 <no-copy>
 <no-validate>
 <re0 | re1>
 <reboot>
 <set [*package-name package-name*]>
 <unlink>
 <upgrade-with-config>
 <upgrade-with-config-format *format*>
 <validate>

Syntax (EX Series Switches) request system software add *package-name*
 <best-effort-load>
 <delay-restart>
 <force>
 <no-copy>
 <no-validate>
 <re0 | re1>
 <reboot>
 <set [*package-name package-name*]>
 <upgrade-with-config>
 <upgrade-with-config-format *format*>
 <validate>

Syntax (TX Matrix Router) request system software add *package-name*
 <best-effort-load>
 <delay-restart>
 <force>
 <lcc *number* | scc>
 <no-copy>
 <no-validate>
 <re0 | re1>
 <reboot>
 <set [*package-name package-name*]>
 <unlink>
 <upgrade-with-config>
 <upgrade-with-config-format *format*>
 <validate>

Syntax (TX Matrix Plus Router) request system software add *package-name*
 <best-effort-load>
 <delay-restart>
 <force>
 <lcc *number* | sfc *number*>
 <no-copy>
 <no-validate>
 <re0 | re1>
 <reboot>
 <set [*package-name package-name*]>

	<div><div><unlink></div><div><upgrade-with-config></div><div><upgrade-with-config-format <i>format</i>></div><div><validate></div></div>
Syntax (MX Series Router)	<div><div>request system software add <i>package-name</i></div><div><best-effort-load></div><div><delay-restart></div><div><force></div><div><member <i>member-id</i>></div><div><no-copy></div><div><no-validate></div><div><re0 re1></div><div><reboot></div><div><set [<i>package-name package-name</i>]></div><div><unlink></div><div><upgrade-with-config></div><div><upgrade-with-config-format <i>format</i>></div><div><validate></div></div>
Syntax (QFX Series)	<div><div>request system software add <i>package-name</i></div><div><best-effort-load></div><div><component all></div><div><delay-restart></div><div><force></div><div><no-copy></div><div><no-validate></div><div><partition></div><div><reboot></div><div><unlink></div><div><upgrade-with-config></div><div><upgrade-with-config-format <i>format</i>></div><div><validate></div></div>
Release Information	<div>Command introduced before Junos OS Release 7.4.</div> <div>best-effort-load and unlink options added in Junos OS Release 7.4.</div> <div>Command introduced in Junos OS Release 9.0 for EX Series switches.</div> <div>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</div> <div>Command introduced in Junos OS Release 11.1 for the QFX Series.</div> <div>set [<i>package-name package-name</i>] option added in Junos OS Release 11.1 for EX Series switches.</div> <div>set [<i>package-name package-name</i>] option added in Junos OS Release 12.2 for M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways.</div>



NOTE: On EX Series switches, the **set [*package-name package-name*]** option allows you to install only two software packages on a mixed EX4200 and EX4500 Virtual Chassis, whereas, on M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways, the **set [*package-name package-name*]** option allows you to install multiple software packages and software add-on packages at the same time.

upgrade-with-config and **upgrade-with-config-format** *format* options added in Junos OS Release 12.3 for M Series routers, MX Series routers, T Series routers, EX Series Ethernet switches, and QFX Series devices.

Description



NOTE: We recommend that you always download the software image to `/var/tmp` only. On EX Series and QFX Series switches, you must use the `/var/tmp` directory. Other directories are not supported.

Install a software package or bundle on the router or switch.

Options

package-name—Location from which the software package or bundle is to be installed.

For example:

- **/var/tmp/package-name**—For a software package or bundle that is being installed from a local directory on the router or switch.
- **protocol://hostname/pathname/package-name**—For a software package or bundle that is to be downloaded and installed from a remote location. Replace **protocol** with one of the following:
 - **ftp**—File Transfer Protocol.
Use **ftp://hostname/pathname/package-name**. To specify authentication credentials, use **ftp://<username>:<password>@hostname/pathname/package-name**. To have the system prompt you for the password, specify **prompt** in place of the password. If a password is required, and you do not specify the password or **prompt**, an error message is displayed.
 - **http**—Hypertext Transfer Protocol.
Use **http://hostname/pathname/package-name**. To specify authentication credentials, use **http://<username>:<password>@hostname/pathname/package-name**. If a password is required and you omit it, you are prompted for it.
 - **scp**—Secure copy (available only for Canada and U.S. version).
Use **scp://hostname/pathname/package-name**. To specify authentication credentials, use **scp://<username>:<password>@hostname/pathname/package-name**.



NOTE:

- The *pathname* in the protocol is the relative path to the user's home directory on the remote system and not the root directory.
- Do not use the `scp` protocol in the `request system software add` command to download and install a software package or bundle from a remote location. The previous statement does not apply to the QFabric switch. The software upgrade is handled by the MGD process which does not support `scp`.
Use the `file copy` command to copy the software package or bundle from the remote location to the `/var/tmp` directory on the hard disk:
`file copy scp://source/package-name /var/tmp`
Then install the software package or bundle using the `request system software add` command:
`request system software add /var/tmp/package-name`
- On a J Series Services Router, when you install the software from a remote location, the package is removed at the earliest opportunity in order to make room for the installation to be completed. If you copy the software to a local directory on the router and then install the new package, use the `unlink` option to achieve the same effect and allow the installation to be completed.

best-effort-load—(Optional) Activate a partial load and treat parsing errors as warnings instead of errors.

component all—(QFabric systems only) (Optional) Install software package on all of the QFabric components.

delay-restart—(Optional) Install a software package or bundle, but do not restart software processes.

force—(Optional) Force the addition of the software package or bundle (ignore warnings).

lcc number—(TX Matrix and TX Matrix Plus routers only) (Optional) In a routing matrix based on the TX Matrix router, install a software package or bundle on a T640 router (or line-card chassis) that is connected to the TX Matrix router. In a routing matrix based on the TX Matrix Plus router, install a software package or bundle on a T1600 router (or line-card chassis) that is connected to the TX Matrix Plus router. Replace *number* with a value from 0 through 3.

member member-id—(MX Series routers only) (Optional) Install a software package on the specified Virtual Chassis member. Replace *member-id* with a value of 0 or 1.

partition—(QFX3500 switches only) (Optional) Format and repartition the media before installation.

scc—(TX Matrix routers only) (Optional) Install a software package or bundle on a Routing Engine on a TX Matrix router (or switch-card chassis).

sfc number—(TX Matrix Plus routers only) (Optional) Install a software package or bundle on a Routing Engine on a TX Matrix Plus router (or switch-fabric chassis). Replace *number* with **0**.

no-copy—(Optional) Install a software package or bundle, but do not save copies of the package or bundle files.

no-validate—(Optional) When loading a software package or bundle with a different release, suppress the default behavior of the **validate** option.

re0 | re1—(Optional) On routers or switches that support dual or redundant Routing Engines, load a software package or bundle on the Routing Engine in slot 0 (**re0**) or the Routing Engine in slot 1 (**re1**).

reboot—(Optional) After adding the software package or bundle, reboot the system. On a QFabric switch, the software installation is not complete until you reboot the component for which you have installed the software.

set [package-name package-name]—(Mixed EX4200 and EX4500 Virtual Chassis only) (Optional) Install two software packages—a package for an EX4200 switch and the same release of the package for an EX4500 switch—to upgrade all member switches in a mixed EX4200 and EX4500 Virtual Chassis.

set [package-name package-name]—(M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways only) (Optional) Install multiple software packages and software add-on packages at the same time.

unlink—(Optional) On J Series Services Routers, this option ensures that the software package is removed at the earliest opportunity in order to make room for the installation to be completed. On M Series, T Series, and MX Series routers, use the **unlink** option to remove the software package from this directory after a successful upgrade is completed.

upgrade-with-config—(Optional) Install one or more configuration files.

upgrade-with-config-format format—(Optional) Specify the configuration file format, **text** or **xml**. The default format is **text**.



NOTE: The **upgrade-with-config** and **upgrade-with-config-format** options are only available locally on the router or switch. In a routing matrix, the configuration is applied only to the local router and is not propagated to other routers.

The options are validated during the validation process and applied to the router or switch during the upgrade process. If the upgrade process is successful, the options are removed from the configuration. If the upgrade process fails, the configuration file is renamed with the **.failed** suffix.

validate—(Optional) Validate the software package or bundle against the current configuration as a prerequisite to adding the software package or bundle. This is the default behavior when the software package or bundle being added is a different release.

Additional Information



NOTE: The **request system snapshot** command is currently not supported on the QFabric system. Also, you cannot add or install multiple packages on a QFabric system.

Before upgrading the software on the router or switch, when you have a known stable system, issue the **request system snapshot** command to back up the software, including the configuration, to the **/altroot** and **/altconfig** file systems. After you have upgraded the software on the router or switch and are satisfied that the new package or bundle is successfully installed and running, issue the **request system snapshot** command again to back up the new software to the **/altroot** and **/altconfig** file systems.

After you run the **request system snapshot** command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

If you are upgrading more than one package at the same time, delete the operating system package, **jkernl**, last. Add the operating system package, **jkernl**, first and the routing software package, **jroute**, last. If you are upgrading all packages at once, delete and add them in the following order:

```
user@host> request system software add /var/tmp/jbase
user@host> request system software add /var/tmp/jkernl
user@host> request system software add /var/tmp/jpfe
user@host> request system software add /var/tmp/jdocs
user@host> request system software add /var/tmp/jroute
user@host> request system software add /var/tmp/jcrypto
```

By default, when you issue the **request system software add package-name** command on a TX Matrix master Routing Engine, all the T640 master Routing Engines that are connected to it are upgraded to the same version of software. If you issue the same command on the TX Matrix backup Routing Engine, all the T640 backup Routing Engines that are connected to it are upgraded to the same version of software.

Likewise, when you issue the **request system software add package-name** command on a TX Matrix Plus master Routing Engine, all the T1600 master Routing Engines that are connected to it are upgraded to the same version of software. If you issue the same command on the TX Matrix Plus backup Routing Engine, all the T1600 backup Routing Engines that are connected to it are upgraded to the same version of software.

Required Privilege Level maintenance

Related Documentation	<ul style="list-style-type: none"> • request system software delete on page 489 • request system software rollback on page 492 • <i>request system storage cleanup</i> • <i>Upgrading Software on QFX3500 and QFX3600 Standalone Switches</i> • <i>Upgrading Software on a QFabric System</i>
List of Sample Output	request system software add validate on page 487 request system software add (Mixed EX4200 and EX4500 Virtual Chassis) on page 488 request system software add component all (QFabric Systems) on page 488
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system software add validate

```

user@host> request system software add validate /var/tmp/jinstall-7.2R1.7-domestic-signed.tgz
Checking compatibility with configuration
Initializing...
Using jbase-7.1R2.2
Using /var/tmp/jinstall-7.2R1.7-domestic-signed.tgz
Verified jinstall-7.2R1.7-domestic.tgz signed by PackageProduction_7_2_0
Using /var/validate/tmp/jinstall-signed/jinstall-7.2R1.7-domestic.tgz
Using /var/validate/tmp/jinstall/jbundle-7.2R1.7-domestic.tgz
Checking jbundle requirements on /
Using /var/validate/tmp/jbundle/jbase-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jkernel-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jcrypto-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jpfe-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jdocs-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jroute-7.2R1.7.tgz
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Validating against /config/rescue.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-7.2R1.7-domestic-signed.tgz' ...
Verified jinstall-7.2R1.7-domestic.tgz signed by PackageProduction_7_2_0
Adding jinstall...

WARNING: This package will load JUNOS 7.2R1.7 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'

```

WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-7.2R1.7-domestic-signed.tgz ...
Saving state for rollback ...

Sample Output


request system software add (Mixed EX4200 and EX4500 Virtual Chassis)

```
user@switch> request system software add set  
[/var/tmp/jinstall-ex-4200-11.1R1.1-domestic-signed.tgz  
/var/tmp/jinstall-ex-4500-11.1R1.1-domestic-signed.tgz]  
...
```

request system software add component all (QFabric Systems)

```
user@switch> request system software add /pbdata/packages/jinstall-qfabric-12.2X50-D1.3.rpm  
component all  
...
```

request system software delete

Syntax	request system software delete <i>software-package</i> <force> <reboot> <set [<i>package-name package-name</i>]>
Syntax (TX Matrix Router)	request system software delete <i>software-package</i> <force> <lcc <i>number</i> scc> <reboot> <set [<i>package-name package-name</i>]>
Syntax (TX Matrix Plus Router)	request system software delete <i>software-package</i> <force> <lcc <i>number</i> sfc <i>number</i> > <reboot> <set [<i>package-name package-name</i>]>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option sfc introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option set [<i>package-name package-name</i>] added in Junos OS Release 12.2 for M Series, MX Series, T Series routers, and Branch SRX Services Gateways.</p> <p>Option reboot introduced in Junos OS Release 12.3.</p>
Description	Remove a software package or bundle from the router or switch.
<div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p>CAUTION: Before removing a software package or bundle, make sure that you have already placed the new software package or bundle that you intend to load onto the router or switch.</p> </div> </div>	
Options	<p><i>software-package</i>—Software package or bundle name. You can delete any or all of the following software bundles or packages:</p> <ul style="list-style-type: none"> • jbase—(Optional) Junos base software suite • jcrypto—(Optional, in domestic version only) Junos security software • jdocs—(Optional) Junos online documentation file • jkernel—(Optional) Junos kernel software suite • jpfe—(Optional) Junos Packet Forwarding Engine support • jroute—(Optional) Junos routing software suite • junos—(Optional) Junos base software



NOTE: On EX Series switches, some of the package names are different than those listed. To see the list of packages that you can delete on an EX Series switch, enter the command **show system software**.

force—(Optional) Ignore warnings and force removal of the software.

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, remove an extension or upgrade package from a specific T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, remove an extension or upgrade package from a specific T1600 router (or line-card chassis) that is connected to the TX Matrix Plus router. Replace ***number*** with a value from 0 through 3.

reboot—As of Junos OS 12.3 and greater, automatically reboot upon completing the **request system software delete** command.

scc—(TX Matrix routers only) (Optional) Remove an extension or upgrade package from the TX Matrix router (or switch-card chassis).

set [*package-name package-name*]—(M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways only) (Optional) Install multiple software packages or software add-on packages at the same time.

sfc *number*—(TX Matrix Plus routers only) (Optional) Remove an extension or upgrade package from the TX Matrix Plus router (or switch-fabric chassis). Replace ***number*** with 0.

Additional Information Before upgrading the software on the router or switch, when you have a known stable system, issue the **request system snapshot** command to back up the software, including the configuration, to the /altroot and /altconfig file systems (on routers) or the /, /altroot, /config, /var, and /var/tmp file systems (on switches). After you have upgraded the software on the router or switch and are satisfied that the new packages are successfully installed and running, issue the **request system snapshot** command again to back up the new software to the /altroot and /altconfig file systems (on routers) or the /, /altroot, /config, /var, and /var/tmp file systems (on switches). After you run the **request system snapshot** command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

Required Privilege Level maintenance

Related Documentation

- [request system software add on page 481](#)
- [request system software rollback on page 492](#)
- [request system software validate on page 496](#)

List of Sample Output [request system software delete jdocs on page 491](#)

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

Sample Output

request system software delete jdocs

The following example displays the system software packages before and after the **jdocs** package is deleted through the **request system software delete** command:

```
user@host> show system software
Information for jbase:

Comment:
JUNOS Base OS Software Suite [7.2R1.7]

Information for jcrypto:

Comment:
JUNOS Crypto Software Suite [7.2R1.7]

Information for jdocs:

Comment:
JUNOS Online Documentation [7.2R1.7]

Information for jkernel:

Comment:
JUNOS Kernel Software Suite [7.2R1.7]

...

user@host> request system software delete jdocs
Removing package 'jdocs' ...

user@host> show system software
Information for jbase:

Comment:
JUNOS Base OS Software Suite [7.2R1.7]

Information for jcrypto:

Comment:
JUNOS Crypto Software Suite [7.2R1.7]

Information for jkernel:

Comment:
JUNOS Kernel Software Suite [7.2R1.7]

...
```

request system software rollback

Syntax	request system software rollback
Syntax (EX Series Switches)	request system software rollback <all-members> <local> <member <i>member-id</i> > <reboot>
Syntax (TX Matrix Router)	request system software rollback <lcc <i>number</i> scc> <reboot>
Syntax (TX Matrix Plus Router)	request system software rollback <lcc <i>number</i> sfc <i>number</i> > <reboot>
Syntax (MX Series Router)	request system software rollback <all-members> <local> <member <i>member-id</i> > <reboot>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option sfc introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command behavior changed in Junos OS Release 12.1.</p> <p>Option reboot introduced in Junos OS Release 12.3.</p>
Description	<p>For all versions of Junos OS up to and including Junos OS 11.4, revert to the software that was loaded at the last successful request system software add command.</p> <p>As of Junos OS 12.1 and greater, revert to the last known good state before the most recent request system software (add delete) command. For example, using rollback in Junos OS 12.1 after using request system software add restores the system to a known good state prior to using the add command. Similarly, using rollback in Junos OS 12.1 after using request system software delete restores the system to a known good state prior to using the delete command.</p> <p>A software rollback fails if any required package (or a jbundle package containing the required package) cannot be found in /var/sw/pkg.</p> <p><i>Additional Information</i></p> <ul style="list-style-type: none">On M Series and T Series routers, if request system software add <jinstall> reboot was used for the previous installation, then request system software rollback has no effect. In this case, use jinstall to reinstall the required package.

- On M Series and T Series routers, if **request system software add <sdk1>** was used for the previous installation, then **request system software rollback** removes the last installed SDK package (**sdk1** in this example).
- On SRX Series devices with dual root systems, when **request system software rollback** is run, the system switches to the alternate root. Each root can have a different version of Junos OS. Rollback takes each root back to the previously installed image.

Options **all-members**—(EX4200 switches and MX Series routers only) (Optional) Attempt to roll back to the previous set of packages on all members of the Virtual Chassis configuration.

lcc number—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, attempt to roll back to the previous set of packages on a T640 router (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, attempt to roll back to the previous set of packages on a T1600 router (or line-card chassis) connected to the TX Matrix Plus router. Replace **number** with a value from 0 through 3.

local—(EX4200 switches and MX Series routers only) (Optional) Attempt to roll back to the previous set of packages on the local Virtual Chassis member.

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

none—For all versions of Junos OS up to and including Junos OS 11.4, revert to the set of software as of the last successful **request system software add**. As of Junos OS 12.1 and greater, revert to the last known good state before the most recent **request system software (add | delete)** command.

reboot—As of Junos OS 12.3 and greater, automatically reboot upon completing the **request system software rollback** command.

scc—(TX Matrix routers only) (Optional) Attempt to roll back to the previous set of packages on the TX Matrix router (or switch-card chassis).

sfc number—(TX Matrix Plus routers only) (Optional) Attempt to roll back to the previous set of packages on the TX Matrix Plus router (or switch-fabric chassis). Replace **number** with 0.

Required Privilege Level maintenance

Related Documentation

- *request system software abort*
- [request system software add on page 481](#)
- [request system software delete on page 489](#)

- [request system software validate on page 496](#)
- *request system configuration rescue delete*
- *request system configuration rescue save*

List of Sample Output [request system software rollback on page 495](#)

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

Sample Output

request system software rollback

```

user@host> request system software rollback
Verified SHA1 checksum of ./jbase-7.2R1.7.tgz
Verified SHA1 checksum of ./jdocs-7.2R1.7.tgz
Verified SHA1 checksum of ./jroute-7.2R1.7.tgz
Installing package './jbase-7.2R1.7.tgz' ...
Available space: 35495 require: 7335
Installing package './jdocs-7.2R1.7.tgz' ...
Available space: 35339 require: 3497
Installing package './jroute-7.2R1.7.tgz' ...
Available space: 35238 require: 6976
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Reloading /config/juniper.conf.gz ...
Activating /config/juniper.conf.gz ...
mgd: commit complete
Restarting mgd ...
Restarting aprobed ...
Restarting apsd ...
Restarting cosd ...
Restarting fsad ...
Restarting fud ...
Restarting gcdrd ...
Restarting ilmid ...
Restarting irsd ...
Restarting l2tpd ...
Restarting mib2d ...
Restarting nasd ...
Restarting pppoed ...
Restarting rdd ...
Restarting rmopd ...
Restarting rtspd ...
Restarting sampled ...
Restarting serviced ...
Restarting snmpd ...
Restarting spd ...
Restarting vrrpd ...

WARNING: cli has been replaced by an updated version:
CLI release 7.2R1.7 built by builder on 2005-04-22 02:03:44 UTC
Restart cli using the new version ? [yes,no] (yes) yes

Restarting cli ...
user@host

```

request system software validate

Syntax	request system software validate <i>package-name</i> <set [<i>package-name package-name</i>]> <upgrade-with-config> <upgrade-with-config-format <i>format</i> >
Syntax (TX Matrix Router)	request system software validate <i>package-name</i> <lcc <i>number</i> scc> <set [<i>package-name package-name</i>]> <upgrade-with-config> <upgrade-with-config-format <i>format</i> >
Syntax (TX Matrix Plus Router)	request system software validate <i>package-name</i> <lcc <i>number</i> sfc <i>number</i> > <set [<i>package-name package-name</i>]> <upgrade-with-config> <upgrade-with-config-format <i>format</i> >
Syntax (MX Series Router)	request system software validate <i>package-name</i> <member <i>member-id</i> > <set [<i>package-name package-name</i>]> <upgrade-with-config> <upgrade-with-config-format <i>format</i> >
Release Information	Command introduced before Junos OS Release 7.4. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. set [<i>package-name package-name</i>] option added in Junos OS Release 12.2 for M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways. upgrade-with-config and upgrade-with-config-format <i>format</i> options added in Junos OS Release 12.3 for M Series routers, MX Series routers, and T Series routers.
Description	Validate candidate software against the current configuration of the router.
Options	<p>lcc <i>number</i>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, validate the software bundle or package on a specific T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, validate the software bundle or package on a specific T1600 router (or line-card chassis) that is connected to the TX Matrix Plus router. Replace <i>number</i> with a value from 0 through 3.</p> <p>member <i>member-id</i>—(MX Series routers only) (Optional) Validate the software bundle or package on the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace <i>member-id</i> with a value of 0 or 1.</p> <p><i>package-name</i>—Name of the software bundle or package to test.</p> <p>scc—(TX Matrix routers only) (Optional) Validate the software bundle or package for the TX Matrix router (or switch-card chassis).</p>

set [*package-name package-name*]**—**(M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways only) (Optional) Install multiple software packages or software add-on packages at the same time.

sfc number**—**(TX Matrix Plus routers only) (Optional) Validate the software bundle or package for the TX Matrix Plus router (or switch-fabric chassis).

upgrade-with-config**—**(Optional) Install one or more configuration files.

upgrade-with-config-format *format***—**(Optional) Specify the configuration file format, **text** or **xml**. The default format is **text**.



NOTE: The **upgrade-with-config** and **upgrade-with-config-format** options are only available locally on the router or switch. In a routing matrix, the configuration is applied only to the local router and is not propagated to other routers.

The options are validated during the validation process and applied to the router or switch during the upgrade process. If the upgrade process is successful, the options are removed from the configuration. If the upgrade process fails, the configuration file is renamed with the **.failed** suffix.

Additional Information By default, when you issue the **request system software validate** command on a TX Matrix master Routing Engine, all the T640 master Routing Engines that are connected to it are validated. If you issue the same command on the TX Matrix backup Routing Engine, all the T640 backup Routing Engines that are connected to it are upgraded to the same version of software.

Likewise, if you issue the **request system software validate** command on a TX Matrix Plus master Routing Engine, all the T1600 master Routing Engines that are connected to it are validated. If you issue the same command on a TX Matrix Plus backup Routing Engine, all the T1600 backup Routing Engines that are connected to it are upgraded to the same version of software.

Required Privilege Level maintenance

Related Documentation

- [request system software abort](#)
- [request system software add on page 481](#)
- [request system software delete on page 489](#)
- [request system software rollback on page 492](#)

List of Sample Output [request system software validate \(Successful Case\) on page 498](#)
[request system software validate \(Failure Case\) on page 498](#)

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

Sample Output

request system software validate (Successful Case)

```
user@host> request system software validate /var/sw/pkg/jbundle-5.3I20020124_0520_sjg.tgz
Checking compatibility with configuration
Initializing...
Using /packages/jbase-5.3I20020122_1901_sjg
Using /var/sw/pkg/jbundle-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jbase-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jkernel-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jcrypto-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jpfe-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jdocs-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jroute-5.3I20020124_0520_sjg.tgz
Validating against /config/juniper.conf.gz
mgd: commit complete

WARNING: cli has been replaced by an updated version:
CLI release 5.3I0 built by sjg on 2002-01-24 05:23:53 UTC
Restart cli using the new version ? [yes,no] (yes)
```


request system software validate (Failure Case)

```
user@host> request system software validate 6.3/
Pushing bundle to lcc0-re0
error: Failed to transfer package to lcc0-re0

user@host> request system software validate test
Pushing bundle to lcc0-re0
Pushing bundle to lcc2-re0

lcc0-re0:
gzip: stdin: not in gzip format
tar: child returned status 1
ERROR: Not a valid package: /var/tmp/test
```


request system zeroize

Syntax	request system zeroize <media>
Release Information	<p>Command introduced before Junos OS Release 9.0.</p> <p>Command introduced in Junos OS Release 11.2 for EX Series switches.</p> <p>Option media added in Junos OS Release 11.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.2 for MX Series devices.</p> <p>Command introduced in Junos OS Release 12.3 for the QFX Series.</p>
Description	<div>  <p>NOTE: The media option is not available on the QFX Series.</p> </div> <p>Remove all configuration information on the Routing Engines and reset all key values. If the device has dual Routing Engines, the command is broadcast to all Routing Engines on the device. The command removes all data files, including customized configuration and log files, by unlinking the files from their directories. The command removes all user-created files from the system including all plain-text passwords, secrets, and private keys for SSH, local encryption, local authentication, IPsec, RADIUS, TACACS+, and SNMP.</p> <p>This command reboots the device and sets it to the factory default configuration. After the reboot, you cannot access the device through the management Ethernet interface. Log in through the console as root and start the Junos OS command-line interface (CLI) by typing cli at the prompt.</p> <p>To completely erase user-created data so that it is unrecoverable, use the media option.</p>
Options	<p>media—(Optional) In addition to removing all configuration and log files, the media option causes memory and the media to be scrubbed, removing all traces of any user-created files. Every storage device attached to the system is scrubbed, including disks, flash drives, removable USBs, and the like. The duration of the scrubbing process is dependent on the size of the media being erased. As a result, the request system zeroize media operation can take considerably more time than the request system zeroize operation. However, the critical security parameters are all removed at the beginning of the process.</p>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • request system snapshot on page 474 • <i>request system snapshot</i> • <i>Reverting to the Default Factory Configuration for the EX Series Switch</i> • <i>Reverting to the Rescue Configuration for the EX Series Switch</i> • <i>Reverting to the Default Factory Configuration</i>

- *Reverting to the Rescue Configuration*
- *Reverting to the Default Factory Configuration by Using the request system zeroize Command*

List of Sample Output [request system zeroize on page 500](#)
[request system zeroize media on page 501](#)

Sample Output

request system zeroize

```
user@host> request system zeroize
warning: System will be rebooted and may not boot without configuration
Erase all data, including configuration and log files? [yes,no] (no) yes

0 1 1 0 0 0 done

syncing disks... All buffers synced.
Uptime: 5d19h20m26s
recorded reboot as normal shutdown
Rebooting...

U-Boot 1.1.6 (Mar 11 2011 - 04:39:06)

Board: EX4200-24T 2.11
EPLD: Version 6.0 (0x85)
DRAM: Initializing (1024 MB)
FLASH: 8 MB

Firmware Version: --- 01.00.00 ---
USB: scanning bus for devices... 2 USB Device(s) found
      scanning bus for storage devices... 1 Storage Device(s) found

ELF file is 32 bit
Consoles: U-Boot console

FreeBSD/PowerPC U-Boot bootstrap loader, Revision 2.4
(user@juniper.net, Fri Mar 11 03:03:36 UTC 2011)
Memory: 1024MB
bootsequencing is enabled
bootsuccess is set
new boot device = disk0s1:
Loading /boot/defaults/loader.conf
/kernel data=0x915c84+0xa1260 syms=[0x4+0x7cbd0+0x4+0xb1c19]

Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [/kernel]...
Kernel entry at 0x800000e0 ...
GDB: no debug ports present
KDB: debugger backends: ddb
KDB: current backend: ddb
Copyright (c) 1996-2011, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
The Regents of the University of California. All rights reserved.
JUNOS 11.1R1.8 #0: 2011-03-09 20:14:25 UTC
```

```

user@juniper.net:/volume/build/junos/11.1/release/11.1R1.8/obj-powerpc/bsd/kernels/
  JUNIPER-EX/kernel
Timecounter "decrementer" frequency 50000000 Hz quality 0
cpu0: Freescale e500v2 core revision 2.2
cpu0: HID0 80004080
...

```

request system zeroize media

```

user@host> request system zeroize media
warning: System will be rebooted and may not boot without configuration
Erase all data, including configuration and log files? [yes,no] (no) yes

warning: ipsec-key-management subsystem not running - not needed by configuration.
warning: zeroizing fpc0

{master:0}
root> Waiting (max 60 seconds) for system process `vnlru' to stop...done
...
Syncing disks, vnodes remaining...2 4 2 4 3 2 1 1 0 0 0 done

syncing disks... All buffers synced.
Uptime: 14m50s
recorded reboot as normal shutdown
Rebooting...

U-Boot 1.1.6 (Apr 21 2011 - 13:58:42)

Board: EX4200-48PX 1.1
EPLD: Version 8.0 (0x82)
DRAM: Initializing (512 MB)
FLASH: 8 MB
NAND: No NAND device found!!!
0 MiB

Firmware Version: --- 01.00.00 ---
USB: scanning bus for devices... 2 USB Device(s) found
      scanning bus for storage devices... 1 Storage Device(s) found

ELF file is 32 bit
Consoles: U-Boot console

FreeBSD/PowerPC U-Boot bootstrap loader, Revision 2.2
(vtseng@svl-junos-pool27.juniper.net, Fri Feb 26 17:48:51 PST 2010)
Memory: 512MB
Loading /boot/defaults/loader.conf
/kernel data=0x9abfdc+0xb06e4 syms=[0x4+0x83b30+0x4+0xbd7c6]

Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [/kernel] in 1 second... Booting [/kernel]...
Kernel entry at 0x800000e0 ...
GDB: no debug ports present
KDB: debugger backends: ddb
KDB: current backend: ddb
Copyright (c) 1996-2011, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
The Regents of the University of California. All rights reserved.
JUNOS 11.4R1.2 #0: 2011-10-27 18:05:39 UTC

```

```
user@juniper.net:/volume/build/junos/11.4/release/11.4R1.2/obj-powerpc/
bsd/kernels/JUNIPER-EX/kernel
can't re-use a leaf (all_slot_serialid)!
Timecounter "decrementer" frequency 50000000 Hz quality 0
cpu0: Freescale e500v2 core revision 2.2
cpu0: HID0 80004080<EMCP,TBEN,EN_MAS7_UPDATE>
real memory = 511705088 (488 MB)
avail memory = 500260864 (477 MB)
ETHERNET SOCKET BRIDGE initialising
Initializing EXSERIES platform properties ...
. . .
Automatic reboot in progress...
Media check on da0 on ex platforms
** /dev/da0s2a
FILE SYSTEM CLEAN; SKIPPING CHECKS
clean, 20055 free (31 frags, 2503 blocks, 0.0% fragmentation)
zeroizing /dev/da0s1a ...
. . .
zeroizing /dev/da0s3d ...
. . .
zeroizing /dev/da0s3e ...
. . .
zeroizing /dev/da0s4d ...
. . .
zeroizing /dev/da0s4e ...
. . .

syncing disks... All buffers synced.
Uptime: 3m40s
Rebooting...

U-Boot 1.1.6 (Apr 21 2011 - 13:58:42)

Board: EX4200-48PX 1.1
EPLD: Version 8.0 (0x82)
DRAM: Initializing (512 MB)
FLASH: 8 MB
NAND: No NAND device found!!!
0 MiB

Firmware Version: --- 01.00.00 ---
USB: scanning bus for devices... 2 USB Device(s) found
      scanning bus for storage devices... 1 Storage Device(s) found

ELF file is 32 bit
Consoles: U-Boot console

FreeBSD/PowerPC U-Boot bootstrap loader, Revision 2.2
(vtseng@svl-junos-pool27.juniper.net, Fri Feb 26 17:48:51 PST 2010)
Memory: 512MB
Loading /boot/defaults/loader.conf
/kernel data=0x9abfdc+0xb06e4 syms=[0x4+0x83b30+0x4+0xbd7c6]

Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [/kernel] in 1 second... Booting [/kernel]...
Kernel entry at 0x800000e0 ...
GDB: no debug ports present
KDB: debugger backends: ddb
KDB: current backend: ddb
Copyright (c) 1996-2011, Juniper Networks, Inc.
All rights reserved.
```

```

Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
The Regents of the University of California. All rights reserved.
JUNOS 11.4R1.2 #0: 2011-10-27 18:05:39 UTC
user@juniper.net:/volume/build/junos/11.4/release/11.4R1.2/obj-powerpc/
bsd/kernels/JUNIPER-EX/kernel
can't re-use a leaf (all_slot_serialid)!
Timecounter "decrementer" frequency 50000000 Hz quality 0
cpu0: Freescale e500v2 core revision 2.2
cpu0: HID0 80004080 <EMCP,TBEN,EN_MAS7_UPDATE>
real memory = 511705088 (488 MB)
avail memory = 500260864 (477 MB)
ETHERNET SOCKET BRIDGE initialising
Initializing EXSERIES platform properties ...
. . .
Automatic reboot in progress...
Media check on da0 on ex platforms
** /dev/da0s1a
FILE SYSTEM CLEAN; SKIPPING CHECKS
clean, 20064 free (48 frags, 2502 blocks, 0.1% fragmentation)
zeroizing /dev/da0s2a ...
. . .
Creating initial configuration...mgd: error: Cannot open configuration file:
/config/juniper.conf
mgd: warning: activating factory configuration
mgd: commit complete
mgd: -----
mgd: Please login as 'root'. No password is required.
mgd: To start Initial Setup, type 'ezsetup' at the JUNOS prompt.
mgd: To start JUNOS CLI, type 'cli' at the JUNOS prompt.
mgd: -----
Setting initial options: debugger_on_panic=NO debugger_on_break=NO.
Starting optional daemons: .
Doing initial network setup:
. . .

Amnesiac (ttyu0)

```

show system boot-messages

Syntax	show system boot-messages
Syntax (EX Series Switches)	show system boot-messages <all-members> <local> <member <i>member-id</i> >
Syntax (TX Matrix Router)	show system boot-messages <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show system boot-messages <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Router)	show system boot-messages <all-members> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show system boot-messages infrastructure <i>name</i> interconnect-device <i>name</i> node-group <i>name</i>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display initial messages generated by the system kernel upon startup. These messages are the contents of <code>/var/run/dmesg.boot</code> .
Options	none —Display all boot time messages. all-chassis —(TX Matrix and TX Matrix Plus routers only) (Optional) Display boot time messages for all of the chassis. all-lcc —(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display boot time messages for all T640 routers (or line-card chassis) connected to a TX Matrix router. On a TX Matrix Plus router, display boot time messages for all T1600 routers (or line-card chassis) connected to a TX Matrix Plus router. all-members —(EX4200 switches and MX Series routers only) (Optional) Display boot time messages on all members of the Virtual Chassis configuration. infrastructure <i>name</i> —(QFabric systems only) (Optional) Display boot time messages on the fabric control Routing Engine or fabric manager Routing engines. interconnect-device <i>name</i> —(QFabric systems only) (Optional) Display boot time messages on the Interconnect device. lcc <i>number</i> —(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display boot time messages for a specific T640 router connected to a TX Matrix

router. On a TX Matrix Plus router, display boot time messages for a specific T1600 router connected to a TX Matrix Plus router. Replace *number* with a value from 0 through 3.

local—(EX4200 switches and MX Series routers only) (Optional) Display boot time messages on the local Virtual Chassis member.

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

node-group *name*—(QFabric systems only) (Optional) Display boot time messages on the Node group.

scc—(TX Matrix routers only) (Optional) Display boot time messages for the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Display boot time messages for the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

Additional Information By default, when you issue the **show system boot-messages** command on a TX Matrix or TX Matrix Plus master Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on a TX Matrix router) master Routing Engines or T1600 (in a routing matrix based on a TX Matrix Plus router) master Routing Engines connected to it. Likewise, if you issue the same command on the TX Matrix or TX Matrix Plus backup Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on a TX Matrix router) backup Routing Engines or T1600 (routing matrix based on a TX Matrix Plus router) backup Routing Engines that are connected to it.

Required Privilege Level view

List of Sample Output [show system boot-messages \(TX Matrix Router\) on page 505](#)
[show system boot-messages lcc \(TX Matrix Router\) on page 507](#)
[show system boot-messages \(TX Matrix Plus Router\) on page 507](#)
[show system boot-messages \(QFX3500 Switch\) on page 508](#)

Sample Output

show system boot-messages (TX Matrix Router)

```
user@host> show system boot-messages
Copyright (c) 1992-1998 FreeBSD Inc.
Copyright (c) 1996-2000 Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1982, 1986, 1989, 1991, 1993
    The Regents of the University of California. All rights reserved.

JUNOS 4.1-20000216-Zf8469 #0: 2000-02-16 12:57:28 UTC
    tlim@single.juniper.net:/p/build/20000216-0905/4.1/release_kernel/sys/compile/GENERIC
CPU: Pentium Pro (332.55-MHz 686-class CPU)
    Origin = "GenuineIntel" Id = 0x66a Stepping=10
```

```

Features=0x183f9ff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,SEP,MTRR,PGE,MCA,CMOV,<b
16>,<b17>,MMX,<b24>>
Teknor CPU Card Recognized
real memory = 805306368 (786432K bytes)
avail memory = 786280448 (767852K bytes)
Probing for devices on PCI bus 0:
chip0 <generic PCI bridge (vendor=8086 device=7192 subclass=0)> rev 3 class 6000
0 on pci0:0:0
chip1 <Intel 82371AB PCI-ISA bridge> rev 1 class 60100 on pci0:7:0
chip2 <Intel 82371AB IDE interface> rev 1 class 10180 on pci0:7:1
chip3 <Intel 82371AB USB interface> rev 1 class c0300 int d irq 11 on pci0:7:2
smb0 <Intel 82371AB SMB controller> rev 1 class 68000 on pci0:7:3
pcic0 <TI PCI-1131 PCI-CardBus Bridge> rev 1 class 60700 int a irq 15 on pci0:13
:0
TI1131 PCI Config Reg: [pci only][FUNC0 pci int]
pcic1 <TI PCI-1131 PCI-CardBus Bridge> rev 1 class 60700 int b irq 12 on pci0:13
:1
TI1131 PCI Config Reg: [pci only][FUNC1 pci int]
fxp0 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 12 on

pci0:16:0
chip4 <generic PCI bridge (vendor=1011 device=0022 subclass=4)> rev 4 class 6040
0 on pci0:17:0
fxp1 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 10 on

pci0:19:0
Probing for devices on PCI bus 1:
mcs0 <Miscellaneous Control Subsystem> rev 12 class ff0000 int a irq 12 on pci1:
13:0
fxp2 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 10 on

pci1:14:0
Probing for devices on the ISA bus:
sc0 at 0x60-0x6f irq 1 on motherboard
sc0: EGA color <16 virtual consoles, flags=0x0>
ed0 not found at 0x300
ed1 not found at 0x280
ed2 not found at 0x340
psm0 not found at 0x60
sio0 at 0x3f8-0x3ff irq 4 flags 0x20010 on isa
sio0: type 16550A, console
sio1 at 0x3e8-0x3ef irq 5 flags 0x20000 on isa
sio1: type 16550A
sio2 at 0x2f8-0x2ff irq 3 flags 0x20000 on isa
sio2: type 16550A
pcic0 at 0x3e0-0x3e1 on isa
PC-Card ctrlr(0) TI PCI-1131 [CardBus bridge mode] (5 mem & 2 I/O windows)
pcic0: slot 0 controller I/O address 0x3e0
npx0 flags 0x1 on motherboard
npx0: INT 16 interface
fdc0: direction bit not set
fdc0: cmd 3 failed at out byte 1 of 3
fdc0 not found at 0x3f0
wdc0 at 0x1f0-0x1f7 irq 14 on isa
wdc0: unit 0 (wd0): <SunDisk SQFXB-80>, single-sector-i/o
wd0: 76MB (156672 sectors), 612 cyls, 8 heads, 32 S/T, 512 B/S
wdc0: unit 1 (wd1): <IBM-DCXA-210000>
wd1: 8063MB (16514064 sectors), 16383 cyls, 16 heads, 63 S/T, 512 B/S
wdc1 not found at 0x170
wdc2 not found at 0x180
ep0 not found at 0x300

```



```

fxp0: Ethernet address 00:a0:a5:12:05:5a
fxp1: Ethernet address 00:a0:a5:12:05:59
fxp2: Ethernet address 02:00:00:00:00:01
swapon: adding /dev/wd1s1b as swap device
Automatic reboot in progress...
/dev/rwd0s1a: clean, 16599 free (95 frags, 2063 blocks, 0.1% fragmentation)
/dev/rwd0s1e: clean, 9233 free (9 frags, 1153 blocks, 0.1% fragmentation)
/dev/rwd0s1a: clean, 16599 free (95 frags, 2063 blocks, 0.1% fragmentation)
/dev/rwd1s1f: clean, 4301055 free (335 frags, 537590 blocks, 0.0% fragmentation)

```

show system boot-messages lcc (TX Matrix Router)

```

user@host> show system boot-messages lcc 2
lcc2-re0:
-----
Copyright (c) 1996-2001, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2001 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
JUNOS 7.0-20040912.0 #0: 2004-09-12 09:16:32 UTC

builder@benten.juniper.net:/build/benten-b/7.0/20040912.0/obj-i386/sys/compile/JUNIPER
Timecounter "i8254" frequency 1193182 Hz
Timecounter "TSC" frequency 601368936 Hz
CPU: Pentium III/Pentium III Xeon/Celeron (601.37-MHz 686-class CPU)
    Origin = "GenuineIntel" Id = 0x68a Stepping = 10

Features=0x387f9ff<FPU,WE,DE,PSE,TSC,MSR,PAE,MCE,CX8,SEP,MTRR,PGE,MCA,CMOV,PAT,PSE36,PN,MMX,FXSR,SSE>
real memory = 2147467264 (2097136K bytes)
sio0: gdb debugging port
avail memory = 2084040704 (2035196K bytes)
Preloaded elf kernel "kernel" at 0xc06d9000.
DEVFS: ready for devices
Pentium Pro MTRR support enabled
md0: Malloc disk
DRAM Data Integrity Mode: ECC Mode with h/w scrubbing
npx0: <math processor> on motherboard
npx0: INT 16 interface
pcib0: <ServerWorks NB6635 3.0LE host to PCI bridge> on motherboard
pci0: <PCI bus> on pcib0
pcic-pci0: <TI PCI-1410 PCI-CardBus Bridge> irq 15 at device 1.0 on pci0
pcic-pci0: TI12XX PCI Config Reg: [pwr save][pci only]
fxp0: <Intel Embedded 10/100 Ethernet> port 0x1000-0x103f mem
0xfb800000-0xfb81ffff,0xfb820000-0xfb820fff irq 9 at device 3.0 on pci0
fxp1: <Intel Embedded 10/100 Ethernet> port 0x1040-0x107f mem
0xfb840000-0xfb85ffff,0xfb821000-0xfb821fff irq 11 at device 4.0 on pci0
...

```

show system boot-messages (TX Matrix Plus Router)

```

user@host> show system boot-messages
sfc0-re0:
-----
Copyright (c) 1996-2009, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
JUNOS 9.6B3.3 #0: 2009-06-17 19:52:08 UTC

```

```

builder@lanath.juniper.net:/volume/build/junos/9.6/release/9.6B3.3/obj-i386/bsd/sys/compile/JUNIPER
MPTable: Timecounter "i8254" frequency 1193182 Hz quality 0 CPU: Intel(R) Xeon(R)
CPU          L5238 @ 2.66GHz (2660.01-MHz 686-class CPU)   Origin =
"GenuineIntel" Id = 0x1067a Stepping = 10   Features=0xbfebfbff
...
lcc1-re0:
-----
Copyright (c) 1996-2009, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
The Regents of the University of California. All rights reserved.
JUNOS 9.6-20090617.0 #0: 2009-06-17 04:15:14 UTC

builder@lanath.juniper.net:/volume/build/junos/9.6/production/20090617.0/obj-i386/bsd/sys/compile/JUNIPER
Timecounter "i8254" frequency 1193182 Hz quality 0
CPU: Intel(R) Xeon(R) CPU @ 1.86GHz (1862.01-MHz 686-class CPU)

Origin = "GenuineIntel" Id = 0x1067a Stepping = 10
Features=0xbfebfbff
...

```

show system boot-messages (QFX3500 Switch)

```

user@switch> show sytem boot-messages
getmemsize: msgbufp[size=32768] = 0x81d07fe4

System physical memory distribution:
-----
Total physical memory: 4160749568 (3968 MB)
Physical memory used: 3472883712 (3312 MB)
Physical memory allocated to kernel: 2130706432 (2032 MB)
Physical memory allocated to user BTLB: 1342177280 (1280 MB)
-----

Copyright (c) 1996-2010, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
The Regents of the University of California. All rights reserved.
JUNOS 11.1I #0: 2010-09-17 19:18:07 UTC

ssiano@svl-junos-pool125.juniper.net:/c/ssiano/DEV_QFX_SI_BRANCH/03/20100917.399988/
obj-xlr/bsd/sys/compile/JUNIPER-DCTOR
WARNING: debug.mpsafenet forced to 0 as ipsec requires Giant
JUNOS 11.1I #0: 2010-09-17 19:18:07 UTC

ssiano@svl-junos-pool125.juniper.net:/c/ssiano/DEV_QFX_SI_BRANCH/03/20100917.399988/
obj-xlr/bsd/sys/compile/JUNIPER-DCTOR
real memory = 3472883712 (3312MB)
avail memory = 1708171264 (1629MB)
cpuid: 0, btlb_cpumap:0xffffffff8
FreeBSD/SMP: Multiprocessor System Detected: 12 CPUs
ETHERNET SOCKET BRIDGE initialising
Initializing QFX platform properties ..
cpu0 on motherboard
: RMI's XLR CPU Rev. 0.3 with no FPU implemented
L1 Cache: I size 32kb(32 line), D size 32kb(32 line), eight way.
L2 Cache: Size 1024kb, eight way
pic_lbus0: <XLR Local Bus>
pic_lbus0: <XLR Local Bus> on motherboard

```

```

Enter qfx control ethernet probe addr:0xc5eeec00
gmac4: <XLR GMAC GE Ethernet> on pic_lbus0
me0: Ethernet address 00:1d:b5:f7:68:40
Enter qfx control ethernet probe addr:0xc5eeeb40
gmac5: <XLR GMAC GE Ethernet> on pic_lbus0
me1: Ethernet address 00:1d:b5:f7:68:41
Enter qfx control ethernet probe addr:0xc5eeea80
gmac6: <XLR GMAC GE Ethernet> on pic_lbus0
me1: Ethernet address 00:1d:b5:f7:68:42
sio0 on pic_lbus0
Entering sioattach
sio0: type 16550A, console
xls_setup_intr: skip irq 3, xlr regs are set up somewhere else.
gblmem0 on pic_lbus0
ehci0: <RMI XLS USB 2.0 controller> on pic_lbus0
ehci_bus_attach: allocated resource. tag=1, base=bef24000
xls_ehci_init: endian hardware swapping NOT enabled.
usb0: EHCI version 1.0
usb0 on ehci0
usb0: USB revision 2.0
uhub0: vendor 0x0000 EHCI root hub, class 9/0, rev 2.00/1.00, addr 1
uhub0: 2 ports with 2 removable, self powered
umass0: USB USBFlashDrive, rev 2.00/11.00, addr 2
pcib0: PCIe link 0 up
pcib0: PCIe link 2 up
pcib0: PCIe link 3 up
pcib0: <XLS PCI Host Controller> on pic_lbus0
pci0: <PCI bus> on pcib0
pcib1: <PCI-PCI bridge> at device 0.0 on pci0
pci1: <PCI bus> on pcib1
pci1: <network, ethernet> at device 0.0 (no driver attached)
pcib2: <PCI-PCI bridge> at device 1.0 on pci0
pcib3: <PCI-PCI bridge> at device 2.0 on pci0
pci2: <PCI bus> on pcib3
pci2: <network, ethernet> at device 0.0 (no driver attached)
pcib4: <PCI-PCI bridge> at device 3.0 on pci0
pci3: <PCI bus> on pcib4
pci3: <network, ethernet> at device 0.0 (no driver attached)
cfi device address space at 0xbc000000
cfi0: <AMD/Fujitsu - 8MB> on pic_lbus0
cfi device address space at 0xbc000000
i2c0: <I2C bus controller> on pic_lbus0
i2c1: <I2C bus controller> on pic_lbus0
qfx_fmn0 on pic_lbus0
pool offset 1503776768
xlr_lbus0: <XLR Local Bus Controller> on motherboard
qfx_bcpld_probe[124]
qfx_bcpld_probe[138]: dev_type=0x0
qfx_bcpld_probe[124]
qfx_bcpld0: QFX BCPLD probe success
qfx_bcpld0qfx_bcpld_attach[174]
qfx_bcpld_attach[207] : bus_space_tag=0x0, bus_space_handle=0xbd900000
qfx_bcpld_probe[124]
qfx_bcpld1: QFX BCPLD probe success
qfx_bcpld1qfx_bcpld_attach[174]
tor_bcpld_slave_attach[1245] : bus_space_tag=0x0, bus_space_handle=0xbda00000
Initializing product: 96 ..
bmeb: bmeb_lib_init done 0xc60a5000, addr 0x809c99a0
bme0:Virtual BME driver initializing
Timecounter "mips" frequency 1200000000 Hz quality 0
Timecounter "xlr_pic_timer" frequency 66666666 Hz quality 1

```

```
Timecounters tick every 1.000 msec
Loading the NETPFE fc module
IPsec: Initialized Security Association Processing.
SMP: AP CPU #3 Launched!
SMP: AP CPU #1 Launched!
SMP: AP CPU #2 Launched!
SMP: AP CPU #4 Launched!
SMP: AP CPU #5 Launched!
SMP: AP CPU #7 Launched!
SMP: AP CPU #6 Launched!
SMP: AP CPU #11 Launched!
SMP: AP CPU #10 Launched!
SMP: AP CPU #9 Launched!
SMP: AP CPU #8 Launched!
da0 at umass-sim0 bus 0 target 0 lun 0
da0: <USB USBFlashDrive 1100> Removable Direct Access SCSI-0 device
da0: 40.000MB/s transfers
da0: 3920MB (8028160 512 byte sectors: 255H 63S/T 499C)
Trying to mount root from ufs:/dev/da0s1a
```

show system license

Syntax	show system license <installed keys usage>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display licenses and information about how they are used.
Options	<p>none—Display all license information.</p> <p>installed—(Optional) Display installed licenses only.</p> <p>keys—(Optional) Display a list of license keys. Use this information to verify that each expected license key is present.</p> <p>usage—(Optional) Display the state of licensed features.</p>
Required Privilege Level	maintenance
List of Sample Output	show system license on page 512 show system license installed on page 512 show system license keys on page 513 show system license usage on page 513 show system license (QFX Series) on page 513
Output Fields	Table 50 on page 511 lists the output fields for the show system license command. Output fields are listed in the approximate order in which they appear.

Table 50: show system license Output Fields

Field Name	Field Description
Feature name	Name assigned to the configured feature. You use this information to verify that all the features for which you installed licenses are present.
Licenses used	<p>Number of licenses used by a router or switch. You use this information to verify that the number of licenses used matches the number configured. If a licensed feature is configured, the feature is considered used.</p> <p>NOTE: In Junos OS Release 10.1 and later, the Licenses used column displays the actual usage count based on the number of active sessions or connections as reported by the corresponding feature daemons. This is applicable for scalable license-based features such as Subscriber Access (scale-subscriber), L2TP (scale-l2tp), Mobile IP (scale-mobile-ip), and so on.</p>

Table 50: show system license Output Fields (*continued*)

Field Name	Field Description
Licenses installed	<p>Information about the installed license key:</p> <ul style="list-style-type: none"> License identifier—Identifier associated with a license key. State—State of the license key: valid or invalid. An invalid state indicates that the key was entered incorrectly or is not valid for the specific device. License version—Version of a license. The version indicates how the license is validated, the type of signature, and the signer of the license key. Valid for device—Device that can use a license key. Group defined—Group membership of a device. Features—Feature associated with a license, such as data link switching (DLSw).
Licenses needed	Number of licenses required for features being used but not yet properly licensed.
Expiry	Amount of time left within the grace period before a license is required for a feature being used.

Sample Output

show system license

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

License usage:

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
subscriber-accounting	2	2	0	permanent
subscriber-authentication	1	2	0	permanent
subscriber-address-assignment	2	2	0	permanent
subscriber-vlan	2	2	0	permanent
subscriber-ip	0	2	0	permanent
scale-subscriber	2	3	0	permanent
scale-l2tp	4	5	0	permanent
scale-mobile-ip	1	2	0	permanent

Licenses installed:

License identifier: XXXXXXXXXX

License version: 2

Features:

```
subscriber-accounting - Per Subscriber Radius Accounting
permanent
subscriber-authentication - Per Subscriber Radius Authentication
permanent
subscriber-address-assignment - Radius/SRC Address Pool Assignment
permanent
subscriber-vlan - Dynamic Auto-sensed Vlan
permanent
subscriber-ip - Dynamic and Static IP
permanent
```

show system license installed

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

License identifier: XXXXXXXXXX

License version: 2

Features:

```

subscriber-accounting - Per Subscriber Radius Accounting
    permanent
subscriber-authentication - Per Subscriber Radius Authentication
    permanent
subscriber-address-assignment - Radius/SRC Address Pool Assignment
    permanent
subscriber-vlan - Dynamic Auto-sensed Vlan
    permanent
subscriber-ip - Dynamic and Static IP
    permanent

```

show system license keys

```

user@host> show system license keys
XXXXXXXXXX xxxxxxx xxxxxxx xxxxxxx xxxxxxx xxxxxxx
          xxxxxxx xxxxxxx xxxxxxx xxxxxxx xxxxxxx
          xxxxxxx xxxxxxx xxx

```

show system license usage

```

user@host> show system license usage
License usage:

```

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
subscriber-accounting	2	2	0	permanent
subscriber-authentication	1	2	0	permanent
subscriber-address-assignment	2	2	0	permanent
subscriber-vlan	2	2	0	permanent
subscriber-ip	0	2	0	permanent
scale-subscriber	2	3	0	permanent
scale-l2tp	4	5	0	permanent
scale-mobile-ip	1	2	0	permanent

show system license (QFX Series)

```

user@switch> show system license
License usage:

```


Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
qfx-edge-fab	1	1	1	permanent

```

Licenses installed:
License identifier: JUNOS417988
License version: 1
Features:
  qfx-edge-fab - QFX3000 Series QF/Node feature license
                permanent

```

show system snapshot

Syntax	show system snapshot
Syntax (EX Series Switches)	show system snapshot <all-members local member <i>member-id</i> > <media (external internal)>
Release Information	Command introduced in Junos OS Release 7.6. Command introduced in Junos OS Release 10.0 for EX Series switches.
Description	<p>Display information about the backup software:</p> <ul style="list-style-type: none"> On the routers, display information about the backup software, which is located in the /altroot, and /altconfig file systems or on the alternate media. On the switches, display information about the backup of the root file system (/) and directories /altroot, /config, /var, and /var/tmp, which are located either on an external USB flash drive or in internal flash memory.
	<div>  <p>NOTE: To back up software, use the request system snapshot command.</p> </div>
Options	<p>none—Display information about the backup software.</p> <p>all-members local member <i>member-id</i>—(EX Series switch Virtual Chassis only) (Optional) Display the snapshot in a Virtual Chassis:</p> <ul style="list-style-type: none"> all-members—Display the snapshot for all members of the Virtual Chassis. local—Display the snapshot on the member of the Virtual Chassis that you are currently logged into. member <i>member-id</i>—Display the snapshot for the specified member of the Virtual Chassis. <p>media (external internal)—(EX Series switch only) (Optional) Display the destination media location for the snapshot. The external option specifies the snapshot on an external mass storage device, such as a USB flash drive. The internal option specifies the snapshot on an internal memory source, such as internal flash memory. If no additional options are specified, the command displays the snapshot stored in both slices.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> request system snapshot on page 474
List of Sample Output	show system snapshot (Router) on page 515

[show system snapshot media external \(Switch\) on page 515](#)

[show system snapshot media internal \(Switch\) on page 515](#)

Output Fields [Table 51 on page 515](#) lists the output fields for the **show system snapshot** command. Output fields are listed in the approximate order in which they appear.

Table 51: show system snapshot Output Fields

Field Name	Field Description
Creation date	Date and time of the last snapshot.
JUNOS version on snapshot	Junos OS release number of individual software packages.

Sample Output

show system snapshot (Router)

```
user@host> show system snapshot
Information for snapshot on hard-disk
Creation date: Oct 5 13:53:29 2005
JUNOS version on snapshot:
  jbase   : 7.3R2.5
  jcrypto: 7.3R2.5
  jdocs   : 7.3R2.5
  jkernel: 7.3R2.5
  jpfe    : M40-7.3R2.5
  jroute  : 7.3R2.5
```

show system snapshot media external (Switch)

```
user@switch> show system snapshot media external
Information for snapshot on      external (/dev/da1s1a) (backup)
Creation date: Mar 19 03:37:18 2012
JUNOS version on snapshot:
  jbase   : ex-12.1I20120111_0048_user
  jcrypto-ex: 12.1I20120111_0048_user
  jdocs-ex: 12.1I20120111_0048_user
  jroute-ex: 12.1I20120111_0048_user
  jswitch-ex: 12.1I20120111_0048_user
  jweb-ex: 12.1I20120111_0048_user
Information for snapshot on      external (/dev/da1s2a) (primary)
Creation date: Mar 19 03:38:25 2012
JUNOS version on snapshot:
  jbase   : ex-12.2I20120305_2240_user
  jcrypto-ex: 12.2I20120305_2240_user
  jdocs-ex: 12.2I20120305_2240_user
  jroute-ex: 12.2I20120305_2240_user
  jswitch-ex: 12.2I20120305_2240_user
  jweb-ex: 12.2I20120305_2240_user
```

show system snapshot media internal (Switch)

```
user@switch> show system snapshot media internal
Information for snapshot on internal (/dev/da0s1a) (backup)
Creation date: Mar 14 05:01:02 2011
JUNOS version on snapshot:
  jbase   : 11.1R1.9
```

```
jcrypto-ex: 11.1R1.9
jdocs-ex: 11.1R1.9
jkernel-ex: 11.1R1.9
jroute-ex: 11.1R1.9
jswitch-ex: 11.1R1.9
jweb-ex: 11.1R1.9
jpfe-ex42x: 11.1R1.9
Information for snapshot on internal (/dev/da0s2a) (primary)
Creation date: Mar 30 08:46:27 2011
JUNOS version on snapshot:
jbase : 11.2-20110330.0
jcrypto-ex: 11.2-20110330.0
jdocs-ex: 11.2-20110330.0
jkernel-ex: 11.2-20110330.0
jroute-ex: 11.2-20110330.0
jswitch-ex: 11.2-20110330.0
jweb-ex: 11.2-20110330.0
jpfe-ex42x: 11.2-20110330.0
```

Operational Commands: System Monitoring

show chassis alarms

Syntax	show chassis alarms
Syntax (TX Matrix Routers)	show chassis alarms <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Routers)	show chassis alarms <lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Routers)	show chassis alarms <all-members> <local> <member <i>member-id</i> >
Syntax (MX2010 3D Universal Edge Routers)	show chassis alarms
Syntax (MX2020 3D Universal Edge Routers)	show chassis alarms
Syntax (QFX Series)	show chassis alarms <interconnect-device <i>name</i> > <node-device <i>name</i> >
Syntax (PTX Series Packet Transport Switches)	show chassis alarms
Syntax (ACX Series Universal Access Routers)	show chassis alarms
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option for the TX Matrix Plus router introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Switches.</p> <p>Command introduced in Junos OS Release 12.2 for the ACX Series Universal Access Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>
Description	Display information about the conditions that have been configured to trigger alarms.
Options	none —Display information about the conditions that have been configured to trigger alarms.

all-members—(MX Series routers only) (Optional) Display information about alarm conditions for all the member routers of the Virtual Chassis configuration.

interconnect-device *name*—(QFabric systems only) (Optional) Display information about alarm conditions for the Interconnect device.

lcc *number* — (TX Matrix and TX Matrix Plus routers only) (Optional) On the TX Matrix router, show information about a specified T640 router (or line-card chassis) that is connected to the TX Matrix router. On the TX Matrix Plus router, show information about a specified T1600 router (or line-card chassis) that is connected to the TX Matrix Plus router. Replace ***number*** with a value from 0 through 3.

local—(MX Series routers only) (Optional) Display information about alarm conditions for the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display information about alarm conditions for the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value of 0 or 1.

node-device *name*—(QFabric systems only) (Optional) Display information about alarm conditions for the Node device.

scc—(TX Matrix router only) (Optional) Show information about the TX Matrix router (or switch-card chassis).

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

Additional Information You cannot clear the alarms for chassis components. Instead, you must remedy the cause of the alarm. When a chassis alarm is lit, it indicates that you are running the router or switch in a manner that we do not recommend.

On routers, you can manually silence external devices connected to the alarm relay contacts by pressing the alarm cutoff button, located on the craft interface. Silencing the device does not remove the alarm messages from the display (if present on the router) or extinguish the alarm LEDs. In addition, new alarms that occur after you silence an external device reactivate the external device.

In Junos OS release 11.1 and later, alarms for fans also show the slot number of the fans in the CLI output.

In Junos OS Release 11.2 and later, the command output on EX8200 switches shows the detailed location (**Plane/FPC/PFE**) for link errors in the chassis.

In Junos OS Release 10.2 and later, an alarm is shown on T Series routers for a standby sonic clock generator (SCG) that is offline or absent.

You may often see the following error messages, in which only the error code is shown and no other information is provided:

```
Apr 12 08:04:10 send: red alarm set, device FPC 6, reason FPC 6 Major Errors - Error code: 257
```

Apr 12 08:04:19 send: red alarm set, device FPC 1, reason FPC 1 Major Errors - Error code: 559

To understand what CM_ALARM error codes mean, you need to first identify the structure of the CM_ALARM code. A CM_ALARM code has the following structure:

Bits:	Error type:
1-31	Major (1)
0	Minor (0)

As per the above table, the LSB (bit 0) identifies the **Error Type** (major alarm, if the bit is set and minor alarm if the bit is unset). The rest of the bits (1 - 31) identify the actual error code.

Take an example of the following error code, which was logged on a T1600:

Apr 12 08:04:10 send: red alarm set, device FPC 1, reason FPC 1 Major Errors - Error code: 559

First, you have to convert 559 to binary; that is **100010111**. The LSB in this case is 1, which means that this is a major alarm. After removing the LSB, you are left with **10001011**, which is equal to 279 in decimal. This is the actual error code, its meaning can be found from the following list:

Chip Type: L Chip	Code
CMALARM_LCHIP_LOUT_DESRD_PARITY_ERR	1
CMALARM_LCHIP_LOUT_DESRD_UNINIT_ERR	2
CMALARM_LCHIP_LOUT_DESRD_ILLEGALLINK_ERR	3
CMALARM_LCHIP_LOUT_DESRD_ILLEGALSIZERR	4
CMALARM_LCHIP_LOUT_HDRF_TOERR_ERR	5
CMALARM_LCHIP_LOUT_HDRF_PARITY_ERR	6
CMALARM_LCHIP_LOUT_HDRF_UCERR_ERR	7
CMALARM_LCHIP_LOUT_NLIF_CRCDROP_ERR	8
CMALARM_LCHIP_LOUT_NLIF_CRCERR_ERR	9
CMALARM_LCHIP_UCODE_TIMEOUT_ERR	10
CMALARM_LCHIP_LIN_SRCTL_ACCT_DROP_ERR	11
CMALARM_LCHIP_LIN_SRCTL_ACCT_ADDR_SIZE_ERR	12

CMALARM_LCHIP_SRAM_PARITY_ERR	13
CMALARM_LCHIP_UCODE_OVFLW_ERR	14
CMALARM_LCHIP_LOUT_HDRF_MTU_ERR	15
<hr/>	
Chip Type: M Chip	Code
CMALARM_MCHIP_ECC_UNCORRECT_ERR	128
<hr/>	
Chip Type: N Chip	Code
CMALARM_NCHIP_RDDMA_JBUS_TIMEOUT_ERR	256
CMALARM_NCHIP_RDDMA_FIFO_OVFLW_ERR	257
CMALARM_NCHIP_RDDMA_FIFO_UNFLW_ERR	258
CMALARM_NCHIP_RDDMA_SIZE_ERR	259
CMALARM_NCHIP_RDDMA_JBUS_CRC_ERR	260
CMALARM_NCHIP_WRDMA_PKTR_ERR	261
CMALARM_NCHIP_WRDMA_PKT_CRC_ERR	262
CMALARM_NCHIP_WRDMA_JBUS_TIMEOUT_ERR	263
CMALARM_NCHIP_WRDMA_FIFO_OVFLW_ERR	264
CMALARM_NCHIP_WRDMA_FIFO_UNFLW_ERR	265
CMALARM_NCHIP_WRDMA_PKT_LEN_ERR	266
CMALARM_NCHIP_WRDMA_JBUS_CRC_ERR	267
CMALARM_NCHIP_PKTR_DMA_AGE_ERR	268
CMALARM_NCHIP_PKTR_ICELLSIG_ERR	269
CMALARM_NCHIP_PKTR_FTTL_ERR	270
CMALARM_NCHIP_RODR_OFFSET_OVFLW_ERR	271
CMALARM_NCHIP_PKTR_TMO_CELL_ERR	272
CMALARM_NCHIP_PKTR_TMO_OUTRANGE_ERR	273
CMALARM_NCHIP_PKTR_MD_REQUEST_Q_OVFLW_ERR	274

CMALARM_NCHIP_PKTR_DMA_BUFFER_OVFLW_ERR	275
CMALARM_NCHIP_PKTR_GRT_OVFLW_ERR	276
CMALARM_NCHIP_FRQ_ERR	277
CMALARM_NCHIP_RODR_IN_Q_OVFLW_ERR	278
CMALARM_NCHIP_DBUF_CRC_ERR	279
<hr/>	
Chip Type: R Chip	Code
CMALARM_RCHIP_SRAM_PARITY_ERR	512
<hr/>	
Chip Type: R Chip	Code
CMALARM_ICHIP_WO_DESRD_ID_ERR	601
CMALARM_ICHIP_WO_DESRD_DATA_ERR	602
CMALARM_ICHIP_WO_DESRD_OFLOW_ERR	603
CMALARM_ICHIP_WO_HDRF_UCERR_ERR	604
CMALARM_ICHIP_WO_HDRF_MTUERR_ERR	605
CMALARM_ICHIP_WO_HDRF_PARITY_ERR	606
CMALARM_ICHIP_WO_HDRF_TOERR_ERR	607
CMALARM_ICHIP_WO_IP_CRC_ERR	608
CMALARM_ICHIP_WO_IP_INTER_ERR	609
CMALARM_ICHIP_WI_WAN_TIMEOUT_ERR	625
CMALARM_ICHIP_WI_FAB_TIMEOUT_ERR	626
CMALARM_ICHIP_RLDRAM_BIST_ERR	630
CMALARM_ICHIP_SDRAM_BIST_ERR	631
CMALARM_ICHIP_RLDRAM_PARITY_ERR	632
CMALARM_ICHIP_SDRAM_UNCORRECT_ERR	633
CMALARM_ICHIP_SDRAM_CORRECT_ERR	634
CMALARM_ICHIP_FUSE_DONE_ERR	635

According to the table above, the **279** error code corresponds to **CMALARM_NCHIP_DBUF_CRC_ERR**; this means that new CRC errors were seen on the NCHIP of this particular FPC, which is FPC as per the logs.

If you do not want to convert decimal to binary and vice-versa, you may use the following shortcut:

For major alarms, the **Actual Error Code = (Error Code - 1)/2**, where **Error Code** is the code that you get in the log message. For example, if you get the following log:

Apr 12 08:04:10 send: red alarm set, device FPC 6, reason FPC 6 Major Errors - Error code: 257

Actual Error Code = $(257-1)/2 = 128$. Similarly, for minor alarms, Actual Error Code = $(\text{Error Code})/2$

Required Privilege Level

view

Related Documentation

- *Configuring an Alarm Entry and Its Attributes*
- *Chassis Conditions That Trigger Alarms*

List of Sample Output

[show chassis alarms \(Alarms Active\) on page 523](#)
[show chassis alarms \(No Alarms Active\) on page 523](#)
[show chassis alarms \(Fan Tray\) on page 523](#)
[show chassis alarms \(MX2020 Router\) on page 523](#)
[show chassis alarms \(MX2010 Router\) on page 523](#)
[show chassis alarms \(T4000 Router\) on page 523](#)
[show chassis alarms \(Unreachable Destinations Present on a T Series Router\) on page 524](#)
[show chassis alarms \(FPC Offline Due to Unreachable Destinations on a T Series Router\) on page 524](#)
[show chassis alarms \(SCG Absent on a T Series Router\) on page 524](#)
[show chassis alarms \(Alarms Active on a TX Matrix Router\) on page 524](#)
[show chassis alarms \(Alarms on a T4000 Router After the enhanced-mode Statement is Enabled\) on page 525](#)
[show chassis alarms \(Backup Routing Engine\) on page 525](#)
[show chassis alarms \(Alarms Active on the QFX Series\) on page 525](#)
[show chassis alarms node-device \(Alarms Active on the QFabric System\) on page 525](#)
[show chassis alarms \(Alarms Active on the QFabric System\) on page 526](#)
[show chassis alarms \(Alarms Active on an EX8200 Switch\) on page 526](#)
[show chassis alarms \(Alarms Active on a PTX5000 Packet Transport Switch\) on page 526](#)
[show chassis alarms \(Alarms Active on an ACX2000 Universal Access Router\) on page 527](#)

Output Fields

[Table 52 on page 523](#) lists the output fields for the **show chassis alarms** command. Output fields are listed in the approximate order in which they appear.

Table 52: show chassis alarms Output Fields

Field Name	Field Description
Alarm time	Date and time the alarm was first recorded.
Class	Severity class for this alarm: Minor or Major .
Description	Information about the alarm.

Sample Output

show chassis alarms (Alarms Active)

```

user@host> show chassis alarms
3 alarms are currently active
Alarm time      Class  Description
2000-02-07 10:12:22 UTC Major fxp0: ethernet link down
2000-02-07 10:11:54 UTC Minor YELLOW ALARM - PEM 1 Removed
2000-02-07 10:11:03 UTC Minor YELLOW ALARM - Lower Fan Tray Removed

```

show chassis alarms (No Alarms Active)

```

user@host> show chassis alarms
No alarms are currently active

```

show chassis alarms (Fan Tray)

```

user@host> show chassis alarms
4 alarms currently active
Alarm time      Class  Description
2010-11-11 20:27:38 UTC Major Side Fan Tray 7 Failure
2010-11-11 20:27:13 UTC Minor Side Fan Tray 7 Overspeed
2010-11-11 20:27:13 UTC Major Side Fan Tray 5 Failure
2010-11-11 20:27:13 UTC Major Side Fan Tray 0 Failure

```

show chassis alarms (MX2020 Router)

```

user@host> show chassis alarms
1 alarms currently active
Alarm time Class Description
2012-10-03 12:14:59 PDT Minor Plane 0 not online

```

show chassis alarms (MX2010 Router)

```

user@host> show chassis alarms
7 alarms currently active
Alarm time      Class  Description
2012-08-07 00:46:06 PDT Major Fan Tray 2 Failure
2012-08-06 18:24:36 PDT Minor Redundant feed missing for PSM 6
2012-08-06 07:41:04 PDT Minor Redundant feed missing for PSM 8
2012-08-04 02:42:06 PDT Minor Redundant feed missing for PSM 5
2012-08-03 21:14:24 PDT Minor Loss of communication with Backup RE
2012-08-03 12:26:03 PDT Minor Redundant feed missing for PSM 4
2012-08-03 10:40:18 PDT Minor Redundant feed missing for PSM 7

```

show chassis alarms (T4000 Router)

```

user@host> show chassis alarms

```

```

9 alarms currently active
Alarm time      Class Description
2007-06-02 01:41:10 UTC Minor RE 0 Not Supported
2007-06-02 01:41:10 UTC Minor CB 0 Not Supported
2007-06-02 01:41:10 UTC Minor Mixed Master and Backup RE types
2007-05-30 19:37:33 UTC Major SPMB 1 not online
2007-05-30 19:37:29 UTC Minor Front Bottom Fan Tray Absent
2007-05-30 19:37:13 UTC Major PEM 1 Input Failure
2007-05-30 19:37:13 UTC Major PEM 0 Not OK
2007-05-30 19:37:03 UTC Major PEM 0 Improper for Platform
2007-05-30 19:37:03 UTC Minor Backup RE Active

```

show chassis alarms (Unreachable Destinations Present on a T Series Router)

```

user@host> show chassis alarms
10 alarms currently active
Alarm time      Class Description
2011-08-30 18:43:53 PDT Major FPC 7 has unreachable destinations
2011-08-30 18:43:53 PDT Major FPC 5 has unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 3 has unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 2 has unreachable destinations
2011-08-30 18:43:52 PDT Minor SIB 0 Not Online
2011-08-30 18:43:33 PDT Minor SIB 4 Not Online
2011-08-30 18:43:28 PDT Minor SIB 3 Not Online
2011-08-30 18:43:05 PDT Minor SIB 2 Not Online
2011-08-30 18:43:28 PDT Minor SIB 1 Not Online
2011-08-30 18:43:05 PDT Major PEM 1 Not Ok

```

show chassis alarms (FPC Offline Due to Unreachable Destinations on a T Series Router)

```

user@host> show chassis alarms
10 alarms currently active
Alarm time      Class Description
2011-08-30 18:43:53 PDT Major FPC 7 offline due to unreachable destinations
2011-08-30 18:43:53 PDT Major FPC 5 offline due to unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 3 offline due to unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 2 offline due to unreachable destinations
2011-08-30 18:43:52 PDT Minor SIB 0 Not Online
2011-08-30 18:43:33 PDT Minor SIB 4 Not Online
2011-08-30 18:43:28 PDT Minor SIB 3 Not Online
2011-08-30 18:43:05 PDT Minor SIB 2 Not Online
2011-08-30 18:43:28 PDT Minor SIB 1 Not Online
2011-08-30 18:43:05 PDT Major PEM 1 Not Ok

```

show chassis alarms (SCG Absent on a T Series Router)

```

user@host> show chassis alarms
4 alarms currently active
Alarm time      Class Description
2011-01-23 21:42:46 PST Major SCG 0 NO EXT CLK MEAS-BKUP SCG ABS

```

show chassis alarms (Alarms Active on a TX Matrix Router)

```

user@host> show chassis alarms
scc-re0:
-----
8 alarms currently active
Alarm time      Class Description
2004-08-05 18:43:53 PDT Minor LCC 0 Minor Errors
2004-08-05 18:43:53 PDT Minor SIB 3 Not Online
2004-08-05 18:43:52 PDT Major SIB 2 Absent
2004-08-05 18:43:52 PDT Major SIB 1 Absent

```

```

2004-08-05 18:43:52 PDT Major SIB 0 Absent
2004-08-05 18:43:33 PDT Major LCC 2 Major Errors
2004-08-05 18:43:28 PDT Major LCC 0 Major Errors
2004-08-05 18:43:05 PDT Minor LCC 2 Minor Errors
lcc0-re0:
-----

```

```

5 alarms currently active
Alarm time          Class Description
2004-08-05 18:43:53 PDT Minor SIB 3 Not Online
2004-08-05 18:43:49 PDT Major SIB 2 Absent
2004-08-05 18:43:49 PDT Major SIB 1 Absent
2004-08-05 18:43:49 PDT Major SIB 0 Absent
2004-08-05 18:43:28 PDT Major PEM 0 Not OK
lcc2-re0:
-----

```

```

5 alarms currently active
Alarm time          Class Description
2004-08-05 18:43:35 PDT Minor SIB 3 Not Online
2004-08-05 18:43:33 PDT Major SIB 2 Absent
2004-08-05 18:43:33 PDT Major SIB 1 Absent
2004-08-05 18:43:33 PDT Major SIB 0 Absent
2004-08-05 18:43:05 PDT Minor PEM 1 Absent

```

show chassis alarms (Alarms on a T4000 Router After the enhanced-mode Statement is Enabled)

On T4000 routers, when you include the **enhanced-mode** statement at the **[edit chassis network-services]** hierarchy level and reboot the system, only the T4000 Type 5 FPCs present on the router are online while the remaining FPCs are offline, and FPC misconfiguration alarms are generated. The **show chassis alarm** command output displays FPC misconfiguration (**FPC *fpc-slot* misconfig**) as the reason for the generation of the alarms.

```

user@host> show chassis alarms
2 alarms currently active
Alarm time          Class Description
2011-10-22 10:10:47 PDT Major FPC 1 misconfig
2011-10-22 10:10:46 PDT Major FPC 0 misconfig

```

show chassis alarms (Backup Routing Engine)

```

user@host> show chassis alarms
2 alarms are currently active
Alarm time          Class Description
2005-04-07 10:12:22 PDT Minor Host 1 Boot from alternate media
2005-04-07 10:11:54 PDT Major Host 1 compact-flash missing in Boot List

```

show chassis alarms (Alarms Active on the QFX Series)

```

user@switch> show chassis alarms
1 alarms currently active
Alarm time          Class Description
2012-03-05 2:10:24 UTC Major FPC 0 PEM 0 Airflow not matching Chassis Airflow

```

show chassis alarms node-device (Alarms Active on the QFabric System)

```

user@switch> show chassis alarms node-device ED3691
node-device ED3694
3 alarms currently active
Alarm time          Class Description
2011-08-24 16:04:15 UTC Major ED3694:fte-0/1/2: Link down

```

```

2011-08-24 16:04:14 UTC Major ED3694:fte-0/1/0: Link down
2011-08-24 14:21:14 UTC Major ED3694 PEM 0 is not supported/powered

```

show chassis alarms (Alarms Active on the QFabric System)

```

user@switch> show chassis alarms
IC-A0001:
-----
1 alarms currently active
Alarm time          Class  Description
2011-08-24 16:04:15 UTC Minor Backup RE Active

ED3694:
-----
3 alarms currently active
Alarm time          Class  Description
2011-08-24 16:04:15 UTC Major ED3694:fte-0/1/2: Link down
2011-08-24 16:04:14 UTC Major ED3694:fte-0/1/0: Link down
2011-08-24 14:21:14 UTC Major ED3694 PEM 0 is not supported/powered

SNG-0:
-----

NW-NG-0:
-----
1 alarms currently active
Alarm time          Class  Description
2011-08-24 15:49:27 UTC Major ED3691 PEM 0 is not supported/powered

```

show chassis alarms (Alarms Active on an EX8200 Switch)

```

user@switch> show chassis alarms

6 alarms currently active
Alarm time          Class  Description
2010-12-02 19:15:22 UTC Major Fan Tray Failure
2010-12-02 19:15:22 UTC Major Fan Tray Failure
2010-12-02 19:15:14 UTC Minor Check CB 0 Fabric Chip 1 on Plane/FPC/PFE: 1/5/0,
1/5/1, 1/5/2, 1/5/3, 1/7/0, 1/7/1, 1/7/2, 1/7/3, 2/5/0, 2/5/1, ...
2010-12-02 19:15:14 UTC Minor Check CB 0 Fabric Chip 0 on Plane/FPC/PFE: 1/5/0,
1/5/1, 1/5/2, 1/5/3, 1/7/0, 1/7/1, 1/7/2, 1/7/3, 2/5/0, 2/5/1, ...
2010-12-02 19:14:18 UTC Major PSU 1 Output Failure
2010-12-02 19:14:18 UTC Minor Loss of communication with Backup RE

```

show chassis alarms (Alarms Active on a PTX5000 Packet Transport Switch)

```

user@switch> show chassis alarms

23 alarms currently active
Alarm time          Class  Description
2011-07-12 16:22:05 PDT Minor No Redundant Power for Rear Chassis
2011-07-12 16:22:05 PDT Major PDU 0 PSM 1 Not OK
2011-07-12 16:21:57 PDT Minor No Redundant Power for Fan 0-2
2011-07-12 16:21:57 PDT Major PDU 0 PSM 0 Not OK
2011-07-12 15:56:06 PDT Major PDU 1 PSM 2 Not OK
2011-07-12 15:56:06 PDT Minor No Redundant Power for FPC 0-7
2011-07-12 15:56:06 PDT Major PDU 0 PSM 3 Not OK
2011-07-12 15:28:20 PDT Major PDU 0 PSM 2 Not OK
2011-07-12 15:19:14 PDT Minor Backup RE Active

```

show chassis alarms (Alarms Active on an ACX2000 Universal Access Router)

```
user@host> show chassis alarms
7 alarms currently active
Alarm time          Class  Description
2012-05-22 11:19:09 UTC Major  xe-0/3/1: Link down
2012-05-22 11:19:09 UTC Major  xe-0/3/0: Link down
2012-05-22 11:19:09 UTC Major  ge-0/1/7: Link down
2012-05-22 11:19:09 UTC Major  ge-0/1/6: Link down
2012-05-22 11:19:09 UTC Major  ge-0/1/3: Link down
2012-05-22 11:19:09 UTC Major  ge-0/1/2: Link down
2012-05-22 11:19:09 UTC Major  ge-0/1/1: Link down
```

show chassis environment

Syntax	show chassis environment
Syntax (T320, T640, T1600, and T4000 Routers)	show chassis environment <cb <i>cb-slot-number</i> > <fpc <i>fpc-slot-number</i> > <fpm> <pem <i>pem-slot-number</i> > <routing-engine <i>re-slot-number</i> > <scg <i>scg-slot-number</i> > <sib <i>sib-slot-number</i> >
Syntax (TX Matrix Routers)	show chassis environment <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Routers)	show chassis environment <lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Routers)	show chassis environment <all-members> <local> <member <i>member-id</i> >
Syntax (MX2020 3D Universal Edge Routers)	show chassis environment <adc <i>adc-slot-number</i> > <cb <i>cb-slot-number</i> > <fpc <i>fpc-slot-number</i> > <fpm> <monitored> <psm <i>psm-slot-number</i> > <routing-engine <i>re-slot-number</i> > <sfb <i>sfb-slot-number</i> >
Syntax (MX2010 3D Universal Edge Routers)	show chassis environment <adc <i>adc-slot-number</i> > <cb <i>cb-slot-number</i> > <fpc <i>fpc-slot-number</i> > <fpm> <monitored> <psm <i>psm-slot-number</i> > <routing-engine <i>re-slot-number</i> > <sfb <i>sfb-slot-number</i> >
Syntax (EX Series Switch)	show chassis environment <all-members> <cb <i>cb-slot-number</i> > <fpc <i>fpc-slot-number</i> > <local> <member <i>member-id</i> > <routing-engine <i>re-slot-number</i> >

Syntax (EX Series Switch)	<pre>show chassis environment <all-members> <cb <i>cb-slot-number</i>> <fpc <i>fpc-slot-number</i>> <local> <member <i>member-id</i>> <power-supply-unit <i>psu-slot-number</i>> <routing-engine <i>slot-number</i>></pre>
Syntax (QFX Series)	<pre>show chassis environment <cb <i>slot-number</i> <interconnect-device <i>name</i>>> <fpc <i>slot-number</i> <interconnect-device <i>name</i>>> <interconnect-device <i>name</i> <slot-number> <node-device <i>name</i>> <pem <i>slot-number</i> (interconnect-device <i>name</i> <i>slot-number</i>) (node-device <i>name</i>)> <routing-engine <i>name</i> <interconnect-device <i>name</i> <i>slot-number</i>>></pre>
Syntax (PTX Series Packet Transport Switches)	<pre>show chassis environment <cb <i>cb-slot-number</i>> <ccg <i>ccg-slot-number</i>> <fpc <i>fpc-slot-number</i>> <fpm> <monitored> <pdu <i>pdu-slot-number</i>> <routing-engine <i>re-slot-number</i>> <sib <i>sib-slot-number</i>></pre>
Syntax (ACX Series Universal Access Routers)	<pre>show chassis environment <cb <i>cb-slot-number</i>> <pem <i>pem-slot-number</i>> <routing-engine <i>re-slot-number</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches.</p> <p>monitored option added in Junos OS Release 12.1 for PTX Series Packet Transport Switches.</p> <p>Command introduced in Junos OS Release 12.1 for T4000 Core Routers.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p> <p>pem option introduced in Junos OS Release 12.3 for ACX4000 Universal Access Routers.</p>
Description	<p>Display environmental information about the router or switch chassis, including the temperature and information about the fans, power supplies, and Routing Engine.</p> <p>In addition on ACX4000 routers, display temperature information about the different channels of a Modular Interface Card (MIC). The number of channels displayed depends on the type of MIC installed.</p>

Options **none**—Display environmental information about the router or switch chassis. On a TX Matrix router, display environmental information about the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display environmental information about the TX Matrix Plus router and its attached T1600 routers.

all-members—(MX Series routers and EX Series switches only) (Optional) Display chassis environmental information for all the members of the Virtual Chassis configuration.

adc **adc-slot-number**—(MX2020 and MX2010 routers only) (Optional) Display chassis environmental information for the adapter cards. For MX2020 routers, replace **adc-slot-number** with a value from 0 through 19. For MX2010 routers, replace **adc-slot-number** with a value from 0 through 9.

cb **cb-slot-number**—(ACX Series Universal Access Routers, EX Series switches, M120, M320, and M40e routers, MX Series routers, MX2020 routers, MX2010 routers, PTX Series Packet Transport Switches, QFX Series, and T Series routers only) (Optional) Display chassis environmental information for the Control Board. On devices other than EX Series switches, replace **cb-slot** with 0 or 1. For the EX Series switches, see *EX Series Switches Hardware and CLI Terminology Mapping* for information on CB slot numbering.

cb interconnect-device name—(QFabric systems only) (Optional) Display chassis environmental information for the Control Board on an Interconnect device.

ccg **ccg-slot-number**—(PTX Series only) (Optional) Display chassis environmental information for the Centralized Clock Generator. Replace **cb-slot** with a value of 0 or 1.

fpc **fpc-slot**—(EX Series switches, M120, M320, and M40e routers, MX Series routers, MX2010 routers, MX2020 routers, PTX Series Packet Transport Switches, QFX Series, QFX3500 switches, QFabric systems, and T Series routers) (Optional) Display chassis environmental information for a specified Flexible PIC Concentrator. For MX2010 routers, replace **fpc-slot** with a value from 0 through 9. For MX2020 routers, replace **fpc-slot** with a value from 0 through 19. For information about FPC numbering, see [show chassis environment fpc](#). On a QFabric system, display chassis environmental information for a specified Flexible PIC Concentrator on an Interconnect device. On an EX Series switch, display chassis environmental information for a specified Flexible PIC Concentrator; see *EX Series Switches Hardware and CLI Terminology Mapping* for information on FPC numbering.

fpm—(M120, M320, and M40e routers, MX2010 routers, MX2020 routers, PTX Series, Packet Transport Switches, and T Series routers only) (Optional) Display chassis environmental information for the craft interface (FPM).

interconnect-device name—(QFabric systems only) (Optional) Display chassis environmental information for the Interconnect device.

monitored—(MX2020 routers and PTX Series Packet Transport Switches only) (Optional) Display chassis environmental information for monitored temperatures only. Temperatures that are not included in temperature alarm computations are not displayed.

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display chassis environmental information for a specified T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display chassis environmental information for a specified T1600 router (or line-card chassis) that is connected to the TX Matrix Plus router. Replace ***number*** with a value from 0 through 3.

local—(MX Series routers and EX Series switches) (Optional) Display chassis environmental information for the local Virtual Chassis member.

member *member-id*—(MX Series routers and EX Series switches only) (Optional) Display chassis environmental information for the specified member of the Virtual Chassis configuration. On MX Series routers, replace ***member-id*** with a value of 0 or 1. For EX Series switches, see *member* for member ID values.

node-device *name*—(QFabric systems only) (Optional) Display chassis environmental information for the Node device.

pdu *pdu-slot-number*—(PTX Series only) (Optional) Display chassis environmental information for the specified power distribution unit.

pem—(QFX3500 switches and QFabric systems only) (Optional) Display chassis environmental information for the Power Entry Module on the specified Interconnect device or Node device.

pem *pem-slot-number*—(ACX Series Universal Access Routers, M120, M320, and M40e routers, MX Series routers, QFX Series, and T Series routers only) (Optional) Display chassis environmental information for the Power Entry Module on the specified Power Entry Module. For information about the options, see [show chassis environment pem](#).

psm *psm-slot-number*—(MX2020 and MX2010 routers only) (Optional) Display chassis environmental information for the power supply module. For MX2020 routers, replace ***psm-slot-number*** with a value from 0 through 17. For MX2010 routers, replace ***psm-slot-number*** with a value from 0 through 8.

psu *psu-slot-number*—(EX Series switches only) (Optional) Display chassis environmental information for a specified power supply. See *EX Series Switches Hardware and CLI Terminology Mapping* for detailed information.

routing-engine—(QFX3500 switches and QFabric systems only) (Optional) Display chassis environmental information for the Routing Engine on the specified Interconnect device.

routing-engine *re-slot-number*—(Optional) Display chassis environmental information for the specified Routing Engine. For information about the options, see [show chassis environment routing-engine](#).

scg—(T Series routers only) (Optional) Display chassis environmental information about the SONET Clock Generator.

scc—(TX Matrix routers only) (Optional) Display chassis environmental information about the TX Matrix router (or switch-card chassis).

sfb *sfb-slot-number*—(MX2020 and MX2010 routers only) (Optional) Display chassis environmental information for the power supply module. Replace ***sfb-slot-number*** with a value from 0 through 7.

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

sib *sib-slot-number*—(M320 routers, PTX Series Packet Transport Switches, and T Series routers only) (Optional) Display chassis environmental information about the specified switch interface board. For information about the options, see *show chassis environment sib*.

Required Privilege Level

view

Related Documentation

- *show chassis environment adc*
- [show chassis environment cb on page 578](#)
- *show chassis environment ccg*
- *show chassis environment cip*
- [show chassis environment fpc on page 595](#)
- *show chassis environment fpm*
- *show chassis environment mcs*
- *show chassis environment monitored*
- *show chassis environment pcg*
- *show chassis environment pdu*
- [show chassis environment pem on page 620](#)
- *show chassis environment psm*
- *show chassis environment psu*
- [show chassis environment routing-engine on page 628](#)
- *show chassis environment scg*
- *show chassis environment sfb*
- *show chassis environment sib*

List of Sample Output

[show chassis environment \(J2300 Router\) on page 535](#)
[show chassis environment \(J4300 or J6300 Router\) on page 535](#)
[show chassis environment \(M5 Router\) on page 535](#)
[show chassis environment \(M7i Router\) on page 536](#)
[show chassis environment \(M10 Router\) on page 536](#)

[show chassis environment \(M10i Router\) on page 536](#)
[show chassis environment \(M20 Router\) on page 537](#)
[show chassis environment \(M40 Router\) on page 537](#)
[show chassis environment \(M40e Router\) on page 537](#)
[show chassis environment \(M120 Router\) on page 538](#)
[show chassis environment \(M160 Router\) on page 539](#)
[show chassis environment \(M320 Router\) on page 539](#)
[show chassis environment \(MX240 Router\) on page 540](#)
[show chassis environment \(MX240 Router with Enhanced MX SCB\) on page 541](#)
[show chassis environment \(MX480 Router\) on page 542](#)
[show chassis environment \(MX480 Router with Enhanced MX SCB\) on page 543](#)
[show chassis environment \(MX960 Router\) on page 544](#)
[show chassis environment \(MX960 Router with Enhanced MX SCB\) on page 545](#)
[show chassis environment \(MX2020 Router\) on page 547](#)
[show chassis environment \(MX2010 Router\) on page 557](#)
[show chassis environment \(T320 Router\) on page 562](#)
[show chassis environment \(T640 Router\) on page 563](#)
[show chassis environment \(T4000 Router\) on page 564](#)
[show chassis environment \(TX Matrix Router\) on page 565](#)
[show chassis environment \(T1600 Router\) on page 567](#)
[show chassis environment \(TX Matrix Plus Router\) on page 568](#)
[show chassis environment \(EX4200 Standalone Switch\) on page 570](#)
[show chassis environment \(EX8216 Switch\) on page 570](#)
[show chassis environment \(QFX Series\) on page 571](#)
[show chassis environment interconnect-device \(QFabric System\) on page 571](#)
[show chassis environment node-device \(QFabric System\) on page 573](#)
[show chassis environment pem node-device \(QFabric System\) on page 574](#)
[show chassis environment \(PTX5000 Packet Transport Switch\) on page 574](#)
[show chassis environment \(ACX2000 Universal Access Router\) on page 576](#)

Output Fields [Table 53 on page 534](#) lists the output fields for the **show chassis environment** command. Output fields are listed in the approximate order in which they appear.

Table 53: show chassis environment Output Fields

Field Name	Field Description
Class	<p>Information about the category or class of chassis component:</p> <ul style="list-style-type: none"> • Power: Power information: <ul style="list-style-type: none"> • (M5, M10, M20, and M40 routers and EX Series switches only) Power supply status: OK, Testing, (during initial power-on), Failed, or Absent. • (M7i, M10i, M40e, M120, M160, M320, and T Series routers and EX Series switches only) Power Entry Modules status: OK, Testing, (during initial power-on), Check, Failed, or Absent. • (PTX Series only) Power information is reported in PDU or PSM combinations. The status is: OK, Testing, (during initial power-on), Check, Failed, or Absent. • Temp: Temperature of air flowing through the chassis in degrees Celsius (C) and Fahrenheit (F). <ul style="list-style-type: none"> • On PTX Series Packet Transport Switches and MX2010 and MX2020 Routers, multiple cooling zones are supported. FRU temperatures in each zone are coordinated with the fan speed of fan trays in those zones. • EX2200 switches have a side-to-rear cooling system. The Local Intake temperature is measured by the sensor on the right side of the chassis, and the Remote Intake temperature is measured by the sensor on the left side of the chassis. • Pic: On ACX4000 Routers, multiple temperature channels on a MIC. The status is: OK and the Measurement is in degrees Celsius (C) and Fahrenheit (F). • Fan: Fan status: OK, Testing (during initial power-on), Failed, or Absent. On PTX Series Packet Transport Switches and MX2010 and MX2020 Routers, multiple fan trays are supported. Fan status is reported in Fan Tray or Fan combinations. Measurement indicates actual fan RPM (PTX and MX2010 and MX2020 Routers only). • Misc: Information about other components of the chassis. <ul style="list-style-type: none"> • On some routers, this field indicates the status of one or more additional components. • On the M40e, M160, and M320 router, Misc includes CIP (Connector Interface Panel). OK indicates that the CIP is present. Absent indicates that the CIP is not present. • On T Series routers, Misc includes CIP and SPMB (Switch Processor Mezzanine Board). OK indicates that the CIP or SPMB is present. Absent indicates that the CIP or SPMB is not present. • On PTX Series Packet Transport Switches, Misc includes the SPMB (Switch Processor Mezzanine Board). The SPMB is located on the control boards. OK indicates that the control board is present. Absent indicates that the control board is not present.
Item	(MX2010 and MX2020 Routers) Information about the chassis component: Routing Engines, Controls Boards (CBs), Switch Fabric Boards (SFBs), PICs, Flexible PIC Concentrators (FPCs), and Adapter Cards (ADCs).

Table 53: show chassis environment Output Fields (*continued*)

Field Name	Field Description
Status	<p>(MX2010 and MX2020 Routers) Status of the specified chassis component. For example, if the Class is Fan, the fan status can be:</p> <ul style="list-style-type: none"> • OK: The fans are operational. • Testing: The fans are being tested during initial power-on. • Failed: The fans have failed or the fans are not spinning. • Absent: The fan tray is not installed. <p>If the Class is Power, the power supply status can be:</p> <ul style="list-style-type: none"> • OK: The power component is operational. • Testing: The power component is being tested during initial power-on. • Check: There is insufficient power---that is, fewer than the minimum required feeds are connected. • Failed: The inputs leads have failed. • Absent: The power component is not installed.
Measurement	<p>(MX2010 and MX2020 Routers) Dependant on the Class. For example, if the Class is Temp, indicates the temperature in degree Celsius and degrees Fahrenheit. If the Class is Fan, indicates actual fan RPM.</p>

Sample Output

show chassis environment (J2300 Router)

```

user@host> show chassis environment
Class Item                Status    Measurement
Temp  Routing Engine         OK        40 degrees C / 104 degrees F
Fan   Fan                    OK

```

show chassis environment (J4300 or J6300 Router)

```

user@host> show chassis environment
Class Item                Status    Measurement
Temp  Routing Engine         OK        41 degrees C / 105 degrees F
Fan   Fan 0                  OK
      Fan 1                  OK

```

show chassis environment (M5 Router)

```

user@host> show chassis environment
Class Item                Status    Measurement
Power  Power Supply A          OK
      Power Supply B      Absent
Temp  FPC 0                  OK        30 degrees C / 86 degrees F
      FEB                  OK        33 degrees C / 91 degrees F
      PS Intake            OK        27 degrees C / 80 degrees F
      PS Exhaust           OK        27 degrees C / 80 degrees F
      Routing Engine       OK        34 degrees C / 93 degrees F
Fans  Left Fan 1           OK        Spinning at normal speed
      Left Fan 2           OK        Spinning at normal speed
      Left Fan 3           OK        Spinning at normal speed
      Left Fan 4           OK        Spinning at normal speed

```

```
Misc  Craft Interface      OK
```

show chassis environment (M7i Router)

```
user@host> show chassis environment
Class Item                Status      Measurement
Power Power Supply 0        OK
      Power Supply 1        Absent
Temp  Intake                OK          22 degrees C / 71 degrees F
      FPC 0                 OK          23 degrees C / 73 degrees F
      Power Supplies        OK          23 degrees C / 73 degrees F
      CFEB Intake           OK          24 degrees C / 75 degrees F
      CFEB Exhaust          OK          29 degrees C / 84 degrees F
      Routing Engine        OK          26 degrees C / 78 degrees F
Fans  Fan 1                 OK          Spinning at normal speed
      Fan 2                 OK          Spinning at normal speed
      Fan 3                 OK          Spinning at normal speed
      Fan 4                 OK          Spinning at normal speed
```

show chassis environment (M10 Router)

```
user@host> show chassis environment
Class Item                Status      Measurement
Power Power Supply A        OK
      Power Supply B        Failed
Temp  FPC 0                 OK          36 degrees C / 96 degrees F
      FPC 1                 OK          35 degrees C / 95 degrees F
      FEB                  OK          34 degrees C / 93 degrees F
      PS Intake            OK          31 degrees C / 87 degrees F
      PS Exhaust           OK          34 degrees C / 93 degrees F
      Routing Engine        OK          35 degrees C / 95 degrees F
Fans  Left Fan 1            OK          Spinning at normal speed
      Left Fan 2            OK          Spinning at normal speed
      Left Fan 3            OK          Spinning at normal speed
      Left Fan 4            OK          Spinning at normal speed
Misc  Craft Interface        OK
```

show chassis environment (M10i Router)

```
user@host> show chassis environment
Class Item                Status      Measurement
Power Power Supply 0        OK
      Power Supply 1        OK
      Power Supply 2        Absent
      Power Supply 3        Absent
Temp  Intake                OK          26 degrees C / 78 degrees F
      FPC 0                 OK          27 degrees C / 80 degrees F
      FPC 1                 OK          28 degrees C / 82 degrees F
      Lower Power Supplies  OK          29 degrees C / 84 degrees F
      Upper Power Supplies  OK          28 degrees C / 82 degrees F
      CFEB Intake           OK          27 degrees C / 80 degrees F
      CFEB Exhaust          OK          36 degrees C / 96 degrees F
      Routing Engine 0       OK          31 degrees C / 87 degrees F
      Routing Engine 1       OK          27 degrees C / 80 degrees F
Fans  Fan Tray 0 Fan 1       OK          Spinning at normal speed
      Fan Tray 0 Fan 2       OK          Spinning at normal speed
      Fan Tray 0 Fan 3       OK          Spinning at normal speed
      Fan Tray 0 Fan 4       OK          Spinning at normal speed
      Fan Tray 0 Fan 5       OK          Spinning at normal speed
```

Fan Tray 0 Fan 6	OK	Spinning at normal speed
Fan Tray 0 Fan 7	OK	Spinning at normal speed
Fan Tray 0 Fan 8	OK	Spinning at normal speed
Fan Tray 1 Fan 1	Absent	
Fan Tray 1 Fan 2	Absent	
Fan Tray 1 Fan 3	Absent	
Fan Tray 1 Fan 4	Absent	
Fan Tray 1 Fan 5	Absent	
Fan Tray 1 Fan 6	Absent	
Fan Tray 1 Fan 7	Absent	
Fan Tray 1 Fan 8	Absent	

show chassis environment (M20 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	Power Supply A	OK	
	Power Supply B	Absent	
Temp	FPC 0	OK	28 degrees C / 82 degrees F
	FPC 1	OK	27 degrees C / 80 degrees F
	Power Supply A	OK	22 degrees C / 71 degrees F
	Power Supply B	Absent	
	SSB 0	OK	30 degrees C / 86 degrees F
	Backplane	OK	22 degrees C / 71 degrees F
Fans	Routing Engine 0	OK	26 degrees C / 78 degrees F
	Routing Engine 1	Testing	
	Rear Fan	OK	Spinning at normal speed
	Front Upper Fan	OK	Spinning at normal speed
	Front Middle Fan	OK	Spinning at normal speed
	Front Bottom Fan	OK	Spinning at normal speed
Misc	Craft Interface	OK	

show chassis environment (M40 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	Power Supply A	OK	
	Power Supply B	Absent	
Temp	FPC 3	OK	24 degrees C / 75 degrees F
	FPC 6	OK	26 degrees C / 78 degrees F
	SCB	OK	26 degrees C / 78 degrees F
	Backplane @ A1	OK	28 degrees C / 82 degrees F
	Backplane @ A2	OK	23 degrees C / 73 degrees F
	Routing Engine	OK	26 degrees C / 78 degrees F
Fans	Top Impeller	OK	Spinning at normal speed
	Bottom impeller	OK	Spinning at normal speed
	Rear Left Fan	OK	Spinning at normal speed
	Rear Center Fan	OK	Spinning at normal speed
	Rear Right Fan	OK	Spinning at normal speed
Misc	Craft Interface	OK	

show chassis environment (M40e Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	PEM 0	OK	
	PEM 1	Absent	
Temp	PCG 0	OK	44 degrees C / 111 degrees F
	PCG 1	OK	47 degrees C / 116 degrees F

	Routing Engine 0	OK	40 degrees C / 104 degrees F
	Routing Engine 1	OK	37 degrees C / 98 degrees F
	MCS 0	OK	45 degrees C / 113 degrees F
	MCS 1	OK	42 degrees C / 107 degrees F
	SFM 0 SPP	OK	40 degrees C / 104 degrees F
	SFM 0 SPR	OK	44 degrees C / 111 degrees F
	SFM 1 SPP	OK	43 degrees C / 109 degrees F
	SFM 1 SPR	OK	45 degrees C / 113 degrees F
	FPC 0	OK	38 degrees C / 100 degrees F
	FPC 1	OK	40 degrees C / 104 degrees F
	FPC 2	OK	38 degrees C / 100 degrees F
	FPC 4	OK	34 degrees C / 93 degrees F
	FPC 5	OK	43 degrees C / 109 degrees F
	FPC 6	OK	41 degrees C / 105 degrees F
	FPC 7	OK	43 degrees C / 109 degrees F
	FPM CMB	OK	28 degrees C / 82 degrees F
	FPM Display	OK	28 degrees C / 82 degrees F
Fans	Rear Bottom Blower	OK	Spinning at normal speed
	Rear Top Blower	OK	Spinning at normal speed
	Front Top Blower	OK	Spinning at normal speed
	Fan Tray Rear Left	OK	Spinning at normal speed
	Fan Tray Rear Right	OK	Spinning at normal speed
	Fan Tray Front Left	OK	Spinning at normal speed
	Fan Tray Front Right	OK	Spinning at normal speed
Misc	CIP	OK	

show chassis environment (M120 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	
	PEM 1	OK	
	Routing Engine 0	OK	43 degrees C / 109 degrees F
	Routing Engine 1	OK	44 degrees C / 111 degrees F
	CB 0 Intake	OK	33 degrees C / 91 degrees F
	CB 0 Exhaust A	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust B	OK	35 degrees C / 95 degrees F
	CB 1 Intake	OK	34 degrees C / 93 degrees F
	CB 1 Exhaust A	OK	38 degrees C / 100 degrees F
	CB 1 Exhaust B	OK	35 degrees C / 95 degrees F
	FEB 3 Intake	OK	35 degrees C / 95 degrees F
	FEB 3 Exhaust A	OK	37 degrees C / 98 degrees F
	FEB 3 Exhaust B	OK	39 degrees C / 102 degrees F
	FEB 4 Intake	OK	33 degrees C / 91 degrees F
	FEB 4 Exhaust A	OK	39 degrees C / 102 degrees F
	FEB 4 Exhaust B	OK	36 degrees C / 96 degrees F
	FPC 2 Exhaust A	OK	32 degrees C / 89 degrees F
	FPC 2 Exhaust B	OK	31 degrees C / 87 degrees F
	FPC 3 Exhaust A	OK	32 degrees C / 89 degrees F
	FPC 3 Exhaust B	OK	33 degrees C / 91 degrees F
	FPC 4 Exhaust A	OK	32 degrees C / 89 degrees F
	FPC 4 Exhaust B	OK	30 degrees C / 86 degrees F
Fans	Front Top Tray Fan 1	OK	Spinning at normal speed
	Front Top Tray Fan 2	OK	Spinning at normal speed
	Front Top Tray Fan 3	OK	Spinning at normal speed
	Front Top Tray Fan 4	OK	Spinning at normal speed
	Front Top Tray Fan 5	OK	Spinning at normal speed
	Front Top Tray Fan 6	OK	Spinning at normal speed
	Front Top Tray Fan 7	OK	Spinning at normal speed
	Front Top Tray Fan 8	OK	Spinning at normal speed

Front Bottom Tray Fan 1	OK	Spinning at normal speed
Front Bottom Tray Fan 2	OK	Spinning at normal speed
Front Bottom Tray Fan 3	OK	Spinning at normal speed
Front Bottom Tray Fan 4	OK	Spinning at normal speed
Front Bottom Tray Fan 5	OK	Spinning at normal speed
Front Bottom Tray Fan 6	OK	Spinning at normal speed
Front Bottom Tray Fan 7	OK	Spinning at normal speed
Front Bottom Tray Fan 8	OK	Spinning at normal speed
Rear Top Tray Fan 1	OK	Spinning at normal speed
Rear Top Tray Fan 2	OK	Spinning at normal speed
Rear Top Tray Fan 3	OK	Spinning at normal speed
Rear Top Tray Fan 4	OK	Spinning at normal speed
Rear Top Tray Fan 5	OK	Spinning at normal speed
Rear Top Tray Fan 6	OK	Spinning at normal speed
Rear Top Tray Fan 7	OK	Spinning at normal speed
Rear Top Tray Fan 8	OK	Spinning at normal speed
Rear Bottom Tray Fan 1	OK	Spinning at normal speed
Rear Bottom Tray Fan 2	OK	Spinning at normal speed
Rear Bottom Tray Fan 3	OK	Spinning at normal speed
Rear Bottom Tray Fan 4	OK	Spinning at normal speed
Rear Bottom Tray Fan 5	OK	Spinning at normal speed
Rear Bottom Tray Fan 6	OK	Spinning at normal speed
Rear Bottom Tray Fan 7	OK	Spinning at normal speed
Rear Bottom Tray Fan 8	OK	Spinning at normal speed

show chassis environment (M160 Router)

```

user@host> show chassis environment

```

Class	Item	Status	Measurement
Power	PEM 0	OK	PEM 1
			Absent
Temp	PCG 0	OK	45 degrees C / 113 degrees F
	PCG 1	Absent	
	Routing Engine 0	OK	35 degrees C / 95 degrees F
	Routing Engine 1	Absent	
	MCS 0	OK	50 degrees C / 122 degrees F
	SFM 0 SPP	OK	47 degrees C / 116 degrees F
	SFM 0 SPR	OK	49 degrees C / 120 degrees F
	SFM 1 SPP	OK	50 degrees C / 122 degrees F
	SFM 1 SPR	OK	50 degrees C / 122 degrees F
	SFM 2 SPP	OK	51 degrees C / 123 degrees F
	SFM 2 SPR	OK	52 degrees C / 125 degrees F
	SFM 3 SPP	OK	52 degrees C / 125 degrees F
	SFM 3 SPR	OK	48 degrees C / 118 degrees F
	FPC 0	OK	45 degrees C / 113 degrees F
	FPC 6	OK	43 degrees C / 109 degrees F
	FPM CMB	OK	31 degrees C / 87 degrees F
	FPM Display	OK	33 degrees C / 91 degrees F
Fans	Rear Bottom Blower	OK	Spinning at normal speed
	Rear Top Blower	OK	Spinning at normal speed
	Front Top Blower	OK	Spinning at normal speed
	Fan Tray Rear Left	OK	Spinning at normal speed
	Fan Tray Rear Right	OK	Spinning at normal speed
	Fan Tray Front Left	OK	Spinning at normal speed
	Fan Tray Front Right	OK	Spinning at normal speed
Misc	CIP	OK	

show chassis environment (M320 Router)

```

user@host> show chassis environment

```

Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	Absent	
	PEM 2	OK	
	PEM 3	OK	
	Routing Engine 0	OK	33 degrees C / 91 degrees F
	Routing Engine 1	OK	32 degrees C / 89 degrees F
	CB 0	OK	36 degrees C / 96 degrees F
	CB 1	OK	36 degrees C / 96 degrees F
	SIB 0	OK	38 degrees C / 100 degrees F
	SIB 1	OK	29 degrees C / 84 degrees F
	SIB 2	OK	38 degrees C / 100 degrees F
	SIB 3	OK	41 degrees C / 105 degrees F
	FPC 0 Intake	OK	28 degrees C / 82 degrees F
	FPC 0 Exhaust	OK	40 degrees C / 104 degrees F
	FPC 1 Intake	OK	29 degrees C / 84 degrees F
	FPC 1 Exhaust	OK	39 degrees C / 102 degrees F
	FPC 2 Intake	OK	28 degrees C / 82 degrees F
	FPC 2 Exhaust	OK	38 degrees C / 100 degrees F
	FPC 3 Intake	OK	28 degrees C / 82 degrees F
	FPC 3 Exhaust	OK	39 degrees C / 102 degrees F
	FPC 6 Intake	OK	27 degrees C / 80 degrees F
	FPC 6 Exhaust	OK	39 degrees C / 102 degrees F
	FPC 7 Intake	OK	27 degrees C / 80 degrees F
	FPC 7 Exhaust	OK	42 degrees C / 107 degrees F
	FPM GBUS	OK	30 degrees C / 86 degrees F
Fan	Top Left Front fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Rear Fan 1 (TOP)	OK	Spinning at normal speed
	Rear Fan 2	OK	Spinning at normal speed
	Rear Fan 3	OK	Spinning at normal speed
	Rear Fan 4	OK	Spinning at normal speed
	Rear Fan 5	OK	Spinning at normal speed
	Rear Fan 6	OK	Spinning at normal speed
	Rear Fan 7 (Bottom)	OK	Spinning at normal speed
Misc	CIP	OK	

show chassis environment (MX240 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	40 degrees C / 104 degrees F
	PEM 1	OK	45 degrees C / 113 degrees F
	PEM 2	Absent	
	PEM 3	Absent	
	Routing Engine 0	OK	39 degrees C / 102 degrees F
	Routing Engine 1	OK	37 degrees C / 98 degrees F
	CB 0 Intake	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust A	OK	34 degrees C / 93 degrees F
	CB 0 Exhaust B	OK	38 degrees C / 100 degrees F
	CB 0 ACBC	OK	37 degrees C / 98 degrees F
	CB 0 SF A	OK	49 degrees C / 120 degrees F
	CB 0 SF B	OK	41 degrees C / 105 degrees F
	CB 1 Intake	OK	37 degrees C / 98 degrees F

CB 1 Exhaust A	OK	34 degrees C / 93 degrees F
CB 1 Exhaust B	OK	39 degrees C / 102 degrees F
CB 1 ACBC	OK	38 degrees C / 100 degrees F
CB 1 SF A	OK	47 degrees C / 116 degrees F
CB 1 SF B	OK	41 degrees C / 105 degrees F
FPC 1 Intake	OK	33 degrees C / 91 degrees F
FPC 1 Exhaust A	OK	38 degrees C / 100 degrees F
FPC 1 Exhaust B	OK	53 degrees C / 127 degrees F
FPC 1 I3 0 TSensor	OK	50 degrees C / 122 degrees F
FPC 1 I3 0 Chip	OK	53 degrees C / 127 degrees F
FPC 1 I3 1 TSensor	OK	49 degrees C / 120 degrees F
FPC 1 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 1 I3 2 TSensor	OK	47 degrees C / 116 degrees F
FPC 1 I3 2 Chip	OK	49 degrees C / 120 degrees F
FPC 1 I3 3 TSensor	OK	44 degrees C / 111 degrees F
FPC 1 I3 3 Chip	OK	46 degrees C / 114 degrees F
FPC 1 IA 0 TSensor	OK	45 degrees C / 113 degrees F
FPC 1 IA 0 Chip	OK	44 degrees C / 111 degrees F
FPC 1 IA 1 TSensor	OK	44 degrees C / 111 degrees F
FPC 1 IA 1 Chip	OK	48 degrees C / 118 degrees F
FPC 2 Intake	OK	32 degrees C / 89 degrees F
FPC 2 Exhaust A	OK	40 degrees C / 104 degrees F
FPC 2 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 2 I3 0 TSensor	OK	52 degrees C / 125 degrees F
FPC 2 I3 0 Chip	OK	56 degrees C / 132 degrees F
FPC 2 I3 1 TSensor	OK	52 degrees C / 125 degrees F
FPC 2 I3 1 Chip	OK	55 degrees C / 131 degrees F
FPC 2 I3 2 TSensor	OK	49 degrees C / 120 degrees F
FPC 2 I3 2 Chip	OK	52 degrees C / 125 degrees F
FPC 2 I3 3 TSensor	OK	44 degrees C / 111 degrees F
FPC 2 I3 3 Chip	OK	48 degrees C / 118 degrees F
FPC 2 IA 0 TSensor	OK	50 degrees C / 122 degrees F
FPC 2 IA 0 Chip	OK	48 degrees C / 118 degrees F
FPC 2 IA 1 TSensor	OK	47 degrees C / 116 degrees F
FPC 2 IA 1 Chip	OK	53 degrees C / 127 degrees F
Fans Front Fan	OK	Spinning at normal speed
Middle Fan	OK	Spinning at normal speed
Rear Fan	OK	Spinning at normal speed

show chassis environment (MX240 Router with Enhanced MX SCB)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	40 degrees C / 104 degrees F
	PEM 1	OK	45 degrees C / 113 degrees F
	PEM 2	Absent	
	PEM 3	Absent	
Routing Engine	Engine 0	OK	39 degrees C / 102 degrees F
	Engine 1	OK	37 degrees C / 98 degrees F
	CB 0 Intake	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust A	OK	34 degrees C / 93 degrees F
	CB 0 Exhaust B	OK	38 degrees C / 100 degrees F
	CB 0 ACBC	OK	37 degrees C / 98 degrees F
	CB 0 XF A	OK	49 degrees C / 120 degrees F
	CB 0 XF B	OK	41 degrees C / 105 degrees F
	CB 1 Intake	OK	37 degrees C / 98 degrees F
	CB 1 Exhaust A	OK	34 degrees C / 93 degrees F
	CB 1 Exhaust B	OK	39 degrees C / 102 degrees F
	CB 1 ACBC	OK	38 degrees C / 100 degrees F
	CB 1 XF A	OK	47 degrees C / 116 degrees F
	CB 1 XF B	OK	41 degrees C / 105 degrees F

	FPC 1 Intake	OK	33 degrees C / 91 degrees F
	FPC 1 Exhaust A	OK	38 degrees C / 100 degrees F
	FPC 1 Exhaust B	OK	53 degrees C / 127 degrees F
	FPC 1 I3 0 TSensor	OK	50 degrees C / 122 degrees F
	FPC 1 I3 0 Chip	OK	53 degrees C / 127 degrees F
	FPC 1 I3 1 TSensor	OK	49 degrees C / 120 degrees F
	FPC 1 I3 1 Chip	OK	52 degrees C / 125 degrees F
	FPC 1 I3 2 TSensor	OK	47 degrees C / 116 degrees F
	FPC 1 I3 2 Chip	OK	49 degrees C / 120 degrees F
	FPC 1 I3 3 TSensor	OK	44 degrees C / 111 degrees F
	FPC 1 I3 3 Chip	OK	46 degrees C / 114 degrees F
	FPC 1 IA 0 TSensor	OK	45 degrees C / 113 degrees F
	FPC 1 IA 0 Chip	OK	44 degrees C / 111 degrees F
	FPC 1 IA 1 TSensor	OK	44 degrees C / 111 degrees F
	FPC 1 IA 1 Chip	OK	48 degrees C / 118 degrees F
	FPC 2 Intake	OK	32 degrees C / 89 degrees F
	FPC 2 Exhaust A	OK	40 degrees C / 104 degrees F
	FPC 2 Exhaust B	OK	52 degrees C / 125 degrees F
	FPC 2 I3 0 TSensor	OK	52 degrees C / 125 degrees F
	FPC 2 I3 0 Chip	OK	56 degrees C / 132 degrees F
	FPC 2 I3 1 TSensor	OK	52 degrees C / 125 degrees F
	FPC 2 I3 1 Chip	OK	55 degrees C / 131 degrees F
	FPC 2 I3 2 TSensor	OK	49 degrees C / 120 degrees F
	FPC 2 I3 2 Chip	OK	52 degrees C / 125 degrees F
	FPC 2 I3 3 TSensor	OK	44 degrees C / 111 degrees F
	FPC 2 I3 3 Chip	OK	48 degrees C / 118 degrees F
	FPC 2 IA 0 TSensor	OK	50 degrees C / 122 degrees F
	FPC 2 IA 0 Chip	OK	48 degrees C / 118 degrees F
	FPC 2 IA 1 TSensor	OK	47 degrees C / 116 degrees F
	FPC 2 IA 1 Chip	OK	53 degrees C / 127 degrees F
Fans	Front Fan	OK	Spinning at normal speed
	Middle Fan	OK	Spinning at normal speed
	Rear Fan	OK	Spinning at normal speed

show chassis environment (MX480 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	35 degrees C / 95 degrees F
	PEM 1	OK	40 degrees C / 104 degrees F
	PEM 2	Absent	
	PEM 3	Absent	
	Routing Engine 0	OK	44 degrees C / 111 degrees F
	Routing Engine 1	OK	45 degrees C / 113 degrees F
	CB 0 Intake	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust A	OK	38 degrees C / 100 degrees F
	CB 0 Exhaust B	OK	39 degrees C / 102 degrees F
	CB 0 ACBC	OK	37 degrees C / 98 degrees F
	CB 0 SF A	OK	51 degrees C / 123 degrees F
	CB 0 SF B	OK	44 degrees C / 111 degrees F
	CB 1 Intake	OK	36 degrees C / 96 degrees F
	CB 1 Exhaust A	OK	39 degrees C / 102 degrees F
	CB 1 Exhaust B	OK	40 degrees C / 104 degrees F
	CB 1 ACBC	OK	37 degrees C / 98 degrees F
	CB 1 SF A	OK	50 degrees C / 122 degrees F
	CB 1 SF B	OK	43 degrees C / 109 degrees F
	FPC 0 Intake	OK	36 degrees C / 96 degrees F
	FPC 0 Exhaust A	OK	39 degrees C / 102 degrees F
	FPC 0 Exhaust B	OK	51 degrees C / 123 degrees F
	FPC 0 I3 0 TSensor	OK	49 degrees C / 120 degrees F
	FPC 0 I3 0 Chip	OK	56 degrees C / 132 degrees F

	FPC 0 I3 1 TSensor	OK	47 degrees C / 116 degrees F
	FPC 0 I3 1 Chip	OK	52 degrees C / 125 degrees F
	FPC 0 I3 2 TSensor	OK	46 degrees C / 114 degrees F
	FPC 0 I3 2 Chip	OK	48 degrees C / 118 degrees F
	FPC 0 I3 3 TSensor	OK	42 degrees C / 107 degrees F
	FPC 0 I3 3 Chip	OK	45 degrees C / 113 degrees F
	FPC 0 IA 0 TSensor	OK	45 degrees C / 113 degrees F
	FPC 0 IA 0 Chip	OK	45 degrees C / 113 degrees F
	FPC 0 IA 1 TSensor	OK	44 degrees C / 111 degrees F
	FPC 0 IA 1 Chip	OK	48 degrees C / 118 degrees F
	FPC 1 Intake	OK	37 degrees C / 98 degrees F
	FPC 1 Exhaust A	OK	41 degrees C / 105 degrees F
	FPC 1 Exhaust B	OK	52 degrees C / 125 degrees F
	FPC 1 I3 0 TSensor	OK	51 degrees C / 123 degrees F
	FPC 1 I3 0 Chip	OK	57 degrees C / 134 degrees F
	FPC 1 I3 1 TSensor	OK	48 degrees C / 118 degrees F
	FPC 1 I3 1 Chip	OK	52 degrees C / 125 degrees F
	FPC 1 I3 2 TSensor	OK	46 degrees C / 114 degrees F
	FPC 1 I3 2 Chip	OK	50 degrees C / 122 degrees F
	FPC 1 I3 3 TSensor	OK	42 degrees C / 107 degrees F
	FPC 1 I3 3 Chip	OK	46 degrees C / 114 degrees F
	FPC 1 IA 0 TSensor	OK	49 degrees C / 120 degrees F
	FPC 1 IA 0 Chip	OK	48 degrees C / 118 degrees F
	FPC 1 IA 1 TSensor	OK	46 degrees C / 114 degrees F
	FPC 1 IA 1 Chip	OK	50 degrees C / 122 degrees F
Fans	Top Rear Fan	OK	Spinning at normal speed
	Bottom Rear Fan	OK	Spinning at normal speed
	Top Middle Fan	OK	Spinning at normal speed
	Bottom Middle Fan	OK	Spinning at normal speed
	Top Front Fan	OK	Spinning at normal speed
	Bottom Front Fan	OK	Spinning at normal speed

show chassis environment (MX480 Router with Enhanced MX SCB)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	35 degrees C / 95 degrees F
	PEM 1	OK	40 degrees C / 104 degrees F
	PEM 2	Absent	
	PEM 3	Absent	
	Routing Engine 0	OK	44 degrees C / 111 degrees F
	Routing Engine 1	OK	45 degrees C / 113 degrees F
	CB 0 Intake	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust A	OK	38 degrees C / 100 degrees F
	CB 0 Exhaust B	OK	39 degrees C / 102 degrees F
	CB 0 ACBC	OK	37 degrees C / 98 degrees F
	CB 0 XF A	OK	51 degrees C / 123 degrees F
	CB 0 XF B	OK	44 degrees C / 111 degrees F
	CB 1 Intake	OK	36 degrees C / 96 degrees F
	CB 1 Exhaust A	OK	39 degrees C / 102 degrees F
	CB 1 Exhaust B	OK	40 degrees C / 104 degrees F
	CB 1 ACBC	OK	37 degrees C / 98 degrees F
	CB 1 XF A	OK	50 degrees C / 122 degrees F
	CB 1 XF B	OK	43 degrees C / 109 degrees F
	FPC 0 Intake	OK	36 degrees C / 96 degrees F
	FPC 0 Exhaust A	OK	39 degrees C / 102 degrees F
	FPC 0 Exhaust B	OK	51 degrees C / 123 degrees F
	FPC 0 I3 0 TSensor	OK	49 degrees C / 120 degrees F
	FPC 0 I3 0 Chip	OK	56 degrees C / 132 degrees F
	FPC 0 I3 1 TSensor	OK	47 degrees C / 116 degrees F
	FPC 0 I3 1 Chip	OK	52 degrees C / 125 degrees F

	FPC 0 I3 2 TSensor	OK	46 degrees C / 114 degrees F
	FPC 0 I3 2 Chip	OK	48 degrees C / 118 degrees F
	FPC 0 I3 3 TSensor	OK	42 degrees C / 107 degrees F
	FPC 0 I3 3 Chip	OK	45 degrees C / 113 degrees F
	FPC 0 IA 0 TSensor	OK	45 degrees C / 113 degrees F
	FPC 0 IA 0 Chip	OK	45 degrees C / 113 degrees F
	FPC 0 IA 1 TSensor	OK	44 degrees C / 111 degrees F
	FPC 0 IA 1 Chip	OK	48 degrees C / 118 degrees F
	FPC 1 Intake	OK	37 degrees C / 98 degrees F
	FPC 1 Exhaust A	OK	41 degrees C / 105 degrees F
	FPC 1 Exhaust B	OK	52 degrees C / 125 degrees F
	FPC 1 I3 0 TSensor	OK	51 degrees C / 123 degrees F
	FPC 1 I3 0 Chip	OK	57 degrees C / 134 degrees F
	FPC 1 I3 1 TSensor	OK	48 degrees C / 118 degrees F
	FPC 1 I3 1 Chip	OK	52 degrees C / 125 degrees F
	FPC 1 I3 2 TSensor	OK	46 degrees C / 114 degrees F
	FPC 1 I3 2 Chip	OK	50 degrees C / 122 degrees F
	FPC 1 I3 3 TSensor	OK	42 degrees C / 107 degrees F
	FPC 1 I3 3 Chip	OK	46 degrees C / 114 degrees F
	FPC 1 IA 0 TSensor	OK	49 degrees C / 120 degrees F
	FPC 1 IA 0 Chip	OK	48 degrees C / 118 degrees F
	FPC 1 IA 1 TSensor	OK	46 degrees C / 114 degrees F
	FPC 1 IA 1 Chip	OK	50 degrees C / 122 degrees F
Fans	Top Rear Fan	OK	Spinning at normal speed
	Bottom Rear Fan	OK	Spinning at normal speed
	Top Middle Fan	OK	Spinning at normal speed
	Bottom Middle Fan	OK	Spinning at normal speed
	Top Front Fan	OK	Spinning at normal speed
	Bottom Front Fan	OK	Spinning at normal speed

show chassis environment (MX960 Router)

user@host> show chassis environment			
Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	Absent	
	PEM 2	Check	
	PEM 3	OK	35 degrees C / 95 degrees F
	Routing Engine 0	OK	37 degrees C / 98 degrees F
	Routing Engine 1	Absent	
	CB 0 Intake	OK	24 degrees C / 75 degrees F
	CB 0 Exhaust A	OK	30 degrees C / 86 degrees F
	CB 0 Exhaust B	OK	27 degrees C / 80 degrees F
	CB 1 Intake	Absent	
	CB 1 Exhaust A	Absent	
	CB 1 Exhaust B	Absent	
	CB 1 ACBC	Absent	
	CB 1 SF A	Absent	
	CB 1 SF B	Absent	
	CB 2 Intake	Absent	
	CB 2 Exhaust A	Absent	
	CB 2 Exhaust B	Absent	
	CB 2 ACBC	Absent	
	CB 2 SF A	Absent	
	CB 2 SF B	Absent	
	FPC 4 Intake	OK	24 degrees C / 75 degrees F
	FPC 4 Exhaust A	OK	36 degrees C / 96 degrees F
	FPC 4 Exhaust B	OK	38 degrees C / 100 degrees F
	FPC 7 Intake	OK	24 degrees C / 75 degrees F
	FPC 7 Exhaust A	OK	36 degrees C / 96 degrees F
	FPC 7 Exhaust B	OK	42 degrees C / 107 degrees F

Fans	Top Fan Tray Temp	Failed	
	Top Tray Fan 1	OK	Spinning at normal speed
	Top Tray Fan 2	OK	Spinning at normal speed
	Top Tray Fan 3	OK	Spinning at normal speed
	Top Tray Fan 4	OK	Spinning at normal speed
	Top Tray Fan 5	OK	Spinning at normal speed
	Top Tray Fan 6	OK	Spinning at normal speed
	Bottom Fan Tray Temp	Failed	
	Bottom Tray Fan 1	OK	Spinning at normal speed
	Bottom Tray Fan 2	OK	Spinning at normal speed
	Bottom Tray Fan 3	OK	Spinning at normal speed
	Bottom Tray Fan 4	OK	Spinning at normal speed
	Bottom Tray Fan 5	OK	Spinning at normal speed
	Bottom Tray Fan 6	OK	Spinning at normal speed

show chassis environment (MX960 Router with Enhanced MX SCB)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	OK	50 degrees C / 122 degrees F
	PEM 2	OK	50 degrees C / 122 degrees F
	PEM 3	OK	50 degrees C / 122 degrees F
	Routing Engine 0	OK	42 degrees C / 107 degrees F
	Routing Engine 0 CPU	OK	51 degrees C / 123 degrees F
	Routing Engine 1	OK	39 degrees C / 102 degrees F
	Routing Engine 1 CPU	OK	44 degrees C / 111 degrees F
	CB 0 Intake	OK	35 degrees C / 95 degrees F
	CB 0 Exhaust A	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust B	OK	43 degrees C / 109 degrees F
	CB 0 ACBC	OK	38 degrees C / 100 degrees F
	CB 0 XF A	OK	53 degrees C / 127 degrees F
	CB 0 XF B	OK	47 degrees C / 116 degrees F
	CB 1 Intake	OK	35 degrees C / 95 degrees F
	CB 1 Exhaust A	OK	35 degrees C / 95 degrees F
	CB 1 Exhaust B	OK	41 degrees C / 105 degrees F
	CB 1 ACBC	OK	38 degrees C / 100 degrees F
	CB 1 XF A	OK	52 degrees C / 125 degrees F
	CB 1 XF B	OK	47 degrees C / 116 degrees F
	CB 2 Intake	OK	32 degrees C / 89 degrees F
	CB 2 Exhaust A	OK	30 degrees C / 86 degrees F
	CB 2 Exhaust B	OK	35 degrees C / 95 degrees F
	CB 2 ACBC	OK	33 degrees C / 91 degrees F
	CB 2 XF A	OK	51 degrees C / 123 degrees F
	CB 2 XF B	OK	50 degrees C / 122 degrees F
	FPC 0 Intake	OK	35 degrees C / 95 degrees F
	FPC 0 Exhaust A	OK	39 degrees C / 102 degrees F
	FPC 0 Exhaust B	OK	50 degrees C / 122 degrees F
	FPC 0 I3 0 TSensor	OK	50 degrees C / 122 degrees F
	FPC 0 I3 0 Chip	OK	56 degrees C / 132 degrees F
	FPC 0 I3 1 TSensor	OK	47 degrees C / 116 degrees F
	FPC 0 I3 1 Chip	OK	50 degrees C / 122 degrees F
	FPC 0 I3 2 TSensor	OK	45 degrees C / 113 degrees F
	FPC 0 I3 2 Chip	OK	48 degrees C / 118 degrees F
	FPC 0 I3 3 TSensor	OK	41 degrees C / 105 degrees F
	FPC 0 I3 3 Chip	OK	44 degrees C / 111 degrees F
	FPC 0 IA 0 TSensor	OK	45 degrees C / 113 degrees F
	FPC 0 IA 0 Chip	OK	45 degrees C / 113 degrees F
	FPC 0 IA 1 TSensor	OK	44 degrees C / 111 degrees F
	FPC 0 IA 1 Chip	OK	48 degrees C / 118 degrees F
	FPC 1 Intake	OK	36 degrees C / 96 degrees F

FPC 1 Exhaust A	OK	47 degrees C / 116 degrees F
FPC 1 Exhaust B	OK	43 degrees C / 109 degrees F
FPC 1 LU 0 TCAM TSensor	OK	53 degrees C / 127 degrees F
FPC 1 LU 0 TCAM Chip	OK	57 degrees C / 134 degrees F
FPC 1 LU 0 TSensor	OK	53 degrees C / 127 degrees F
FPC 1 LU 0 Chip	OK	60 degrees C / 140 degrees F
FPC 1 MQ 0 TSensor	OK	53 degrees C / 127 degrees F
FPC 1 MQ 0 Chip	OK	56 degrees C / 132 degrees F
FPC 1 LU 1 TCAM TSensor	OK	51 degrees C / 123 degrees F
FPC 1 LU 1 TCAM Chip	OK	52 degrees C / 125 degrees F
FPC 1 LU 1 TSensor	OK	51 degrees C / 123 degrees F
FPC 1 LU 1 Chip	OK	53 degrees C / 127 degrees F
FPC 1 MQ 1 TSensor	OK	51 degrees C / 123 degrees F
FPC 1 MQ 1 Chip	OK	58 degrees C / 136 degrees F
FPC 2 Intake	OK	35 degrees C / 95 degrees F
FPC 2 Exhaust A	OK	39 degrees C / 102 degrees F
FPC 2 Exhaust B	OK	54 degrees C / 129 degrees F
FPC 2 I3 0 TSensor	OK	52 degrees C / 125 degrees F
FPC 2 I3 0 Chip	OK	59 degrees C / 138 degrees F
FPC 2 I3 1 TSensor	OK	48 degrees C / 118 degrees F
FPC 2 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 2 I3 2 TSensor	OK	47 degrees C / 116 degrees F
FPC 2 I3 2 Chip	OK	49 degrees C / 120 degrees F
FPC 2 I3 3 TSensor	OK	41 degrees C / 105 degrees F
FPC 2 I3 3 Chip	OK	44 degrees C / 111 degrees F
FPC 2 IA 0 TSensor	OK	47 degrees C / 116 degrees F
FPC 2 IA 0 Chip	OK	46 degrees C / 114 degrees F
FPC 2 IA 1 TSensor	OK	45 degrees C / 113 degrees F
FPC 2 IA 1 Chip	OK	49 degrees C / 120 degrees F
FPC 3 Intake	OK	34 degrees C / 93 degrees F
FPC 3 Exhaust A	OK	34 degrees C / 93 degrees F
FPC 3 Exhaust B	OK	47 degrees C / 116 degrees F
FPC 3 I3 0 TSensor	OK	48 degrees C / 118 degrees F
FPC 3 I3 0 Chip	OK	52 degrees C / 125 degrees F
FPC 3 I3 1 TSensor	OK	46 degrees C / 114 degrees F
FPC 3 I3 1 Chip	OK	48 degrees C / 118 degrees F
FPC 3 IA 0 TSensor	OK	41 degrees C / 105 degrees F
FPC 3 IA 0 Chip	OK	40 degrees C / 104 degrees F
FPC 5 Intake	OK	42 degrees C / 107 degrees F
FPC 5 Exhaust A	OK	42 degrees C / 107 degrees F
FPC 5 Exhaust B	OK	53 degrees C / 127 degrees F
FPC 5 LU 0 TSensor	OK	53 degrees C / 127 degrees F
FPC 5 LU 0 Chip	OK	54 degrees C / 129 degrees F
FPC 5 LU 1 TSensor	OK	53 degrees C / 127 degrees F
FPC 5 LU 1 Chip	OK	61 degrees C / 141 degrees F
FPC 5 LU 2 TSensor	OK	53 degrees C / 127 degrees F
FPC 5 LU 2 Chip	OK	51 degrees C / 123 degrees F
FPC 5 LU 3 TSensor	OK	53 degrees C / 127 degrees F
FPC 5 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 5 MQ 0 TSensor	OK	47 degrees C / 116 degrees F
FPC 5 MQ 0 Chip	OK	52 degrees C / 125 degrees F
FPC 5 MQ 1 TSensor	OK	47 degrees C / 116 degrees F
FPC 5 MQ 1 Chip	OK	52 degrees C / 125 degrees F
FPC 5 MQ 2 TSensor	OK	47 degrees C / 116 degrees F
FPC 5 MQ 2 Chip	OK	46 degrees C / 114 degrees F
FPC 5 MQ 3 TSensor	OK	47 degrees C / 116 degrees F
FPC 5 MQ 3 Chip	OK	45 degrees C / 113 degrees F
FPC 7 Intake	OK	36 degrees C / 96 degrees F
FPC 7 Exhaust A	OK	35 degrees C / 95 degrees F
FPC 7 Exhaust B	OK	33 degrees C / 91 degrees F
FPC 7 QX 0 TSensor	OK	42 degrees C / 107 degrees F

	FPC 7 QX 0 Chip	OK	47 degrees C / 116 degrees F
	FPC 7 LU 0 TCAM TSensor	OK	42 degrees C / 107 degrees F
	FPC 7 LU 0 TCAM Chip	OK	44 degrees C / 111 degrees F
	FPC 7 LU 0 TSensor	OK	42 degrees C / 107 degrees F
	FPC 7 LU 0 Chip	OK	46 degrees C / 114 degrees F
	FPC 7 MQ 0 TSensor	OK	42 degrees C / 107 degrees F
	FPC 7 MQ 0 Chip	OK	45 degrees C / 113 degrees F
	FPC 8 Intake	OK	33 degrees C / 91 degrees F
	FPC 8 Exhaust A	OK	33 degrees C / 91 degrees F
	FPC 8 Exhaust B	OK	36 degrees C / 96 degrees F
	FPC 8 I3 0 TSensor	OK	38 degrees C / 100 degrees F
	FPC 8 I3 0 Chip	OK	43 degrees C / 109 degrees F
	FPC 8 BDS 0 TSensor	OK	37 degrees C / 98 degrees F
	FPC 8 BDS 0 Chip	OK	36 degrees C / 96 degrees F
	FPC 8 IA 0 TSensor	OK	37 degrees C / 98 degrees F
	FPC 8 IA 0 Chip	OK	37 degrees C / 98 degrees F
	FPC 10 Intake	OK	38 degrees C / 100 degrees F
	FPC 10 Exhaust A	OK	36 degrees C / 96 degrees F
	FPC 10 Exhaust B	OK	41 degrees C / 105 degrees F
	FPC 10 I3 0 TSensor	OK	40 degrees C / 104 degrees F
	FPC 10 I3 0 Chip	OK	42 degrees C / 107 degrees F
	FPC 10 I3 1 TSensor	OK	40 degrees C / 104 degrees F
	FPC 10 I3 1 Chip	OK	44 degrees C / 111 degrees F
	FPC 10 I3 2 TSensor	OK	42 degrees C / 107 degrees F
	FPC 10 I3 2 Chip	OK	43 degrees C / 109 degrees F
	FPC 10 I3 3 TSensor	OK	39 degrees C / 102 degrees F
	FPC 10 I3 3 Chip	OK	44 degrees C / 111 degrees F
	FPC 10 IA 0 TSensor	OK	36 degrees C / 96 degrees F
	FPC 10 IA 0 Chip	OK	36 degrees C / 96 degrees F
	FPC 10 IA 1 TSensor	OK	43 degrees C / 109 degrees F
	FPC 10 IA 1 Chip	OK	42 degrees C / 107 degrees F
Fans	Top Fan Tray Temp	OK	37 degrees C / 98 degrees F
	Top Tray Fan 1	OK	Spinning at normal speed
	Top Tray Fan 2	OK	Spinning at normal speed
	Top Tray Fan 3	OK	Spinning at normal speed
	Top Tray Fan 4	OK	Spinning at normal speed
	Top Tray Fan 5	OK	Spinning at normal speed
	Top Tray Fan 6	OK	Spinning at normal speed
	Bottom Fan Tray Temp	OK	28 degrees C / 82 degrees F
	Bottom Tray Fan 1	OK	Spinning at normal speed
	Bottom Tray Fan 2	OK	Spinning at normal speed
	Bottom Tray Fan 3	OK	Spinning at normal speed
	Bottom Tray Fan 4	OK	Spinning at normal speed
	Bottom Tray Fan 5	OK	Spinning at normal speed
	Bottom Tray Fan 6	OK	Spinning at normal speed

show chassis environment (MX2020 Router)

user@host> show chassis environment			
Class	Item	Status	Measurement
Temp	PSM 0	Absent	
	PSM 1	Absent	
	PSM 2	OK	41 degrees C / 105 degrees F
	PSM 3	OK	39 degrees C / 102 degrees F
	PSM 4	OK	39 degrees C / 102 degrees F
	PSM 5	OK	38 degrees C / 100 degrees F
	PSM 6	OK	38 degrees C / 100 degrees F
	PSM 7	OK	38 degrees C / 100 degrees F
	PSM 8	OK	37 degrees C / 98 degrees F
	PSM 9	Absent	
	PSM 10	Absent	

PSM 11	OK	47 degrees C / 116 degrees F
PSM 12	OK	45 degrees C / 113 degrees F
PSM 13	OK	44 degrees C / 111 degrees F
PSM 14	OK	44 degrees C / 111 degrees F
PSM 15	OK	43 degrees C / 109 degrees F
PSM 16	OK	42 degrees C / 107 degrees F
PSM 17	OK	41 degrees C / 105 degrees F
PDM 0	OK	
PDM 1	Absent	
PDM 2	Absent	
PDM 3	OK	
CB 0 IntakeA-Zone0	OK	45 degrees C / 113 degrees F
CB 0 IntakeB-Zone1	OK	34 degrees C / 93 degrees F
CB 0 IntakeC-Zone0	OK	48 degrees C / 118 degrees F
CB 0 ExhaustA-Zone0	OK	45 degrees C / 113 degrees F
CB 0 ExhaustB-Zone1	OK	37 degrees C / 98 degrees F
CB 0 TCBC-Zone0	OK	41 degrees C / 105 degrees F
CB 1 IntakeA-Zone0	OK	46 degrees C / 114 degrees F
CB 1 IntakeB-Zone1	OK	42 degrees C / 107 degrees F
CB 1 IntakeC-Zone0	OK	49 degrees C / 120 degrees F
CB 1 ExhaustA-Zone0	OK	46 degrees C / 114 degrees F
CB 1 ExhaustB-Zone1	OK	41 degrees C / 105 degrees F
CB 1 TCBC-Zone0	OK	46 degrees C / 114 degrees F
SPMB 0 Intake	OK	33 degrees C / 91 degrees F
SPMB 1 Intake	OK	42 degrees C / 107 degrees F
Routing Engine 0	OK	35 degrees C / 95 degrees F
Routing Engine 0 CPU	OK	34 degrees C / 93 degrees F
Routing Engine 1	OK	44 degrees C / 111 degrees F
Routing Engine 1 CPU	OK	42 degrees C / 107 degrees F
SFB 0 Intake-Zone0	OK	55 degrees C / 131 degrees F
SFB 0 Exhaust-Zone1	OK	48 degrees C / 118 degrees F
SFB 0 IntakeA-Zone0	OK	50 degrees C / 122 degrees F
SFB 0 IntakeB-Zone1	OK	40 degrees C / 104 degrees F
SFB 0 Exhaust-Zone0	OK	52 degrees C / 125 degrees F
SFB 0 SFB-XF2-Zone1	OK	61 degrees C / 141 degrees F
SFB 0 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 0 SFB-XF0-Zone0	OK	68 degrees C / 154 degrees F
SFB 1 Intake-Zone0	OK	56 degrees C / 132 degrees F
SFB 1 Exhaust-Zone1	OK	47 degrees C / 116 degrees F
SFB 1 IntakeA-Zone0	OK	51 degrees C / 123 degrees F
SFB 1 IntakeB-Zone1	OK	40 degrees C / 104 degrees F
SFB 1 Exhaust-Zone0	OK	51 degrees C / 123 degrees F
SFB 1 SFB-XF2-Zone1	OK	62 degrees C / 143 degrees F
SFB 1 SFB-XF1-Zone0	OK	67 degrees C / 152 degrees F
SFB 1 SFB-XF0-Zone0	OK	69 degrees C / 156 degrees F
SFB 2 Intake-Zone0	OK	56 degrees C / 132 degrees F
SFB 2 Exhaust-Zone1	OK	47 degrees C / 116 degrees F
SFB 2 IntakeA-Zone0	OK	51 degrees C / 123 degrees F
SFB 2 IntakeB-Zone1	OK	40 degrees C / 104 degrees F
SFB 2 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 2 SFB-XF2-Zone1	OK	65 degrees C / 149 degrees F
SFB 2 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 2 SFB-XF0-Zone0	OK	70 degrees C / 158 degrees F
SFB 3 Intake-Zone0	OK	57 degrees C / 134 degrees F
SFB 3 Exhaust-Zone1	OK	48 degrees C / 118 degrees F
SFB 3 IntakeA-Zone0	OK	52 degrees C / 125 degrees F
SFB 3 IntakeB-Zone1	OK	41 degrees C / 105 degrees F
SFB 3 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 3 SFB-XF2-Zone1	OK	66 degrees C / 150 degrees F
SFB 3 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 3 SFB-XF0-Zone0	OK	71 degrees C / 159 degrees F

SFB 4 Intake-Zone0	OK	58 degrees C / 136 degrees F
SFB 4 Exhaust-Zone1	OK	49 degrees C / 120 degrees F
SFB 4 IntakeA-Zone0	OK	54 degrees C / 129 degrees F
SFB 4 IntakeB-Zone1	OK	42 degrees C / 107 degrees F
SFB 4 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 4 SFB-XF2-Zone1	OK	64 degrees C / 147 degrees F
SFB 4 SFB-XF1-Zone0	OK	68 degrees C / 154 degrees F
SFB 4 SFB-XF0-Zone0	OK	71 degrees C / 159 degrees F
SFB 5 Intake-Zone0	OK	58 degrees C / 136 degrees F
SFB 5 Exhaust-Zone1	OK	50 degrees C / 122 degrees F
SFB 5 IntakeA-Zone0	OK	53 degrees C / 127 degrees F
SFB 5 IntakeB-Zone1	OK	43 degrees C / 109 degrees F
SFB 5 Exhaust-Zone0	OK	54 degrees C / 129 degrees F
SFB 5 SFB-XF2-Zone1	OK	66 degrees C / 150 degrees F
SFB 5 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 5 SFB-XF0-Zone0	OK	74 degrees C / 165 degrees F
SFB 6 Intake-Zone0	OK	58 degrees C / 136 degrees F
SFB 6 Exhaust-Zone1	OK	49 degrees C / 120 degrees F
SFB 6 IntakeA-Zone0	OK	53 degrees C / 127 degrees F
SFB 6 IntakeB-Zone1	OK	43 degrees C / 109 degrees F
SFB 6 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 6 SFB-XF2-Zone1	OK	65 degrees C / 149 degrees F
SFB 6 SFB-XF1-Zone0	OK	68 degrees C / 154 degrees F
SFB 6 SFB-XF0-Zone0	OK	72 degrees C / 161 degrees F
SFB 7 Intake-Zone0	OK	57 degrees C / 134 degrees F
SFB 7 Exhaust-Zone1	OK	50 degrees C / 122 degrees F
SFB 7 IntakeA-Zone0	OK	53 degrees C / 127 degrees F
SFB 7 IntakeB-Zone1	OK	43 degrees C / 109 degrees F
SFB 7 Exhaust-Zone0	OK	54 degrees C / 129 degrees F
SFB 7 SFB-XF2-Zone1	OK	68 degrees C / 154 degrees F
SFB 7 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 7 SFB-XF0-Zone0	OK	73 degrees C / 163 degrees F
FPC 0 Intake	OK	41 degrees C / 105 degrees F
FPC 0 Exhaust A	OK	48 degrees C / 118 degrees F
FPC 0 Exhaust B	OK	62 degrees C / 143 degrees F
FPC 0 LU 0 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 0 Chip	OK	62 degrees C / 143 degrees F
FPC 0 LU 1 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 1 Chip	OK	64 degrees C / 147 degrees F
FPC 0 LU 2 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 0 LU 3 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 0 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 0 Chip	OK	49 degrees C / 120 degrees F
FPC 0 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 1 Chip	OK	51 degrees C / 123 degrees F
FPC 0 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 0 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 3 Chip	OK	45 degrees C / 113 degrees F
FPC 1 Intake	OK	40 degrees C / 104 degrees F
FPC 1 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 1 Exhaust B	OK	58 degrees C / 136 degrees F
FPC 1 LU 0 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 0 Chip	OK	56 degrees C / 132 degrees F
FPC 1 LU 1 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 1 Chip	OK	58 degrees C / 136 degrees F
FPC 1 LU 2 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 2 Chip	OK	49 degrees C / 120 degrees F
FPC 1 LU 3 TSen	OK	55 degrees C / 131 degrees F

FPC 1 LU 3 Chip	OK	51 degrees C / 123 degrees F
FPC 1 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 0 Chip	OK	48 degrees C / 118 degrees F
FPC 1 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 1 Chip	OK	50 degrees C / 122 degrees F
FPC 1 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 1 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 3 Chip	OK	44 degrees C / 111 degrees F
FPC 2 Intake	OK	39 degrees C / 102 degrees F
FPC 2 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 2 Exhaust B	OK	61 degrees C / 141 degrees F
FPC 2 LU 0 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 0 Chip	OK	60 degrees C / 140 degrees F
FPC 2 LU 1 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 1 Chip	OK	65 degrees C / 149 degrees F
FPC 2 LU 2 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 2 Chip	OK	51 degrees C / 123 degrees F
FPC 2 LU 3 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 2 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 0 Chip	OK	50 degrees C / 122 degrees F
FPC 2 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 1 Chip	OK	52 degrees C / 125 degrees F
FPC 2 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 2 Chip	OK	45 degrees C / 113 degrees F
FPC 2 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 3 Chip	OK	46 degrees C / 114 degrees F
FPC 3 Intake	OK	40 degrees C / 104 degrees F
FPC 3 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 3 Exhaust B	OK	61 degrees C / 141 degrees F
FPC 3 LU 0 TSen	OK	58 degrees C / 136 degrees F
FPC 3 LU 0 Chip	OK	61 degrees C / 141 degrees F
FPC 3 LU 1 TSen	OK	58 degrees C / 136 degrees F
FPC 3 LU 1 Chip	OK	62 degrees C / 143 degrees F
FPC 3 LU 2 TSen	OK	58 degrees C / 136 degrees F
FPC 3 LU 2 Chip	OK	51 degrees C / 123 degrees F
FPC 3 LU 3 TSen	OK	58 degrees C / 136 degrees F
FPC 3 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 3 MQ 0 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 0 Chip	OK	50 degrees C / 122 degrees F
FPC 3 MQ 1 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 1 Chip	OK	54 degrees C / 129 degrees F
FPC 3 MQ 2 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 2 Chip	OK	45 degrees C / 113 degrees F
FPC 3 MQ 3 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 3 Chip	OK	48 degrees C / 118 degrees F
FPC 4 Intake	OK	40 degrees C / 104 degrees F
FPC 4 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 4 Exhaust B	OK	62 degrees C / 143 degrees F
FPC 4 LU 0 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 0 Chip	OK	62 degrees C / 143 degrees F
FPC 4 LU 1 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 1 Chip	OK	65 degrees C / 149 degrees F
FPC 4 LU 2 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 2 Chip	OK	51 degrees C / 123 degrees F
FPC 4 LU 3 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 4 MQ 0 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 0 Chip	OK	52 degrees C / 125 degrees F
FPC 4 MQ 1 TSen	OK	48 degrees C / 118 degrees F

FPC 4 MQ 1 Chip	OK	53 degrees C / 127 degrees F
FPC 4 MQ 2 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 2 Chip	OK	46 degrees C / 114 degrees F
FPC 4 MQ 3 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 5 Intake	OK	41 degrees C / 105 degrees F
FPC 5 Exhaust A	OK	50 degrees C / 122 degrees F
FPC 5 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 5 LU 0 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 0 Chip	OK	63 degrees C / 145 degrees F
FPC 5 LU 1 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 1 Chip	OK	66 degrees C / 150 degrees F
FPC 5 LU 2 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 2 Chip	OK	56 degrees C / 132 degrees F
FPC 5 LU 3 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 5 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 0 Chip	OK	52 degrees C / 125 degrees F
FPC 5 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 1 Chip	OK	53 degrees C / 127 degrees F
FPC 5 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 2 Chip	OK	48 degrees C / 118 degrees F
FPC 5 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 6 Intake	OK	42 degrees C / 107 degrees F
FPC 6 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 6 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 6 LU 0 TSen	OK	61 degrees C / 141 degrees F
FPC 6 LU 0 Chip	OK	64 degrees C / 147 degrees F
FPC 6 LU 1 TSen	OK	61 degrees C / 141 degrees F
FPC 6 LU 1 Chip	OK	66 degrees C / 150 degrees F
FPC 6 LU 2 TSen	OK	61 degrees C / 141 degrees F
FPC 6 LU 2 Chip	OK	56 degrees C / 132 degrees F
FPC 6 LU 3 TSen	OK	61 degrees C / 141 degrees F
FPC 6 LU 3 Chip	OK	56 degrees C / 132 degrees F
FPC 6 MQ 0 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 0 Chip	OK	56 degrees C / 132 degrees F
FPC 6 MQ 1 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 1 Chip	OK	59 degrees C / 138 degrees F
FPC 6 MQ 2 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 2 Chip	OK	49 degrees C / 120 degrees F
FPC 6 MQ 3 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 3 Chip	OK	49 degrees C / 120 degrees F
FPC 7 Intake	OK	41 degrees C / 105 degrees F
FPC 7 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 7 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 7 LU 0 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 0 Chip	OK	61 degrees C / 141 degrees F
FPC 7 LU 1 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 1 Chip	OK	65 degrees C / 149 degrees F
FPC 7 LU 2 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 2 Chip	OK	54 degrees C / 129 degrees F
FPC 7 LU 3 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 7 MQ 0 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 0 Chip	OK	53 degrees C / 127 degrees F
FPC 7 MQ 1 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 1 Chip	OK	54 degrees C / 129 degrees F
FPC 7 MQ 2 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 2 Chip	OK	47 degrees C / 116 degrees F
FPC 7 MQ 3 TSen	OK	50 degrees C / 122 degrees F

FPC 7 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 8 Intake	OK	41 degrees C / 105 degrees F
FPC 8 Exhaust A	OK	50 degrees C / 122 degrees F
FPC 8 Exhaust B	OK	62 degrees C / 143 degrees F
FPC 8 LU 0 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 0 Chip	OK	62 degrees C / 143 degrees F
FPC 8 LU 1 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 1 Chip	OK	64 degrees C / 147 degrees F
FPC 8 LU 2 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 2 Chip	OK	55 degrees C / 131 degrees F
FPC 8 LU 3 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 8 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 0 Chip	OK	51 degrees C / 123 degrees F
FPC 8 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 1 Chip	OK	52 degrees C / 125 degrees F
FPC 8 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 2 Chip	OK	46 degrees C / 114 degrees F
FPC 8 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 9 Intake	OK	42 degrees C / 107 degrees F
FPC 9 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 9 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 9 LU 0 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 0 Chip	OK	65 degrees C / 149 degrees F
FPC 9 LU 1 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 1 Chip	OK	67 degrees C / 152 degrees F
FPC 9 LU 2 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 2 Chip	OK	54 degrees C / 129 degrees F
FPC 9 LU 3 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 9 MQ 0 TSen	OK	51 degrees C / 123 degrees F
FPC 9 MQ 0 Chip	OK	55 degrees C / 131 degrees F
FPC 9 MQ 1 TSen	OK	51 degrees C / 123 degrees F
FPC 9 MQ 1 Chip	OK	59 degrees C / 138 degrees F
FPC 9 MQ 2 TSen	OK	51 degrees C / 123 degrees F
FPC 9 MQ 2 Chip	OK	49 degrees C / 120 degrees F
FPC 9 MQ 3 TSen	OK	51 degrees C / 123 degrees F
FPC 9 MQ 3 Chip	OK	49 degrees C / 120 degrees F
FPC 10 Intake	OK	44 degrees C / 111 degrees F
FPC 10 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 10 Exhaust B	OK	55 degrees C / 131 degrees F
FPC 10 LU 0 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 0 Chip	OK	55 degrees C / 131 degrees F
FPC 10 LU 1 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 1 Chip	OK	59 degrees C / 138 degrees F
FPC 10 LU 2 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 2 Chip	OK	52 degrees C / 125 degrees F
FPC 10 LU 3 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 3 Chip	OK	51 degrees C / 123 degrees F
FPC 10 MQ 0 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 0 Chip	OK	49 degrees C / 120 degrees F
FPC 10 MQ 1 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 1 Chip	OK	52 degrees C / 125 degrees F
FPC 10 MQ 2 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 2 Chip	OK	47 degrees C / 116 degrees F
FPC 10 MQ 3 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 11 Intake	OK	30 degrees C / 86 degrees F
FPC 11 Exhaust A	OK	35 degrees C / 95 degrees F
FPC 11 Exhaust B	OK	30 degrees C / 86 degrees F

FPC 11 LU 0 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 0 Chip	OK	58 degrees C / 136 degrees F
FPC 11 LU 1 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 1 Chip	OK	62 degrees C / 143 degrees F
FPC 11 LU 2 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 11 LU 3 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 11 MQ 0 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 0 Chip	OK	52 degrees C / 125 degrees F
FPC 11 MQ 1 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 1 Chip	OK	57 degrees C / 134 degrees F
FPC 11 MQ 2 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 2 Chip	OK	48 degrees C / 118 degrees F
FPC 11 MQ 3 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 3 Chip	OK	52 degrees C / 125 degrees F
FPC 12 Intake	OK	40 degrees C / 104 degrees F
FPC 12 Exhaust A	OK	47 degrees C / 116 degrees F
FPC 12 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 12 LU 0 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 0 Chip	OK	52 degrees C / 125 degrees F
FPC 12 LU 1 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 1 Chip	OK	55 degrees C / 131 degrees F
FPC 12 LU 2 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 2 Chip	OK	47 degrees C / 116 degrees F
FPC 12 LU 3 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 3 Chip	OK	50 degrees C / 122 degrees F
FPC 12 MQ 0 TSen	OK	46 degrees C / 114 degrees F
FPC 12 MQ 0 Chip	OK	46 degrees C / 114 degrees F
FPC 12 MQ 1 TSen	OK	46 degrees C / 114 degrees F
FPC 12 MQ 1 Chip	OK	50 degrees C / 122 degrees F
FPC 12 MQ 2 TSen	OK	46 degrees C / 114 degrees F
FPC 12 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 12 MQ 3 TSen	OK	46 degrees C / 114 degrees F
FPC 12 MQ 3 Chip	OK	46 degrees C / 114 degrees F
FPC 13 Intake	OK	40 degrees C / 104 degrees F
FPC 13 Exhaust A	OK	48 degrees C / 118 degrees F
FPC 13 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 13 LU 0 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 0 Chip	OK	52 degrees C / 125 degrees F
FPC 13 LU 1 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 1 Chip	OK	55 degrees C / 131 degrees F
FPC 13 LU 2 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 2 Chip	OK	48 degrees C / 118 degrees F
FPC 13 LU 3 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 3 Chip	OK	48 degrees C / 118 degrees F
FPC 13 MQ 0 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 0 Chip	OK	46 degrees C / 114 degrees F
FPC 13 MQ 1 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 1 Chip	OK	50 degrees C / 122 degrees F
FPC 13 MQ 2 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 13 MQ 3 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 3 Chip	OK	46 degrees C / 114 degrees F
FPC 14 Intake	OK	40 degrees C / 104 degrees F
FPC 14 Exhaust A	OK	50 degrees C / 122 degrees F
FPC 14 Exhaust B	OK	51 degrees C / 123 degrees F
FPC 14 LU 0 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 0 Chip	OK	50 degrees C / 122 degrees F
FPC 14 LU 1 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 1 Chip	OK	54 degrees C / 129 degrees F

FPC 14 LU 2 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 2 Chip	OK	47 degrees C / 116 degrees F
FPC 14 LU 3 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 3 Chip	OK	49 degrees C / 120 degrees F
FPC 14 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 0 Chip	OK	46 degrees C / 114 degrees F
FPC 14 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 1 Chip	OK	51 degrees C / 123 degrees F
FPC 14 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 2 Chip	OK	45 degrees C / 113 degrees F
FPC 14 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 3 Chip	OK	48 degrees C / 118 degrees F
FPC 15 Intake	OK	44 degrees C / 111 degrees F
FPC 15 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 15 Exhaust B	OK	60 degrees C / 140 degrees F
FPC 15 LU 0 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 0 Chip	OK	56 degrees C / 132 degrees F
FPC 15 LU 1 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 1 Chip	OK	50 degrees C / 122 degrees F
FPC 15 LU 2 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 2 Chip	OK	58 degrees C / 136 degrees F
FPC 15 LU 3 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 3 Chip	OK	63 degrees C / 145 degrees F
FPC 15 XM 0 TSen	OK	50 degrees C / 122 degrees F
FPC 15 XM 0 Chip	OK	56 degrees C / 132 degrees F
FPC 15 XF 0 TSen	OK	50 degrees C / 122 degrees F
FPC 15 XF 0 Chip	OK	68 degrees C / 154 degrees F
FPC 15 PLX Switch TSen	OK	50 degrees C / 122 degrees F
FPC 15 PLX Switch Chip	OK	56 degrees C / 132 degrees F
FPC 16 Intake	OK	42 degrees C / 107 degrees F
FPC 16 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 16 Exhaust B	OK	53 degrees C / 127 degrees F
FPC 16 LU 0 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 0 Chip	OK	52 degrees C / 125 degrees F
FPC 16 LU 1 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 1 Chip	OK	55 degrees C / 131 degrees F
FPC 16 LU 2 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 2 Chip	OK	48 degrees C / 118 degrees F
FPC 16 LU 3 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 3 Chip	OK	49 degrees C / 120 degrees F
FPC 16 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 0 Chip	OK	48 degrees C / 118 degrees F
FPC 16 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 1 Chip	OK	53 degrees C / 127 degrees F
FPC 16 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 2 Chip	OK	46 degrees C / 114 degrees F
FPC 16 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 3 Chip	OK	49 degrees C / 120 degrees F
FPC 17 Intake	OK	43 degrees C / 109 degrees F
FPC 17 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 17 Exhaust B	OK	55 degrees C / 131 degrees F
FPC 17 LU 0 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 0 Chip	OK	57 degrees C / 134 degrees F
FPC 17 LU 1 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 1 Chip	OK	60 degrees C / 140 degrees F
FPC 17 LU 2 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 17 LU 3 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 17 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 0 Chip	OK	50 degrees C / 122 degrees F

FPC 17 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 1 Chip	OK	54 degrees C / 129 degrees F
FPC 17 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 2 Chip	OK	47 degrees C / 116 degrees F
FPC 17 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 3 Chip	OK	51 degrees C / 123 degrees F
FPC 18 Intake	OK	44 degrees C / 111 degrees F
FPC 18 Exhaust A	OK	53 degrees C / 127 degrees F
FPC 18 Exhaust B	OK	57 degrees C / 134 degrees F
FPC 18 LU 0 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 0 Chip	OK	57 degrees C / 134 degrees F
FPC 18 LU 1 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 1 Chip	OK	62 degrees C / 143 degrees F
FPC 18 LU 2 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 18 LU 3 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 3 Chip	OK	55 degrees C / 131 degrees F
FPC 18 MQ 0 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 0 Chip	OK	54 degrees C / 129 degrees F
FPC 18 MQ 1 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 1 Chip	OK	58 degrees C / 136 degrees F
FPC 18 MQ 2 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 2 Chip	OK	50 degrees C / 122 degrees F
FPC 18 MQ 3 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 3 Chip	OK	53 degrees C / 127 degrees F
FPC 19 Intake	OK	48 degrees C / 118 degrees F
FPC 19 Exhaust A	OK	56 degrees C / 132 degrees F
FPC 19 Exhaust B	OK	64 degrees C / 147 degrees F
FPC 19 LU 0 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 0 Chip	OK	64 degrees C / 147 degrees F
FPC 19 LU 1 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 1 Chip	OK	70 degrees C / 158 degrees F
FPC 19 LU 2 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 2 Chip	OK	61 degrees C / 141 degrees F
FPC 19 LU 3 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 3 Chip	OK	62 degrees C / 143 degrees F
FPC 19 MQ 0 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 0 Chip	OK	60 degrees C / 140 degrees F
FPC 19 MQ 1 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 1 Chip	OK	62 degrees C / 143 degrees F
FPC 19 MQ 2 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 2 Chip	OK	56 degrees C / 132 degrees F
FPC 19 MQ 3 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 3 Chip	OK	57 degrees C / 134 degrees F
ADC 0 Intake	OK	40 degrees C / 104 degrees F
ADC 0 Exhaust	OK	52 degrees C / 125 degrees F
ADC 0 ADC-XF1	OK	59 degrees C / 138 degrees F
ADC 0 ADC-XF0	OK	66 degrees C / 150 degrees F
ADC 1 Intake	OK	38 degrees C / 100 degrees F
ADC 1 Exhaust	OK	50 degrees C / 122 degrees F
ADC 1 ADC-XF1	OK	59 degrees C / 138 degrees F
ADC 1 ADC-XF0	OK	63 degrees C / 145 degrees F
ADC 2 Intake	OK	37 degrees C / 98 degrees F
ADC 2 Exhaust	OK	52 degrees C / 125 degrees F
ADC 2 ADC-XF1	OK	53 degrees C / 127 degrees F
ADC 2 ADC-XF0	OK	61 degrees C / 141 degrees F
ADC 3 Intake	OK	40 degrees C / 104 degrees F
ADC 3 Exhaust	OK	51 degrees C / 123 degrees F
ADC 3 ADC-XF1	OK	61 degrees C / 141 degrees F
ADC 3 ADC-XF0	OK	64 degrees C / 147 degrees F
ADC 4 Intake	OK	39 degrees C / 102 degrees F

ADC 4 Exhaust	OK	51 degrees C / 123 degrees F
ADC 4 ADC-XF1	OK	60 degrees C / 140 degrees F
ADC 4 ADC-XF0	OK	63 degrees C / 145 degrees F
ADC 5 Intake	OK	38 degrees C / 100 degrees F
ADC 5 Exhaust	OK	54 degrees C / 129 degrees F
ADC 5 ADC-XF1	OK	56 degrees C / 132 degrees F
ADC 5 ADC-XF0	OK	67 degrees C / 152 degrees F
ADC 6 Intake	OK	39 degrees C / 102 degrees F
ADC 6 Exhaust	OK	52 degrees C / 125 degrees F
ADC 6 ADC-XF1	OK	59 degrees C / 138 degrees F
ADC 6 ADC-XF0	OK	66 degrees C / 150 degrees F
ADC 7 Intake	OK	39 degrees C / 102 degrees F
ADC 7 Exhaust	OK	54 degrees C / 129 degrees F
ADC 7 ADC-XF1	OK	62 degrees C / 143 degrees F
ADC 7 ADC-XF0	OK	70 degrees C / 158 degrees F
ADC 8 Intake	OK	39 degrees C / 102 degrees F
ADC 8 Exhaust	OK	52 degrees C / 125 degrees F
ADC 8 ADC-XF1	OK	61 degrees C / 141 degrees F
ADC 8 ADC-XF0	OK	65 degrees C / 149 degrees F
ADC 9 Intake	OK	41 degrees C / 105 degrees F
ADC 9 Exhaust	OK	51 degrees C / 123 degrees F
ADC 9 ADC-XF1	OK	63 degrees C / 145 degrees F
ADC 9 ADC-XF0	OK	63 degrees C / 145 degrees F
ADC 10 Intake	OK	48 degrees C / 118 degrees F
ADC 10 Exhaust	OK	53 degrees C / 127 degrees F
ADC 10 ADC-XF1	OK	67 degrees C / 152 degrees F
ADC 10 ADC-XF0	OK	66 degrees C / 150 degrees F
ADC 12 Intake	OK	49 degrees C / 120 degrees F
ADC 12 Exhaust	OK	54 degrees C / 129 degrees F
ADC 12 ADC-XF1	OK	67 degrees C / 152 degrees F
ADC 12 ADC-XF0	OK	67 degrees C / 152 degrees F
ADC 13 Intake	OK	49 degrees C / 120 degrees F
ADC 13 Exhaust	OK	57 degrees C / 134 degrees F
ADC 13 ADC-XF1	OK	66 degrees C / 150 degrees F
ADC 13 ADC-XF0	OK	69 degrees C / 156 degrees F
ADC 14 Intake	OK	51 degrees C / 123 degrees F
ADC 14 Exhaust	OK	59 degrees C / 138 degrees F
ADC 14 ADC-XF1	OK	69 degrees C / 156 degrees F
ADC 14 ADC-XF0	OK	74 degrees C / 165 degrees F
ADC 15 Intake	OK	50 degrees C / 122 degrees F
ADC 15 Exhaust	OK	59 degrees C / 138 degrees F
ADC 15 ADC-XF1	OK	68 degrees C / 154 degrees F
ADC 15 ADC-XF0	OK	69 degrees C / 156 degrees F
ADC 16 Intake	OK	52 degrees C / 125 degrees F
ADC 16 Exhaust	OK	58 degrees C / 136 degrees F
ADC 16 ADC-XF1	OK	68 degrees C / 154 degrees F
ADC 16 ADC-XF0	OK	70 degrees C / 158 degrees F
ADC 17 Intake	OK	52 degrees C / 125 degrees F
ADC 17 Exhaust	OK	59 degrees C / 138 degrees F
ADC 17 ADC-XF1	OK	69 degrees C / 156 degrees F
ADC 17 ADC-XF0	OK	71 degrees C / 159 degrees F
ADC 18 Intake	OK	53 degrees C / 127 degrees F
ADC 18 Exhaust	OK	59 degrees C / 138 degrees F
ADC 18 ADC-XF1	OK	68 degrees C / 154 degrees F
ADC 18 ADC-XF0	OK	73 degrees C / 163 degrees F
ADC 19 Intake	OK	50 degrees C / 122 degrees F
ADC 19 Exhaust	OK	59 degrees C / 138 degrees F
ADC 19 ADC-XF1	OK	68 degrees C / 154 degrees F
ADC 19 ADC-XF0	OK	72 degrees C / 161 degrees F
Fans Fan Tray 0 Fan 1	OK	7440 RPM
Fans Fan Tray 0 Fan 2	OK	7200 RPM

Fan Tray 0 Fan 3	OK	6960 RPM
Fan Tray 0 Fan 4	OK	7200 RPM
Fan Tray 0 Fan 5	OK	7080 RPM
Fan Tray 0 Fan 6	OK	6840 RPM
Fan Tray 1 Fan 1	OK	6840 RPM
Fan Tray 1 Fan 2	OK	6960 RPM
Fan Tray 1 Fan 3	OK	6960 RPM
Fan Tray 1 Fan 4	OK	7080 RPM
Fan Tray 1 Fan 5	OK	6960 RPM
Fan Tray 1 Fan 6	OK	6960 RPM
Fan Tray 2 Fan 1	OK	8640 RPM
Fan Tray 2 Fan 2	OK	8640 RPM
Fan Tray 2 Fan 3	OK	8760 RPM
Fan Tray 2 Fan 4	OK	8760 RPM
Fan Tray 2 Fan 5	OK	8640 RPM
Fan Tray 2 Fan 6	OK	8640 RPM
Fan Tray 3 Fan 1	OK	8520 RPM
Fan Tray 3 Fan 2	OK	8520 RPM
Fan Tray 3 Fan 3	OK	8640 RPM
Fan Tray 3 Fan 4	OK	8640 RPM
Fan Tray 3 Fan 5	OK	8520 RPM
Fan Tray 3 Fan 6	OK	8520 RPM

show chassis environment (MX2010 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PSM 0	OK	7 degrees C / 44 degrees F
	PSM 1	OK	7 degrees C / 44 degrees F
	PSM 2	OK	7 degrees C / 44 degrees F
	PSM 3	OK	6 degrees C / 42 degrees F
	PSM 4	OK	6 degrees C / 42 degrees F
	PSM 5	OK	6 degrees C / 42 degrees F
	PSM 6	OK	6 degrees C / 42 degrees F
	PSM 7	OK	7 degrees C / 44 degrees F
	PSM 8	OK	7 degrees C / 44 degrees F
	PDM 0	OK	
	PDM 1	Absent	
	CB 0 IntakeA-Zone0	OK	14 degrees C / 57 degrees F
	CB 0 IntakeB-Zone1	OK	7 degrees C / 44 degrees F
	CB 0 IntakeC-Zone0	OK	22 degrees C / 71 degrees F
	CB 0 ExhaustA-Zone0	OK	14 degrees C / 57 degrees F
	CB 0 ExhaustB-Zone1	OK	9 degrees C / 48 degrees F
	CB 0 TCBC-Zone0	OK	11 degrees C / 51 degrees F
	CB 1 IntakeA-Zone0	OK	9 degrees C / 48 degrees F
	CB 1 IntakeB-Zone1	OK	5 degrees C / 41 degrees F
	CB 1 IntakeC-Zone0	OK	20 degrees C / 68 degrees F
	CB 1 ExhaustA-Zone0	OK	12 degrees C / 53 degrees F
	CB 1 ExhaustB-Zone1	OK	7 degrees C / 44 degrees F
	CB 1 TCBC-Zone0	OK	10 degrees C / 50 degrees F
	SPMB 0 Intake	OK	5 degrees C / 41 degrees F
	SPMB 1 Intake	OK	4 degrees C / 39 degrees F
	Routing Engine 0	OK	9 degrees C / 48 degrees F
	Routing Engine 0 CPU	OK	9 degrees C / 48 degrees F
	Routing Engine 1	OK	6 degrees C / 42 degrees F
	Routing Engine 1 CPU	OK	6 degrees C / 42 degrees F
	SFB 0 Intake-Zone0	OK	26 degrees C / 78 degrees F
	SFB 0 Exhaust-Zone1	OK	17 degrees C / 62 degrees F
	SFB 0 IntakeA-Zone0	OK	16 degrees C / 60 degrees F
	SFB 0 IntakeB-Zone1	OK	11 degrees C / 51 degrees F
	SFB 0 Exhaust-Zone0	OK	18 degrees C / 64 degrees F

SFB 0	SFB-XF2-Zone1	OK	25 degrees C / 77 degrees F
SFB 0	SFB-XF1-Zone0	OK	23 degrees C / 73 degrees F
SFB 0	SFB-XF0-Zone0	OK	33 degrees C / 91 degrees F
SFB 1	Intake-Zone0	OK	27 degrees C / 80 degrees F
SFB 1	Exhaust-Zone1	OK	15 degrees C / 59 degrees F
SFB 1	IntakeA-Zone0	OK	20 degrees C / 68 degrees F
SFB 1	IntakeB-Zone1	OK	10 degrees C / 50 degrees F
SFB 1	Exhaust-Zone0	OK	19 degrees C / 66 degrees F
SFB 1	SFB-XF2-Zone1	OK	26 degrees C / 78 degrees F
SFB 1	SFB-XF1-Zone0	OK	27 degrees C / 80 degrees F
SFB 1	SFB-XF0-Zone0	OK	32 degrees C / 89 degrees F
SFB 2	Intake-Zone0	OK	21 degrees C / 69 degrees F
SFB 2	Exhaust-Zone1	OK	13 degrees C / 55 degrees F
SFB 2	IntakeA-Zone0	OK	18 degrees C / 64 degrees F
SFB 2	IntakeB-Zone1	OK	9 degrees C / 48 degrees F
SFB 2	Exhaust-Zone0	OK	16 degrees C / 60 degrees F
SFB 2	SFB-XF2-Zone1	OK	24 degrees C / 75 degrees F
SFB 2	SFB-XF1-Zone0	OK	21 degrees C / 69 degrees F
SFB 2	SFB-XF0-Zone0	OK	26 degrees C / 78 degrees F
SFB 4	Intake-Zone0	OK	28 degrees C / 82 degrees F
SFB 4	Exhaust-Zone1	OK	16 degrees C / 60 degrees F
SFB 4	IntakeA-Zone0	OK	18 degrees C / 64 degrees F
SFB 4	IntakeB-Zone1	OK	11 degrees C / 51 degrees F
SFB 4	Exhaust-Zone0	OK	19 degrees C / 66 degrees F
SFB 4	SFB-XF2-Zone1	OK	27 degrees C / 80 degrees F
SFB 4	SFB-XF1-Zone0	OK	27 degrees C / 80 degrees F
SFB 4	SFB-XF0-Zone0	OK	32 degrees C / 89 degrees F
SFB 5	Intake-Zone0	OK	22 degrees C / 71 degrees F
SFB 5	Exhaust-Zone1	OK	14 degrees C / 57 degrees F
SFB 5	IntakeA-Zone0	OK	18 degrees C / 64 degrees F
SFB 5	IntakeB-Zone1	OK	10 degrees C / 50 degrees F
SFB 5	Exhaust-Zone0	OK	17 degrees C / 62 degrees F
SFB 5	SFB-XF2-Zone1	OK	22 degrees C / 71 degrees F
SFB 5	SFB-XF1-Zone0	OK	29 degrees C / 84 degrees F
SFB 5	SFB-XF0-Zone0	OK	27 degrees C / 80 degrees F
SFB 6	Intake-Zone0	OK	27 degrees C / 80 degrees F
SFB 6	Exhaust-Zone1	OK	13 degrees C / 55 degrees F
SFB 6	IntakeA-Zone0	OK	19 degrees C / 66 degrees F
SFB 6	IntakeB-Zone1	OK	10 degrees C / 50 degrees F
SFB 6	Exhaust-Zone0	OK	20 degrees C / 68 degrees F
SFB 6	SFB-XF2-Zone1	OK	24 degrees C / 75 degrees F
SFB 6	SFB-XF1-Zone0	OK	32 degrees C / 89 degrees F
SFB 6	SFB-XF0-Zone0	OK	33 degrees C / 91 degrees F
SFB 7	Intake-Zone0	OK	25 degrees C / 77 degrees F
SFB 7	Exhaust-Zone1	OK	13 degrees C / 55 degrees F
SFB 7	IntakeA-Zone0	OK	14 degrees C / 57 degrees F
SFB 7	IntakeB-Zone1	OK	8 degrees C / 46 degrees F
SFB 7	Exhaust-Zone0	OK	17 degrees C / 62 degrees F
SFB 7	SFB-XF2-Zone1	OK	21 degrees C / 69 degrees F
SFB 7	SFB-XF1-Zone0	OK	21 degrees C / 69 degrees F
SFB 7	SFB-XF0-Zone0	OK	33 degrees C / 91 degrees F
FPC 0	Intake	OK	13 degrees C / 55 degrees F
FPC 0	Exhaust A	OK	13 degrees C / 55 degrees F
FPC 0	Exhaust B	OK	14 degrees C / 57 degrees F
FPC 0	LU 0 TSen	OK	28 degrees C / 82 degrees F
FPC 0	LU 0 Chip	OK	25 degrees C / 77 degrees F
FPC 0	LU 1 TSen	OK	28 degrees C / 82 degrees F
FPC 0	LU 1 Chip	OK	27 degrees C / 80 degrees F
FPC 0	LU 2 TSen	OK	28 degrees C / 82 degrees F
FPC 0	LU 2 Chip	OK	19 degrees C / 66 degrees F
FPC 0	LU 3 TSen	OK	28 degrees C / 82 degrees F

FPC 0 LU 3 Chip	OK	23 degrees C / 73 degrees F
FPC 0 XM 0 TSen	OK	28 degrees C / 82 degrees F
FPC 0 XM 0 Chip	OK	33 degrees C / 91 degrees F
FPC 0 XM 1 TSen	OK	28 degrees C / 82 degrees F
FPC 0 XM 1 Chip	OK	26 degrees C / 78 degrees F
FPC 0 PLX Switch TSen	OK	28 degrees C / 82 degrees F
FPC 0 PLX Switch Chip	OK	26 degrees C / 78 degrees F
FPC 1 Intake	OK	10 degrees C / 50 degrees F
FPC 1 Exhaust A	OK	24 degrees C / 75 degrees F
FPC 1 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 1 LU 0 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 0 Chip	OK	31 degrees C / 87 degrees F
FPC 1 LU 1 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 1 Chip	OK	21 degrees C / 69 degrees F
FPC 1 LU 2 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 2 Chip	OK	25 degrees C / 77 degrees F
FPC 1 LU 3 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 3 Chip	OK	33 degrees C / 91 degrees F
FPC 1 XM 0 TSen	OK	22 degrees C / 71 degrees F
FPC 1 XM 0 Chip	OK	30 degrees C / 86 degrees F
FPC 1 XF 0 TSen	OK	22 degrees C / 71 degrees F
FPC 1 XF 0 Chip	OK	37 degrees C / 98 degrees F
FPC 1 PLX Switch TSen	OK	22 degrees C / 71 degrees F
FPC 1 PLX Switch Chip	OK	22 degrees C / 71 degrees F
FPC 2 Intake	OK	9 degrees C / 48 degrees F
FPC 2 Exhaust A	OK	10 degrees C / 50 degrees F
FPC 2 Exhaust B	OK	10 degrees C / 50 degrees F
FPC 2 LU 0 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 0 Chip	OK	25 degrees C / 77 degrees F
FPC 2 LU 1 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 1 Chip	OK	26 degrees C / 78 degrees F
FPC 2 LU 2 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 2 Chip	OK	17 degrees C / 62 degrees F
FPC 2 LU 3 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 3 Chip	OK	22 degrees C / 71 degrees F
FPC 2 XM 0 TSen	OK	26 degrees C / 78 degrees F
FPC 2 XM 0 Chip	OK	34 degrees C / 93 degrees F
FPC 2 XM 1 TSen	OK	26 degrees C / 78 degrees F
FPC 2 XM 1 Chip	OK	26 degrees C / 78 degrees F
FPC 2 PLX Switch TSen	OK	26 degrees C / 78 degrees F
FPC 2 PLX Switch Chip	OK	20 degrees C / 68 degrees F
FPC 3 Intake	OK	12 degrees C / 53 degrees F
FPC 3 Exhaust A	OK	16 degrees C / 60 degrees F
FPC 3 Exhaust B	OK	26 degrees C / 78 degrees F
FPC 3 LU 0 TSen	OK	23 degrees C / 73 degrees F
FPC 3 LU 0 Chip	OK	26 degrees C / 78 degrees F
FPC 3 LU 1 TSen	OK	23 degrees C / 73 degrees F
FPC 3 LU 1 Chip	OK	27 degrees C / 80 degrees F
FPC 3 LU 2 TSen	OK	23 degrees C / 73 degrees F
FPC 3 LU 2 Chip	OK	22 degrees C / 71 degrees F
FPC 3 LU 3 TSen	OK	23 degrees C / 73 degrees F
FPC 3 LU 3 Chip	OK	21 degrees C / 69 degrees F
FPC 3 MQ 0 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 0 Chip	OK	18 degrees C / 64 degrees F
FPC 3 MQ 1 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 1 Chip	OK	20 degrees C / 68 degrees F
FPC 3 MQ 2 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 2 Chip	OK	17 degrees C / 62 degrees F
FPC 3 MQ 3 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 3 Chip	OK	16 degrees C / 60 degrees F
FPC 4 Intake	OK	11 degrees C / 51 degrees F

FPC 4 Exhaust A	OK	22 degrees C / 71 degrees F
FPC 4 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 4 LU 0 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 0 Chip	OK	33 degrees C / 91 degrees F
FPC 4 LU 1 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 1 Chip	OK	21 degrees C / 69 degrees F
FPC 4 LU 2 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 2 Chip	OK	26 degrees C / 78 degrees F
FPC 4 LU 3 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 3 Chip	OK	33 degrees C / 91 degrees F
FPC 4 XM 0 TSen	OK	22 degrees C / 71 degrees F
FPC 4 XM 0 Chip	OK	30 degrees C / 86 degrees F
FPC 4 XF 0 TSen	OK	22 degrees C / 71 degrees F
FPC 4 XF 0 Chip	OK	37 degrees C / 98 degrees F
FPC 4 PLX Switch TSen	OK	22 degrees C / 71 degrees F
FPC 4 PLX Switch Chip	OK	23 degrees C / 73 degrees F
FPC 5 Intake	OK	12 degrees C / 53 degrees F
FPC 5 Exhaust A	OK	12 degrees C / 53 degrees F
FPC 5 Exhaust B	OK	12 degrees C / 53 degrees F
FPC 5 LU 0 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 0 Chip	OK	28 degrees C / 82 degrees F
FPC 5 LU 1 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 1 Chip	OK	27 degrees C / 80 degrees F
FPC 5 LU 2 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 2 Chip	OK	19 degrees C / 66 degrees F
FPC 5 LU 3 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 3 Chip	OK	22 degrees C / 71 degrees F
FPC 5 XM 0 TSen	OK	27 degrees C / 80 degrees F
FPC 5 XM 0 Chip	OK	36 degrees C / 96 degrees F
FPC 5 XM 1 TSen	OK	27 degrees C / 80 degrees F
FPC 5 XM 1 Chip	OK	26 degrees C / 78 degrees F
FPC 5 PLX Switch TSen	OK	27 degrees C / 80 degrees F
FPC 5 PLX Switch Chip	OK	24 degrees C / 75 degrees F
FPC 6 Intake	OK	12 degrees C / 53 degrees F
FPC 6 Exhaust A	OK	17 degrees C / 62 degrees F
FPC 6 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 6 LU 0 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 0 Chip	OK	29 degrees C / 84 degrees F
FPC 6 LU 1 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 1 Chip	OK	30 degrees C / 86 degrees F
FPC 6 LU 2 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 2 Chip	OK	24 degrees C / 75 degrees F
FPC 6 LU 3 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 3 Chip	OK	22 degrees C / 71 degrees F
FPC 6 MQ 0 TSen	OK	16 degrees C / 60 degrees F
FPC 6 MQ 0 Chip	OK	19 degrees C / 66 degrees F
FPC 6 MQ 1 TSen	OK	16 degrees C / 60 degrees F
FPC 6 MQ 1 Chip	OK	20 degrees C / 68 degrees F
FPC 6 MQ 2 TSen	OK	16 degrees C / 60 degrees F
FPC 6 MQ 2 Chip	OK	17 degrees C / 62 degrees F
FPC 6 MQ 3 TSen	OK	16 degrees C / 60 degrees F
FPC 6 MQ 3 Chip	OK	16 degrees C / 60 degrees F
FPC 7 Intake	OK	10 degrees C / 50 degrees F
FPC 7 Exhaust A	OK	10 degrees C / 50 degrees F
FPC 7 Exhaust B	OK	11 degrees C / 51 degrees F
FPC 7 LU 0 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 0 Chip	OK	26 degrees C / 78 degrees F
FPC 7 LU 1 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 1 Chip	OK	29 degrees C / 84 degrees F
FPC 7 LU 2 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 2 Chip	OK	19 degrees C / 66 degrees F

FPC 7 LU 3 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 3 Chip	OK	24 degrees C / 75 degrees F
FPC 7 XM 0 TSen	OK	26 degrees C / 78 degrees F
FPC 7 XM 0 Chip	OK	34 degrees C / 93 degrees F
FPC 7 XM 1 TSen	OK	26 degrees C / 78 degrees F
FPC 7 XM 1 Chip	OK	32 degrees C / 89 degrees F
FPC 7 PLX Switch TSen	OK	26 degrees C / 78 degrees F
FPC 7 PLX Switch Chip	OK	22 degrees C / 71 degrees F
FPC 8 Intake	OK	10 degrees C / 50 degrees F
FPC 8 Exhaust A	OK	22 degrees C / 71 degrees F
FPC 8 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 8 LU 0 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 0 Chip	OK	33 degrees C / 91 degrees F
FPC 8 LU 1 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 1 Chip	OK	23 degrees C / 73 degrees F
FPC 8 LU 2 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 2 Chip	OK	26 degrees C / 78 degrees F
FPC 8 LU 3 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 3 Chip	OK	33 degrees C / 91 degrees F
FPC 8 XM 0 TSen	OK	20 degrees C / 68 degrees F
FPC 8 XM 0 Chip	OK	29 degrees C / 84 degrees F
FPC 8 XF 0 TSen	OK	20 degrees C / 68 degrees F
FPC 8 XF 0 Chip	OK	38 degrees C / 100 degrees F
FPC 8 PLX Switch TSen	OK	20 degrees C / 68 degrees F
FPC 8 PLX Switch Chip	OK	24 degrees C / 75 degrees F
FPC 9 Intake	OK	11 degrees C / 51 degrees F
FPC 9 Exhaust A	OK	11 degrees C / 51 degrees F
FPC 9 Exhaust B	OK	11 degrees C / 51 degrees F
FPC 9 LU 0 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 0 Chip	OK	24 degrees C / 75 degrees F
FPC 9 LU 1 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 1 Chip	OK	26 degrees C / 78 degrees F
FPC 9 LU 2 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 2 Chip	OK	16 degrees C / 60 degrees F
FPC 9 LU 3 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 3 Chip	OK	21 degrees C / 69 degrees F
FPC 9 XM 0 TSen	OK	25 degrees C / 77 degrees F
FPC 9 XM 0 Chip	OK	32 degrees C / 89 degrees F
FPC 9 XM 1 TSen	OK	25 degrees C / 77 degrees F
FPC 9 XM 1 Chip	OK	25 degrees C / 77 degrees F
FPC 9 PLX Switch TSen	OK	25 degrees C / 77 degrees F
FPC 9 PLX Switch Chip	OK	21 degrees C / 69 degrees F
ADC 0 Intake	OK	12 degrees C / 53 degrees F
ADC 0 Exhaust	OK	20 degrees C / 68 degrees F
ADC 0 ADC-XF1	OK	26 degrees C / 78 degrees F
ADC 0 ADC-XF0	OK	32 degrees C / 89 degrees F
ADC 1 Intake	OK	11 degrees C / 51 degrees F
ADC 1 Exhaust	OK	21 degrees C / 69 degrees F
ADC 1 ADC-XF1	OK	24 degrees C / 75 degrees F
ADC 1 ADC-XF0	OK	31 degrees C / 87 degrees F
ADC 2 Intake	OK	14 degrees C / 57 degrees F
ADC 2 Exhaust	OK	21 degrees C / 69 degrees F
ADC 2 ADC-XF1	OK	28 degrees C / 82 degrees F
ADC 2 ADC-XF0	OK	34 degrees C / 93 degrees F
ADC 3 Intake	OK	13 degrees C / 55 degrees F
ADC 3 Exhaust	OK	19 degrees C / 66 degrees F
ADC 3 ADC-XF1	OK	24 degrees C / 75 degrees F
ADC 3 ADC-XF0	OK	31 degrees C / 87 degrees F
ADC 4 Intake	OK	9 degrees C / 48 degrees F
ADC 4 Exhaust	OK	22 degrees C / 71 degrees F
ADC 4 ADC-XF1	OK	28 degrees C / 82 degrees F

	ADC 4 ADC-XF0	OK	35 degrees C / 95 degrees F
	ADC 5 Intake	OK	12 degrees C / 53 degrees F
	ADC 5 Exhaust	OK	22 degrees C / 71 degrees F
	ADC 5 ADC-XF1	OK	28 degrees C / 82 degrees F
	ADC 5 ADC-XF0	OK	34 degrees C / 93 degrees F
	ADC 6 Intake	OK	11 degrees C / 51 degrees F
	ADC 6 Exhaust	OK	21 degrees C / 69 degrees F
	ADC 6 ADC-XF1	OK	26 degrees C / 78 degrees F
ADC 6	ADC-XF0	OK	35 degrees C / 95 degrees F
	ADC 7 Intake	OK	14 degrees C / 57 degrees F
	ADC 7 Exhaust	OK	22 degrees C / 71 degrees F
	ADC 7 ADC-XF1	OK	26 degrees C / 78 degrees F
	ADC 7 ADC-XF0	OK	34 degrees C / 93 degrees F
	ADC 8 Intake	OK	14 degrees C / 57 degrees F
	ADC 8 Exhaust	OK	21 degrees C / 69 degrees F
	ADC 8 ADC-XF1	OK	24 degrees C / 75 degrees F
	ADC 8 ADC-XF0	OK	31 degrees C / 87 degrees F
	ADC 9 Intake	OK	10 degrees C / 50 degrees F
	ADC 9 Exhaust	OK	22 degrees C / 71 degrees F
	ADC 9 ADC-XF1	OK	28 degrees C / 82 degrees F
	ADC 9 ADC-XF0	OK	36 degrees C / 96 degrees F
Fans	Fan Tray 0 Fan 1	OK	3480 RPM
	Fan Tray 0 Fan 2	OK	3480 RPM
	Fan Tray 0 Fan 3	OK	3480 RPM
	Fan Tray 0 Fan 4	OK	3360 RPM
	Fan Tray 0 Fan 5	OK	3360 RPM
	Fan Tray 0 Fan 6	OK	3480 RPM
	Fan Tray 1 Fan 1	OK	3360 RPM
	Fan Tray 1 Fan 2	OK	3360 RPM
	Fan Tray 1 Fan 3	OK	3360 RPM
	Fan Tray 1 Fan 4	OK	3480 RPM
	Fan Tray 1 Fan 5	OK	3480 RPM
	Fan Tray 1 Fan 6	OK	3480 RPM
	Fan Tray 2 Fan 1	OK	3360 RPM
	Fan Tray 2 Fan 2	OK	3360 RPM
	Fan Tray 2 Fan 3	OK	3480 RPM
	Fan Tray 2 Fan 4	OK	3480 RPM
	Fan Tray 2 Fan 5	OK	3360 RPM
	Fan Tray 2 Fan 6	OK	3480 RPM
	Fan Tray 3 Fan 1	OK	3360 RPM
	Fan Tray 3 Fan 2	OK	3360 RPM
	Fan Tray 3 Fan 3	OK	3480 RPM
	Fan Tray 3 Fan 4	OK	3480 RPM
	Fan Tray 3 Fan 5	OK	3480 RPM
	Fan Tray 3 Fan 6	OK	3360 RPM

show chassis environment (T320 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	PEM 0	OK	
	PEM 1	Absent	
Temp	SCG 0	OK	28 degrees C / 82 degrees F
	SCG 1	OK	28 degrees C / 82 degrees F
	Routing Engine 0	OK	31 degrees C / 87 degrees F
	Routing Engine 1	OK	30 degrees C / 86 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	32 degrees C / 89 degrees F
	SIB 0	OK	33 degrees C / 91 degrees F
	SIB 1	OK	33 degrees C / 91 degrees F
	SIB 2	OK	34 degrees C / 93 degrees F

	FPC 0 Top	OK	38 degrees C / 100 degrees F
	FPC 0 Bottom	OK	32 degrees C / 89 degrees F
	FPC 1 Top	OK	38 degrees C / 100 degrees F
	FPC 1 Bottom	OK	33 degrees C / 91 degrees F
	FPC 2 Top	OK	36 degrees C / 96 degrees F
	FPC 2 Bottom	OK	31 degrees C / 87 degrees F
	FPM GBUS	OK	26 degrees C / 78 degrees F
	FPM Display	OK	29 degrees C / 84 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed
	Rear Tray Middle fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
	Rear Tray Bottom fan	OK	Spinning at normal speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

show chassis environment (T640 Router)

user@host> show chassis environment			
Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	OK	22 degrees C / 71 degrees F
	SCG 0	OK	30 degrees C / 86 degrees F
	SCG 1	OK	30 degrees C / 86 degrees F
	Routing Engine 0	Present	
	Routing Engine 1	OK	27 degrees C / 80 degrees F
	CB 0	Present	
	CB 1	OK	33 degrees C / 91 degrees F
	SIB 0	Absent	
	SIB 1	Absent	
Fans	SIB 2	Absent	
	SIB 3	Absent	
	SIB 4	Absent	
	FPC 4 Top	Testing	
	FPC 4 Bottom	Testing	
	FPC 5 Top	Testing	
	FPC 5 Bottom	Testing	
	FPC 6 Top	Testing	
	FPC 6 Bottom	Testing	
	FPM GBUS	OK	23 degrees C / 73 degrees F
	FPM Display	Absent	
	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed

Bottom Left Front fan	OK	Spinning at normal speed
Bottom Left Middle fan	OK	Spinning at normal speed
Bottom Left Rear fan	OK	Spinning at normal speed
Bottom Right Front fan	OK	Spinning at normal speed
Bottom Right Middle fan	OK	Spinning at normal speed
Bottom Right Rear fan	OK	Spinning at normal speed
Fourth Blower from top	OK	Spinning at normal speed
Bottom Blower	OK	Spinning at normal speed
Middle Blower	OK	Spinning at normal speed
Top Blower	OK	Spinning at normal speed
Second Blower from top	OK	Spinning at normal speed
Misc CIP	OK	
SPMB 0	OK	
SPMB 1	OK	

show chassis environment (T4000 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	33 degrees C / 91 degrees F
	PEM 1	Absent	
	SCG 0	OK	33 degrees C / 91 degrees F
	SCG 1	OK	33 degrees C / 91 degrees F
	Routing Engine 0	OK	33 degrees C / 91 degrees F
	Routing Engine 0 CPU	OK	50 degrees C / 122 degrees F
	Routing Engine 1	OK	32 degrees C / 89 degrees F
	Routing Engine 1 CPU	OK	46 degrees C / 114 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	33 degrees C / 91 degrees F
	SIB 0	OK	42 degrees C / 107 degrees F
	SIB 1	OK	42 degrees C / 107 degrees F
	SIB 2	OK	42 degrees C / 107 degrees F
	SIB 3	OK	43 degrees C / 109 degrees F
	SIB 4	OK	45 degrees C / 113 degrees F
	FPC 0 Fan Intake	OK	34 degrees C / 93 degrees F
	FPC 0 Fan Exhaust	OK	48 degrees C / 118 degrees F
	FPC 0 PMB	OK	47 degrees C / 116 degrees F
	FPC 0 LMB0	OK	50 degrees C / 122 degrees F
	FPC 0 LMB1	OK	41 degrees C / 105 degrees F
	FPC 0 LMB2	OK	35 degrees C / 95 degrees F
	FPC 0 PFE1 LU2	OK	46 degrees C / 114 degrees F
	FPC 0 PFE1 LU0	OK	41 degrees C / 105 degrees F
	FPC 0 PFE0 LU0	OK	57 degrees C / 134 degrees F
	FPC 0 XF1	OK	46 degrees C / 114 degrees F
	FPC 0 XF0	OK	52 degrees C / 125 degrees F
	FPC 0 XM1	OK	41 degrees C / 105 degrees F
	FPC 0 XM0	OK	50 degrees C / 122 degrees F
	FPC 0 PFE0 LU1	OK	56 degrees C / 132 degrees F
	FPC 0 PFE0 LU2	OK	45 degrees C / 113 degrees F
	FPC 0 PFE1 LU1	OK	37 degrees C / 98 degrees F
	FPC 3 Fan Intake	OK	36 degrees C / 96 degrees F
	FPC 3 Fan Exhaust	OK	51 degrees C / 123 degrees F
	FPC 3 PMB	OK	43 degrees C / 109 degrees F
	FPC 3 LMB0	OK	57 degrees C / 134 degrees F
	FPC 3 LMB1	OK	54 degrees C / 129 degrees F
	FPC 3 LMB2	OK	38 degrees C / 100 degrees F
	FPC 3 PFE1 LU2	OK	63 degrees C / 145 degrees F
	FPC 3 PFE1 LU0	OK	45 degrees C / 113 degrees F
	FPC 3 PFE0 LU0	OK	69 degrees C / 156 degrees F
	FPC 3 XF1	OK	62 degrees C / 143 degrees F

	FPC 3 XF0	OK	63 degrees C / 145 degrees F
	FPC 3 XM1	OK	43 degrees C / 109 degrees F
	FPC 3 XM0	OK	67 degrees C / 152 degrees F
	FPC 3 PFE0 LU1	OK	63 degrees C / 145 degrees F
	FPC 3 PFE0 LU2	OK	66 degrees C / 150 degrees F
	FPC 3 PFE1 LU1	OK	41 degrees C / 105 degrees F
	FPC 5 Top	OK	39 degrees C / 102 degrees F
	FPC 5 Bottom	OK	38 degrees C / 100 degrees F
	FPC 6 Fan Intake	OK	33 degrees C / 91 degrees F
	FPC 6 Fan Exhaust	OK	49 degrees C / 120 degrees F
	FPC 6 PMB	OK	40 degrees C / 104 degrees F
	FPC 6 LMB0	OK	60 degrees C / 140 degrees F
	FPC 6 LMB1	OK	58 degrees C / 136 degrees F
	FPC 6 LMB2	OK	40 degrees C / 104 degrees F
	FPC 6 PFE1 LU2	OK	69 degrees C / 156 degrees F
	FPC 6 PFE1 LU0	OK	45 degrees C / 113 degrees F
	FPC 6 PFE0 LU0	OK	71 degrees C / 159 degrees F
	FPC 6 XF1	OK	58 degrees C / 136 degrees F
	FPC 6 XF0	OK	65 degrees C / 149 degrees F
	FPC 6 XM1	OK	39 degrees C / 102 degrees F
	FPC 6 XM0	OK	66 degrees C / 150 degrees F
	FPC 6 PFE0 LU1	OK	69 degrees C / 156 degrees F
	FPC 6 PFE0 LU2	OK	69 degrees C / 156 degrees F
	FPC 6 PFE1 LU1	OK	42 degrees C / 107 degrees F
	FPM GBUS	OK	24 degrees C / 75 degrees F
	FPM Display	OK	27 degrees C / 80 degrees F
Fans	Top Left Front fan	OK	Spinning at high speed
	Top Left Middle fan	OK	Spinning at high speed
	Top Left Rear fan	OK	Spinning at high speed
	Top Right Front fan	OK	Spinning at high speed
	Top Right Middle fan	OK	Spinning at high speed
	Top Right Rear fan	OK	Spinning at high speed
	Bottom Left Front fan	OK	Spinning at high speed
	Bottom Left Middle fan	OK	Spinning at high speed
	Bottom Left Rear fan	OK	Spinning at high speed
	Bottom Right Front fan	OK	Spinning at high speed
	Bottom Right Middle fan	OK	Spinning at high speed
	Bottom Right Rear fan	OK	Spinning at high speed
	Rear Tray Top fan	OK	Spinning at high speed
	Rear Tray Second fan	OK	Spinning at high speed
	Rear Tray Third fan	OK	Spinning at high speed
	Rear Tray Fourth fan	OK	Spinning at high speed
	Rear Tray Fifth fan	OK	Spinning at high speed
	Rear Tray Sixth fan	OK	Spinning at high speed
	Rear Tray Seventh fan	OK	Spinning at high speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

show chassis environment (TX Matrix Router)

```
user@host> show chassis environment
scc-re0:
```

Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	OK	29 degrees C / 84 degrees F
	Routing Engine 0	OK	34 degrees C / 93 degrees F
	Routing Engine 1	OK	34 degrees C / 93 degrees F
	CB 0	OK	32 degrees C / 89 degrees F

	CB 1	OK	32 degrees C / 89 degrees F
	SIB 0	OK	44 degrees C / 111 degrees F
	SIB 0 (B)	OK	44 degrees C / 111 degrees F
	FPM GBUS	OK	27 degrees C / 80 degrees F
	FPM Display	OK	32 degrees C / 89 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed
	Rear Tray Third fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
	Rear Tray Fifth fan	OK	Spinning at normal speed
	Rear Tray Sixth fan	OK	Spinning at normal speed
	Rear Tray Seventh fan	OK	Spinning at normal speed
	Rear Tray Bottom fan	OK	Spinning at normal speed
Misc	CIP 0	OK	
	CIP 1	OK	
	SPMB 0	OK	
	SPMB 1	OK	

lcc0-re0:

Class	Item	Status	Measurement
Temp	PEM 0	OK	29 degrees C / 84 degrees F
	PEM 1	Absent	
	SCG 0	OK	35 degrees C / 95 degrees F
	SCG 1	Absent	
	Routing Engine 0	OK	39 degrees C / 102 degrees F
	Routing Engine 1	OK	36 degrees C / 96 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	32 degrees C / 89 degrees F
	SIB 0	OK	40 degrees C / 104 degrees F
	SIB 0 (B)	OK	51 degrees C / 123 degrees F
	FPC 0 Top	OK	45 degrees C / 113 degrees F
	FPC 0 Bottom	OK	31 degrees C / 87 degrees F
	FPC 1 Top	OK	34 degrees C / 93 degrees F
	FPC 1 Bottom	OK	31 degrees C / 87 degrees F
	FPM GBUS	OK	30 degrees C / 86 degrees F
	FPM Display	OK	34 degrees C / 93 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed

```

Rear Tray Top fan      OK      Spinning at normal speed
Rear Tray Second fan   OK      Spinning at normal speed
Rear Tray Third fan    OK      Spinning at normal speed
Rear Tray Fourth fan   OK      Spinning at normal speed
Rear Tray Fifth fan    OK      Spinning at normal speed
Rear Tray Sixth fan    OK      Spinning at normal speed
Rear Tray Seventh fan  OK      Spinning at normal speed
Rear Tray Bottom fan   OK      Spinning at normal speed
Misc CIP               OK
SPMB 0                 OK
SPMB 1                 OK

```

```
lcc2-re0:
```

```

-----
Class Item              Status      Measurement
Temp PEM 0              OK        29 degrees C / 84 degrees F
      PEM 1              Absent
      SCG 0              OK        32 degrees C / 89 degrees F
      SCG 1              Absent
      Routing Engine 0    OK        31 degrees C / 87 degrees F
      Routing Engine 1    OK        32 degrees C / 89 degrees F
      CB 0               OK        30 degrees C / 86 degrees F
      SIB 0              OK        38 degrees C / 100 degrees F
      SIB 0 (B)          OK        49 degrees C / 120 degrees F
      FPC 0 Top           OK        45 degrees C / 113 degrees F
      FPC 0 Bottom        OK        33 degrees C / 91 degrees F
      FPC 1 Top           OK        37 degrees C / 98 degrees F
      FPC 1 Bottom        OK        33 degrees C / 91 degrees F
      FPM GBUS            OK        30 degrees C / 86 degrees F
      FPM Display         OK        34 degrees C / 93 degrees F
Fans  Top Left Front fan OK        Spinning at normal speed
      Top Left Middle fan OK        Spinning at normal speed
...

```

show chassis environment (T1600 Router)

```

user@host> show chassis environment
Class Item              Status      Measurement
Temp PEM 0              OK        27 degrees C / 80 degrees F
      PEM 1              Absent
      SCG 0              OK        31 degrees C / 87 degrees F
      SCG 1              OK        35 degrees C / 95 degrees F
      Routing Engine 0    OK        30 degrees C / 86 degrees F
      Routing Engine 1    OK        30 degrees C / 86 degrees F
      CB 0               OK        31 degrees C / 87 degrees F
      CB 1               OK        31 degrees C / 87 degrees F
      SIB 0              OK        41 degrees C / 105 degrees F
      SIB 0 (B)          OK        34 degrees C / 93 degrees F
      SIB 1              OK        0 degrees C / 32 degrees F
      SIB 1 (B)          OK        0 degrees C / 32 degrees F
      SIB 2              OK        0 degrees C / 32 degrees F
      SIB 2 (B)          OK        0 degrees C / 32 degrees F
      SIB 3              OK        0 degrees C / 32 degrees F
      SIB 3 (B)          OK        0 degrees C / 32 degrees F
      SIB 4              OK        0 degrees C / 32 degrees F
      SIB 4 (B)          OK        0 degrees C / 32 degrees F
      FPC 0 Top           OK        49 degrees C / 120 degrees F
      FPC 0 Bottom        OK        50 degrees C / 122 degrees F
      FPC 1 Top           OK        48 degrees C / 118 degrees F
      FPC 1 Bottom        OK        49 degrees C / 120 degrees F
      FPM GBUS            OK        27 degrees C / 80 degrees F

```

	FPM Display	OK	30 degrees C / 86 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed
	Rear Tray Third fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
	Rear Tray Fifth fan	OK	Spinning at normal speed
	Rear Tray Sixth fan	OK	Spinning at normal speed
	Rear Tray Seventh fan	OK	Spinning at normal speed
	Rear Tray Bottom fan	OK	Spinning at normal speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

show chassis environment (TX Matrix Plus Router)

```
user@host> show chassis environment
sfc0-re0:
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	28 degrees C / 82 degrees F
	PEM 1	Absent	
	Routing Engine 0	OK	27 degrees C / 80 degrees F
	Routing Engine 1	OK	29 degrees C / 84 degrees F
	CB 0 Intake	OK	26 degrees C / 78 degrees F
	CB 0 Exhaust A	OK	25 degrees C / 77 degrees F
	CB 0 Exhaust B	OK	25 degrees C / 77 degrees F
	CB 1 Intake	OK	26 degrees C / 78 degrees F
	CB 1 Exhaust A	OK	26 degrees C / 78 degrees F
	CB 1 Exhaust B	OK	26 degrees C / 78 degrees F
	SIB F13 0	OK	47 degrees C / 116 degrees F
	SIB F13 0 (B)	OK	48 degrees C / 118 degrees F
	SIB F13 1	OK	38 degrees C / 100 degrees F
	SIB F13 1 (B)	OK	37 degrees C / 98 degrees F
	SIB F2S 0/0	OK	27 degrees C / 80 degrees F
	SIB F2S 0/2	OK	28 degrees C / 82 degrees F
	SIB F2S 0/4	OK	27 degrees C / 80 degrees F
	SIB F2S 0/6	OK	28 degrees C / 82 degrees F
	SIB F2S 1/0	OK	26 degrees C / 78 degrees F
	SIB F2S 1/2	OK	26 degrees C / 78 degrees F
	SIB F2S 1/4	OK	26 degrees C / 78 degrees F
	SIB F2S 1/6	OK	26 degrees C / 78 degrees F
	SIB F2S 2/0	OK	25 degrees C / 77 degrees F
	SIB F2S 2/2	OK	25 degrees C / 77 degrees F
	SIB F2S 2/4	OK	23 degrees C / 73 degrees F
	CIP 0 Intake	OK	23 degrees C / 73 degrees F
	CIP 0 Exhaust A	OK	24 degrees C / 75 degrees F
	CIP 0 Exhaust B	OK	24 degrees C / 75 degrees F
	CIP 1 Intake	OK	24 degrees C / 75 degrees F
	CIP 1 Exhaust A	OK	25 degrees C / 77 degrees F

	CIP 1 Exhaust B	OK	25 degrees C / 77 degrees F
Fans	Fan Tray 0 Fan 1	OK	Spinning at normal speed
	Fan Tray 0 Fan 2	OK	Spinning at normal speed
	Fan Tray 0 Fan 3	OK	Spinning at normal speed
	Fan Tray 0 Fan 4	OK	Spinning at normal speed
	Fan Tray 0 Fan 5	OK	Spinning at normal speed
	Fan Tray 0 Fan 6	OK	Spinning at normal speed
	Fan Tray 1 Fan 1	OK	Spinning at normal speed
	Fan Tray 1 Fan 2	OK	Spinning at normal speed
	Fan Tray 1 Fan 3	OK	Spinning at normal speed
	Fan Tray 1 Fan 4	OK	Spinning at normal speed
	Fan Tray 1 Fan 5	OK	Spinning at normal speed
	Fan Tray 1 Fan 6	OK	Spinning at normal speed
	Fan Tray 2 Fan 1	OK	Spinning at normal speed
	Fan Tray 2 Fan 2	OK	Spinning at normal speed
	Fan Tray 2 Fan 3	OK	Spinning at normal speed
	Fan Tray 2 Fan 4	OK	Spinning at normal speed
	Fan Tray 2 Fan 5	OK	Spinning at normal speed
	Fan Tray 2 Fan 6	OK	Spinning at normal speed
	Fan Tray 2 Fan 7	OK	Spinning at normal speed
	Fan Tray 2 Fan 8	OK	Spinning at normal speed
	Fan Tray 2 Fan 9	OK	Spinning at normal speed
	Fan Tray 3 Fan 1	OK	Spinning at normal speed
	Fan Tray 3 Fan 2	OK	Spinning at normal speed
	Fan Tray 3 Fan 3	OK	Spinning at normal speed
	Fan Tray 3 Fan 4	OK	Spinning at normal speed
	Fan Tray 3 Fan 5	OK	Spinning at normal speed
	Fan Tray 3 Fan 6	OK	Spinning at normal speed
	Fan Tray 3 Fan 7	OK	Spinning at normal speed
	Fan Tray 3 Fan 8	OK	Spinning at normal speed
	Fan Tray 3 Fan 9	OK	Spinning at normal speed
	Fan Tray 4 Fan 1	OK	Spinning at normal speed
	Fan Tray 4 Fan 2	OK	Spinning at normal speed
	Fan Tray 4 Fan 3	OK	Spinning at normal speed
	Fan Tray 4 Fan 4	OK	Spinning at normal speed
	Fan Tray 4 Fan 5	OK	Spinning at normal speed
	Fan Tray 4 Fan 6	OK	Spinning at normal speed
	Fan Tray 4 Fan 7	OK	Spinning at normal speed
	Fan Tray 4 Fan 8	OK	Spinning at normal speed
	Fan Tray 4 Fan 9	OK	Spinning at normal speed
	Fan Tray 5 Fan 1	OK	Spinning at normal speed
	Fan Tray 5 Fan 2	OK	Spinning at normal speed
	Fan Tray 5 Fan 3	OK	Spinning at normal speed
	Fan Tray 5 Fan 4	OK	Spinning at normal speed
	Fan Tray 5 Fan 5	OK	Spinning at normal speed
	Fan Tray 5 Fan 6	OK	Spinning at normal speed
	Fan Tray 5 Fan 7	OK	Spinning at normal speed
	Fan Tray 5 Fan 8	OK	Spinning at normal speed
	Fan Tray 5 Fan 9	OK	Spinning at normal speed
Misc	SPMB 0	OK	
	SPMB 1	OK	

1cc0-re0:

Class	Item	Status	Measurement
Temp	PEM 0	OK	27 degrees C / 80 degrees F
	PEM 1	Absent	
	SCG 0	OK	31 degrees C / 87 degrees F
	SCG 1	OK	35 degrees C / 95 degrees F
	Routing Engine 0	OK	30 degrees C / 86 degrees F
	Routing Engine 1	OK	30 degrees C / 86 degrees F

	CB 0	OK	31 degrees C / 87 degrees F
	CB 1	OK	31 degrees C / 87 degrees F
	SIB 0	OK	41 degrees C / 105 degrees F
	SIB 0 (B)	OK	34 degrees C / 93 degrees F
	SIB 1	OK	0 degrees C / 32 degrees F
	SIB 1 (B)	OK	0 degrees C / 32 degrees F
	SIB 2	OK	0 degrees C / 32 degrees F
	SIB 2 (B)	OK	0 degrees C / 32 degrees F
	SIB 3	OK	0 degrees C / 32 degrees F
	SIB 3 (B)	OK	0 degrees C / 32 degrees F
	SIB 4	OK	0 degrees C / 32 degrees F
	SIB 4 (B)	OK	0 degrees C / 32 degrees F
	FPC 0 Top	OK	49 degrees C / 120 degrees F
	FPC 0 Bottom	OK	50 degrees C / 122 degrees F
	FPC 1 Top	OK	48 degrees C / 118 degrees F
	FPC 1 Bottom	OK	49 degrees C / 120 degrees F
	FPM GBUS	OK	27 degrees C / 80 degrees F
	FPM Display	OK	30 degrees C / 86 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed
	Rear Tray Third fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
	Rear Tray Fifth fan	OK	Spinning at normal speed
	Rear Tray Sixth fan	OK	Spinning at normal speed
	Rear Tray Seventh fan	OK	Spinning at normal speed
	Rear Tray Bottom fan	OK	Spinning at normal speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

show chassis environment (EX4200 Standalone Switch)

user@switch> show chassis environment			
Class	Item	Status	Measurement
Power	FPC 0 Power Supply 0	OK	
	FPC 0 Power Supply 1	Absent	
Temp	FPC 0 CPU	OK	41 degrees C / 105 degrees F
	FPC 0 EX-PFE1	OK	42 degrees C / 107 degrees F
	FPC 0 EX-PFE2	OK	46 degrees C / 114 degrees F
	FPC 0 GEPHY Front Left	OK	25 degrees C / 77 degrees F
	FPC 0 GEPHY Front Right	OK	27 degrees C / 80 degrees F
	FPC 0 Uplink Conn	OK	29 degrees C / 84 degrees F
Fans	FPC 0 Fan 1	OK	Spinning at normal speed
	FPC 0 Fan 2	OK	Spinning at normal speed
	FPC 0 Fan 3	OK	Spinning at normal speed

show chassis environment (EX8216 Switch)

```
user@switch> show chassis environment
```


Class	Item	Status	Measurement
Power	PSU 0	OK	
	PSU 1	OK	
	PSU 2	OK	
	PSU 3	Check	
	PSU 4	Absent	
	PSU 5	Absent	
Temp	CB 0 Intake	OK	23 degrees C / 73 degrees F
	CB 0 Exhaust	OK	26 degrees C / 78 degrees F
	CB 1 Intake	OK	22 degrees C / 71 degrees F
	CB 1 Exhaust	OK	25 degrees C / 77 degrees F
	FPC 4 Intake	OK	49 degrees C / 120 degrees F
	FPC 4 Exhaust	OK	59 degrees C / 138 degrees F
	SIB 5 Intake	OK	25 degrees C / 77 degrees F
	SIB 5 Exhaust	OK	35 degrees C / 95 degrees F
	SIB 6 Intake	OK	25 degrees C / 77 degrees F
Fans	SIB 6 Exhaust	OK	38 degrees C / 100 degrees F
	Top Fan 1	OK	Spinning at normal speed
	Top Fan 2	OK	Spinning at normal speed
	Top Fan 3	OK	Spinning at normal speed
	Top Fan 4	OK	Spinning at normal speed
	Top Fan 5	OK	Spinning at normal speed
	Top Fan 6	OK	Spinning at normal speed
	Top Fan 7	OK	Spinning at normal speed
	Top Fan 8	OK	Spinning at normal speed
	Top Fan 9	OK	Spinning at normal speed
	Bottom Fan 1	OK	Spinning at normal speed
	Bottom Fan 2	OK	Spinning at normal speed
	Bottom Fan 3	OK	Spinning at normal speed
	Bottom Fan 4	OK	Spinning at normal speed
	Bottom Fan 5	OK	Spinning at normal speed
	Bottom Fan 6	OK	Spinning at normal speed
	Bottom Fan 7	OK	Spinning at normal speed
	Bottom Fan 8	OK	Spinning at normal speed
	Bottom Fan 9	OK	Spinning at normal speed

show chassis environment (QFX Series)

```
user@switch> show chassis environment
```

Class	Item	Status	Measurement
Power	FPC 0 Power Supply 0	OK	
	FPC 0 Power Supply 1	OK	
Temp	FPC 0 Sensor TopLeft I	OK	26 degrees C / 78 degrees F
	FPC 0 Sensor TopRight I	OK	24 degrees C / 75 degrees F
	FPC 0 Sensor TopLeft E	OK	30 degrees C / 86 degrees F
	FPC 0 Sensor TopRight E	OK	30 degrees C / 86 degrees F
	FPC 0 Sensor TopMiddle I	OK	30 degrees C / 86 degrees F
	FPC 0 Sensor TopMiddle E	OK	38 degrees C / 100 degrees F
	FPC 0 Sensor Bottom I	OK	34 degrees C / 93 degrees F
	FPC 0 Sensor Bottom E	OK	38 degrees C / 100 degrees F
	FPC 0 Sensor Die Temp	OK	38 degrees C / 100 degrees F
	FPC 0 Sensor Mgmt Brd I	OK	24 degrees C / 75 degrees F
Fans	FPC 0 Sensor Switch I	OK	28 degrees C / 82 degrees F
	FPC 0 Fan 1 (left)	Failed	
	FPC 0 Fan 2 (right)	OK	Spinning at normal speed
	FPC 0 Fan 3 (middle)	OK	Spinning at normal speed

show chassis environment interconnect-device (QFabric System)

```
user@switch> show chassis environment interconnect-device IC-A0004
```

Class	Item	Status	Measurement
	CB 0		
	CB 0 L Intake	OK	30 degrees C / 86 degrees F
	CB 0 R Intake	OK	31 degrees C / 87 degrees F
	CB 0 L Exhaust	OK	32 degrees C / 89 degrees F
	CB 0 R Exhaust	OK	33 degrees C / 91 degrees F
	Routing Engine 0 CPU temp	OK	51 degrees C / 123 degrees F
	CB 1		
	CB 1 L Intake	OK	27 degrees C / 80 degrees F
	CB 1 R Intake	OK	29 degrees C / 84 degrees F
	CB 1 L Exhaust	OK	31 degrees C / 87 degrees F
	CB 1 R Exhaust	OK	32 degrees C / 89 degrees F
	Routing Engine 1 CPU temp	OK	40 degrees C / 104 degrees F
	FC 0 FPC 0		
	FPC 0 L Intake	OK	25 degrees C / 77 degrees F
	FPC 0 R Intake	OK	28 degrees C / 82 degrees F
	FPC 0 L Exhaust	OK	28 degrees C / 82 degrees F
	FPC 0 R Exhaust	OK	29 degrees C / 84 degrees F
	FC 7 FPC 7		
	FPC 7 L Intake	OK	25 degrees C / 77 degrees F
	FPC 7 R Intake	OK	26 degrees C / 78 degrees F
	FPC 7 L Exhaust	OK	28 degrees C / 82 degrees F
	FPC 7 R Exhaust	OK	29 degrees C / 84 degrees F
	RC 0 FPC 8		
	FPC 8 L Intake	OK	25 degrees C / 77 degrees F
	FPC 8 R Intake	OK	26 degrees C / 78 degrees F
	FPC 8 L Exhaust	OK	32 degrees C / 89 degrees F
	FPC 8 R Exhaust	OK	30 degrees C / 86 degrees F
	RC 7 FPC 15		
	FPC 15 L Intake	OK	24 degrees C / 75 degrees F
	FPC 15 R Intake	OK	25 degrees C / 77 degrees F
	FPC 15 L Exhaust	OK	33 degrees C / 91 degrees F
	FPC 15 R Exhaust	OK	31 degrees C / 87 degrees F
Fans	TFT 0 Fan 0	OK	Spinning at normal speed
Fans	TFT 0 Fan 1	OK	Spinning at normal speed
Fans	TFT 0 Fan 2	OK	Spinning at normal speed
Fans	TFT 0 Fan 3	OK	Spinning at normal speed
Fans	TFT 0 Fan 4	OK	Spinning at normal speed
Fans	TFT 0 Fan 5	OK	Spinning at normal speed
Fans	BFT 1 Fan 0	OK	Spinning at normal speed
Fans	BFT 1 Fan 1	OK	Spinning at normal speed
Fans	BFT 1 Fan 2	OK	Spinning at normal speed
Fans	BFT 1 Fan 3	Check	
Fans	BFT 1 Fan 4	OK	Spinning at normal speed
Fans	BFT 1 Fan 5	OK	Spinning at normal speed
Fans	SFT 0 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans	SFT 0 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans	SFT 0 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans	SFT 0 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans	SFT 0 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans	SFT 0 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans	SFT 0 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans	SFT 0 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans	SFT 1 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans	SFT 1 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans	SFT 1 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans	SFT 1 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans	SFT 1 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans	SFT 1 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans	SFT 1 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans	SFT 1 Fan 3 Rotor 1	OK	Spinning at normal speed

Fans	SFT 2 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans	SFT 2 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans	SFT 2 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans	SFT 2 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans	SFT 2 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans	SFT 2 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans	SFT 2 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans	SFT 2 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans	SFT 3 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans	SFT 3 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans	SFT 3 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans	SFT 3 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans	SFT 3 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans	SFT 3 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans	SFT 3 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans	SFT 3 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans	SFT 4 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans	SFT 4 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans	SFT 4 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans	SFT 4 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans	SFT 4 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans	SFT 4 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans	SFT 4 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans	SFT 4 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans	SFT 5 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans	SFT 5 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans	SFT 5 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans	SFT 5 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans	SFT 5 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans	SFT 5 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans	SFT 5 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans	SFT 5 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans	SFT 6 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans	SFT 6 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans	SFT 6 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans	SFT 6 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans	SFT 6 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans	SFT 6 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans	SFT 6 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans	SFT 6 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans	SFT 7 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans	SFT 7 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans	SFT 7 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans	SFT 7 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans	SFT 7 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans	SFT 7 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans	SFT 7 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans	SFT 7 Fan 3 Rotor 1	OK	Spinning at normal speed
Power	PEM 0	OK	30 degrees C / 86 degrees F
Power	PEM 1	OK	30 degrees C / 86 degrees F
Power	PEM 2	OK	30 degrees C / 86 degrees F
Power	PEM 3	Absent	
Power	PEM 4	Absent	
Power	PEM 5	Absent	

show chassis environment node-device (QFabric System)

```

user@switch> show chassis environment node-device node1
Class Item                               Status Measurement
Power node1 Power Supply 0              Absent
      node1 Power Supply 1              Absent

```

```

Fans  node1 Fan Tray 0      Testing
      node1 Fan Tray 1      Testing
      node1 Fan Tray 2      Testing

```

show chassis environment pem node-device (QFabric System)

```

user@switch> show chassis environment pem node-device node1
FPC 0 PEM 0 status:
  State          Check
  Airflow        Front to Back
  Temperature     OK
  AC Input:      OK
  DC Output      Voltage(V) Current(A) Power(W) Load(%)
                  12         10        120     18

FPC 0 PEM 1 status:
  State          Online
  Airflow        Back to Front
  Temperature     OK
  AC Input:      OK
  DC Output      Voltage(V) Current(A) Power(W) Load(%)
                  11         10        110     17

```

show chassis environment (PTX5000 Packet Transport Switch)

```

user@switch> show chassis environment
Class Item                               Status      Measurement
Temp  PDU 0                               OK
      PDU 0 PSM 0                         OK          36 degrees C / 96 degrees F
      PDU 0 PSM 1                         OK          38 degrees C / 100 degrees F
      PDU 0 PSM 2                         OK          38 degrees C / 100 degrees F
      PDU 0 PSM 3                         OK          37 degrees C / 98 degrees F
      PDU 1                               Absent
      CCG 0                               OK          44 degrees C / 111 degrees F
      CCG 1                               OK          44 degrees C / 111 degrees F
      Routing Engine 0                     OK          62 degrees C / 143 degrees F
      Routing Engine 0 CPU                 OK          75 degrees C / 167 degrees F
      Routing Engine 1                     OK          51 degrees C / 123 degrees F
      Routing Engine 1 CPU                 OK          64 degrees C / 147 degrees F
      CB 0 Intake                          OK          38 degrees C / 100 degrees F
      CB 0 Exhaust A                      OK          46 degrees C / 114 degrees F
      CB 0 Exhaust B                      OK          42 degrees C / 107 degrees F
      CB 1 Intake                          OK          35 degrees C / 95 degrees F
      CB 1 Exhaust A                      OK          39 degrees C / 102 degrees F
      CB 1 Exhaust B                      OK          36 degrees C / 96 degrees F
      SIB 0 Intake                        OK          39 degrees C / 102 degrees F
      SIB 0 Exhaust                      OK          37 degrees C / 98 degrees F
      SIB 0 Junction                     OK          43 degrees C / 109 degrees F
      SIB 1 Intake                        OK          39 degrees C / 102 degrees F
      SIB 1 Exhaust                      OK          36 degrees C / 96 degrees F
      SIB 1 Junction                     OK          46 degrees C / 114 degrees F
      SIB 2 Intake                        OK          37 degrees C / 98 degrees F
      SIB 2 Exhaust                      OK          37 degrees C / 98 degrees F
      SIB 2 Junction                     OK          42 degrees C / 107 degrees F
      SIB 3 Intake                        OK          40 degrees C / 104 degrees F
      SIB 3 Exhaust                      OK          40 degrees C / 104 degrees F
      SIB 3 Junction                     OK          45 degrees C / 113 degrees F
      SIB 4 Intake                        OK          47 degrees C / 116 degrees F
      SIB 4 Exhaust                      OK          44 degrees C / 111 degrees F
      SIB 4 Junction                     OK          58 degrees C / 136 degrees F
      SIB 5 Intake                        OK          58 degrees C / 136 degrees F
      SIB 5 Exhaust                      OK          43 degrees C / 109 degrees F

```

SIB 5 Junction	OK	71 degrees C / 159 degrees F
SIB 6 Intake	OK	57 degrees C / 134 degrees F
SIB 6 Exhaust	OK	42 degrees C / 107 degrees F
SIB 6 Junction	OK	65 degrees C / 149 degrees F
SIB 7 Intake	OK	58 degrees C / 136 degrees F
SIB 7 Exhaust	OK	42 degrees C / 107 degrees F
SIB 7 Junction	OK	66 degrees C / 150 degrees F
SIB 8 Intake	OK	57 degrees C / 134 degrees F
SIB 8 Exhaust	OK	42 degrees C / 107 degrees F
SIB 8 Junction	OK	70 degrees C / 158 degrees F
FPC 0 PMB	OK	35 degrees C / 95 degrees F
FPC 0 Intake	OK	33 degrees C / 91 degrees F
FPC 0 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 0 Exhaust B	OK	43 degrees C / 109 degrees F
FPC 0 TL0	OK	48 degrees C / 118 degrees F
FPC 0 TQ0	OK	53 degrees C / 127 degrees F
FPC 0 TL1	OK	56 degrees C / 132 degrees F
FPC 0 TQ1	OK	58 degrees C / 136 degrees F
FPC 0 TL2	OK	55 degrees C / 131 degrees F
FPC 0 TQ2	OK	56 degrees C / 132 degrees F
FPC 0 TL3	OK	59 degrees C / 138 degrees F
FPC 0 TQ3	OK	59 degrees C / 138 degrees F
FPC 2 PMB	OK	35 degrees C / 95 degrees F
FPC 2 Intake	OK	34 degrees C / 93 degrees F
FPC 2 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 2 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 2 TL0	OK	53 degrees C / 127 degrees F
FPC 2 TQ0	OK	53 degrees C / 127 degrees F
FPC 2 TL1	OK	57 degrees C / 134 degrees F
FPC 2 TQ1	OK	58 degrees C / 136 degrees F
FPC 2 TL2	OK	54 degrees C / 129 degrees F
FPC 2 TQ2	OK	59 degrees C / 138 degrees F
FPC 2 TL3	OK	60 degrees C / 140 degrees F
FPC 2 TQ3	OK	64 degrees C / 147 degrees F
PIC 2/0 Ambient	OK	49 degrees C / 120 degrees F
FPC 3 PMB	OK	34 degrees C / 93 degrees F
FPC 3 Intake	OK	35 degrees C / 95 degrees F
FPC 3 Exhaust A	OK	54 degrees C / 129 degrees F
FPC 3 Exhaust B	OK	49 degrees C / 120 degrees F
FPC 3 TL0	OK	49 degrees C / 120 degrees F
FPC 3 TQ0	OK	55 degrees C / 131 degrees F
FPC 3 TL1	OK	56 degrees C / 132 degrees F
FPC 3 TQ1	OK	58 degrees C / 136 degrees F
FPC 3 TL2	OK	56 degrees C / 132 degrees F
FPC 3 TQ2	OK	59 degrees C / 138 degrees F
FPC 3 TL3	OK	62 degrees C / 143 degrees F
FPC 3 TQ3	OK	63 degrees C / 145 degrees F
PIC 3/1	Absent	
FPC 5 PMB	OK	35 degrees C / 95 degrees F
FPC 5 Intake	OK	34 degrees C / 93 degrees F
FPC 5 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 5 Exhaust B	OK	53 degrees C / 127 degrees F
FPC 5 TL0	OK	54 degrees C / 129 degrees F
FPC 5 TQ0	OK	52 degrees C / 125 degrees F
FPC 5 TL1	OK	61 degrees C / 141 degrees F
FPC 5 TQ1	OK	60 degrees C / 140 degrees F
FPC 5 TL2	OK	55 degrees C / 131 degrees F
FPC 5 TQ2	OK	55 degrees C / 131 degrees F
FPC 5 TL3	OK	59 degrees C / 138 degrees F
FPC 5 TQ3	OK	58 degrees C / 136 degrees F
PIC 5/0 Ambient	OK	51 degrees C / 123 degrees F

PIC 5/1 Ambient	OK	34 degrees C / 93 degrees F
PIC 5/1 cfp-5/1/0	OK	34 degrees C / 93 degrees F
PIC 5/1 cfp-5/1/1	OK	36 degrees C / 96 degrees F
FPC 6 PMB	OK	36 degrees C / 96 degrees F
FPC 6 Intake	OK	33 degrees C / 91 degrees F
FPC 6 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 6 Exhaust B	OK	39 degrees C / 102 degrees F
FPC 6 TL0	OK	44 degrees C / 111 degrees F
FPC 6 TQ0	OK	54 degrees C / 129 degrees F
FPC 6 TL1	OK	59 degrees C / 138 degrees F
FPC 6 TQ1	OK	58 degrees C / 136 degrees F
FPC 6 TL2	OK	60 degrees C / 140 degrees F
FPC 6 TQ2	OK	57 degrees C / 134 degrees F
FPC 6 TL3	OK	65 degrees C / 149 degrees F
FPC 6 TQ3	OK	60 degrees C / 140 degrees F
FPC 7 PMB	OK	35 degrees C / 95 degrees F
FPC 7 Intake	OK	33 degrees C / 91 degrees F
FPC 7 Exhaust A	OK	53 degrees C / 127 degrees F
FPC 7 Exhaust B	OK	40 degrees C / 104 degrees F
FPC 7 TL0	OK	46 degrees C / 114 degrees F
FPC 7 TQ0	OK	58 degrees C / 136 degrees F
FPC 7 TL1	OK	53 degrees C / 127 degrees F
FPC 7 TQ1	OK	59 degrees C / 138 degrees F
FPC 7 TL2	OK	56 degrees C / 132 degrees F
FPC 7 TQ2	OK	61 degrees C / 141 degrees F
FPC 7 TL3	OK	63 degrees C / 145 degrees F
FPC 7 TQ3	OK	63 degrees C / 145 degrees F
FPM I2CS	OK	37 degrees C / 98 degrees F
Fans Fan Tray 0 Fan 1	OK	3042 RPM
Fans Fan Tray 0 Fan 2	OK	3042 RPM
Fans Fan Tray 0 Fan 3	OK	3000 RPM
Fans Fan Tray 0 Fan 4	OK	3042 RPM
Fans Fan Tray 0 Fan 5	OK	3000 RPM
Fans Fan Tray 0 Fan 6	OK	3042 RPM
Fans Fan Tray 0 Fan 7	OK	3085 RPM
Fans Fan Tray 0 Fan 8	OK	3042 RPM
Fans Fan Tray 0 Fan 9	OK	3042 RPM
Fans Fan Tray 0 Fan 10	OK	3085 RPM
Fans Fan Tray 0 Fan 11	OK	3085 RPM
Fans Fan Tray 0 Fan 12	OK	3128 RPM
Fans Fan Tray 0 Fan 13	OK	3128 RPM
Fans Fan Tray 0 Fan 14	OK	3042 RPM
Fans Fan Tray 1 Fan 1	OK	2299 RPM
Fans Fan Tray 1 Fan 2	OK	2399 RPM
Fans Fan Tray 1 Fan 3	OK	2299 RPM
Fans Fan Tray 1 Fan 4	OK	2266 RPM
Fans Fan Tray 1 Fan 5	OK	2266 RPM
Fans Fan Tray 1 Fan 6	OK	2366 RPM
Fans Fan Tray 2 Fan 1	OK	2199 RPM
Fans Fan Tray 2 Fan 2	OK	2133 RPM
Fans Fan Tray 2 Fan 3	OK	2366 RPM
Fans Fan Tray 2 Fan 4	OK	2233 RPM
Fans Fan Tray 2 Fan 5	OK	2399 RPM
Fans Fan Tray 2 Fan 6	OK	2233 RPM
Misc SPMB 0 Intake	OK	50 degrees C / 122 degrees F
Misc SPMB 1 Intake	OK	40 degrees C / 104 degrees F

show chassis environment (ACX2000 Universal Access Router)

user@host> show chassis environment

Class	Item	Status	Measurement
	PCB Left	OK	44 degrees C / 111 degrees F
	SFP+ Xcvr	OK	50 degrees C / 122 degrees F
	FEB	OK	70 degrees C / 158 degrees F
	PCB Up	OK	63 degrees C / 145 degrees F
	PCB Mid	OK	66 degrees C / 150 degrees F
	Telecom Mod	OK	65 degrees C / 149 degrees F
	Routing Engine	OK	54 degrees C / 129 degrees F
	Heater off		

show chassis environment cb

Syntax	show chassis environment cb <slot>
Syntax (TX Matrix Routers)	show chassis environment cb <lcc number scc> <slot>
Syntax (TX Matrix Plus Routers)	show chassis environment cb <lcc number sfc number > <slot>
Syntax (MX Series Routers)	show chassis environment cb <slot> <all-members> <local> <member member-id>
Syntax (MX2010 3D Universal Edge Routers)	show chassis environment cb <slot>
Syntax (MX2020 3D Universal Edge Routers)	show chassis environment cb <slot>
Syntax (QFabric System)	show chassis environment cb <slot interconnect-device interconnect-device-name> < interconnect-device interconnect-device-name slot>
Release Information	<p>Command introduced before Junos Release 7.4.</p> <p>Command introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches.</p> <p>Command introduced in Junos OS Release 12.1 for T4000 Core Routers.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos Release 9.6.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>
Description	(M120, M320, MX Series, and T Series routers, EX8200 switches, and PTX Series Packet Transport Switches only) Display environmental information about the Control Boards (CBs). For information about the meaning of “CBs” on the switches, see <i>EX Series Switches Hardware and CLI Terminology Mapping</i> .
Options	none —Display environmental information about all CBs. For a TX Matrix router, display environmental information about all CBs on the TX Matrix router and its attached T640 routers. For a TX Matrix Plus router, display environmental information about all CBs on the TX Matrix Plus router and its attached T1600 routers.

all-members—(MX Series routers only) (Optional) Display environmental information about the CBs on all the members of the Virtual Chassis configuration.

interconnect-device—(QFabric systems only) Display environmental information about the CBs on the Interconnect device.

lcc-number—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, display environmental information about the CBs in a specified T640 router (or line-card chassis) that is connected to the TX Matrix router. For a TX Matrix Plus router, display environmental information about the CBs in a specified T1600 router (or line-card chassis) that is connected to the TX Matrix Plus router. Replace **number** with a value from 0 through 3.

local—(MX Series routers only) (Optional) Display environmental information about the CBs on the local Virtual Chassis member.

member member-id—(MX Series routers only) (Optional) Display environmental information about the CBs on the specified member of the Virtual Chassis configuration. Replace **member-id** with a value of 0 or 1.

scc—(TX Matrix router only) (Optional) Display environmental information about the CBs in the TX Matrix router (or switch-card chassis).

sfc-number—(TX Matrix Plus router only) (Optional) Display environmental information about the CBs in the TX Matrix Plus router (or switch-fabric chassis).

slot—(Optional) Display environmental information about the specified CB. On routers and PTX Series Packet Transport Switches, replace **slot** with 0 or 1. On EX Series switches replace **slot** with 0, 1, or 2. On QFX Series switches, replace **slot** with 0 or 1.

Required Privilege Level

view

Related Documentation

- *request chassis cb*
- *Switching Control Board Redundancy*
- *Routing Engine and Switching Control Board Redundancy Configuration Statements*

List of Sample Output

[show chassis environment cb \(M120 Router\) on page 580](#)
[show chassis environment cb \(M320 Router\) on page 581](#)
[show chassis environment cb \(MX80 Router\) on page 581](#)
[show chassis environment cb \(MX240 Router\) on page 582](#)
[show chassis environment cb \(MX240 Router with Enhanced MX SCB\) on page 582](#)
[show chassis environment cb \(MX480 Router\) on page 583](#)
[show chassis environment cb \(MX480 Router with Enhanced MX SCB\) on page 583](#)
[show chassis environment cb \(MX960 Router\) on page 584](#)
[show chassis environment cb \(MX960 Router with Enhanced MX SCB\) on page 584](#)
[show chassis environment cb \(MX2020 Router\) on page 585](#)
[show chassis environment cb \(MX2010 Router\) on page 585](#)
[show chassis environment cb \(T4000 Core Router\) on page 586](#)

[show chassis environment cb \(TX Matrix Router\) on page 587](#)
[show chassis environment cb \(TX Matrix Plus Router\) on page 588](#)
[show chassis environment cb \(EX8200 Switch\) on page 591](#)
[show chassis environment cb \(EX8208 Switch\) on page 592](#)
[show chassis environment cb \(PTX5000 Packet Transport Switch\) on page 593](#)
[show chassis environment cb \(QFabric System\) on page 594](#)

Output Fields [Table 54 on page 580](#) lists the output fields for the **show chassis environment cb** command. Output fields are listed in the approximate order in which they appear.

Table 54: show chassis environment cb Output Fields

Field Name	Field Description
State	<p>Status of the CB. If two CBs are installed and online, one is functioning as the master, and the other is the standby.</p> <ul style="list-style-type: none"> • Online—CB is online and running. • Offline—CB is powered down. <p>NOTE: On the EX8208 switch, the installation can include three CBs. See <i>EX Series Switches Hardware and CLI Terminology Mapping</i>.</p>
Temperature	<p>Temperature in Celsius (C) and Fahrenheit (F) of the air flowing past the CB.</p> <ul style="list-style-type: none"> • Temperature Intake—Measures the temperature of the air intake to cool the power supplies. • Temperature Exhaust—Measures the temperature of the hot air exhaust. <p>NOTE: On the MX2010 and MX2020 routers, the intake temperature measures the temperature of the air intake to cool the Control Board (CB). The MX2010 and MX2020 routers include intake and exhaust temperatures for multiple zones (Intake A, Intake B, Intake C, Exhaust A, Exhaust B, and TCBC).</p>
Power	<p>Power required and measured on the CB. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.</p>
BUS Revision	<p>Revision level of the generic bus device. (Not on switches.)</p>
FPGA Revision	<p>Revision level of the field-programmable gate array (FPGA). (Not on switches.)</p>
PMBus device (on MX240, MX480, and MX960 routers with Enhanced MX SCB)	<p>Enhanced SCB on MX 240, MX480, and MX960 routers allows the system to save power by supplying only the amount of voltage that is required. Configurable PMBus devices are used to provide the voltage for each individual device. There is one PMBus device for each XF ASIC so that the output can be customized to each device. The following PMBus device information is displayed for routers with Enhanced MX SCB:</p> <ul style="list-style-type: none"> • Expected voltage • Measured voltage • Measured current • Calculated power

Sample Output

show chassis environment cb (M120 Router)

```
user@host> show chassis environment cb
```

```

CB 0 status:
State                Online Master
Temperature           33 degrees C / 91 degrees F
Power
  1.2 V              1214 mV
  1.5 V              1495 mV
  2.5 V              2494 mV
  3.3 V              3319 mV
  5.0 V              5085 mV
  3.3 V bias         3296 mV
Bus Revision         12
FPGA Revision        17

CB 1 status:
State                Online Standby
Temperature           34 degrees C / 93 degrees F
Power
  1.2 V              1195 mV
  1.5 V              1495 mV
  2.5 V              2504 mV
  3.3 V              3312 mV
  5.0 V              5111 mV
  3.3 V bias         3296 mV
Bus Revision         12
FPGA Revision        17

```

show chassis environment cb (M320 Router)

```

user@host> show chassis environment cb

CB 0 status:
State                Online Master
Temperature           29 degrees C / 84 degrees F
Power:
  1.8 V              1805 mV
  2.5 V              2501 mV
  3.3 V              3293 mV
  4.6 V              4725 mV
  5.0 V              5032 mV
  12.0 V             11975 mV
  3.3 V bias         3286 mV
  8.0 V bias         7589 mV
BUS Revision         40
FPGA Revision        7

CB 1 status:
State                Online Standby
Temperature           32 degrees C / 89 degrees F
Power:
  1.8 V              1802 mV
  2.5 V              2482 mV
  3.3 V              3289 mV
  4.6 V              4720 mV
  5.0 V              5001 mV
  12.0 V             11946 mV
  3.3 V bias         3274 mV
  8.0 V bias         7562 mV
BUS Revision         40
FPGA Revision        7

```

show chassis environment cb (MX80 Router)

```

user@host> show chassis environment cb

```

```
CB 0 status:
State                Online Master
Temperature          36 degrees C / 96 degrees F
Power 1
  1.0 V              1034 mV
  1.0 V MQ           1037 mV
  1.0 V LU           1005 mV
  1.2 V              1218 mV
  1.5 V              1524 mV
  1.8 V              1814 mV
  2.5 V              2558 mV
  3.3 V              3296 mV
  5.0 V              5233 mV
  5.0 V bias         5207 mV
  12.0 V             12162 mV
```

show chassis environment cb (MX240 Router)

```
user@host> show chassis environment cb
CB 0 status:
State                Online Standby
Temperature          37 degrees C / 98 degrees F
Power 1
  1.2 V              1208 mV
  1.5 V              1521 mV
  1.8 V              1811 mV
  2.5 V              2513 mV
  3.3 V              3332 mV
  5.0 V              5059 mV
  12.0 V             12162 mV
  1.25 V             1260 mV
  3.3 V SM3          3306 mV
  5.0 V RE           5085 mV
  12.0 V RE          11872 mV
Power 2
  11.3 V bias PEM    11272 mV
  4.6 V bias MidPlane 4827 mV
  11.3 V bias FPD    11272 mV
  11.3 V bias POE 0   11292 mV
  11.3 V bias POE 1   11253 mV
Bus Revision         42
FPGA Revision        1
```

show chassis environment cb (MX240 Router with Enhanced MX SCB)

```
user@host> show chassis environment cb
CB 0 status:
State                Online Standby
Temperature          37 degrees C / 98 degrees F
Power 1
  1.2 V              1208 mV
  1.5 V              1521 mV
  1.8 V              1811 mV
  2.5 V              2513 mV
  3.3 V              3332 mV
  5.0 V              5059 mV
  12.0 V             12162 mV
  1.25 V             1260 mV
  3.3 V SM3          3306 mV
  5.0 V RE           5085 mV
  12.0 V RE          11872 mV
```

```

Power 2
  11.3 V bias PEM          11272 mV
  4.6 V bias MidPlane      4827 mV
  11.3 V bias FPD          11272 mV
  11.3 V bias POE 0        11292 mV
  11.3 V bias POE 1        11253 mV
Bus Revision               42
FPGA Revision              1
PMBus
device      Expected    Measured    Measured    Calculated
            voltage     voltage     current     power
XF ASIC A   1000 mV      997 mV      11031 mA    10997 mW
XF ASIC B   1000 mV      996 mV      12125 mA    12076 mW

```

show chassis environment cb (MX480 Router)

```

user@host> show chassis environment cb
CB 0 status:
State                Online Master
Temperature           41 degrees C / 105 degrees F
Power 1
  1.2 V               1202 mV
  1.5 V               1511 mV
  1.8 V               1798 mV
  2.5 V               2507 mV
  3.3 V               3312 mV
  5.0 V               5027 mV
  12.0 V              12200 mV
  1.25 V              1260 mV
  3.3 V SM3           3293 mV
  5 V RE              5040 mV
  12 V RE             11910 mV
Power 2
  11.3 V bias PEM     11156 mV
  4.6 V bias MidPlane 4801 mV
  11.3 V bias FPD     11214 mV
  11.3 V bias POE 0   11098 mV
  11.3 V bias POE 1   11330 mV
Bus Revision          42
FPGA Revision         1

```

show chassis environment cb (MX480 Router with Enhanced MX SCB)

```

user@host> show chassis environment cb
CB 0 status:
State                Online Master
Temperature           41 degrees C / 105 degrees F
Power 1
  1.2 V               1202 mV
  1.5 V               1511 mV
  1.8 V               1798 mV
  2.5 V               2507 mV
  3.3 V               3312 mV
  5.0 V               5027 mV
  12.0 V              12200 mV
  1.25 V              1260 mV
  3.3 V SM3           3293 mV
  5 V RE              5040 mV
  12 V RE             11910 mV
Power 2
  11.3 V bias PEM     11156 mV
  4.6 V bias MidPlane 4801 mV

```

11.3 V bias FPD	11214 mV			
11.3 V bias POE 0	11098 mV			
11.3 V bias POE 1	11330 mV			
Bus Revision	42			
FPGA Revision	1			
PMBus	Expected	Measured	Measured	Calculated
device	voltage	voltage	current	power
XF ASIC A	1000 mV	997 mV	11031 mA	10997 mW
XF ASIC B	1000 mV	996 mV	12125 mA	12076 mW

show chassis environment cb (MX960 Router)

```

user@host> show chassis environment cb
CB 0 status:
State                               Online Master
Temperature                         24 degrees C / 75 degrees F
Power 1
  1.2 V                             1965 mV
  1.5 V                             2465 mV
  1.8 V                             2990 mV
  2.5 V                             3296 mV
  3.3 V                             3296 mV
  5.0 V                             6593 mV
  12.0 V                            13187 mV
  3.3 V bias                        3296 mV
  1.25 V                            1994 mV
  3.3 V SM3                         3296 mV
  5 V RE                            6593 mV
  12 V RE                           13174 mV
Power 2                             Sensor failure
Bus Revision                        4
FPGA Revision                       3

```

show chassis environment cb (MX960 Router with Enhanced MX SCB)

```

user@host> show chassis environment cb
CB 0 status:
State                               Online Master
Temperature                         24 degrees C / 75 degrees F
Power 1
  1.2 V                             1965 mV
  1.5 V                             2465 mV
  1.8 V                             2990 mV
  2.5 V                             3296 mV
  3.3 V                             3296 mV
  5.0 V                             6593 mV
  12.0 V                            13187 mV
  3.3 V bias                        3296 mV
  1.25 V                            1994 mV
  3.3 V SM3                         3296 mV
  5 V RE                            6593 mV
  12 V RE                           13174 mV
Power 2                             Sensor failure
Bus Revision                        4
FPGA Revision                       3
PMBus
device                             Expected    Measured    Measured    Calculated
                                voltage      voltage      current      power
  XF ASIC A                      1000 mV      997 mV      11031 mA     10997 mW
  XF ASIC B                      1000 mV      996 mV      12125 mA     12076 mW

```

show chassis environment cb (MX2020 Router)

```

user@host> show chassis environment cb
CB 0 status:
  State Online Master
  IntakeA-Zone0 Temperature 44 degrees C / 111 degrees F
  IntakeB-Zone1 Temperature 34 degrees C / 93 degrees F
  IntakeC-Zone0 Temperature 45 degrees C / 113 degrees F
  ExhaustA-Zone0 Temperature 43 degrees C / 109 degrees F
  ExhaustB-Zone1 Temperature 36 degrees C / 96 degrees F
  TCBC-Zone0 Temperature 39 degrees C / 102 degrees F
  Power 1
    1.0 V 1011 mV
    1.2 V 1208 mV
    1.8 V 1801 mV
    2.5 V 2552 mV
    3.3 V 3312 mV
    5.0 V 5040 mV
    5.0 V RE 4988 mV
    12.0 V 12065 mV
    12.0 V RE 12046 mV
  Bus Revision 99
  FPGA Revision 270
CB 1 status:
  State Online Standby
  IntakeA-Zone0 Temperature 45 degrees C / 113 degrees F
  IntakeB-Zone1 Temperature 41 degrees C / 105 degrees F
  IntakeC-Zone0 Temperature 46 degrees C / 114 degrees F
  ExhaustA-Zone0 Temperature 44 degrees C / 111 degrees F
  ExhaustB-Zone1 Temperature 41 degrees C / 105 degrees F
  TCBC-Zone0 Temperature 45 degrees C / 113 degrees F
  Power 1
    1.0 V 1008 mV
    1.2 V 1208 mV
    1.8 V 1798 mV
    2.5 V 2539 mV
    3.3 V 3325 mV
    5.0 V 5033 mV
    5.0 V RE 4950 mV
    12.0 V 12046 mV
    12.0 V RE 11968 mV
  Bus Revision 99
  FPGA Revision 0

```

show chassis environment cb (MX2010 Router)

```

user@host> show chassis environment cb
CB 0 status:
  State Online Master
  IntakeA-Zone0 Temperature 36 degrees C / 96 degrees F
  IntakeB-Zone1 Temperature 30 degrees C / 86 degrees F
  IntakeC-Zone0 Temperature 38 degrees C / 100 degrees F
  ExhaustA-Zone0 Temperature 36 degrees C / 96 degrees F
  ExhaustB-Zone1 Temperature 32 degrees C / 89 degrees F
  TCBC-Zone0 Temperature 34 degrees C / 93 degrees F
  Power 1
    1.0 V 1015 mV
    1.2 V 1205 mV
    1.8 V 1804 mV
    2.5 V 2552 mV
    3.3 V 3325 mV

```

```

5.0 V          5020 mV
5.0 V RE       4988 mV
12.0 V         12104 mV
12.0 V RE      12026 mV
Bus Revision    100
FPGA Revision   270
CB 1 status:
State           Online
IntakeA-Zone0 Temperature 35 degrees C / 95 degrees F
IntakeB-Zone1 Temperature 28 degrees C / 82 degrees F
IntakeC-Zone0 Temperature 37 degrees C / 98 degrees F
ExhaustA-Zone0 Temperature 34 degrees C / 93 degrees F
ExhaustB-Zone1 Temperature 29 degrees C / 84 degrees F
TCBC-Zone0 Temperature 32 degrees C / 89 degrees F
Power 1
1.0 V          1011 mV
1.2 V          1208 mV
1.8 V          1788 mV
2.5 V          2526 mV
3.3 V          3319 mV
5.0 V          5046 mV
5.0 V RE       4975 mV
12.0 V         12046 mV
12.0 V RE      12007 mV
Bus Revision    100
FPGA Revision   0

```

show chassis environment cb (T4000 Core Router)

```

user@host> show chassis environment cb
CB 0 status:
State           Online Master
Temperature      33 degrees C / 91 degrees F
Power 1
1.8 V           1805 mV
2.5 V           2523 mV
3.3 V           3324 mV
3.3 V bias      3296 mV
4.6 V           4680 mV
5.0 V           4893 mV
8.0 V bias      7572 mV
12.0 V          11916 mV
Power 2
1.0 V           993 mV
1.2 V           1210 mV
3.3 V RE        3330 mV
Bus Revision     51
FPGA Revision    5
CB 1 status:
State           Online Standby
Temperature      33 degrees C / 91 degrees F
Power 1
1.8 V           1810 mV
2.5 V           2496 mV
3.3 V           3308 mV
3.3 V bias      3286 mV
4.6 V           4692 mV
5.0 V           4954 mV
8.0 V bias      7282 mV
12.0 V          11926 mV
Power 2

```


1.0 V	993 mV
1.2 V	1185 mV
3.3 V RE	3316 mV
Bus Revision	51
FPGA Revision	5

show chassis environment cb (TX Matrix Router)

```
user@host> show chassis environment cb
```

```
-----
CB 0 status:
  State                Online Master
  Temperature          32 degrees C / 89 degrees F
  Power:
    1.8 V              1797 mV
    2.5 V              2477 mV
    3.3 V              3311 mV
    4.6 V              4727 mV
    5.0 V              5015 mV
    12.0 V             12185 mV
    3.3 V bias         3304 mV
    8.0 V bias         7870 mV
  BUS Revision         40
  FPGA Revision        1
CB 1 status:
  State                Online Standby
...
```

```
lcc0-re0:
```

```
-----
CB 0 status:
  State                Online Master
  Temperature          32 degrees C / 89 degrees F
  Power:
    1.8 V              1787 mV
    2.5 V              2473 mV
    3.3 V              3306 mV
    4.6 V              4793 mV
    5.0 V              5025 mV
    12.0 V             12156 mV
    3.3 V bias         3289 mV
    8.0 V bias         7609 mV
  BUS Revision         40
  FPGA Revision        5
CB 1 status:
  State                Online Standby
....
  BUS Revision         40
  FPGA Revision        5
```

```
lcc2-re0:
```

```
-----
CB 0 status:
  State                Online Master
...
CB 1 status:
  State                Online Standby
...
```

show chassis environment cb (TX Matrix Plus Router)

```
user@host> show chassis environment cb
sfc0-re0:
```

CB 0 status:

State	Online Master
Temperature	38 degrees C / 100 degrees F
Power 1	
1.0 V	1005 mV
1.1 V	1108 mV
1.2 V	1205 mV
1.25 V	1269 mV
1.5 V	1508 mV
1.8 V	1814 mV
2.5 V	2507 mV
3.3 V	3306 mV
3.3 V bias	3300 mV
9.0 V	9058 mV
9.0 V RE	9107 mV
Power 2	
3.9 V	3963 mV
5.0 V	5020 mV
9.0 V	9087 mV
Bus Revision	79
FPGA Revision	23

CB 1 status:

State	Online Standby
Temperature	39 degrees C / 102 degrees F
Power 1	
1.0 V	1002 mV
1.1 V	1105 mV
1.2 V	1198 mV
1.25 V	1276 mV
1.5 V	1504 mV
1.8 V	1804 mV
2.5 V	2507 mV
3.3 V	3300 mV
3.3 V bias	3293 mV
9.0 V	9039 mV
9.0 V RE	9049 mV
Power 2	
3.9 V	3892 mV
5.0 V	5040 mV
9.0 V	9058 mV
Bus Revision	79
FPGA Revision	23

```
lcc0-re0:
```

CB 0 status:

State	Online Master
Temperature	39 degrees C / 102 degrees F
Power 1	
1.8 V	1799 mV
2.5 V	2499 mV
3.3 V	3327 mV
3.3 V bias	3299 mV
4.6 V	4673 mV
5.0 V	4918 mV
8.0 V bias	7308 mV

```

    12.0 V          11887 mV
Power 2
    1.0 V          996 mV
    1.2 V          1199 mV
    3.3 V RE       3319 mV
Bus Revision      51
FPGA Revision     3
CB 1 status:
State             Online Standby
Temperature       40 degrees C / 104 degrees F
Power 1
    1.8 V          1800 mV
    2.5 V          2496 mV
    3.3 V          3322 mV
    3.3 V bias     3284 mV
    4.6 V          4680 mV
    5.0 V          4954 mV
    8.0 V bias     7284 mV
    12.0 V         11902 mV
Power 2
    1.0 V          998 mV
    1.2 V          1205 mV
    3.3 V RE       3327 mV
Bus Revision      51
FPGA Revision     3

```

```
lcc1-re0:
```

```

-----
CB 0 status:
State             Online Master
Temperature       41 degrees C / 105 degrees F
Power 1
    1.8 V          1804 mV
    2.5 V          2517 mV
    3.3 V          3300 mV
    3.3 V bias     3284 mV
    4.6 V          4681 mV
    5.0 V          4927 mV
    8.0 V bias     7357 mV
    12.0 V         11907 mV
Power 2
    1.0 V          991 mV
    1.2 V          1202 mV
    3.3 V RE       3301 mV
Bus Revision      51
FPGA Revision     3
CB 1 status:
State             Online Standby
Temperature       40 degrees C / 104 degrees F
Power 1
    1.8 V          1805 mV
    2.5 V          2528 mV
    3.3 V          3324 mV
    3.3 V bias     3289 mV
    4.6 V          4694 mV
    5.0 V          4959 mV
    8.0 V bias     7311 mV
    12.0 V         11926 mV
Power 2
    1.0 V          998 mV
    1.2 V          1200 mV

```

3.3 V RE	3313 mV
Bus Revision	51
FPGA Revision	3

1cc2-re0:

CB 0 status:

State	Online Master
Temperature	41 degrees C / 105 degrees F
Power 1	
1.8 V	1805 mV
2.5 V	2494 mV
3.3 V	3333 mV
3.3 V bias	3296 mV
4.6 V	4673 mV
5.0 V	4901 mV
8.0 V bias	7343 mV
12.0 V	11916 mV
Power 2	
1.0 V	993 mV
1.2 V	1213 mV
3.3 V RE	3328 mV
Bus Revision	51
FPGA Revision	3

CB 1 status:

State	Online Standby
Temperature	41 degrees C / 105 degrees F
Power 1	
1.8 V	1804 mV
2.5 V	2523 mV
3.3 V	3334 mV
3.3 V bias	3291 mV
4.6 V	4697 mV
5.0 V	4969 mV
8.0 V bias	7308 mV
12.0 V	11936 mV
Power 2	
1.0 V	996 mV
1.2 V	1200 mV
3.3 V RE	3328 mV
Bus Revision	51
FPGA Revision	3

1cc3-re0:

CB 0 status:

State	Online Master
Temperature	37 degrees C / 98 degrees F
Power 1	
1.8 V	1809 mV
2.5 V	2510 mV
3.3 V	3296 mV
3.3 V bias	3291 mV
4.6 V	4670 mV
5.0 V	4905 mV
8.0 V bias	7211 mV
12.0 V	11882 mV
Power 2	
1.0 V	996 mV
1.2 V	1188 mV
3.3 V RE	3326 mV

```

Bus Revision          51
FPGA Revision         5
CB 1 status:
State                 Online Standby
Temperature           38 degrees C / 100 degrees F
Power 1
  1.8 V               1813 mV
  2.5 V               2510 mV
  3.3 V               3322 mV
  3.3 V bias          3289 mV
  4.6 V               4692 mV
  5.0 V               4967 mV
  8.0 V bias          7194 mV
  12.0 V              11916 mV
Power 2
  1.0 V               996 mV
  1.2 V               1205 mV
  3.3 V RE            3273 mV
Bus Revision          51
FPGA Revision         5

```

show chassis environment cb (EX8200 Switch)

```
user@host> show chassis environment cb
```

```

CB 0 status:
State                 Online Master
Temperature Intake     20 degrees C / 68 degrees F
Temperature Exhaust    24 degrees C / 75 degrees F
Power 1
  1.1 V               1086 mV
  1.2 V               1179 mV
  1.2 V *             1182 mV
  1.2 V *             1182 mV
  1.25 V              1211 mV
  1.5 V               1472 mV
  1.8 V               1756 mV
  2.5 V               2449 mV
  3.3 V               3254 mV
  3.3 V bias          3300 mV
  5.0 V               4911 mV
  12.0 V              11891 mV
Power 2
  3.3 V bias *        3615 mV
  3.3 V bias *        3615 mV
  3.3 V bias *        3567 mV
  3.3 V bias *        3664 mV
  4.3 V bias *        4224 mV
  4.3 V bias *        4215 mV
  4.3 V bias *        4224 mV
  4.3 V bias *        4205 mV
  4.3 V bias *        4195 mV
  4.3 V bias *        4215 mV
  5.0 V bias          4920 mV
CB 1 status:
State                 Online Standby
Temperature Intake     19 degrees C / 66 degrees F
Temperature Exhaust    23 degrees C / 73 degrees F
Power 1
  1.1 V               1082 mV
  1.2 V               1169 mV

```

```

1.2 V *          1179 mV
1.2 V *          1179 mV
1.25 V          1214 mV
1.5 V           1482 mV
1.8 V           1759 mV
2.5 V           2481 mV
3.3 V           3248 mV
3.3 V bias      3306 mV
5.0 V           4911 mV
12.0 V          11910 mV
Power 2
3.3 V bias *    3644 mV
3.3 V bias *    3664 mV
3.3 V bias *    3586 mV
3.3 V bias *    3654 mV
4.3 V bias *    4224 mV
4.3 V bias *    4215 mV
4.3 V bias *    4224 mV
4.3 V bias *    4205 mV
4.3 V bias *    4244 mV
4.3 V bias *    4215 mV
5.0 V bias      4930 mV
CB 2 status:
State           Online
Temperature Intake 19 degrees C / 66 degrees F
Temperature Exhaust 23 degrees C / 73 degrees F
Power 1
1.2 V           1195 mV
1.5 V           1511 mV
1.8 V           1804 mV
2.5 V           2526 mV
3.3 V           3300 mV
3.3 V bias      3306 mV
12.0 V          12220 mV

```

show chassis environment cb (EX8208 Switch)

```

user@host> show chassis environment cb
CB 0 status:
State           Online Master
Temperature Intake 20 degrees C / 68 degrees F
Temperature Exhaust 24 degrees C / 75 degrees F
Power 1
1.1 V           1086 mV
1.2 V           1179 mV
1.2 V *         1182 mV
1.2 V *         1182 mV
1.25 V          1211 mV
1.5 V           1466 mV
1.8 V           1759 mV
2.5 V           2455 mV
3.3 V           3261 mV
3.3 V bias      3300 mV
5.0 V           4930 mV
12.0 V          11891 mV
Power 2
3.3 V bias *    3606 mV
3.3 V bias *    3615 mV
3.3 V bias *    3567 mV
3.3 V bias *    3673 mV
4.3 V bias *    4224 mV

```

```

4.3 V bias *      4215 mV
4.3 V bias *      4234 mV
4.3 V bias *      4205 mV
4.3 V bias *      4186 mV
4.3 V bias *      4215 mV
5.0 V bias        4940 mV
CB 1 status:
State              Online Standby
Temperature Intake  19 degrees C / 66 degrees F
Temperature Exhaust 23 degrees C / 73 degrees F
Power 1
1.1 V              1086 mV
1.2 V              1169 mV
1.2 V *            1179 mV
1.2 V *            1179 mV
1.25 V             1211 mV
1.5 V              1479 mV
1.8 V              1759 mV
2.5 V              2475 mV
3.3 V              3235 mV
3.3 V bias         3306 mV
5.0 V              4930 mV
12.0 V             11891 mV
Power 2
3.3 V bias *       3644 mV
3.3 V bias *       3664 mV
3.3 V bias *       3586 mV
3.3 V bias *       3654 mV
4.3 V bias *       4215 mV
4.3 V bias *       4224 mV
4.3 V bias *       4215 mV
4.3 V bias *       4215 mV
4.3 V bias *       4234 mV
4.3 V bias *       4224 mV
5.0 V bias         4920 mV
CB 2 status:
State              Online
Temperature Intake  20 degrees C / 68 degrees F
Temperature Exhaust 24 degrees C / 75 degrees F
Power 1
1.2 V              1202 mV
1.5 V              1508 mV
1.8 V              1804 mV
2.5 V              2520 mV
3.3 V              3300 mV
3.3 V bias         3300 mV
12.0 V             12200 mV

```

show chassis environment cb (PTX5000 Packet Transport Switch)

```

user@host> show chassis environment cb
CB 0 status:
State              Online Master
Intake Temperature 38 degrees C / 100 degrees F
Exhaust A Temperature 45 degrees C / 113 degrees F
Exhaust B Temperature 42 degrees C / 107 degrees F
Power 1
1.2 V              1200 mV
1.25 V             1250 mV
2.5 V              2500 mV
3.3 V              3300 mV

```

```
Power 2
  1.0 V          1000 mV
  3.3 V bias     3293 mV
  3.9 V          3921 mV
Bus Revision     132
FPGA Revision    27
CB 1 status:
State            Online Standby
Intake Temperature 34 degrees C / 93 degrees F
Exhaust A Temperature 39 degrees C / 102 degrees F
Exhaust B Temperature 36 degrees C / 96 degrees F
Power 1
  1.2 V          1199 mV
  1.25 V         1250 mV
  2.5 V          2499 mV
  3.3 V          3299 mV
Power 2
  1.0 V          1000 mV
  3.3 V bias     3312 mV
  3.9 V          3961 mV
Bus Revision     132
FPGA Revision    28
```

show chassis environment cb (QFabric System)

```
user@switch> show chassis environment cb interconnect-device IC-123 0
CB 0 status:
State            Online Master
Left Intake Temperature 33 degrees C / 91 degrees F
Right Intake Temperature 33 degrees C / 91 degrees F
Left Exhaust Temperature 36 degrees C / 96 degrees F
Right Exhaust Temperature 35 degrees C / 95 degrees F
Power            OK
  VDD 3V3        3294 mV
  VDD 2V5        2436 mV
  VDD 1V8        1746 mV
  VDD 1V5        1460 mV
  VDD 1V25       1210 mV
  VDD 1V2        1164 mV
  CPU CORE 1V2   1120 mV
  VDD 1V0        968 mV
  VDD 5V0        5088 mV
  CPU MP BIAS 4V3 4050 mV
  BIAS 3V3       3180 mV
  VTT 0V9        866 mV
```


show chassis environment fpc

Syntax	show chassis environment fpc <slot>
Syntax (TX Matrix and TX Matrix Plus Routers)	show chassis environment fpc <fcc number> <slot>
Syntax (MX Series Routers)	show chassis environment fpc <slot> <all-members> <local> <member <i>member-id</i> >
Syntax (MX2010 3D Universal Edge Routers)	show chassis environment fpc <slot>
Syntax (MX2020 3D Universal Edge Routers)	show chassis environment fpc <slot>
Syntax (QFX Series)	show chassis environment fpc <fpc-slot> interconnect-device <i>name</i>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for QFX Series. Command introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches. Command introduced in Junos OS Release 12.1 for T4000 Core Routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.
Description	(M40e, M120, M160, M320, MX Series, T Series routers, EX Series, QFX Series, and PTX Series switches only) Display environmental information about Flexible PIC Concentrators (FPCs).
Options	<p>none—Display environmental information about all FPCs. On a TX Matrix router, display environmental information about all FPCs on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display environmental information about all FPCs on the TX Matrix Plus router and its attached T1600 routers.</p> <p>all-members—(MX Series routers only) (Optional) Display environmental information for the FPCs in all the members of the Virtual Chassis configuration.</p> <p>interconnect-device <i>name</i>—(QFabric systems only) (Optional) Display chassis environmental information for the Interconnect device.</p> <p>fcc number—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display environmental information about the FPC in a T640 router (or line-card</p>

chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display environmental information about the FPC in a T1600 router (or line-card chassis) that is connected to the TX Matrix Plus router. Replace **number** with a value from 0 through 3.

local—(MX Series routers only) (Optional) Display environmental information for the FPCs in the local Virtual Chassis member.

member member-id—(MX Series routers only) (Optional) Display environmental information for the FPCs in the specified member of the Virtual Chassis configuration. Replace **member-id** with a value of 0 or 1.

slot or fpc-slot—(Optional) Display environmental information about an individual FPC:

- (TX Matrix and TX Matrix Plus routers only) On a TX Matrix router, if you specify the number of the T640 router by using only the **lcc number** option (the recommended method), replace **slot** with a value from 0 through 7. Similarly, on a TX Matrix Plus router, if you specify the number of the T1600 router by using only the **lcc number** option (the recommended method), replace **slot** with a value from 0 through 7. Otherwise, replace **slot** with a value from 0 through 31. For example, the following commands have the same result:

```
user@host> show chassis environment fpc 1 lcc 1
```

```
user@host> show chassis environment fpc 9
```

- M120 router—Replace **slot** with a value from 0 through 5.
- MX240 router—Replace **slot** with a value from 0 through 2.
- MX480 router—Replace **slot** with a value from 0 through 5.
- MX960 router—Replace **slot** with a value from 0 through 11.
- MX2010 router—Replace **slot** with a value from 0 through 9.
- MX2020 router—Replace **slot** with a value from 0 through 19.
- Other routers—Replace **slot** with a value from 0 through 7.
- EX Series switches:
 - EX3200 switches and EX4200 standalone switches—Replace **slot** with 0.
 - EX4200 switches in a Virtual Chassis configuration—Replace **slot** with a value from 0 through 9 (switch's member ID).
 - EX6210 switches—Replace **slot** with a value from 0 through 3 (line card only), 4 or 5 (line card or Switch Fabric and Rotating Engine (SRE) module), or 6 through 9 (line card only).
 - EX8208 switches—Replace **slot** with a value from 0 through 7 (line card).
 - EX8216 switches—Replace **slot** with a value from 0 through 15 (line card).
- QFX3500 switches —Replace **fpc-slot** with 0 through 15.
- PTX5000 Packet Transport Switch—Replace **fpc-slot** with 0 through 7.

Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>request chassis fpc</i> • show chassis fpc on page 4298 • <i>show chassis fpc-feb-connectivity</i> • <i>Configuring the Junos OS to Resynchronize FPC Sequence Numbers with Active FPCs when an FPC Comes Online</i> • <i>MX960 Flexible PIC Concentrator Description</i>
List of Sample Output	show chassis environment fpc (M120 Router) on page 599 show chassis environment fpc (M160 Router) on page 600 show chassis environment fpc (M320 Router) on page 600 show chassis environment fpc (MX2020 Router) on page 601 show chassis environment fpc (MX2010 Router) on page 604 show chassis environment fpc (MX240 Router) on page 606 show chassis environment fpc (MX480 Router) on page 607 show chassis environment fpc (MX960 Router) on page 608 show chassis environment fpc (MX480 Router with 100-Gigabit Ethernet CFP) on page 609 show chassis environment fpc (MX240, MX480, MX960 with Application Services Modular Line Card on page 610 show chassis environment fpc (T320, T640, and T1600 Routers) on page 610 show chassis environment fpc (T4000 Router) on page 611 show chassis environment fpc lcc (TX Matrix Router) on page 616 show chassis environment fpc lcc (TX Matrix Plus Router) on page 617 show chassis environment fpc (QFX Series) on page 618 show chassis environment fpc interconnect-device (QFabric Systems) on page 618 show chassis environment fpc 0 (PTX5000 Packet Transport Switch) on page 618 show chassis environment FPC 1 (MX Routers with Media Services Blade [MSB]) on page 619
Output Fields	<p>Table 55 on page 598 lists the output fields for the show chassis environment fpc command. Output fields are listed in the approximate order in which they appear.</p>

Table 55: show chassis environment fpc Output Fields

Field Name	Field Description
State	<p>Status of the FPC:</p> <ul style="list-style-type: none"> • Unknown—FPC is not detected by the router. • Empty—No FPC is present. • Present—FPC is detected by the chassis daemon but is either not supported by the current version of the Junos OS, or the FPC is coming up but not yet online. • Ready—FPC is in intermediate or transition state. • Announce online—Intermediate state during which the FPC is coming up but not yet online, and the chassis manager acknowledges the chassisd FPC online initiative. • Online—FPC is online and running. • Offline—FPC is powered down. • Diagnostics—FPC is set to operate in diagnostics mode.
Temperature	(M40e and M160 routers and QFX Series only) Temperature of the air flowing past the FPC.
PMB Temperature	(PTX Series only) Temperature of the air flowing past the PMB (bottom of the FPC).
Temperature Intake	(M320 routers, MX2010 routers, MX2020 routers, and PTX Series only) Temperature of the air flowing into the chassis.
Temperature Top	(T Series routers only) Temperature of the air flowing past the top of the FPC.
Temperature Exhaust	<p>(M120 and M320 routers, MX2010 routers, MX2020 routers, and PTX Series only) Temperature of the air flowing out of the chassis.</p> <p>The PTX Series Packet Transport Switches, and the MX2010 and MX2020 routers include exhaust temperatures for multiple zones (Exhaust A and Exhaust B).</p>
Temperature Bottom	(T Series routers only) Temperature of the air flowing past the bottom of the FPC.
TL <i>n</i> Temperature	(PTX Series only) Temperature of the air flowing past the specified TL area of the packet forwarding engine (PFE) on the FPC.
TQ <i>n</i> Temperature	(PTX Series only) Temperature of the air flowing past the specified TQ area of the packet forwarding engine (PFE) on the FPC.
Temperature MMBO	(T640 router only) Temperature of the air flowing past the type 3 FPC.
Temperature MMB1	(M320 and T Series routers only) Temperature of the air flowing past the type 1, type 2, and type 3 FPC.
Power	Information about the voltage supplied to the FPC. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.
CMB Revision or BUS revision	Revision level of the chassis management bus device (M Series router) or bus (T Series routers).

Sample Output

show chassis environment fpc (M120 Router)

```

user@host> show chassis environment fpc
FPC 2 status:
  State                               Online
  Temperature Exhaust A               32 degrees C / 89 degrees F
  Temperature Exhaust B               31 degrees C / 87 degrees F
  Power A-Board
    1.2 V                             1202 mV
    1.5 V                             1508 mV
    1.8 V                             1798 mV
    2.5 V                             2507 mV
    3.3 V                             3351 mV
    5.0 V                             4995 mV
    3.3 V bias                         3296 mV
    1.2 V Rocket IO                   1205 mV
    1.5 V Rocket IO                   1501 mV
  I2C Slave Revision                 12
FPC 3 status:
  State                               Online
  Temperature Exhaust A               31 degrees C / 87 degrees F
  Temperature Exhaust B               33 degrees C / 91 degrees F
  Power A-Board
    1.2 V                             1211 mV
    1.5 V                             1501 mV
    1.8 V                             1798 mV
    2.5 V                             2471 mV
    3.3 V                             3293 mV
    5.0 V                             4930 mV
    3.3 V bias                         3296 mV
    1.2 V Rocket IO                   1205 mV
    1.5 V Rocket IO                   1501 mV
  Power B-Board
    1.2 V                             1214 mV
    1.5 V                             1501 mV
    2.5 V                             2471 mV
    3.3 V                             3300 mV
    5.0 V                             4943 mV
    3.3 V bias                         3296 mV
    1.2 V Rocket IO                   1205 mV
    1.5 V Rocket IO                   1501 mV
  I2C Slave Revision                 12
FPC 4 status:
  State                               Online
  Temperature Exhaust A               32 degrees C / 89 degrees F
  Temperature Exhaust B               30 degrees C / 86 degrees F
  Power A-Board
    1.2 V                             1195 mV
    1.5 V                             1504 mV
    1.8 V                             1801 mV
    2.5 V                             2504 mV
    3.3 V                             3293 mV
    5.0 V                             4917 mV
    3.3 V bias                         3296 mV
    1.2 V Rocket IO                   1202 mV
    1.5 V Rocket IO                   1492 mV
  I2C Slave Revision                 12

```

show chassis environment fpc (M160 Router)

```
user@host> show chassis environment fpc
FPC 0 status:
  State                               Online
  Temperature                         42 degrees C / 107 degrees F
  Power:
    1.5 V                             1500 mV
    2.5 V                             2509 mV
    3.3 V                             3308 mV
    5.0 V                             4991 mV
    5.0 V bias                         4952 mV
    8.0 V bias                         8307 mV
  CMB Revision                         12
FPC 1 status:
  State                               Online
  Temperature                         45 degrees C / 113 degrees F
  Power:
    1.5 V                             1498 mV
    2.5 V                             2501 mV
    3.3 V                             3319 mV
    5.0 V                             5020 mV
    5.0 V bias                         5025 mV
    8.0 V bias                         8307 mV
  CMB Revision                         12
```

show chassis environment fpc (M320 Router)

```
user@host> show chassis environment fpc
FPC 0 status:
  State                               Online
  Temperature Intake                  27 degrees C / 80 degrees F
  Temperature Exhaust                 38 degrees C / 100 degrees F
  Temperature MMB1                    31 degrees C / 87 degrees F
  Power:
    1.5 V                             1487 mV
    1.5 V *                           1494 mV
    1.8 V                             1821 mV
    2.5 V                             2533 mV
    3.3 V                             3323 mV
    5.0 V                             5028 mV
    3.3 V bias                        3296 mV
    5.0 V bias                        4984 mV
  CMB Revision                        16
FPC 1 status:
  State                               Online
  Temperature Intake                  27 degrees C / 80 degrees F
  Temperature Exhaust                 37 degrees C / 98 degrees F
  Temperature MMB1                    32 degrees C / 89 degrees F
  Power:
    1.5 V                             1504 mV
    1.5 V *                           1499 mV
    1.8 V                             1820 mV
    2.5 V                             2529 mV
    3.3 V                             3328 mV
    5.0 V                             5013 mV
    3.3 V bias                        3294 mV
    5.0 V bias                        4984 mV
  CMB Revision                        16
FPC 2 status:
  State                               Online
```

```

Temperature Intake          28 degrees C / 82 degrees F
Temperature Exhaust         38 degrees C / 100 degrees F
Temperature MMB1            32 degrees C / 89 degrees F
Power:
  1.5 V                     1498 mV
  1.5 V *                   1487 mV
  1.8 V                     1816 mV
  2.5 V                     2531 mV
  3.3 V                     3324 mV
  5.0 V                     5025 mV
  3.3 V bias                3277 mV
  5.0 V bias                5013 mV
CMB Revision                17
FPC 3 status:
...

```

show chassis environment fpc (MX2020 Router)

```

user@host> show chassis environment fpc
FPC 0 status:
State                               Online
Temperature Intake                  41 degrees C / 105 degrees F
Temperature Exhaust A               48 degrees C / 118 degrees F
Temperature Exhaust B               60 degrees C / 140 degrees F
Temperature LU 0 TSen               56 degrees C / 132 degrees F
Temperature LU 0 Chip               59 degrees C / 138 degrees F
Temperature LU 1 TSen               56 degrees C / 132 degrees F
Temperature LU 1 Chip               61 degrees C / 141 degrees F
Temperature LU 2 TSen               56 degrees C / 132 degrees F
Temperature LU 2 Chip               52 degrees C / 125 degrees F
Temperature LU 3 TSen               56 degrees C / 132 degrees F
Temperature LU 3 Chip               52 degrees C / 125 degrees F
Temperature MQ 0 TSen               49 degrees C / 120 degrees F
Temperature MQ 0 Chip               49 degrees C / 120 degrees F
Temperature MQ 1 TSen               49 degrees C / 120 degrees F
Temperature MQ 1 Chip               52 degrees C / 125 degrees F
Temperature MQ 2 TSen               49 degrees C / 120 degrees F
Temperature MQ 2 Chip               45 degrees C / 113 degrees F
Temperature MQ 3 TSen               49 degrees C / 120 degrees F
Temperature MQ 3 Chip               46 degrees C / 114 degrees F
Power
  AS-BIAS3V3-z12105                3299 mV
  AS-VDD1V8-z12006                 1807 mV
  AS-VDD2V5-z12006                 2512 mV
  AS-AVDD1V0-z12004                 997 mV
  AS-PCIE_1V0-z12004                 996 mV
  AS-VDD3V3-z12004                 3294 mV
  AS-VDD_1V5A-z12004                1501 mV
  AS-VDD_1V5B-z12004                1498 mV
  AS-LU0_1V0-z12004                 998 mV
  AS-LU1_1V0-z12004                1002 mV
  AS-MQ0_1V0-z12004                 999 mV
  AS-MQ1_1V0-z12004                 994 mV
  AS-LU2_1V0-z12004                1000 mV
  AS-LU3_1V0-z12004                 998 mV
  AS-MQ2_1V0-z12004                1002 mV
  AS-MQ3_1V0-z12004                 999 mV
  AS-PMB_1V1-z12006                1096 mV
I2C Slave Revision                68
FPC 1 status:
State                               Online

```

Temperature Intake	39 degrees C / 102 degrees F
Temperature Exhaust A	48 degrees C / 118 degrees F
Temperature Exhaust B	55 degrees C / 131 degrees F
Temperature LU 0 TSen	52 degrees C / 125 degrees F
Temperature LU 0 Chip	54 degrees C / 129 degrees F
Temperature LU 1 TSen	52 degrees C / 125 degrees F
Temperature LU 1 Chip	56 degrees C / 132 degrees F
Temperature LU 2 TSen	52 degrees C / 125 degrees F
Temperature LU 2 Chip	49 degrees C / 120 degrees F
Temperature LU 3 TSen	52 degrees C / 125 degrees F
Temperature LU 3 Chip	50 degrees C / 122 degrees F
Temperature MQ 0 TSen	48 degrees C / 118 degrees F
Temperature MQ 0 Chip	48 degrees C / 118 degrees F
Temperature MQ 1 TSen	48 degrees C / 118 degrees F
Temperature MQ 1 Chip	51 degrees C / 123 degrees F
Temperature MQ 2 TSen	48 degrees C / 118 degrees F
Temperature MQ 2 Chip	45 degrees C / 113 degrees F
Temperature MQ 3 TSen	48 degrees C / 118 degrees F
Temperature MQ 3 Chip	45 degrees C / 113 degrees F
Power	
AS-BIAS3V3-z12105	3291 mV
AS-VDD1V8-z12006	1786 mV
AS-VDD2V5-z12006	2496 mV
AS-AVDD1V0-z12004	1000 mV
AS-PCIE_1V0-z12004	1000 mV
AS-VDD3V3-z12004	3294 mV
AS-VDD_1V5A-z12004	1500 mV
AS-VDD_1V5B-z12004	1498 mV
AS-LU0_1V0-z12004	1003 mV
AS-LU1_1V0-z12004	1000 mV
AS-MQ0_1V0-z12004	1000 mV
AS-MQ1_1V0-z12004	995 mV
AS-LU2_1V0-z12004	1002 mV
AS-LU3_1V0-z12004	997 mV
AS-MQ2_1V0-z12004	1000 mV
AS-MQ3_1V0-z12004	998 mV
AS-PMB_1V1-z12006	1096 mV
I2C Slave Revision	68
FPC 2 status:	
State	Online
Temperature Intake	39 degrees C / 102 degrees F
Temperature Exhaust A	48 degrees C / 118 degrees F
Temperature Exhaust B	58 degrees C / 136 degrees F
Temperature LU 0 TSen	55 degrees C / 131 degrees F
Temperature LU 0 Chip	57 degrees C / 134 degrees F
Temperature LU 1 TSen	55 degrees C / 131 degrees F
Temperature LU 1 Chip	63 degrees C / 145 degrees F
Temperature LU 2 TSen	55 degrees C / 131 degrees F
Temperature LU 2 Chip	51 degrees C / 123 degrees F
Temperature LU 3 TSen	55 degrees C / 131 degrees F
Temperature LU 3 Chip	52 degrees C / 125 degrees F
Temperature MQ 0 TSen	48 degrees C / 118 degrees F
Temperature MQ 0 Chip	50 degrees C / 122 degrees F
Temperature MQ 1 TSen	48 degrees C / 118 degrees F
Temperature MQ 1 Chip	52 degrees C / 125 degrees F
Temperature MQ 2 TSen	48 degrees C / 118 degrees F
Temperature MQ 2 Chip	47 degrees C / 116 degrees F
Temperature MQ 3 TSen	48 degrees C / 118 degrees F
Temperature MQ 3 Chip	47 degrees C / 116 degrees F
Power	
AS-BIAS3V3-z12105	3299 mV

AS-VDD1V8-z12006	1805 mV
AS-VDD2V5-z12006	2510 mV
AS-AVDD1V0-z12004	999 mV
AS-PCIE_1V0-z12004	998 mV
AS-VDD3V3-z12004	3296 mV
AS-VDD_1V5A-z12004	1492 mV
AS-VDD_1V5B-z12004	1497 mV
AS-LU0_1V0-z12004	997 mV
AS-LU1_1V0-z12004	1000 mV
AS-MQ0_1V0-z12004	998 mV
AS-MQ1_1V0-z12004	1001 mV
AS-LU2_1V0-z12004	996 mV
AS-LU3_1V0-z12004	995 mV
AS-MQ2_1V0-z12004	998 mV
AS-MQ3_1V0-z12004	997 mV
AS-PMB_1V1-z12006	1100 mV
I2C Slave Revision	68
FPC 3 status:	
State	Online
Temperature Intake	41 degrees C / 105 degrees F
Temperature Exhaust A	48 degrees C / 118 degrees F
Temperature Exhaust B	58 degrees C / 136 degrees F
Temperature LU 0 TSen	56 degrees C / 132 degrees F
Temperature LU 0 Chip	59 degrees C / 138 degrees F
Temperature LU 1 TSen	56 degrees C / 132 degrees F
Temperature LU 1 Chip	61 degrees C / 141 degrees F
Temperature LU 2 TSen	56 degrees C / 132 degrees F
Temperature LU 2 Chip	51 degrees C / 123 degrees F
Temperature LU 3 TSen	56 degrees C / 132 degrees F
Temperature LU 3 Chip	53 degrees C / 127 degrees F
Temperature MQ 0 TSen	50 degrees C / 122 degrees F
Temperature MQ 0 Chip	51 degrees C / 123 degrees F
Temperature MQ 1 TSen	50 degrees C / 122 degrees F
Temperature MQ 1 Chip	55 degrees C / 131 degrees F
Temperature MQ 2 TSen	50 degrees C / 122 degrees F
Temperature MQ 2 Chip	47 degrees C / 116 degrees F
Temperature MQ 3 TSen	50 degrees C / 122 degrees F
Temperature MQ 3 Chip	50 degrees C / 122 degrees F
Power	
AS-BIAS3V3-z12105	3305 mV
AS-VDD1V8-z12006	1810 mV
AS-VDD2V5-z12006	2508 mV
AS-AVDD1V0-z12004	999 mV
AS-PCIE_1V0-z12004	1001 mV
AS-VDD3V3-z12004	3294 mV
AS-VDD_1V5A-z12004	1500 mV
AS-VDD_1V5B-z12004	1498 mV
AS-LU0_1V0-z12004	998 mV
AS-LU1_1V0-z12004	998 mV
AS-MQ0_1V0-z12004	999 mV
AS-MQ1_1V0-z12004	998 mV
AS-LU2_1V0-z12004	1000 mV
AS-LU3_1V0-z12004	1001 mV
AS-MQ2_1V0-z12004	996 mV
AS-MQ3_1V0-z12004	998 mV
AS-PMB_1V1-z12006	1098 mV
I2C Slave Revision	68
FPC 4 status:	
...	

show chassis environment fpc (MX2010 Router)

```

user@host> show chassis environment fpc
FPC 0 status:
  State      Online
  Temperature Intake      36 degrees C / 96 degrees F
  Temperature Exhaust A   42 degrees C / 107 degrees F
  Temperature Exhaust B   51 degrees C / 123 degrees F
  Temperature LU 0 TSen    49 degrees C / 120 degrees F
  Temperature LU 0 Chip    50 degrees C / 122 degrees F
  Temperature LU 1 TSen    49 degrees C / 120 degrees F
  Temperature LU 1 Chip    54 degrees C / 129 degrees F
  Temperature LU 2 TSen    49 degrees C / 120 degrees F
  Temperature LU 2 Chip    45 degrees C / 113 degrees F
  Temperature LU 3 TSen    49 degrees C / 120 degrees F
  Temperature LU 3 Chip    46 degrees C / 114 degrees F
  Temperature MQ 0 TSen    40 degrees C / 104 degrees F
  Temperature MQ 0 Chip    41 degrees C / 105 degrees F
  Temperature MQ 1 TSen    40 degrees C / 104 degrees F
  Temperature MQ 1 Chip    44 degrees C / 111 degrees F
  Temperature MQ 2 TSen    40 degrees C / 104 degrees F
  Temperature MQ 2 Chip    38 degrees C / 100 degrees F
  Temperature MQ 3 TSen    40 degrees C / 104 degrees F
  Temperature MQ 3 Chip    41 degrees C / 105 degrees F
  Power
    AS-BIAS3V3-z12105      3300 mV
    AS-VDD1V8-z12006       1805 mV
    AS-VDD2V5-z12006       2505 mV
    AS-AVDD1V0-z12004       998 mV
    AS-PCIE_1V0-z12004       999 mV
    AS-VDD3V3-z12004       3303 mV
    AS-VDD_1V5A-z12004     1497 mV
    AS-VDD_1V5B-z12004     1497 mV
    AS-LU0_1V0-z12004       998 mV
    AS-LU1_1V0-z12004     1003 mV
    AS-MQ0_1V0-z12004       998 mV
    AS-MQ1_1V0-z12004       998 mV
    AS-LU2_1V0-z12004       997 mV
    AS-LU3_1V0-z12004     1001 mV
    AS-MQ2_1V0-z12004       996 mV
    AS-MQ3_1V0-z12004       994 mV
    AS-PMB_1V1-z12006     1097 mV
  I2C Slave Revision      68
FPC 1 status:
  State      Online
  Temperature Intake      34 degrees C / 93 degrees F
  Temperature Exhaust A   46 degrees C / 114 degrees F
  Temperature Exhaust B   54 degrees C / 129 degrees F
  Temperature LU 0 TSen    45 degrees C / 113 degrees F
  Temperature LU 0 Chip    55 degrees C / 131 degrees F
  Temperature LU 1 TSen    45 degrees C / 113 degrees F
  Temperature LU 1 Chip    44 degrees C / 111 degrees F
  Temperature LU 2 TSen    45 degrees C / 113 degrees F
  Temperature LU 2 Chip    50 degrees C / 122 degrees F
  Temperature LU 3 TSen    45 degrees C / 113 degrees F
  Temperature LU 3 Chip    58 degrees C / 136 degrees F
  Temperature XM 0 TSen    45 degrees C / 113 degrees F
  Temperature XM 0 Chip    51 degrees C / 123 degrees F
  Temperature XF 0 TSen    45 degrees C / 113 degrees F
  Temperature XF 0 Chip    63 degrees C / 145 degrees F
  Temperature PLX Switch TSen 45 degrees C / 113 degrees F

```

```

Temperature PLX Switch Chip47 degrees C / 116 degrees F
Power
MPC-BIAS3V3-z12105      3300 mV
MPC-VDD3V3-z16100      3294 mV
MPC-VDD2V5-z16100      2505 mV
MPC-VDD1V8-z12004      1796 mV
MPC-AVDD1V0-z12004      991 mV
MPC-VDD1V2-z16100      1196 mV
MPC-VDD1V5A-z12004      1491 mV
MPC-VDD1V5B-z12004      1492 mV
MPC-XF_0V9-z12004      996 mV
MPC-PCIE_1V0-z16100     1003 mV
MPC-LU0_1V0-z12004      996 mV
MPC-LU1_1V0-z12004      996 mV
MPC-LU2_1V0-z12004      998 mV
MPC-LU3_1V0-z12004      994 mV
MPC-12VA-BMR453         12031 mV
MPC-12VB-BMR453         12003 mV
MPC-PMB_1V1-z12006      1104 mV
MPC-PMB_1V2-z12106      1194 mV
MPC-XM_0V9-vt273m       911 mV
I2C Slave Revision      110
FPC 8 status:
State                    Online
Temperature Intake        32 degrees C / 89 degrees F
Temperature Exhaust A     44 degrees C / 111 degrees F
Temperature Exhaust B     37 degrees C / 98 degrees F
Temperature LU 0 TCAM TSen 41 degrees C / 105 degrees F
Temperature LU 0 TCAM Chip 49 degrees C / 120 degrees F
Temperature LU 0 TSen      41 degrees C / 105 degrees F
Temperature LU 0 Chip      52 degrees C / 125 degrees F
Temperature MQ 0 TSen      41 degrees C / 105 degrees F
Temperature MQ 0 Chip      47 degrees C / 116 degrees F
Temperature LU 1 TCAM TSen 39 degrees C / 102 degrees F
Temperature LU 1 TCAM Chip 42 degrees C / 107 degrees F
Temperature LU 1 TSen      39 degrees C / 102 degrees F
Temperature LU 1 Chip      46 degrees C / 114 degrees F
Temperature MQ 1 TSen      39 degrees C / 102 degrees F
Temperature MQ 1 Chip      45 degrees C / 113 degrees F
Power
MPC-BIAS3V3-z12105      3296 mV
MPC-VDD3V3-z12006      3298 mV
MPC-VDD2V5-z12006      2505 mV
MPC-TCAM_1V0-z12004      997 mV
MPC-AVDD1V0-z12006      1007 mV
MPC-VDD1V8-z12006      1803 mV
MPC-PCIE_1V0-z12006      1004 mV
MPC-LU0_1V0-z12004      1000 mV
MPC-MQ0_1V0-z12004      999 mV
MPC-VDD_1V5-z12004      1498 mV
MPC-PMB_1V1-z12006      1102 mV
MPC-9VA-BMR453          9009 mV
MPC-9VB-BMR453          8960 mV
MPC-PMB_1V2-z12105      1202 mV
MPC-LU1_1V0-z12004      1005 mV
MPC-MQ1_1V0-z12004      1000 mV
I2C Slave Revision      70
FPC 9 status:
State                    Online
Temperature Intake        34 degrees C / 93 degrees F
Temperature Exhaust A     41 degrees C / 105 degrees F

```

Temperature Exhaust B	54 degrees C / 129 degrees F
Temperature LU 0 TSen	51 degrees C / 123 degrees F
Temperature LU 0 Chip	52 degrees C / 125 degrees F
Temperature LU 1 TSen	51 degrees C / 123 degrees F
Temperature LU 1 Chip	55 degrees C / 131 degrees F
Temperature LU 2 TSen	51 degrees C / 123 degrees F
Temperature LU 2 Chip	47 degrees C / 116 degrees F
Temperature LU 3 TSen	51 degrees C / 123 degrees F
Temperature LU 3 Chip	47 degrees C / 116 degrees F
Temperature MQ 0 TSen	40 degrees C / 104 degrees F
Temperature MQ 0 Chip	42 degrees C / 107 degrees F
Temperature MQ 1 TSen	40 degrees C / 104 degrees F
Temperature MQ 1 Chip	44 degrees C / 111 degrees F
Temperature MQ 2 TSen	40 degrees C / 104 degrees F
Temperature MQ 2 Chip	38 degrees C / 100 degrees F
Temperature MQ 3 TSen	40 degrees C / 104 degrees F
Temperature MQ 3 Chip	40 degrees C / 104 degrees F
Power	
AS-BIAS3V3-z12105	3302 mV
AS-VDD1V8-z12006	1808 mV
AS-VDD2V5-z12006	2513 mV
AS-AVDD1V0-z12004	997 mV
AS-PCIE_1V0-z12004	999 mV
AS-VDD3V3-z12004	3294 mV
AS-VDD_1V5A-z12004	1503 mV
AS-VDD_1V5B-z12004	1502 mV
AS-LU0_1V0-z12004	996 mV
AS-LU1_1V0-z12004	999 mV
AS-MQ0_1V0-z12004	997 mV
AS-MQ1_1V0-z12004	999 mV
AS-LU2_1V0-z12004	997 mV
AS-LU3_1V0-z12004	998 mV
AS-MQ2_1V0-z12004	1000 mV
AS-MQ3_1V0-z12004	1000 mV
AS-PMB_1V1-z12006	1102 mV
I2C Slave Revision	68

show chassis environment fpc (MX240 Router)

```

user@host> show chassis environment fpc
FPC 1 status:
State                               Online
Temperature Intake                  34 degrees C / 93 degrees F
Temperature Exhaust A               39 degrees C / 102 degrees F
Temperature Exhaust B               53 degrees C / 127 degrees F
Temperature I3 0 TSensor            51 degrees C / 123 degrees F
Temperature I3 0 Chip               54 degrees C / 129 degrees F
Temperature I3 1 TSensor            50 degrees C / 122 degrees F
Temperature I3 1 Chip               53 degrees C / 127 degrees F
Temperature I3 2 TSensor            48 degrees C / 118 degrees F
Temperature I3 2 Chip               51 degrees C / 123 degrees F
Temperature I3 3 TSensor            45 degrees C / 113 degrees F
Temperature I3 3 Chip               48 degrees C / 118 degrees F
Temperature IA 0 TSensor            45 degrees C / 113 degrees F
Temperature IA 0 Chip               45 degrees C / 113 degrees F
Temperature IA 1 TSensor            45 degrees C / 113 degrees F
Temperature IA 1 Chip               49 degrees C / 120 degrees F
Power
  1.5 V                             1492 mV
  2.5 V                             2507 mV
  3.3 V                             3306 mV

```

```

1.8 V PFE 0          1801 mV
1.8 V PFE 1          1804 mV
1.8 V PFE 2          1798 mV
1.8 V PFE 3          1798 mV
1.2 V PFE 0          1169 mV
1.2 V PFE 1          1189 mV
1.2 V PFE 2          1182 mV
1.2 V PFE 3          1176 mV
I2C Slave Revision   42
FPC 2 status:
State                Online
Temperature Intake    33 degrees C / 91 degrees F
Temperature Exhaust A 41 degrees C / 105 degrees F
Temperature Exhaust B 53 degrees C / 127 degrees F
Temperature I3 0 TSensor 53 degrees C / 127 degrees F
Temperature I3 0 Chip  58 degrees C / 136 degrees F
Temperature I3 1 TSensor 52 degrees C / 125 degrees F
Temperature I3 1 Chip  56 degrees C / 132 degrees F
Temperature I3 2 TSensor 50 degrees C / 122 degrees F
Temperature I3 2 Chip  52 degrees C / 125 degrees F
Temperature I3 3 TSensor 46 degrees C / 114 degrees F
Temperature I3 3 Chip  49 degrees C / 120 degrees F
Temperature IA 0 TSensor 51 degrees C / 123 degrees F
Temperature IA 0 Chip  49 degrees C / 120 degrees F
Temperature IA 1 TSensor 48 degrees C / 118 degrees F
Temperature IA 1 Chip  53 degrees C / 127 degrees F
Power
1.5 V                1492 mV
2.5 V                2445 mV
3.3 V                3293 mV
1.8 V PFE 0          1827 mV
1.8 V PFE 1          1775 mV
1.8 V PFE 2          1788 mV
1.8 V PFE 3          1798 mV
1.2 V PFE 0          1250 mV
1.2 V PFE 1          1234 mV
1.2 V PFE 2          1231 mV
1.2 V PFE 3          1192 mV
I2C Slave Revision   42

```

show chassis environment fpc (MX480 Router)

```

user@host> show chassis environment fpc
FPC 1 status:
State                Online
Temperature Intake    36 degrees C / 96 degrees F
Temperature Exhaust A 41 degrees C / 105 degrees F
Temperature Exhaust B 55 degrees C / 131 degrees F
Temperature I3 0 TSensor 55 degrees C / 131 degrees F
Temperature I3 0 Chip  57 degrees C / 134 degrees F
Temperature I3 1 TSensor 53 degrees C / 127 degrees F
Temperature I3 1 Chip  53 degrees C / 127 degrees F
Temperature I3 2 TSensor 52 degrees C / 125 degrees F
Temperature I3 2 Chip  49 degrees C / 120 degrees F
Temperature I3 3 TSensor 47 degrees C / 116 degrees F
Temperature I3 3 Chip  47 degrees C / 116 degrees F
Temperature IA 0 TSensor 54 degrees C / 129 degrees F
Temperature IA 0 Chip  58 degrees C / 136 degrees F
Temperature IA 1 TSensor 48 degrees C / 118 degrees F
Temperature IA 1 Chip  53 degrees C / 127 degrees F
Power

```

1.5 V	1479 mV
2.5 V	2542 mV
3.3 V	3319 mV
1.8 V PFE 0	1811 mV
1.8 V PFE 1	1804 mV
1.8 V PFE 2	1804 mV
1.8 V PFE 3	1814 mV
1.2 V PFE 0	1192 mV
1.2 V PFE 1	1202 mV
1.2 V PFE 2	1205 mV
1.2 V PFE 3	1189 mV
I2C Slave Revision	40

show chassis environment fpc (MX960 Router)

```
user@host> show chassis environment fpc
```

```
FPC 5 status:
```

State	Online
Temperature Intake	27 degrees C / 80 degrees F
Temperature Exhaust A	34 degrees C / 93 degrees F
Temperature Exhaust B	40 degrees C / 104 degrees F
Temperature I3 0 TSensor	39 degrees C / 102 degrees F
Temperature I3 0 Chip	41 degrees C / 105 degrees F
Temperature I3 1 TSensor	38 degrees C / 100 degrees F
Temperature I3 1 Chip	37 degrees C / 98 degrees F
Temperature I3 2 TSensor	37 degrees C / 98 degrees F
Temperature I3 2 Chip	34 degrees C / 93 degrees F
Temperature I3 3 TSensor	32 degrees C / 89 degrees F
Temperature I3 3 Chip	33 degrees C / 91 degrees F
Temperature IA 0 TSensor	39 degrees C / 102 degrees F
Temperature IA 0 Chip	44 degrees C / 111 degrees F
Temperature IA 1 TSensor	36 degrees C / 96 degrees F
Temperature IA 1 Chip	44 degrees C / 111 degrees F
Power	
1.5 V	1479 mV
2.5 V	2523 mV
3.3 V	3254 mV
1.8 V PFE 0	1798 mV
1.8 V PFE 1	1798 mV
1.8 V PFE 2	1807 mV
1.8 V PFE 3	1791 mV
1.2 V PFE 0	1173 mV
1.2 V PFE 1	1179 mV
1.2 V PFE 2	1179 mV
1.2 V PFE 3	1185 mV
I2C Slave Revision	6

```
FPC 6 status:
```

State	Online
Temperature Intake	25 degrees C / 77 degrees F
Temperature Exhaust A	38 degrees C / 100 degrees F
Temperature Exhaust B	38 degrees C / 100 degrees F
Temperature I3 0 TSensor	40 degrees C / 104 degrees F
Temperature I3 0 Chip	40 degrees C / 104 degrees F
Temperature I3 1 TSensor	40 degrees C / 104 degrees F
Temperature I3 1 Chip	38 degrees C / 100 degrees F
Temperature I3 2 TSensor	37 degrees C / 98 degrees F
Temperature I3 2 Chip	32 degrees C / 89 degrees F
Temperature I3 3 TSensor	34 degrees C / 93 degrees F
Temperature I3 3 Chip	33 degrees C / 91 degrees F
Temperature IA 0 TSensor	45 degrees C / 113 degrees F
Temperature IA 0 Chip	47 degrees C / 116 degrees F

```

Temperature IA 1 TSensor 37 degrees C / 98 degrees F
Temperature IA 1 Chip    42 degrees C / 107 degrees F
Power
  1.5 V                  1485 mV
  2.5 V                  2510 mV
  3.3 V                  3332 mV
  1.8 V PFE 0            1801 mV
  1.8 V PFE 1            1814 mV
  1.8 V PFE 2            1804 mV
  1.8 V PFE 3            1820 mV
  1.2 V PFE 0            1192 mV
  1.2 V PFE 1            1189 mV
  1.2 V PFE 2            1202 mV
  1.2 V PFE 3            1156 mV
I2C Slave Revision      40

```

show chassis environment fpc (MX480 Router with 100-Gigabit Ethernet CFP)

```

user@host> show chassis environment fpc
FPC 0 status:
  State                               Online
  Temperature Intake                  32 degrees C / 89 degrees F
  Temperature Exhaust A               39 degrees C / 102 degrees F
  Temperature Exhaust B               37 degrees C / 98 degrees F
  Temperature QX 0 TSen               44 degrees C / 111 degrees F
  Temperature QX 0 Chip               48 degrees C / 118 degrees F
  Temperature LU 0 TCAM TSen          44 degrees C / 111 degrees F
  Temperature LU 0 TCAM Chip          47 degrees C / 116 degrees F
  Temperature LU 0 TSen               44 degrees C / 111 degrees F
  Temperature LU 0 Chip               48 degrees C / 118 degrees F
  Temperature MQ 0 TSen               44 degrees C / 111 degrees F
  Temperature MQ 0 Chip               47 degrees C / 116 degrees F
  Power
    MPC-BIAS3V3-z12105                3297 mV
    MPC-VDD3V3-z12105                 3306 mV
    MPC-VDD2V5-z12105                 2498 mV
    MPC-TCAM_1V0-z12004                999 mV
    MPC-AVDD1V0-z12006                 999 mV
    MPC-VDD1V8-z12006                 1796 mV
    MPC-PCIE_1V0-z12006                1002 mV
    MPC-LU0_1V0-z12004                 997 mV
    MPC-MQ0_1V0-z12004                 995 mV
    MPC-VDD_1V5-z12004                 1496 mV
    MPC-PMB_1V1-z12006                 1094 mV
    MPC-9VA-BMR453                     9054 mV
    MPC-9VB-BMR453                     9037 mV
    MPC-PMB_1V2-z12106                 1191 mV
    MPC-QXM0_1V0-z12006                1000 mV
  I2C Slave Revision                  66
FPC 1 status:
  State                               Online
  Temperature Intake                  35 degrees C / 95 degrees F
  Temperature Exhaust A               50 degrees C / 122 degrees F
  Temperature Exhaust B               56 degrees C / 132 degrees F
  Temperature LU 0 TSen               46 degrees C / 114 degrees F
  Temperature LU 0 Chip               59 degrees C / 138 degrees F
  Temperature LU 1 TSen               46 degrees C / 114 degrees F
  Temperature LU 1 Chip               45 degrees C / 113 degrees F
  Temperature LU 2 TSen               46 degrees C / 114 degrees F
  Temperature LU 2 Chip               60 degrees C / 140 degrees F
  Temperature LU 3 TSen               46 degrees C / 114 degrees F

```

```

Temperature LU 3 Chip      71 degrees C / 159 degrees F
Temperature XM 0 TSen      46 degrees C / 114 degrees F
Temperature XM 0 Chip      -18 degrees C / 0 degrees F
Temperature XF 0 TSen      46 degrees C / 114 degrees F
Temperature XF 0 Chip      76 degrees C / 168 degrees F
Power
MPC-BIAS3V3-z12105        3292 mV
MPC-VDD3V3-z16100         3303 mV
MPC-VDD2V5-z16100         2501 mV
MPC-VDD1V8-z12004         1801 mV
MPC-AVDD1V0-z12006         996 mV
MPC-VDD1V2-z16100         1199 mV
MPC-VDD1V5A-z12004        1493 mV
MPC-VDD1V5B-z12004        1498 mV
MPC-XF_0V9-z12006         996 mV
MPC-PCIE_1V0-z16100       1000 mV
MPC-LU0_1V0-z12004         994 mV
MPC-LU1_1V0-z12004         994 mV
MPC-LU2_1V0-z12004         992 mV
MPC-LU3_1V0-z12004         993 mV
MPC-12VA-BMR453           12003 mV
MPC-12VB-BMR453           12043 mV
MPC-PMB_1V1-z12006        1091 mV
MPC-PMB_1V2-z12106        1196 mV
MPC-XM_0V9-vt273m         899 mV
I2C Slave Revision        106

```

show chassis environment fpc (MX240, MX480, MX960 with Application Services Modular Line Card)

```

user@host>show chassis environment fpc 1
FPC 1 status:
State                Online
Temperature Intake    36 degrees C / 96 degrees F
Temperature Exhaust A 39 degrees C / 102 degrees F
Temperature LU TSen    52 degrees C / 125 degrees F
Temperature LU Chip    54 degrees C / 129 degrees F
Temperature XM TSen    52 degrees C / 125 degrees F
Temperature XM Chip    60 degrees C / 140 degrees F
Temperature PCIE TSen  52 degrees C / 125 degrees F
Temperature PCIE Chip  69 degrees C / 156 degrees F
Power
MPC-BIAS3V3-z12106    3302 mV
MPC-VDD3V3-z16100     3325 mV
MPC-AVDD1V0-z16100    1007 mV
MPC-PCIE_1V0-z16100    904 mV
MPC-LU0_1V0-z12004     996 mV
MPC-VDD_1V5-z12004    1498 mV
MPC-12VA-BMR453       11733 mV
MPC-12VB-BMR453       11728 mV
MPC-XM_0V9-vt273m     900 mV
I2C Slave Revision    81

```

show chassis environment fpc (T320, T640, and T1600 Routers)

```

user@host> show chassis environment fpc
FPC 0 status:
State                Online
Temperature Top       42 degrees C / 107 degrees F
Temperature Bottom    36 degrees C / 96 degrees F
Temperature MMB1       39 degrees C / 102 degrees F
Power:

```



```

1.8 V          1959 mV
2.5 V          2495 mV
3.3 V          3344 mV
5.0 V          5047 mV
1.8 V bias     1787 mV
3.3 V bias     3291 mV
5.0 V bias     4998 mV
8.0 V bias     7343 mV
BUS Revision   40
FPC 1 status:
State          Online
Temperature Top 42 degrees C / 107 degrees F
Temperature Bottom 39 degrees C / 102 degrees F
Temperature MMB1 40 degrees C / 104 degrees F
Power:
1.8 V          1956 mV
2.5 V          2498 mV
3.3 V          3340 mV
5.0 V          5023 mV
1.8 V bias     1782 mV
3.3 V bias     3277 mV
5.0 V bias     4989 mV
8.0 V bias     7289 mV
BUS Revision   40
FPC 2 status:
State          Online
Temperature Top 43 degrees C / 109 degrees F
Temperature Bottom 39 degrees C / 102 degrees F
Temperature MMB1 41 degrees C / 105 degrees F
Power:
1.8 V          1963 mV
2.5 V          2503 mV
3.3 V          3340 mV
5.0 V          5042 mV
1.8 V bias     1797 mV
3.3 V bias     3311 mV
5.0 V bias     5013 mV
8.0 V bias     7221 mV
BUS Revision   40

```

show chassis environment fpc (T4000 Router)

```

user@host> show chassis environment fpc
FPC 0 status:
State          Online
Fan Intake     34 degrees C / 93 degrees F
Fan Exhaust    48 degrees C / 118 degrees F
PMB            47 degrees C / 116 degrees F
LMB0           50 degrees C / 122 degrees F
LMB1           41 degrees C / 105 degrees F
LMB2           35 degrees C / 95 degrees F
PFE1 LU2       46 degrees C / 114 degrees F
PFE1 LU0       41 degrees C / 105 degrees F
PFE0 LU0       57 degrees C / 134 degrees F
XF1            47 degrees C / 116 degrees F
XF0            52 degrees C / 125 degrees F
XM1            41 degrees C / 105 degrees F
XM0            50 degrees C / 122 degrees F
PFE0 LU1       56 degrees C / 132 degrees F
PFE0 LU2       45 degrees C / 113 degrees F
PFE1 LU1       37 degrees C / 98 degrees F

```

Power 1	
1.0 V	991 mV
1.2 V bias	1195 mV
1.8 V	1788 mV
2.5 V	2483 mV
3.3 V	3289 mV
3.3 V bias	3299 mV
12.0 V A	10608 mV
12.0 V B	10637 mV
Power 2	
0.9 V	881 mV
0.9 V PFE0	916 mV
0.9 V PFE1	903 mV
1.0 V PFE0	1012 mV
1.0 V PFE1	1002 mV
1.1 V	1095 mV
1.5 V_0	1494 mV
1.5 V_1	1479 mV
Power 3	
1.0 V PFE0	1000 mV
1.0 V PFE1	1002 mV
1.0 V PFE0 *	995 mV
1.0 V PFE1 *	995 mV
1.8 V PFE 0	1788 mV
1.8 V PFE 1	1789 mV
2.5 V	2482 mV
12.0 V	11614 mV
Power 4	
1.0 V PFE0 LU0	1003 mV
1.0 V PFE1 LU0	1003 mV
1.0 V PFE1 LU2	1004 mV
1.0 V PFE0 LU0 *	995 mV
1.0 V PFE1 LU0 *	998 mV
1.0 V PFE1 LU2 *	996 mV
12.0 V	11643 mV
12.0 V C	11711 mV
Power (Base/PMB/MMB)	
LMB0 VDD2V5	2488 mV
LMB0 VDD1V8	1788 mV
LMB0 VDD1V5	1496 mV
LMB0 PFE0 LU0 AVDD1V0	1002 mV
LMB0 PFE0 LU0 VDD1V0	1000 mV
LMB0 VDD12V0	10752 mV
LMB1 VDD2V5	2472 mV
LMB1 VDD1V8	1792 mV
LMB1 VDD1V5	1480 mV
LMB1 PFE0 LU2 AVDD1V0	994 mV
LMB1 PFE0 LU2 VDD1V0	1002 mV
LMB1 VDD12V0	10800 mV
LMB2 VDD2V5	2472 mV
LMB2 VDD1V8	1792 mV
LMB2 VDD1V5	1486 mV
LMB2 PFE1 LU1 AVDD1V0	996 mV
LMB2 PFE1 LU1 VDD1V0	998 mV
LMB2 VDD12V0	10704 mV
PMB 1.05v	1049 mV
PMB 1.5v	1500 mV
PMB 2.5v	2500 mV
PMB 3.3v	3299 mV
Bus Revision	113
FPC 3 status:	

State	Online
Fan Intake	37 degrees C / 98 degrees F
Fan Exhaust	51 degrees C / 123 degrees F
PMB	43 degrees C / 109 degrees F
LMB0	57 degrees C / 134 degrees F
LMB1	54 degrees C / 129 degrees F
LMB2	38 degrees C / 100 degrees F
PFE1 LU2	63 degrees C / 145 degrees F
PFE1 LU0	45 degrees C / 113 degrees F
PFE0 LU0	69 degrees C / 156 degrees F
XF1	62 degrees C / 143 degrees F
XF0	63 degrees C / 145 degrees F
XM1	43 degrees C / 109 degrees F
XM0	67 degrees C / 152 degrees F
PFE0 LU1	63 degrees C / 145 degrees F
PFE0 LU2	66 degrees C / 150 degrees F
PFE1 LU1	41 degrees C / 105 degrees F
Power 1	
1.0 V	1002 mV
1.2 V bias	1201 mV
1.8 V	1785 mV
2.5 V	2485 mV
3.3 V	3288 mV
3.3 V bias	3285 mV
12.0 V A	10412 mV
12.0 V B	10515 mV
Power 2	
0.9 V	882 mV
0.9 V PFE0	920 mV
0.9 V PFE1	905 mV
1.0 V PFE0	1015 mV
1.0 V PFE1	1001 mV
1.1 V	1094 mV
1.5 V_0	1495 mV
1.5 V_1	1478 mV
Power 3	
0.92 V PFE1	998 mV
1.0 V PFE0	997 mV
1.0 V PFE0 *	992 mV
1.0 V PFE1 *	991 mV
1.8 V PFE 0	1780 mV
1.8 V PFE 1	1797 mV
2.5 V	2492 mV
12.0 V	11604 mV
Power 4	
1.0 V PFE0 LU0	1003 mV
1.0 V PFE1 LU0	1004 mV
1.0 V PFE1 LU2	1003 mV
1.0 V PFE0 LU0 *	1000 mV
1.0 V PFE1 LU0 *	1001 mV
1.0 V PFE1 LU2 *	1003 mV
12.0 V	11653 mV
12.0 V C	11672 mV
Power (Base/PMB/MMB)	
LMB0 VDD2V5	2512 mV
LMB0 VDD1V8	1790 mV
LMB0 VDD1V5	1500 mV
LMB0 PFE0 LU0 AVDD1V0	1004 mV
LMB0 PFE0 LU0 VDD1V0	1002 mV
LMB0 VDD12V0	10608 mV
LMB1 VDD2V5	2472 mV

LMB1 VDD1V8	1788 mV
LMB1 VDD1V5	1480 mV
LMB1 PFE0 LU2 AVDD1V0	1000 mV
LMB1 PFE0 LU2 VDD1V0	1004 mV
LMB1 VDD12V0	10672 mV
LMB2 VDD2V5	2488 mV
LMB2 VDD1V8	1798 mV
LMB2 VDD1V5	1494 mV
LMB2 PFE1 LU1 AVDD1V0	1000 mV
LMB2 PFE1 LU1 VDD1V0	1004 mV
LMB2 VDD12V0	10528 mV
PMB 1.05v	1050 mV
PMB 1.5v	1500 mV
PMB 2.5v	2499 mV
PMB 3.3v	3299 mV
Bus Revision	113
FPC 5 status:	
State	Online
Temperature Top	39 degrees C / 102 degrees F
Temperature Bottom	38 degrees C / 100 degrees F
Power	
1.8 V	1804 mV
1.8 V bias	1802 mV
3.3 V	3294 mV
3.3 V bias	3277 mV
5.0 V bias	5008 mV
5.0 V TOP	5067 mV
8.0 V bias	6642 mV
Power (Base/PMB/MMB)	
1.2 V	1202 mV
1.5 V	1504 mV
5.0 V BOT	5079 mV
12.0 V TOP Base	11848 mV
12.0 V BOT Base	11780 mV
1.1 V PMB	1111 mV
1.2 V PMB	1189 mV
1.5 V PMB	1494 mV
1.8 V PMB	1819 mV
2.5 V PMB	2503 mV
3.3 V PMB	3294 mV
5.0 V PMB	5035 mV
12.0 V PMB	11788 mV
0.75 MMB TOP	766 mV
1.5 V MMB TOP	1484 mV
1.8 V MMB TOP	1772 mV
2.5 V MMB TOP	2485 mV
1.2 V MMB TOP	1137 mV
5.0 V MMB TOP	4946 mV
12.0 V MMB TOP	11772 mV
3.3 V MMB TOP	3289 mV
0.75 MMB BOT	759 mV
1.5 V MMB BOT	1482 mV
1.8 V MMB BOT	1792 mV
2.5 V MMB BOT	2490 mV
1.2 V MMB BOT	1145 mV
5.0 V MMB BOT	4922 mV
12.0 V MMB BOT	11625 mV
3.3 V MMB BOT	3282 mV
APS 00	2495 mV
APS 01	3308 mV
APS 02	3301 mV

5.0 V PIC 0	4967 mV
APS 10	2512 mV
APS 11	3316 mV
APS 12	3304 mV
5.0 V PIC 1	5081 mV
Bus Revision	49
FPC 6 status:	
State	Online
Fan Intake	34 degrees C / 93 degrees F
Fan Exhaust	49 degrees C / 120 degrees F
PMB	40 degrees C / 104 degrees F
LMB0	60 degrees C / 140 degrees F
LMB1	58 degrees C / 136 degrees F
LMB2	40 degrees C / 104 degrees F
PFE1 LU2	69 degrees C / 156 degrees F
PFE1 LU0	45 degrees C / 113 degrees F
PFE0 LU0	71 degrees C / 159 degrees F
XF1	58 degrees C / 136 degrees F
XF0	65 degrees C / 149 degrees F
XM1	40 degrees C / 104 degrees F
XM0	66 degrees C / 150 degrees F
PFE0 LU1	69 degrees C / 156 degrees F
PFE0 LU2	68 degrees C / 154 degrees F
PFE1 LU1	42 degrees C / 107 degrees F
Power 1	
1.0 V	998 mV
1.2 V bias	1191 mV
1.8 V	1781 mV
2.5 V	2487 mV
3.3 V	3302 mV
3.3 V bias	3300 mV
12.0 V A	10388 mV
12.0 V B	10388 mV
Power 2	
0.9 V	902 mV
0.9 V PFE0	921 mV
0.9 V PFE1	907 mV
1.0 V PFE0	996 mV
1.0 V PFE1	974 mV
1.1 V	1095 mV
1.5 V_0	1495 mV
1.5 V_1	1478 mV
Power 3	
1.0 V PFE0	997 mV
1.0 V PFE1	998 mV
1.0 V PFE0 *	993 mV
1.0 V PFE1 *	991 mV
1.8 V PFE 0	1796 mV
1.8 V PFE 1	1789 mV
2.5 V	2465 mV
12.0 V	11609 mV
Power 4	
1.0 V PFE0 LU0	1003 mV
1.0 V PFE1 LU0	1006 mV
1.0 V PFE1 LU2	1002 mV
1.0 V PFE0 LU0 *	1000 mV
1.0 V PFE1 LU0 *	998 mV
1.0 V PFE1 LU2 *	998 mV
12.0 V	11638 mV
12.0 V C	11702 mV
Power (Base/PMB/MMB)	

LMB0	VDD2V5	2484 mV
LMB0	VDD1V8	1780 mV
LMB0	VDD1V5	1496 mV
LMB0	PFE0 LU0 AVDD1V0	998 mV
LMB0	PFE0 LU0 VDD1V0	1004 mV
LMB0	VDD12V0	10528 mV
LMB1	VDD2V5	2472 mV
LMB1	VDD1V8	1776 mV
LMB1	VDD1V5	1474 mV
LMB1	PFE0 LU2 AVDD1V0	994 mV
LMB1	PFE0 LU2 VDD1V0	1004 mV
LMB1	VDD12V0	10544 mV
LMB2	VDD2V5	2476 mV
LMB2	VDD1V8	1790 mV
LMB2	VDD1V5	1492 mV
LMB2	PFE1 LU1 AVDD1V0	996 mV
LMB2	PFE1 LU1 VDD1V0	1010 mV
LMB2	VDD12V0	10528 mV
PMB	1.05v	1050 mV
PMB	1.5v	1499 mV
PMB	2.5v	2500 mV
PMB	3.3v	3300 mV
Bus Revision		80

show chassis environment fpc lcc (TX Matrix Router)

```
user@host> show chassis environment fpc lcc 0
lcc0-re0:
```

FPC 1 status:

State	Online
Temperature Top	30 degrees C / 86 degrees F
Temperature Bottom	25 degrees C / 77 degrees F
Temperature MMB0	Absent
Temperature MMB1	27 degrees C / 80 degrees F
Power:	
1.8 V	1813 mV
2.5 V	2504 mV
3.3 V	3338 mV
5.0 V	5037 mV
1.8 V bias	1797 mV
3.3 V bias	3301 mV
5.0 V bias	5013 mV
8.0 V bias	7345 mV
BUS Revision	40

FPC 2 status:

State	Online
Temperature Top	37 degrees C / 98 degrees F
Temperature Bottom	26 degrees C / 78 degrees F
Temperature MMB0	32 degrees C / 89 degrees F
Temperature MMB1	27 degrees C / 80 degrees F
Power:	
1.8 V	1791 mV
2.5 V	2517 mV
3.3 V	3308 mV
5.0 V	5052 mV
1.8 V bias	1797 mV
3.3 V bias	3289 mV
5.0 V bias	4991 mV

8.0 V bias	7477 mV
BUS Revision	40

show chassis environment fpc lcc (TX Matrix Plus Router)

```
user@host> show chassis environment fpc lcc 0
lcc0-re0:
```

FPC 1 status:

State	Online
Temperature Top	46 degrees C / 114 degrees F
Temperature Bottom	47 degrees C / 116 degrees F
Power	
1.8 V	1788 mV
1.8 V bias	1787 mV
3.3 V	3321 mV
3.3 V bias	3306 mV
5.0 V bias	5018 mV
5.0 V TOP	5037 mV
8.0 V bias	7223 mV
Power (Base/PMB/MMB)	
1.2 V	1205 mV
1.5 V	1503 mV
5.0 V BOT	5084 mV
12.0 V TOP Base	11775 mV
12.0 V BOT Base	11794 mV
1.1 V PMB	1108 mV
1.2 V PMB	1196 mV
1.5 V PMB	1499 mV
1.8 V PMB	1811 mV
2.5 V PMB	2515 mV
3.3 V PMB	3318 mV
5.0 V PMB	5030 mV
12.0 V PMB	11832 mV
0.75 MMB TOP	752 mV
1.5 V MMB TOP	1489 mV
1.8 V MMB TOP	1782 mV
2.5 V MMB TOP	2498 mV
1.2 V MMB TOP	1155 mV
5.0 V MMB TOP	4902 mV
12.0 V MMB TOP	11721 mV
3.3 V MMB TOP	3316 mV
0.75 MMB BOT	754 mV
1.5 V MMB BOT	1482 mV
1.8 V MMB BOT	1758 mV
2.5 V MMB BOT	2488 mV
1.2 V MMB BOT	1157 mV
5.0 V MMB BOT	4962 mV
12.0 V MMB BOT	11691 mV
3.3 V MMB BOT	3308 mV
APS 00	1484 mV
APS 01	2503 mV
APS 02	3313 mV
5.0 V PIC 0	5025 mV
APS 10	1501 mV
APS 11	2466 mV
APS 12	3311 mV
5.0 V PIC 1	5081 mV
Bus Revision	49

show chassis environment fpc (QFX Series)

```
user@switch> show chassis environment fpc 0
FPC 0 status:
  State                Online
  Temperature          42 degrees C / 107 degrees F
```

show chassis environment fpc interconnect-device (QFabric Systems)

```
user@switch> show chassis environment fpc interconnect-device interconnect1 0
FC 0 FPC 0 status:
  State                Online
  Left Intake Temperature 24 degrees C / 75 degrees F
  Right Intake Temperature 24 degrees C / 75 degrees F
  Left Exhaust Temperature 27 degrees C / 80 degrees F
  Right Exhaust Temperature 27 degrees C / 80 degrees F
  Power
    BIAS 3V3            3330 mV
    VDD 3V3             3300 mV
    VDD 2V5             2502 mV
    VDD 1V5             1496 mV
    VDD 1V2             1194 mV
    VDD 1V0             1000 mV
    SW0 VDD 1V0         1020 mV
    SW0 CVDD 1V025      1032 mV
    SW1 VDD 1V0         1022 mV
    SW1 CVDD 1V025      1030 mV
    VDD 12V0 DIV3_33    3414 mV
```

show chassis environment fpc 0 (PTX5000 Packet Transport Switch)

```
user@switch> show chassis environment fpc 0
FPC 0 status:
  State                Online
  PMB Temperature      35 degrees C / 95 degrees F
  Intake Temperature   33 degrees C / 91 degrees F
  Exhaust A Temperature 51 degrees C / 123 degrees F
  Exhaust B Temperature 43 degrees C / 109 degrees F
  TL0 Temperature      48 degrees C / 118 degrees F
  TQ0 Temperature      53 degrees C / 127 degrees F
  TL1 Temperature      56 degrees C / 132 degrees F
  TQ1 Temperature      58 degrees C / 136 degrees F
  TL2 Temperature      55 degrees C / 131 degrees F
  TQ2 Temperature      57 degrees C / 134 degrees F
  TL3 Temperature      59 degrees C / 138 degrees F
  TQ3 Temperature      59 degrees C / 138 degrees F
  Power
    PMB 1.05v          1049 mV
    PMB 1.5v           1500 mV
    PMB 2.5v           2500 mV
    PMB 3.3v           3299 mV
    PFE0 1.5v          1500 mV
    PFE0 1.0v          999 mV
    TQ0 0.9v           900 mV
    TL0 0.9v           900 mV
    PFE1 1.5v          1499 mV
    PFE1 1.0v          999 mV
    TQ1 0.9v           899 mV
    TL1 0.9v           900 mV
    PFE2 1.5v          1500 mV
    PFE2 1.0v          1000 mV
```


TQ2	0.9v	900 mV
TL2	0.9v	900 mV
PFE3	1.5v	1499 mV
PFE3	1.0v	1000 mV
TQ3	0.9v	900 mV
TL3	0.9v	900 mV
Bias	3.3v	3327 mV
FPC	3.3v	3300 mV
FPC	2.5v	2500 mV
SAM	0.9v	900 mV
A	12.0v	2014 mV
B	12.0v	2030 mV

show chassis environment FPC 1 (MX Routers with Media Services Blade [MSB])

```

user@switch> show chassis environment fpc 1
FPC 1 status:
State                               Online
Temperature Intake                  36 degrees C / 96 degrees F
Temperature Exhaust A               39 degrees C / 102 degrees F
Temperature LU TSen                 52 degrees C / 125 degrees F
Temperature LU Chip                 54 degrees C / 129 degrees F
Temperature XM TSen                 52 degrees C / 125 degrees F
Temperature XM Chip                 60 degrees C / 140 degrees F
Temperature PCIe TSen               52 degrees C / 125 degrees F
Temperature PCIe Chip               69 degrees C / 156 degrees F
Power
MPC-BIAS3V3-z12106                 3302 mV
MPC-VDD3V3-z16100                  3325 mV
MPC-AVDD1V0-z16100                 1007 mV
MPC-PCIE_1V0-z16100                 904 mV
MPC-LU0_1V0-z12004                 996 mV
MPC-VDD_1V5-z12004                 1498 mV
MPC-12VA-BMR453                    11733 mV
MPC-12VB-BMR453                    11728 mV
MPC-XM_0V9-vt273m                  900 mV
I2C Slave Revision                  81

```

show chassis environment pem

Syntax	show chassis environment pem <slot>
Syntax (ACX4000 Router)	show chassis environment pem
Syntax (TX Matrix Routers)	show chassis environment pem <lcc number scc> <slot>
Syntax (TX Matrix Plus Routers)	show chassis environment pem <lcc number sfc number> <slot>
Syntax (MX Series Router)	show chassis environment pem <slot> <all-members> <local> <member member-id>
Syntax (QFX Series)	show chassis environment pem <slot (interconnect-device name slot) (node-device name)>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS 11.3 for the QFX Series. Command introduced in Junos OS 12.3R2 for EX Series switches.
Description	Display Power Entry Module (PEM) environmental status information.



NOTE: The new high-capacity (4100W) enhanced DC PEM on MX960 routers includes a new design that can condition the input voltage. This results in the output voltage differing from the input voltage. The earlier generation of DC PEMs coupled the input power directly to the output, thereby making it safe to assume that the output voltage was equal to the input voltage.

- Options** **none**—Display environmental information about both PEMs. For the TX Matrix router, display environmental information about the PEMs, the TX Matrix router, and its attached T640 routers. For the TX Matrix Plus router, display environmental information about the PEMs, the TX Matrix Plus router, and its attached T1600 routers.
- all-members**—(MX Series routers only) (Optional) Display environmental information about the PEMs in all the member routers of the Virtual Chassis configuration.
- interconnect-device name**—(QFabric systems only) (Optional) Display chassis environmental information about the PEMs in the Interconnect device.

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display environmental information about the PEM in a specified T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display environmental information about the PEM in a specified T1600 router (or line-card chassis) that is connected to a TX Matrix Plus router. Replace ***number*** with a value from **0** through **3**.

local—(MX Series routers only) (Optional) Display environmental information about the PEM in the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display environmental information about the PEM in the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value of 0 or 1.

node-device *name*—(QFabric systems only) (Optional) Display chassis environmental information about the PEMs in the Node device.

scc—(TX Matrix routers only) (Optional) Display environmental information about the PEM in the TX Matrix router (or switch-card chassis).

sfc—(TX Matrix Plus routers only) (Optional) Display environmental information about the PEM in the TX Matrix Plus router (or switch-fabric chassis).

slot —(Optional) Display environmental information about an individual PEM. Replace ***slot*** with **0** or **1**.

Required Privilege Level

view

Related Documentation

- [show chassis hardware on page 644](#)

List of Sample Output

[show chassis environment pem \(M40e Router\) on page 622](#)
[show chassis environment pem \(M120 Router\) on page 622](#)
[show chassis environment pem \(M160 Router\) on page 623](#)
[show chassis environment pem \(M320 Router\) on page 623](#)
[show chassis environment pem \(MX240 Router\) on page 623](#)
[show chassis environment pem \(MX480 Router\) on page 624](#)
[show chassis environment pem \(MX960 Router\) on page 624](#)
[show chassis environment pem \(T320 Router\) on page 624](#)
[show chassis environment pem \(T640 Router\) on page 624](#)
[show chassis environment pem \(T4000 Router\) on page 624](#)
[show chassis environment pem \(T640/T1600/T4000 Routers With Six-Input DC Power Supply\) on page 625](#)
[show chassis environment pem lcc \(TX Matrix Routing Matrix\) on page 625](#)
[show chassis environment pem scc \(TX Matrix Routing Matrix\) on page 626](#)
[show chassis environment pem sfc \(TX Matrix Plus Routing Matrix\) on page 626](#)
[show chassis environment pem lcc \(TX Matrix Plus Routing Matrix\) on page 626](#)
[show chassis environment pem node-device \(QFabric System\) on page 627](#)
[show chassis environment pem \(QFX Series\) on page 627](#)
[show chassis environment pem interconnect-device \(QFabric System\) on page 627](#)

Output Fields Table 56 on page 622 lists the output fields for the **show chassis environment pem** command. Output fields are listed in the approximate order in which they appear.

Table 56: show chassis environment pem Output Fields

Field Name	Field Description
PEM slot status	Number of the PEM slot.
State	Status of the PEM.
Temperature	Temperature of the air flowing past the PEM.
AC Input	Status of the AC input for the specified component
AC Output	Status of the AC output for the specified component.
DC input	Status of the DC input for the specified component.
DC output	Status of the DC output for the specified component.
Load	(Not available on M40e or M160 routers) Information about the load on supply, in percentage of rated current being used.
Voltage	(M120, M160, M320, T640, T1600, TX Matrix, and TX Matrix Plus routers only) Information about voltage supplied to the PEM.
Current	(T640, T1600, TX Matrix, and TX Matrix Plus routers only) Information about the PEM current.
Power	(T640, T1600, TX Matrix, and TX Matrix Plus routers only) Information about the PEM power.
SCG/CB/SIB	(T640, T1600, TX Matrix, and TX Matrix Plus routers only) SONET Clock Generator/Control Board/Switch Interface Board.
FAN	(T640, T1600, and T4000 routers with six-input DC power supply only) Information about the DC output to the fan.

Sample Output

show chassis environment pem (M40e Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State           Online
  Temperature      OK
  AC input         OK
  DC output        OK
```

show chassis environment pem (M120 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State           Online
  Temperature      OK
```

```

DC Input:           OK
DC Output:          OK
Load                Less than 20 percent
Voltage:
  48.0 V input      52864 mV
  48.0 V fan supply 41655 mV
  3.3 V             3399 mV
PEM 1 status:
State               Online
Temperature         OK
DC Input:           OK
DC Output:          OK
Load                Less than 20 percent
Voltage:
  48.0 V input      54537 mV
  48.0 V fan supply 42910 mV
  3.3 V             3506 mV

```

show chassis environment pem (M160 Router)

```

user@host> show chassis environment pem
PEM 0 status:
State               Online
Temperature         OK
DC input            OK
DC output           OK
Load                Less than 20 percent
Voltage:
  48.0 V input      54833 mV
  48.0 V fan supply 50549 mV
  8.0 V bias        8239 mV
  5.0 V bias        5006 mV

```

show chassis environment pem (M320 Router)

```

user@host> show chassis environment pem
PEM 2 status:
State               Online
Temperature         OK
DC input            OK
Load                Less than 40 percent
  48.0 V input      51853 mV
  48.0 V fan supply 48877 mV
  8.0 V bias        8449 mV
  5.0 V bias        4998 mV
PEM 3 status:
State               Online
Temperature         OK
DC input            OK
Load                Less than 40 percent
  48.0 V input      51717 mV
  48.0 V fan supply 49076 mV
  8.0 V bias        8442 mV
  5.0 V bias        4998 mV

```

show chassis environment pem (MX240 Router)

```

user@host> show chassis environment pem
PEM 0 status:
State               Online
Temperature         OK
DC Output:          OK

```

```
PEM 1 status:
State          Online
Temperature    OK
DC Output:     OK
```

show chassis environment pem (MX480 Router)

```
user@host> show chassis environment pem
PEM 0 status:
State          Online
Temperature    OK
DC Input:      OK
DC Output:     OK
Voltage:
PEM 1 status:
State          Online
Temperature    OK
DC Input:      OK
DC Output:     OK
Voltage:
```

show chassis environment pem (MX960 Router)

```
user@host> show chassis environment pem
PEM 2 status:
State          Present
PEM 3 status:
State          Online
Temperature    OK
DC Output:     OK
```

show chassis environment pem (T320 Router)

```
user@host> show chassis environment pem
PEM 0 status:
State          Online
Temperature    OK
DC input:      OK
```

show chassis environment pem (T640 Router)

```
user@host> show chassis environment pem
PEM 0 status:
State          Online
Temperature    22 degrees C / 71 degrees F
AC input: OK
DC output:
Voltage      Current      Power      Load
FPC 0        56875      606        34        4
FPC 1        57016      525        29        3
FPC 2         0         0          0         0
FPC 3         0         0          0         0
FPC 4         0         0          0         0
FPC 5         0         0          0         0
FPC 6        57158     1581       90        12
FPC 7         0         0          0         0
SCG/CB/SIB   56750     1125       63         5
```

show chassis environment pem (T4000 Router)

```
user@host> show chassis environment pem
PEM 0 status:
State          Online
```

```

Temperature          33 degrees C / 91 degrees F
DC Input:            OK
                    Voltage(V) Current(A) Power(W) Load(%)
INPUT 0              54.625    9.812    535    22
INPUT 1              54.625   10.250    559    23
INPUT 2              55.125    0.125     6     0
INPUT 3              54.500   10.062    548    22
INPUT 4              54.750    9.375    513    21
INPUT 5              54.750   10.187    557    23
DC Output            Voltage(V) Current(A) Power(W) Load(%)
FPC 0                55.750   10.125    564    37
FPC 1                51.625    0.000     0     0
FPC 2                52.000    0.000     0     0
FPC 3                55.062   10.437    574    38
FPC 4                52.125    0.000     0     0
FPC 5                55.000    9.375    515    34
FPC 6                55.187    9.687    534    35
FPC 7                51.437    0.000     0     0
SCG/CB/SIB           55.375   15.750    872    35
FAN                  54.562   14.750    804    42

```

show chassis environment pem (T640/T1600/T4000 Routers With Six-Input DC Power Supply)

```

user@host> show chassis environment pem
PEM 1 status:
State          Online
Temperature     36 degrees C / 96 degrees F
DC Input:      OK
              Voltage(V) Current(A) Power(W) Load(%)
INPUT 0        0.000    0.000     0     0
INPUT 1       54.875    3.812    209    27
INPUT 2       55.375    3.937    218    29
INPUT 3       54.625    3.750    204    27
INPUT 4       55.125    3.375    186    24
INPUT 5       55.125    3.375    186    24
DC Output      Voltage(V) Current(A) Power(W) Load(%)
FPC 0         52.312    0.000     0     0
FPC 1         52.687    0.000     0     0
FPC 2         52.812    0.000     0     0
FPC 3         55.812    7.062    394    52
FPC 4         52.625    0.000     0     0
FPC 5         52.625    0.000     0     0
FPC 6         52.750    0.000     0     0
FPC 7         52.750    0.000     0     0
SCG/CB/SIB    55.937   11.937    667    55
FAN           55.812    4.937    275    36

```

show chassis environment pem lcc (TX Matrix Routing Matrix)

```

user@host> show chassis environment pem 0 lcc 0
lcc0-re0:

```

```

-----
PEM 0 status:
State          Present
Temperature     27 degrees C / 80 degrees F
DC input:      Check
DC output:     Voltage Current Power Load
FPC 0          0       0       0       0
FPC 1          0       0       0       0
FPC 2          0       0       0       0
FPC 3          0       0       0       0

```

FPC 4	0	0	0	0
FPC 5	0	0	0	0
FPC 6	0	0	0	0
FPC 7	0	0	0	0
SCG/CB/SIB	0	0	0	0

show chassis environment pem scc (TX Matrix Routing Matrix)

```
user@host> show chassis environment pem scc
scc-re0:
```

```
-----
PEM 1 status:
State                Online
Temperature          24 degrees C / 75 degrees F
DC input:            OK
DC output:           Voltage Current      Power    Load
SIB 0                0         0         0        0
SIB 1                0         0         0        0
SIB 2                0         0         0        0
SIB 3                56550        0         0        0
SIB 4                55958       6912       386       51
```

show chassis environment pem sfc (TX Matrix Plus Routing Matrix)

```
user@host> show chassis environment pem sfc 0
sfc0-re0:
```

```
-----
PEM 0 status:
State                Online
Temperature          35 degrees C / 95 degrees F
DC Input:            OK
DC Output           Voltage Current      Power    Load
Channel 0           53820    14140       761      59
Channel 1           53550    12720       681      53
Channel 2           53840    12930       696      54
Channel 3           53690    14990       804      63
Channel 4           53620    15070       808      63
Channel 5           53900    14820       798      62
Channel 6           54120     5020       271      21
```

show chassis environment pem lcc (TX Matrix Plus Routing Matrix)

```
user@host> show chassis environment lcc 0
```

```
lcc0-re1:
-----
PEM 0 status:
State                Online
Temperature          38 degrees C / 100 degrees F
DC Input:            OK
DC Output           Voltage Current      Power    Load
FPC 0                0         0         0        0
FPC 1                0         0         0        0
FPC 2                0         0         0        0
FPC 3                0         0         0        0
FPC 4                56408     7575       427      56
FPC 5                0         0         0        0
FPC 6                56266     7956       447      59
FPC 7                56283     6100       343      45
SCG/CB/SIB           55916     8950       500      41
PEM 1 status:
State                Present
```



```

Temperature                35 degrees C / 95 degrees F
DC Input:                   Check
DC Output                   Voltage    Current    Power    Load
FPC 0                       0         0         0         0
FPC 1                       0         0         0         0
FPC 2                       0         0         0         0
FPC 3                       0         0         0         0
FPC 4                       0         0         0         0
FPC 5                       0         0         0         0
FPC 6                       0         0         0         0
FPC 7                       0         0         0         0
SCG/CB/SIB                  0         0         0         0

```

show chassis environment pem node-device (QFabric System)

```

user@switch> show chassis environment pem node-device node1
FPC 0 PEM 0 status:
State                Check
Airflow              Front to Back
Temperature          OK
AC Input:            OK
DC Output            Voltage(V) Current(A) Power(W) Load(%)
                   12         10         120      18
FPC 0 PEM 1 status:
State                Online
Airflow              Back to Front
Temperature          OK
AC Input:            OK
DC Output            Voltage(V) Current(A) Power(W) Load(%)
                   11         10         110      17

```

show chassis environment pem (QFX Series)

```

user@switch> show chassis environment pem
FPC 0 PEM 1 status:
State                Online
Airflow              Front to Back
Temperature          OK
AC Input:            OK
DC Output            Voltage(V) Current(A) Power(W) Load(%)
                   12         17         204      31

```

show chassis environment pem interconnect-device (QFabric System)

```

user@switch> show chassis environment pem interconnect-device IC11
IC1 PEM 1 status:
State                Online
Airflow              Front to Back
Temperature          OK
AC Input:            OK
DC Output            Voltage(V) Current(A) Power(W) Load(%)
                   12         18         216      33

```

show chassis environment routing-engine

Syntax	show chassis environment routing-engine <slot>
Syntax (TX Matrix Routers)	show chassis environment routing-engine <lcc number scc> <slot>
Syntax (TX Matrix Plus Routers)	show chassis environment routing-engine <lcc number sfc number> <slot>
Syntax (MX2010 3D Universal Edge Routers)	show chassis environment routing-engine <slot>
Syntax (MX Series Routers)	show chassis environment routing-engine <slot> <all-members> <local> <member member-id>
Syntax (MX2020 3D Universal Edge Routers)	show chassis environment routing-engine <slot>
Syntax (QFX Series)	show chassis environment routing-engine interconnect-device <i>name</i>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Switches. Command introduced in Junos OS Release 12.1 for the T4000 Core Routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.
Description	Display Routing Engine environmental status information.
Options	none —Display environmental information about all Routing Engines. For a TX Matrix router, display environmental information about all Routing Engines on the TX Matrix router and its attached T640 routers. For a TX Matrix Plus router, display environmental information about all Routing Engines on the TX Matrix Plus router and its attached T1600 routers. all-members —(MX Series routers only) (Optional) Display environmental information about the Routing Engines in all member routers in the Virtual Chassis configuration.

interconnect-device *name*—(QFabric systems only) (Optional) Display environmental information about the Routing Engines for the Interconnect device.

lcc *number*—(TX Matrix and TX Matrix routers only) (Optional) On a TX Matrix router, display environmental information about the Routing Engine in a specified T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display environmental information about the Routing Engine in a specified T1600 router (or line-card chassis) that is connected to the TX Matrix Plus router. Replace ***number*** with a value from 0 through 3.

local—(MX Series routers only) (Optional) Display environmental information about the Routing Engines in the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display environmental information about the Routing Engines in the specified member in the Virtual Chassis configuration. Replace ***member-id*** with the value of 0 or 1.

scc—(TX Matrix router only) (Optional) Display environmental information about the Routing Engine in the TX Matrix router (or switch-card chassis).

sfc—(TX Matrix Plus router only) (Optional) Display environmental information about the Routing Engine in the TX Matrix Plus router (or switch-fabric chassis).

slot—(Optional) Display environmental information about an individual Routing Engine. On M10i, M20, M40e, M120, M160, M320, MX Series, MX2020 routers, and T Series routers, replace **slot** with 0 or 1. On M5, M7i, M10, and M40 routers and on the J Series router, replace **slot** with 0. On EX3200 and EX4200 standalone switches, replace **slot** with 0. On EX4200 switches in a Virtual Chassis configuration and on EX8208 and EX8216 switches, replace **slot** with 0 or 1. On the QFX3500 switch, there is only one Routing Engine, so you do not need to specify the slot number. On PTX Series Packet Transport Switches, replace **slot** with 0 or 1.

Required Privilege Level view

Related Documentation

- [request chassis routing-engine master](#)
- [show chassis routing-engine on page 767](#)

List of Sample Output

- [show chassis environment routing-engine \(Nonredundant\) on page 630](#)
- [show chassis environment routing-engine \(Redundant\) on page 630](#)
- [show chassis environment routing-engine \(MX2010 Router\) on page 630](#)
- [show chassis environment routing-engine \(MX2020 Router\) on page 630](#)
- [show chassis environment routing-engine \(TX Matrix Plus Router\) on page 631](#)
- [show chassis environment routing-engine \(T4000 Core Router\) on page 631](#)
- [show chassis environment routing-engine \(QFX Series\) on page 631](#)
- [show chassis environment routing-engine interconnect-device \(QFabric System\) on page 631](#)
- [show chassis environment routing-engine \(PTX5000 Packet Transport Switch\) on page 631](#)

Output Fields Table 57 on page 630 lists the output fields for the **show chassis environment routing-engine** command. Output fields are listed in the approximate order in which they appear.

Table 57: show chassis environment routing-engine Output Fields

Field Name	Field Description
Routing engine slot status	Number of the Routing Engine slot: 0 or 1.
State	Status of the Routing Engine: <ul style="list-style-type: none"> • Online Master—Routing Engine is online, operating as Master. • Online Standby—Routing Engine is online, operating as Standby. • Offline—Routing Engine is offline.
Temperature	Temperature of the air flowing past the Routing Engine.
CPU Temperature	(PTX Series and T4000 Core Routers only) Temperature of the air flowing past the Routing Engine CPU.

Sample Output

show chassis environment routing-engine (Nonredundant)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State                Online Master
  Temperature          27 degrees C / 80 degrees
```

show chassis environment routing-engine (Redundant)

```
user@host> show chassis environment routing-engine
Route Engine 0 status:
  State:                Online Master
  Temperature:          26 degrees C / 78 degrees F
Route Engine 1 status:
  State:                Online Standby
  Temperature:          26 degrees C / 78 degrees F
```

show chassis environment routing-engine (MX2010 Router)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State                Online Master
  Temperature          37 degrees C / 98 degrees F
  CPU Temperature      37 degrees C / 98 degrees F
Routing Engine 1 status:
  State                Online Standby
  Temperature          35 degrees C / 95 degrees F
  CPU Temperature      34 degrees C / 93 degrees F
```

show chassis environment routing-engine (MX2020 Router)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State                Online Master
  Temperature          35 degrees C / 95 degrees F
  CPU Temperature      34 degrees C / 93 degrees F
```

```

Routing Engine 1 status:
  State           Online Standby
  Temperature     44 degrees C / 111 degrees F
  CPU Temperature 43 degrees C / 109 degrees F

```

show chassis environment routing-engine (TX Matrix Plus Router)

```

user@host> show chassis environment routing-engine
sfc0-re0:

```

```

-----
Routing Engine 0 status:
  State           Online Master
  Temperature     26 degrees C / 78 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature     28 degrees C / 82 degrees F

```

```

lcc0-re0:

```

```

-----
Routing Engine 0 status:
  State           Online Master
  Temperature     30 degrees C / 86 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature     29 degrees C / 84 degrees F

```

show chassis environment routing-engine (T4000 Core Router)

```

user@host> show chassis environment routing-engine

```

```

Routing Engine 0 status:
  State           Online Master
  Temperature     33 degrees C / 91 degrees F
  CPU Temperature 50 degrees C / 122 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature     33 degrees C / 91 degrees F
  CPU Temperature 46 degrees C / 114 degrees F

```

show chassis environment routing-engine (QFX Series)

```

user@switch> show chassis environment routing-engine

```

```

Routing Engine 0 status:
  State           Online Master
  Temperature     42 degrees C / 107 degrees F

```

show chassis environment routing-engine interconnect-device (QFabric System)

```

user@switch> show chassis environment routing-engine interconnect-device interconnect1
routing-engine interconnect-device interconnect1

```

```

Routing Engine 0 status:
  State           Online Standby
  Temperature     52 degrees C / 125 degrees F
Routing Engine 1 status:
  State           Online Master
  Temperature     57 degrees C / 134 degrees F

```

show chassis environment routing-engine (PTX5000 Packet Transport Switch)

```

user@switch> show chassis environment routing-engine

```

```

Routing Engine 0 status:
  State           Online Master
  Temperature     55 degrees C / 131 degrees F

```

CPU Temperature	66 degrees C / 150 degrees F
Routing Engine 1 status:	
State	Online Standby
Temperature	52 degrees C / 125 degrees F
CPU Temperature	64 degrees C / 147 degrees F

show chassis fan

Syntax	show chassis fan
Syntax (ACX4000 Series Router)	show chassis fan
Syntax (MX Series Router)	show chassis fan <all-members> <local> <member <i>member-id</i> >
Syntax (T Series Routers)	show chassis fan
Syntax (MX2010 3D Universal Edge Router)	show chassis fan
Syntax (MX2020 3D Universal Edge Router)	show chassis fan
Syntax (QFabric Systems)	show chassis fan <interconnect-device <i>name</i> >
Syntax (TX Matrix Router)	show chassis fan <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show chassis fan <lcc <i>number</i> sfc <i>number</i> >
Release Information	<p>Command introduced in Junos OS Release 10.0 on MX Series 3D Universal Edge Routers, M120 routers, and M320 routers, T320 routers, T640 routers, T1600 routers, TX Matrix Routers, and TX Matrix Plus Routers.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 11.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.3 for PTX5000 Packet Transport Switches.</p> <p>Command introduced in Junos OS Release 12.1 for T4000 routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for ACX Series Routers.</p>
Description	(T Series routers, TX Matrix routers, TX Matrix Plus router, M120 routers, M320 routers, MX2010 routers, MX2020 routers, MX Series 3D Universal Edge Routers, QFX3008-I Interconnect devices, EX Series switches, and PTX Series Packet Transport Switches only) Show information about the fan tray and fans.
Options	<p>all-members—(MX Series routers only) (Optional) Display information about the fan tray and fans for all members of the Virtual Chassis configuration.</p> <p>local—(MX Series routers only) (Optional) Display information about the fan tray and fans for the local Virtual Chassis member.</p>

member *member-id*—(MX Series routers only) (Optional) Display information about the fan tray and fans for the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

interconnect-device *name*—(QFX3000-G QFabric systems only) (Optional) Display information about the fan tray and fans for the specified QFX3008-I Interconnect device.

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display information about the fan tray and fans for the specified T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display information about the fan tray and fans for the specified T1600 router (or line-card chassis) that is connected to a TX Matrix Plus router. Replace number with a value from 0 through 3.

scc—(TX Matrix router only) (Optional) Display information about the fan tray and fans for the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus router only) (Optional) Display information about the fan tray and fans for the TX Matrix Plus router (or switch-fabric chassis). Replace number with 0.

Required Privilege Level view

List of Sample Output

- [show chassis fan on page 635](#)
- [show chassis fan \(QFabric Systems\) on page 636](#)
- [show chassis fan \(EX Series Switches\) on page 637](#)
- [show chassis fan \(T320 Router\) on page 637](#)
- [show chassis fan \(T640 Router\) on page 638](#)
- [show chassis fan \(T1600 Router\) on page 638](#)
- [show chassis fan \(T4000 Core Router\) on page 638](#)
- [show chassis fan \(TX Matrix Router\) on page 639](#)
- [show chassis fan \(TX Matrix Plus Router\) on page 640](#)
- [show chassis fan \(PTX5000 Packet Transport Switch\) on page 641](#)
- [show chassis fan \(MX2010 Router\) on page 641](#)
- [show chassis fan \(MX2020 Router\) on page 642](#)
- [show chassis fan \(ACX4000 Router\) on page 642](#)

Output Fields [Table 58 on page 634](#) lists the output fields for the **show chassis fan** command. Output fields are listed in the approximate order in which they appear.

Table 58: show chassis fan Output Fields

Field Name	Field Description
Item	Fan item identifier.

Table 58: show chassis fan Output Fields (*continued*)

Field Name	Field Description
Status	<p>Status of the fan:</p> <ul style="list-style-type: none"> • OK-Fan is running properly and within the normal range. • Check-Fan is in Check state because of some fault or alarm condition.
RPM	(T Series routers, TX Matrix routers, TX Matrix Plus router, MX Series 3D Universal Edge Routers, QFX3108 Interconnect devices, and EX Series switches only) Fan speed in revolutions per minute (RPM).
% RPM	(MX2010 routers, MX2020 routers, and PTX Series Packet Transport Switches only) Percentage of the fan speed being used.
Measurement	<p>(T Series routers, TX Matrix routers, TX Matrix Plus router, MX Series 3D Universal Edge Routers, QFX3108 Interconnect devices, and EX Series switches only) Fan speed status based on different chassis cooling requirements:</p> <ul style="list-style-type: none"> • Spinning at high speed • Spinning at intermediate speed • Spinning at normal speed • Spinning at low speed (except EX Series switches) <p>(MX2010 routers, MX2020 routers, and PTX Series Packet Transport Switches only) Fan speed in revolutions per minute (RPM) for each fan in the fan tray.</p>

Sample Output

show chassis fan

```

user@host> show chassis fan
user@host> show chassis fan
  Item                Status  RPM    Measurement
  Top Tray Fan 1      OK      3790    Spinning at normal speed
  Top Tray Fan 2      OK      3769    Spinning at normal speed
  Top Tray Fan 3      OK      3769    Spinning at normal speed
  Top Tray Fan 4      OK      3790    Spinning at normal speed
  Top Tray Fan 5      OK      3790    Spinning at normal speed
  Top Tray Fan 6      OK      3769    Spinning at normal speed
  Top Tray Fan 7      OK      3790    Spinning at normal speed
  Top Tray Fan 8      OK      3769    Spinning at normal speed
  Top Tray Fan 9      OK      3769    Spinning at normal speed
  Top Tray Fan 10     OK      3790    Spinning at normal speed
  Top Tray Fan 11     OK      3790    Spinning at normal speed
  Top Tray Fan 12     OK      3769    Spinning at normal speed
  Bottom Tray Fan 1   OK      2880    Spinning at normal speed
  Bottom Tray Fan 2   OK      2912    Spinning at normal speed
  Bottom Tray Fan 3   OK      2928    Spinning at normal speed
  Bottom Tray Fan 4   OK      2896    Spinning at normal speed
  Bottom Tray Fan 5   OK      2896    Spinning at normal speed
  Bottom Tray Fan 6   OK      2928    Spinning at normal speed

```

show chassis fan (QFabric Systems)

user@host> show chassis fan interconnect-device *interconnect1*

Item	Status	RPM	Measurement
TFT 0 Fan 0	OK	2849	Spinning at normal speed
TFT 0 Fan 1	OK	2821	Spinning at normal speed
TFT 0 Fan 2	OK	2735	Spinning at normal speed
TFT 0 Fan 3	OK	2815	Spinning at normal speed
TFT 0 Fan 4	OK	2828	Spinning at normal speed
TFT 0 Fan 5	OK	2863	Spinning at normal speed
BFT 1 Fan 0	OK	2941	Spinning at normal speed
BFT 1 Fan 1	OK	3008	Spinning at normal speed
BFT 1 Fan 2	OK	3073	Spinning at normal speed
BFT 1 Fan 3	OK	2925	Spinning at normal speed
BFT 1 Fan 4	OK	2863	Spinning at normal speed
BFT 1 Fan 5	OK	2933	Spinning at normal speed
SFT 0 Fan 0 Rotor 0	OK	15472	Spinning at normal speed
SFT 0 Fan 0 Rotor 1	OK	14477	Spinning at normal speed
SFT 0 Fan 1 Rotor 0	OK	15561	Spinning at normal speed
SFT 0 Fan 1 Rotor 1	OK	14210	Spinning at normal speed
SFT 0 Fan 2 Rotor 0	OK	16167	Spinning at normal speed
SFT 0 Fan 2 Rotor 1	OK	14248	Spinning at normal speed
SFT 0 Fan 3 Rotor 0	OK	16463	Spinning at normal speed
SFT 0 Fan 3 Rotor 1	OK	14099	Spinning at normal speed
SFT 1 Fan 0 Rotor 0	OK	15083	Spinning at normal speed
SFT 1 Fan 0 Rotor 1	OK	13533	Spinning at normal speed
SFT 1 Fan 1 Rotor 0	OK	16071	Spinning at normal speed
SFT 1 Fan 1 Rotor 1	OK	14400	Spinning at normal speed
SFT 1 Fan 2 Rotor 0	OK	15517	Spinning at normal speed
SFT 1 Fan 2 Rotor 1	OK	14210	Spinning at normal speed
SFT 1 Fan 3 Rotor 0	OK	16413	Spinning at normal speed
SFT 1 Fan 3 Rotor 1	OK	14400	Spinning at normal speed
SFT 2 Fan 0 Rotor 0	OK	15297	Spinning at normal speed
SFT 2 Fan 0 Rotor 1	OK	14634	Spinning at normal speed
SFT 2 Fan 1 Rotor 0	OK	15561	Spinning at normal speed
SFT 2 Fan 1 Rotor 1	OK	14285	Spinning at normal speed
SFT 2 Fan 2 Rotor 0	OK	15835	Spinning at normal speed
SFT 2 Fan 2 Rotor 1	OK	14400	Spinning at normal speed
SFT 2 Fan 3 Rotor 0	OK	15789	Spinning at normal speed
SFT 2 Fan 3 Rotor 1	OK	14323	Spinning at normal speed
SFT 3 Fan 0 Rotor 0	OK	16314	Spinning at normal speed
SFT 3 Fan 0 Rotor 1	OK	14876	Spinning at normal speed
SFT 3 Fan 1 Rotor 0	OK	15835	Spinning at normal speed
SFT 3 Fan 1 Rotor 1	OK	14323	Spinning at normal speed
SFT 3 Fan 2 Rotor 0	OK	16265	Spinning at normal speed
SFT 3 Fan 2 Rotor 1	OK	14594	Spinning at normal speed
SFT 3 Fan 3 Rotor 0	OK	16071	Spinning at normal speed
SFT 3 Fan 3 Rotor 1	OK	14323	Spinning at normal speed
SFT 4 Fan 0 Rotor 0	OK	15652	Spinning at normal speed
SFT 4 Fan 0 Rotor 1	OK	14438	Spinning at normal speed
SFT 4 Fan 1 Rotor 0	OK	16167	Spinning at normal speed
SFT 4 Fan 1 Rotor 1	OK	14555	Spinning at normal speed
SFT 4 Fan 2 Rotor 0	OK	16023	Spinning at normal speed
SFT 4 Fan 2 Rotor 1	OK	14361	Spinning at normal speed
SFT 4 Fan 3 Rotor 0	OK	16216	Spinning at normal speed
SFT 4 Fan 3 Rotor 1	OK	14438	Spinning at normal speed
SFT 5 Fan 0 Rotor 0	OK	15297	Spinning at normal speed
SFT 5 Fan 0 Rotor 1	OK	14173	Spinning at normal speed
SFT 5 Fan 1 Rotor 0	OK	15472	Spinning at normal speed
SFT 5 Fan 1 Rotor 1	OK	13846	Spinning at normal speed

SFT 5 Fan 2 Rotor 0	OK	15340	Spinning at normal speed
SFT 5 Fan 2 Rotor 1	OK	13917	Spinning at normal speed
SFT 5 Fan 3 Rotor 0	OK	15835	Spinning at normal speed
SFT 5 Fan 3 Rotor 1	OK	13917	Spinning at normal speed
SFT 6 Fan 0 Rotor 0	OK	15743	Spinning at normal speed
SFT 6 Fan 0 Rotor 1	OK	14594	Spinning at normal speed
SFT 6 Fan 1 Rotor 0	OK	16167	Spinning at normal speed
SFT 6 Fan 1 Rotor 1	OK	14634	Spinning at normal speed
SFT 6 Fan 2 Rotor 0	OK	16167	Spinning at normal speed
SFT 6 Fan 2 Rotor 1	OK	14516	Spinning at normal speed
SFT 6 Fan 3 Rotor 0	OK	16666	Spinning at normal speed
SFT 6 Fan 3 Rotor 1	OK	14438	Spinning at normal speed
SFT 7 Fan 0 Rotor 0	OK	15517	Spinning at normal speed
SFT 7 Fan 0 Rotor 1	OK	14438	Spinning at normal speed
SFT 7 Fan 1 Rotor 0	OK	15517	Spinning at normal speed
SFT 7 Fan 1 Rotor 1	OK	14361	Spinning at normal speed
SFT 7 Fan 2 Rotor 0	OK	16167	Spinning at normal speed
SFT 7 Fan 2 Rotor 1	OK	14555	Spinning at normal speed
SFT 7 Fan 3 Rotor 0	OK	15697	Spinning at normal speed
SFT 7 Fan 3 Rotor 1	OK	14361	Spinning at normal speed

show chassis fan (EX Series Switches)

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
Fan 1	OK	3477	Spinning at normal speed
Fan 2	OK	3477	Spinning at normal speed
Fan 3	OK	3479	Spinning at normal speed
Fan 4	OK	3508	Spinning at normal speed
Fan 5	OK	3517	Spinning at normal speed
Fan 6	OK	3531	Spinning at normal speed
Fan 7	OK	3439	Spinning at normal speed
Fan 8	OK	3424	Spinning at normal speed
Fan 9	OK	3413	Spinning at normal speed
Fan 10	OK	3439	Spinning at normal speed
Fan 11	OK	3446	Spinning at normal speed
Fan 12	OK	3432	Spinning at normal speed

show chassis fan (T320 Router)

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	2850	Spinning at normal speed
Top Left Middle fan	OK	2820	Spinning at normal speed
Top Left Rear fan	OK	2970	Spinning at normal speed
Top Right Front fan	OK	2790	Spinning at normal speed
Top Right Middle fan	OK	2640	Spinning at normal speed
Top Right Rear fan	OK	2790	Spinning at normal speed
Bottom Left Front fan	OK	2520	Spinning at normal speed
Bottom Left Middle fan	OK	2610	Spinning at normal speed
Bottom Left Rear fan	OK	2550	Spinning at normal speed
Bottom Right Front fan	OK	2610	Spinning at normal speed
Bottom Right Middle fan	OK	2880	Spinning at normal speed
Bottom Right Rear fan	OK	2790	Spinning at normal speed
Rear Tray Top fan	OK	2130	Spinning at normal speed
Rear Tray Second fan	OK	2190	Spinning at normal speed
Rear Tray Middle fan	OK	2250	Spinning at normal speed
Rear Tray Fourth fan	OK	2220	Spinning at normal speed
Rear Tray Bottom fan	OK	2280	Spinning at normal speed

show chassis fan (T640 Router)

user@host> show chassis fan

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3420	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3450	Spinning at normal speed
Bottom Left Front fan	OK	3390	Spinning at normal speed
Bottom Left Middle fan	OK	3420	Spinning at normal speed
Bottom Left Rear fan	OK	3390	Spinning at normal speed
Bottom Right Front fan	OK	3390	Spinning at normal speed
Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3390	Spinning at normal speed
Rear Tray Top fan	OK	5220	Spinning at normal speed
Rear Tray Second fan	OK	5220	Spinning at normal speed
Rear Tray Third fan	OK	5220	Spinning at normal speed
Rear Tray Fourth fan	OK	5220	Spinning at normal speed
Rear Tray Fifth fan	OK	5220	Spinning at normal speed
Rear Tray Sixth fan	OK	5220	Spinning at normal speed
Rear Tray Seventh fan	OK	5220	Spinning at normal speed
Rear Tray Bottom fan	OK	5220	Spinning at normal speed

show chassis fan (T1600 Router)

user@host> show chassis fan

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3420	Spinning at normal speed
Top Left Rear fan	OK	3450	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3390	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3420	Spinning at normal speed
Bottom Left Rear fan	OK	3390	Spinning at normal speed
Bottom Right Front fan	OK	3390	Spinning at normal speed
Bottom Right Middle fan	OK	3420	Spinning at normal speed
Bottom Right Rear fan	OK	3390	Spinning at normal speed
Rear Tray Top fan	OK	5190	Spinning at normal speed
Rear Tray Second fan	OK	5190	Spinning at normal speed
Rear Tray Third fan	OK	5190	Spinning at normal speed
Rear Tray Fourth fan	OK	5190	Spinning at normal speed
Rear Tray Fifth fan	OK	5190	Spinning at normal speed
Rear Tray Sixth fan	OK	5190	Spinning at normal speed
Rear Tray Seventh fan	OK	5190	Spinning at normal speed
Rear Tray Bottom fan	OK	5190	Spinning at normal speed

show chassis fan (T4000 Core Router)

user@host> show chassis fan

Item	Status	RPM	Measurement
Top Left Front fan	OK	5190	Spinning at high speed
Top Left Middle fan	OK	5220	Spinning at high speed
Top Left Rear fan	OK	5190	Spinning at high speed
Top Right Front fan	OK	5160	Spinning at high speed

Top Right Middle fan	OK	5190	Spinning at high speed
Top Right Rear fan	OK	5160	Spinning at high speed
Bottom Left Front fan	OK	6030	Spinning at high speed
Bottom Left Middle fan	OK	6090	Spinning at high speed
Bottom Left Rear fan	OK	6090	Spinning at high speed
Bottom Right Front fan	OK	6030	Spinning at high speed
Bottom Right Middle fan	OK	6060	Spinning at high speed
Bottom Right Rear fan	OK	6060	Spinning at high speed
Rear Tray Top fan	OK	10000	Spinning at high speed
Rear Tray Second fan	OK	10000	Spinning at high speed
Rear Tray Third fan	OK	10000	Spinning at high speed
Rear Tray Fourth fan	OK	10000	Spinning at high speed
Rear Tray Fifth fan	OK	10000	Spinning at high speed
Rear Tray Sixth fan	OK	10000	Spinning at high speed
Rear Tray Seventh fan	OK	10000	Spinning at high speed
Rear Tray Bottom fan	OK	10000	Spinning at high speed

show chassis fan (TX Matrix Router)

```
user@host> show chassis fan
scc-re0:
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3390	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3390	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3390	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3450	Spinning at normal speed
Bottom Left Rear fan	OK	3420	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3420	Spinning at normal speed
Bottom Right Rear fan	OK	3420	Spinning at normal speed
Rear Tray Top fan	OK	3420	Spinning at normal speed
Rear Tray Second fan	OK	5190	Spinning at normal speed
Rear Tray Third fan	OK	5190	Spinning at normal speed
Rear Tray Fourth fan	OK	5190	Spinning at normal speed
Rear Tray Fifth fan	OK	3420	Spinning at normal speed
Rear Tray Sixth fan	OK	3420	Spinning at normal speed
Rear Tray Seventh fan	OK	3420	Spinning at normal speed
Rear Tray Bottom fan	OK	3420	Spinning at normal speed

```
lcc2-re0:
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3420	Spinning at normal speed
Top Left Rear fan	OK	3450	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3450	Spinning at normal speed
Top Right Rear fan	OK	3360	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3480	Spinning at normal speed
Bottom Left Rear fan	OK	3420	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3420	Spinning at normal speed
Rear Tray Top fan	OK	3420	Spinning at normal speed
Rear Tray Second fan	OK	3420	Spinning at normal speed

Rear Tray Third fan	OK	3420	Spinning at normal speed
Rear Tray Fourth fan	OK	3420	Spinning at normal speed
Rear Tray Fifth fan	OK	3420	Spinning at normal speed
Rear Tray Sixth fan	OK	3420	Spinning at normal speed
Rear Tray Seventh fan	OK	3420	Spinning at normal speed
Rear Tray Bottom fan	OK	3420	Spinning at normal speed

show chassis fan (TX Matrix Plus Router)

```
user@host> show chassis fan
sfc0-re0:
```

Item	Status	RPM	Measurement
Fan Tray 0 Fan 1	OK	4350	Spinning at normal speed
Fan Tray 0 Fan 2	OK	4380	Spinning at normal speed
Fan Tray 0 Fan 3	OK	4410	Spinning at normal speed
Fan Tray 0 Fan 4	OK	4380	Spinning at normal speed
Fan Tray 0 Fan 5	OK	4350	Spinning at normal speed
Fan Tray 0 Fan 6	OK	4380	Spinning at normal speed
Fan Tray 1 Fan 1	OK	4410	Spinning at normal speed
Fan Tray 1 Fan 2	OK	4380	Spinning at normal speed
Fan Tray 1 Fan 3	OK	4410	Spinning at normal speed
Fan Tray 1 Fan 4	OK	4380	Spinning at normal speed
Fan Tray 1 Fan 5	OK	4410	Spinning at normal speed
Fan Tray 1 Fan 6	OK	4410	Spinning at normal speed
Fan Tray 2 Fan 1	OK	4380	Spinning at normal speed
Fan Tray 2 Fan 2	OK	4380	Spinning at normal speed
Fan Tray 2 Fan 3	OK	4380	Spinning at normal speed
Fan Tray 2 Fan 4	OK	4410	Spinning at normal speed
Fan Tray 2 Fan 5	OK	4380	Spinning at normal speed
Fan Tray 2 Fan 6	OK	4410	Spinning at normal speed
Fan Tray 2 Fan 7	OK	4410	Spinning at normal speed
Fan Tray 2 Fan 8	OK	4380	Spinning at normal speed
Fan Tray 2 Fan 9	OK	4380	Spinning at normal speed
Fan Tray 3 Fan 1	OK	4350	Spinning at normal speed
Fan Tray 3 Fan 2	OK	4380	Spinning at normal speed
Fan Tray 3 Fan 3	OK	4410	Spinning at normal speed
Fan Tray 3 Fan 4	OK	4440	Spinning at normal speed
Fan Tray 3 Fan 5	OK	4380	Spinning at normal speed
Fan Tray 3 Fan 6	OK	4410	Spinning at normal speed
Fan Tray 3 Fan 7	OK	4410	Spinning at normal speed
Fan Tray 3 Fan 8	OK	4380	Spinning at normal speed
Fan Tray 3 Fan 9	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 1	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 2	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 3	OK	4380	Spinning at normal speed
Fan Tray 4 Fan 4	OK	4380	Spinning at normal speed
Fan Tray 4 Fan 5	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 6	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 7	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 8	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 9	OK	4410	Spinning at normal speed
Fan Tray 5 Fan 1	OK	4350	Spinning at normal speed
Fan Tray 5 Fan 2	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 3	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 4	OK	4350	Spinning at normal speed
Fan Tray 5 Fan 5	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 6	OK	4410	Spinning at normal speed
Fan Tray 5 Fan 7	OK	4410	Spinning at normal speed
Fan Tray 5 Fan 8	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 9	OK	4410	Spinning at normal speed

```
1cc0-re0:
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3420	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3450	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3420	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3420	Spinning at normal speed
Bottom Left Rear fan	OK	3390	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3390	Spinning at normal speed
Rear Tray Top fan	OK	7050	Spinning at normal speed
Rear Tray Second fan	OK	7050	Spinning at normal speed
Rear Tray Third fan	OK	7050	Spinning at normal speed
Rear Tray Fourth fan	OK	7050	Spinning at normal speed
Rear Tray Fifth fan	OK	7050	Spinning at normal speed
Rear Tray Sixth fan	OK	7050	Spinning at normal speed
Rear Tray Seventh fan	OK	7050	Spinning at normal speed
Rear Tray Bottom fan	OK	7050	Spinning at normal speed

show chassis fan (PTX5000 Packet Transport Switch)

```
user@host> show chassis fan
user@host> show chassis fan
```

Item	Status	% RPM	Measurement
Fan Tray 0 Fan 1	OK	29%	2700 RPM
Fan Tray 0 Fan 2	OK	29%	2700 RPM
Fan Tray 0 Fan 3	OK	29%	2742 RPM
Fan Tray 0 Fan 4	OK	29%	2700 RPM
Fan Tray 0 Fan 5	OK	30%	2828 RPM
Fan Tray 0 Fan 6	OK	30%	2828 RPM
Fan Tray 0 Fan 7	OK	29%	2700 RPM
Fan Tray 0 Fan 8	OK	30%	2785 RPM
Fan Tray 0 Fan 9	OK	30%	2828 RPM
Fan Tray 0 Fan 10	OK	30%	2828 RPM
Fan Tray 0 Fan 11	OK	30%	2785 RPM
Fan Tray 0 Fan 12	OK	30%	2828 RPM
Fan Tray 0 Fan 13	OK	31%	2871 RPM
Fan Tray 0 Fan 14	OK	30%	2828 RPM
Fan Tray 1 Fan 1	OK	42%	3033 RPM
Fan Tray 1 Fan 2	OK	42%	3066 RPM
Fan Tray 1 Fan 3	OK	43%	3099 RPM
Fan Tray 1 Fan 4	OK	43%	3166 RPM
Fan Tray 1 Fan 5	OK	45%	3266 RPM
Fan Tray 1 Fan 6	OK	43%	3133 RPM
Fan Tray 2 Fan 1	OK	29%	2099 RPM
Fan Tray 2 Fan 2	OK	30%	2199 RPM
Fan Tray 2 Fan 3	OK	30%	2166 RPM
Fan Tray 2 Fan 4	OK	33%	2399 RPM
Fan Tray 2 Fan 5	OK	29%	2133 RPM
Fan Tray 2 Fan 6	OK	32%	2366 RPM

show chassis fan (MX2010 Router)

```
user@host > show chassis fan
```

Item	Status	% RPM	Measurement
Fan Tray 0 Fan 1	OK	37%	3360 RPM
Fan Tray 0 Fan 2	OK	38%	3480 RPM
Fan Tray 0 Fan 3	OK	37%	3360 RPM
Fan Tray 0 Fan 4	OK	37%	3360 RPM
Fan Tray 0 Fan 5	OK	38%	3480 RPM
Fan Tray 0 Fan 6	OK	37%	3360 RPM
Fan Tray 1 Fan 1	OK	38%	3480 RPM
Fan Tray 1 Fan 2	OK	40%	3600 RPM
Fan Tray 1 Fan 3	OK	38%	3480 RPM
Fan Tray 1 Fan 4	OK	38%	3480 RPM
Fan Tray 1 Fan 5	OK	38%	3480 RPM
Fan Tray 1 Fan 6	OK	38%	3480 RPM
Fan Tray 2 Fan 1	OK	38%	3480 RPM
Fan Tray 2 Fan 2	OK	41%	3720 RPM
Fan Tray 2 Fan 3	OK	38%	3480 RPM
Fan Tray 2 Fan 4	OK	38%	3480 RPM
Fan Tray 2 Fan 5	OK	38%	3480 RPM
Fan Tray 2 Fan 6	OK	38%	3480 RPM
Fan Tray 3 Fan 1	OK	38%	3480 RPM
Fan Tray 3 Fan 2	OK	40%	3600 RPM
Fan Tray 3 Fan 3	OK	40%	3600 RPM
Fan Tray 3 Fan 4	OK	40%	3600 RPM
Fan Tray 3 Fan 5	OK	40%	3600 RPM
Fan Tray 3 Fan 6	OK	38%	3480 RPM

show chassis fan (MX2020 Router)

```
user@host > show chassis fan
```

Item	Status	% RPM	Measurement
Fan Tray 0 Fan 1	OK	37%	3360 RPM
Fan Tray 0 Fan 2	OK	37%	3360 RPM
Fan Tray 0 Fan 3	OK	36%	3240 RPM
Fan Tray 0 Fan 4	OK	37%	3360 RPM
Fan Tray 0 Fan 5	OK	37%	3360 RPM
Fan Tray 0 Fan 6	OK	37%	3360 RPM
Fan Tray 1 Fan 1	OK	37%	3360 RPM
Fan Tray 1 Fan 2	OK	37%	3360 RPM
Fan Tray 1 Fan 3	OK	37%	3360 RPM
Fan Tray 1 Fan 4	OK	37%	3360 RPM
Fan Tray 1 Fan 5	OK	37%	3360 RPM
Fan Tray 1 Fan 6	OK	36%	3240 RPM
Fan Tray 2 Fan 1	OK	37%	3360 RPM
Fan Tray 2 Fan 2	OK	37%	3360 RPM
Fan Tray 2 Fan 3	OK	37%	3360 RPM
Fan Tray 2 Fan 4	OK	37%	3360 RPM
Fan Tray 2 Fan 5	OK	37%	3360 RPM
Fan Tray 2 Fan 6	OK	38%	3480 RPM
Fan Tray 3 Fan 1	OK	38%	3480 RPM
Fan Tray 3 Fan 2	OK	38%	3480 RPM
Fan Tray 3 Fan 3	OK	38%	3480 RPM
Fan Tray 3 Fan 4	OK	37%	3360 RPM
Fan Tray 3 Fan 5	OK	37%	3360 RPM
Fan Tray 3 Fan 6	OK	37%	3360 RPM

show chassis fan (ACX4000 Router)

```
user@host > show chassis fan
```


Item	Status	RPM	Measurement
Fan 1	OK	4140	Spinning at normal speed
Fan 2	OK	4200	Spinning at normal speed

show chassis hardware

Syntax	show chassis hardware <detail extensive> <clei-models> <models>
Syntax (EX Series)	show chassis hardware <clei-models> <detail extensive> <models>
Syntax (T4000 Router)	show chassis hardware <clei-models> <detail extensive> <models>
Syntax (TX Matrix Router)	show chassis hardware <clei-models> <detail extensive> <models> <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show chassis hardware <clei-models> <detail extensive> <models> <lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Routers and EX Series Switches)	show chassis hardware <detail extensive> <clei-models> <models> <all-members> <local> <member <i>member-id</i> >
Syntax (MX2010 and MX2020 3D Universal Edge Routers)	show chassis hardware <clei-models> <detail extensive> <models>
Syntax (QFX Series)	show chassis hardware <detail extensive> <clei-models> <interconnect-device <i>name</i> > <node-device <i>name</i> > <models>
Syntax (PTX Series Packet Transport Switches)	show chassis hardware <detail extensive> <clei-models> <models>

Syntax (ACX Series Universal Access Routers)	<pre>show chassis hardware <detail extensive> <clei-models> <models></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>models option introduced in Junos OS Release 8.2.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>
Description	<p>Display a list of all Flexible PIC Concentrators (FPCs) and PICs installed in the router or switch chassis, including the hardware version level and serial number.</p> <p>In the EX Series switch command output, FPC refers to the following:</p> <ul style="list-style-type: none"> On EX2200 switches, EX3200 switches, EX4200 standalone switches, and EX4500 switches—Refers to the switch; FPC number is always 0. On EX4200 switches in a Virtual Chassis configuration—Refers to the member of a Virtual Chassis; FPC number equals the member ID, from 0 through 9. On EX8208 and EX8216 switches—Refers to a line card; FPC number equals the slot number for the line card. <p>On a QFX3500 standalone switch, both the FPC and FPC number are always 0.</p> <p>On Type 5 FPC on T4000 routers, there are no top temperature sensor or bottom temperature sensor parameters. Instead, fan intake temperature sensor and fan exhaust temperature sensors parameters are displayed.</p>
Options	<p>none—Display information about hardware. For a TX Matrix router, display information about the TX Matrix router and its attached T640 routers. For a TX Matrix Plus router, display information about the TX Matrix Plus router and its attached T1600 routers.</p> <p>clei-models—(Optional) Display Common Language Equipment Identifier (CLEI) barcode and model number for orderable field-replaceable units (FRUs).</p> <p>detail—(Optional) Include RAM and disk information in output.</p> <p>extensive—(Optional) Display ID EEPROM information.</p> <p>all-members—(MX Series routers and EX Series switches only) (Optional) Display hardware-specific information for all the members of the Virtual Chassis configuration.</p> <p>interconnect-device name—(QFabric systems only) (Optional) Display hardware-specific information for the Interconnect device.</p>

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display hardware information for a specified T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display hardware information for a specified T1600 router (or line-card chassis) that is connected to the TX Matrix Plus router. Replace ***number*** with a value from **0** through **3**.

local—(MX Series routers and EX Series switches only) (Optional) Display hardware-specific information for the local Virtual Chassis members.

member *member-id*—(MX Series routers and EX Series switches only) (Optional) Display hardware-specific information for the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value of **0** or **1**.

models—(Optional) Display model numbers and part numbers for orderable FRUs and, for components that use ID EEPROM format v2, the CLEI code.

node-device *name*—(QFabric systems only) (Optional) Display hardware-specific information for the Node device.

scc—(TX Matrix router only) (Optional) Display hardware information for the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus router only) (Optional) Display hardware information for the TX Matrix Plus router (or switch-fabric chassis). Replace ***number*** with **0**.

Additional Information The **show chassis hardware detail** command now displays DIMM information for the following Routing Engines:

Table 59: Routing Engines Displaying DIMM Information

Routing Engines	Routers
RE-S-1800x2 and RE-S-1800x4	MX240, MX480, and MX960 routers
RE-A-1800x2	M120 and M320 routers

Required Privilege Level view

Related Documentation

- *show chassis power*

List of Sample Output

- [show chassis hardware \(EX8216 Switch\) on page 651](#)
- [show chassis hardware clei-models \(EX8216 Switch\) on page 652](#)
- [show chassis hardware clei-models \(T1600 Router\) on page 653](#)
- [show chassis hardware detail \(EX4200 Switch\) on page 654](#)
- [show chassis hardware models \(EX4500 Switch\) on page 654](#)
- [show chassis hardware \(J6350 Router\) on page 654](#)
- [show chassis hardware \(J6300 Router\) on page 654](#)
- [show chassis hardware \(M7i Router\) on page 655](#)
- [show chassis hardware \(M10 Router\) on page 655](#)
- [show chassis hardware models \(M10 Router\) on page 656](#)

[show chassis hardware \(M20 Router\) on page 656](#)
[show chassis hardware models \(M20 Router\) on page 657](#)
[show chassis hardware \(M40 Router\) on page 657](#)
[show chassis hardware \(M40e Router\) on page 658](#)
[show chassis hardware \(M120 Router\) on page 658](#)
[show chassis hardware detail \(M120 Router\) on page 659](#)
[show chassis hardware models \(M120 Router\) on page 660](#)
[show chassis hardware \(M160 Router\) on page 661](#)
[show chassis hardware models \(M160 Router\) on page 661](#)
[show chassis hardware detail \(M160 Router\) on page 662](#)
[show chassis hardware \(M320 Router\) on page 663](#)
[show chassis hardware models \(M320 Router\) on page 664](#)
[show chassis hardware \(MX5 Router\) on page 664](#)
[show chassis hardware \(MX10 Router\) on page 665](#)
[show chassis hardware \(MX40 Router\) on page 666](#)
[show chassis hardware \(Fixed MX80 Router\) on page 666](#)
[show chassis hardware \(Modular MX80 Router\) on page 667](#)
[show chassis hardware \(MX240 Router\) on page 667](#)
[show chassis hardware detail \(MX240 Router with RE Displaying DIMM information\) on page 668](#)
[show chassis hardware \(MX240 Router with Enhanced MX SCB\) on page 668](#)
[show chassis hardware \(MX480 Router\) on page 669](#)
[show chassis hardware \(MX480 Router with Enhanced MX SCB\) on page 670](#)
[show chassis hardware \(MX960 Router\) on page 670](#)
[show chassis hardware \(MX960 Router with Bidirectional Optics\) on page 670](#)
[show chassis hardware \(MX960 Router with Enhanced MX SCB\) on page 671](#)
[show chassis hardware models \(MX960 Router with Enhanced MX SCB\) on page 673](#)
[show chassis hardware detail \(MX960 Router\) on page 674](#)
[show chassis hardware \(MX2010 Router\) on page 674](#)
[show chassis hardware detail \(MX2010 Router\) on page 676](#)
[show chassis hardware extensive \(MX2010 Router\) on page 681](#)
[show chassis hardware models \(MX2010 Router\) on page 687](#)
[show chassis hardware clei-models \(MX2010 Routers\) on page 687](#)
[show chassis hardware \(MX2020 Router\) on page 688](#)
[show chassis hardware detail \(MX2020 Router\) on page 697](#)
[show chassis hardware models \(MX2020 Router\) on page 705](#)
[show chassis hardware clei-models \(MX2020 Router\) on page 707](#)
[show chassis hardware \(MX Series Routers with ATM MIC\) on page 708](#)
[show chassis hardware \(MX240, MX480, MX960 Routers with Application Services Modular Line Card\) on page 709](#)
[show chassis hardware extensive \(MX240, MX480, MX960 Routers with Application Services Modular Line Card\) on page 709](#)
[show chassis hardware \(MX480 Router with MPC4E\) on page 710](#)
[show chassis hardware \(MX2020 Router with MPC4E\) on page 711](#)
[show chassis hardware \(T320 Router\) on page 713](#)
[show chassis hardware \(T640 Router\) on page 714](#)
[show chassis hardware models \(T640 Router\) on page 714](#)
[show chassis hardware extensive \(T640 Router\) on page 715](#)
[show chassis hardware \(T4000 Router\) on page 716](#)

[show chassis hardware \(T4000 Router with 16 GB line card chassis \(LCC\) Routing Engine\) on page 718](#)
[show chassis hardware clei-models \(T4000 Router\) on page 718](#)
[show chassis hardware detail \(T4000 Router\) on page 719](#)
[show chassis hardware models \(T4000 Router\) on page 721](#)
[show chassis hardware lcc \(TX Matrix Router\) on page 721](#)
[show chassis hardware scc \(TX Matrix Router\) on page 722](#)
[show chassis hardware \(T1600 Router\) on page 722](#)
[show chassis hardware \(TX Matrix Plus Router\) on page 725](#)
[show chassis hardware sfc \(TX Matrix Plus Router\) on page 729](#)
[show chassis hardware extensive \(TX Matrix Plus Router\) on page 731](#)
[show chassis hardware clei-models \(TX Matrix Plus Router\) on page 732](#)
[show chassis hardware detail \(TX Matrix Plus Router\) on page 734](#)
[show chassis hardware models \(TX Matrix Plus Router\) on page 736](#)
[show chassis hardware \(16-Port 10-Gigabit Ethernet MPC with SFP+ Optics \[MX Series Routers\]\) on page 739](#)
[show chassis hardware \(MPC3E \[MX Series Routers\]\) on page 739](#)
[show chassis hardware \(QFX3500 Switches\) on page 741](#)
[show chassis hardware detail \(QFX3500 Switches\) on page 741](#)
[show chassis hardware models \(QFX3500 Switches\) on page 742](#)
[show chassis hardware clei-models \(QFX3500 Switches\) on page 742](#)
[show chassis hardware interconnect-device \(QFabric Systems\) on page 742](#)
[show chassis hardware node-device \(QFabric Systems\) on page 743](#)
[show chassis hardware \(PTX5000 Packet Transport Switch\) on page 743](#)
[show chassis hardware clei-models \(PTX5000 Packet Transport Switch\) on page 744](#)
[show chassis hardware detail \(PTX5000 Packet Transport Switch\) on page 745](#)
[show chassis hardware models \(PTX5000 Packet Transport Switch\) on page 746](#)
[show chassis hardware extensive \(PTX5000 Packet Transport Switch\) on page 747](#)
[show chassis hardware \(MX Routers with Media Services Blade \[MSB\]\) on page 748](#)
[show chassis hardware extensive \(MX Routers with Media Services Blade \[MSB\]\) on page 748](#)

Output Fields [Table 60 on page 649](#) lists the output fields for the **show chassis hardware** command. Output fields are listed in the approximate order in which they appear.

Table 60: show chassis hardware Output Fields

Field Name	Field Description	Level of Output
Item	<p>Chassis component:</p> <ul style="list-style-type: none"> (EX Series switches)—Information about the chassis, Routing Engine (SRE and Routing Engine modules in EX8200 switches), power supplies, fan trays, and LCD panel. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs). Information about the backplane, midplane, and SIBs (SF modules) is displayed for EX8200 switches. See <i>EX Series Switches Hardware and CLI Terminology Mapping</i>. (MX Series routers and EX Series switches)—Information about the backplane, Routing Engine, Power Entry Modules (PEMs), and fan trays. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs), Modular Port Concentrators (MPCs) and associated Modular Interface Cards (MICs), or Dense Port Concentrators (DPCs). MX80 routers have a single Routing Engine and a built-in Packet Forwarding Engine that attaches directly to MICs. The Packet Forwarding Engine has two “pseudo” FPCs (FPC 0 and FPC1). MX80 routers also have a Forwarding Engine Board (FEB). (M Series routers, except for the M320 router)—Information about the backplane; power supplies; fan trays; Routing Engine; maxicab (the connection between the Routing Engine and the backplane, for the M40 router only); SCB, SSB, SFM, or FEB; MCS and PCG (for the M160 router only); each FPC and PIC; and each fan, blower, and impeller. (M120, M320, and T Series routers)—Information about the backplane, power supplies, fan trays, midplane, FPM (craft interface), CIP, PEM, SCG, CB, FPC, PIC, SFP, SPMB, and SIB. (QFX Series)—Information about the chassis, Routing Engine, power supplies, fan trays, Interconnect devices, and Node devices. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs). (PTX Series)—Information about the chassis, midplane, craft interface (FPM), power distribution units (PDUs) and Power Supply Modules (PSMs), Centralized Clock Generators (CCGs), Routing Engines, Control Boards (CBs) and Switch Processor Mezzanine Boards (SPMBs), Flexible PIC Concentrators (FPCs), PICs, Switch Interface Boards (SIBs), and fan trays (vertical and horizontal). (MX2010 and MX2020 routers)—Information about the chassis, midplane, craft interface (FPM), power midplane (PMP), Power Supply Modules (PSMs), Power Distribution Modules (PDMs), Routing Engines, Control Boards (CBs) and Switch Processor Mezzanine Boards (SPMBs), Switch Fabric Boards (SFBs), Flexible PIC Concentrators (FPCs), PICs, adapter cards (ADCs) and fan trays. 	All levels
Version	Revision level of the chassis component.	All levels
Part number	Part number of the chassis component.	All levels
Serial number	Serial number of the chassis component. The serial number of the backplane is also the serial number of the router chassis. Use this serial number when you need to contact Juniper Networks Customer Support about the router or switch chassis.	All levels
Assb ID or Assembly ID	(extensive keyword only) Identification number that describes the FRU hardware.	extensive

Table 60: show chassis hardware Output Fields (*continued*)

Field Name	Field Description	Level of Output
Assembly Version	(extensive keyword only) Version number of the FRU hardware.	extensive
Assembly Flags	(extensive keyword only) Flags.	extensive
FRU model number	(clei-models , extensive , and models keyword only) Model number of the FRU hardware component.	none specified
CLEI code	(clei-models and extensive keyword only) Common Language Equipment Identifier code. This value is displayed only for hardware components that use ID EEPROM format v2. This value is not displayed for components that use ID EEPROM format v1.	none specified
EEPROM Version	ID EEPROM version used by the hardware component: 0x00 (version 0), 0x01 (version 1), or 0x02 (version 2).	extensive
Description	<p>Brief description of the hardware item:</p> <ul style="list-style-type: none"> • Type of power supply. • Type of PIC. If the PIC type is not supported on the current software release, the output states Hardware Not Supported. • Type of FPC: FPC Type 1, FPC Type 2, FPC Type 3, FPC Type 4, or FPC TypeOC192. <p>On EX Series switches, a brief description of the FPC.</p> <p>On the J Series routers, the FPC type corresponds to the Physical Interface Module (PIM). The following list shows the PIM abbreviation in the output and the corresponding PIM name.</p> <ul style="list-style-type: none"> • 2x FE—Either two built-in Fast Ethernet interfaces (fixed PIM) or dual-port Fast Ethernet PIM • 4x FE—4-port Fast Ethernet ePIM • 1x GE Copper—Copper Gigabit Ethernet ePIM (one 10-Mbps, 100-Mbps, or 1000-Mbps port) • 1x GE SFP—SFP Gigabit Ethernet ePIM (one fiber port) • 4x GE Base PIC—Four built-in Gigabit Ethernet ports on a J4350 or J6350 chassis (fixed PIM) • 2x Serial—Dual-port serial PIM • 2x T1—Dual-port T1 PIM • 2x E1—Dual-port E1 PIM • 2x CTIE1—Dual-port channelized T1/E1 PIM • 1x T3—T3 PIM (one port) • 1x E3—E3 PIM (one port) • 4x BRI S/T—4-port ISDN BRI S/T PIM • 4x BRI U—4-port ISDN BRI U PIM • 1x ADSL Annex A—ADSL 2/2+ Annex A PIM (one port, for POTS) • 1x ADSL Annex B—ADSL 2/2+ Annex B PIM (one port, for ISDN) • 1x TGM550—TGM550 Telephony Gateway Module (Avaya VoIP gateway module with one console port, two analog LINE ports, and two analog TRUNK ports) 	All levels

Table 60: show chassis hardware Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<ul style="list-style-type: none"> • 1x DS1 TIM510—TIM510 E1/T1 Telephony Interface Module (Avaya VoIP media module with one E1 or T1 trunk termination port and ISDN PRI backup) • 4x FXS, 4x FX0, TIM514—TIM514 Analog Telephony Interface Module (Avaya VoIP media module with four analog LINE ports and four analog TRUNK ports) • 4x BRI TIM521—TIM521 BRI Telephony Interface Module (Avaya VoIP media module with four ISDN BRI ports) • Crypto Accelerator Module—For enhanced performance of cryptographic algorithms used in IP Security (IPsec) services • MPC M 16x 10GE—16-port 10-Gigabit Module Port Concentrator that supports SFP+ optical transceivers. (Not on EX Series switches.) • For hosts, the Routing Engine type. • For small form-factor pluggable transceiver (SFP) modules, the type of fiber: LX, SX, LH, or T. • LCD description for EX Series switches (except EX2200 switches). • MPC2—1-port MPC2 that supports two separate slots for MICs. • MPC3E—1-port MPC3E that supports two separate slots for MICs (MIC-3D-1X100GE-CFP and MIC-3D-20GE-SFP) on MX960, MX480, and MX240 routers. The MPC3E maps one MIC to one PIC (1 MIC, 1 PIC), which differs from the mapping of legacy MPCs. • 100GBASE-LR4, pluggable CFP optics • Supports the Enhanced MX Switch Control Board with fabric redundancy and existing SCBs without fabric redundancy. • Interoperates with existing MX Series line cards, including Flexible Port Concentrators (FPC), Dense Port Concentrators (DPCs), and Modular Port Concentrators (MPCs). • MPC4E—Fixed configuration MPC4E that is available in two flavors: MPC4E-3D-32XGE-SFP and MPC4E-3D-2CGE-8XGE) on MX2020, MX960, MX480, and MX240 routers. • LCD description for MX Series routers and EX Series switches 	

Sample Output

show chassis hardware (EX8216 Switch)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis       REV 06                CY0109220035  EX8216
Midplane      REV 06    710-016845  BA0909120112  EX8216-MP
CB 0          REV 22    710-020771  AX0109197723  EX8216-RE320
CB 1          REV 22    710-020771  AX0109197726  EX8216-RE320
Routing Engine 1    BUILTIN    BUILTIN      RE-EX8216
FPC 3         REV 19    710-020683  BC0109083125  EX8200-48F
CPU           REV 13    710-020598  BF0109144549  EX8200-CPU
FPC 4         REV 17    710-020683  BC0108500127  EX8200-48F
CPU           REV 10    710-020598  BF0108460510  EX8200-CPU
PIC 0                    BUILTIN    BUILTIN      48x 100 Base-QFX/1000
Base-X
Xcvr 1        REV 01    740-011613  PE70V89       SFP-SX

```

Xcvr 11	REV 01	740-011613	PE70YCE	SFP-SX
Xcvr 12	REV 01	740-011613	PE70VSH	SFP-SX
Xcvr 13	REV 01	740-011613	E08C02063	SFP-SX
Xcvr 14	REV 01	740-011613	PE70VKU	SFP-SX
Xcvr 15	REV 01	740-011613	E08E03372	SFP-SX
Xcvr 21	REV 01	740-011613	PE70VAD	SFP-SX
Xcvr 22	REV 01	740-011613	E08E01228	SFP-SX
Xcvr 23	REV 01	740-011613	PE70VSL	SFP-SX
Xcvr 24	REV 01	740-011613	E08E03409	SFP-SX
Xcvr 25	REV 01	740-011613	PE70VL4	SFP-SX
Xcvr 26	REV 01	740-011613	PDQ4L2Z	SFP-SX
Xcvr 27	REV 01	740-011613	PE70WFK	SFP-SX
Xcvr 28	REV 01	740-011782	PBD2B5U	SFP-SX
Xcvr 29	REV 01	740-011613	PE70UQX	SFP-SX
Xcvr 30	REV 01	740-011613	PE70VL5	SFP-SX
Xcvr 31	REV 01	740-011613	PE70V0F	SFP-SX
Xcvr 32	REV 01	740-011613	E08C02052	SFP-SX
Xcvr 33	REV 01	740-011613	E08C02197	SFP-SX
Xcvr 34	REV 01	740-011613	PE70V0L	SFP-SX
Xcvr 35	REV 01	740-011613	E08E03390	SFP-SX
Xcvr 36	REV 01	740-011613	PDQ4VL9	SFP-SX
Xcvr 37	REV 01	740-011613	E08E03370	SFP-SX
Xcvr 38	REV 01	740-011613	E08E03362	SFP-SX
Xcvr 39	REV 01	740-011613	E08C02065	SFP-SX
Xcvr 40	REV 01	740-011613	E08E03405	SFP-SX
Xcvr 41	REV 01	740-011613	E08E03411	SFP-SX
Xcvr 43	REV 01	740-011613	E08C02171	SFP-SX
Xcvr 45	REV 01	740-011613	E08E03410	SFP-SX
FPC 13	REV 16	710-016837	BB0109051344	EX8200-8XS
CPU				
SIB 0	REV 10	710-021613	AY0109166244	EX8216-SF320
SIB 1	REV 10	710-021613	AY0109166357	EX8216-SF320
SIB 2	REV 10	710-021613	AY0109166362	EX8216-SF320
SIB 3	REV 10	710-021613	AY0109166338	EX8216-SF320
SIB 4	REV 10	710-021613	AY0109166350	EX8216-SF320
SIB 5	REV 10	710-021613	AY0109166365	EX8216-SF320
SIB 6	REV 10	710-021613	AY0109166361	EX8216-SF320
SIB 7	REV 10	710-021613	AY0109166399	EX8216-SF320
PSU 0	REV 17	740-021466	BG0709170003	EX8200-AC2K
PSU 1	REV 17	740-021466	BG0709170004	EX8200-AC2K
PSU 2	REV 17	740-021466	BG0709170020	EX8200-AC2K
PSU 3	REV 17	740-021466	BG0709170017	EX8200-AC2K
PSU 4	REV 17	740-021466	BG0709170008	EX8200-AC2K
PSU 5	REV 17	740-021466	BG0709170018	EX8200-AC2K
Top Fan Tray				
FTC 0	REV 4	760-022620	CX1209140212	EX8216-FT
FTC 1	REV 4	760-022620	CX1209140212	EX8216-FT
Bottom Fan Tray				
FTC 0	REV 4	760-022620	CX1209140211	EX8216-FT
FTC 1	REV 4	760-022620	CX1209140211	EX8216-FT
LCD 0	REV 04	710-025742	CE0109186919	EX8200 LCD

show chassis hardware clei-models (EX8216 Switch)

```

user@host> show chassis hardware clei-models
Hardware inventory:
Item          Version  Part number  CLEI code      FRU model number
Midplane      REV 08   710-016845
PSU 0         REV 05   740-023002   COUPAEAEAA     EX8200-PWR-AC3KR
PSU 1         REV 05   740-023002   COUPAEAEAA     EX8200-PWR-AC3KR
PSU 2         REV 05   740-023002   COUPAEAEAA     EX8200-PWR-AC3KR

```

```

PSU 3          REV 05  740-023002  COUPAEAEAA  EX8200-PWR-AC3KR
PSU 4          REV 05  740-023002  COUPAEAEAA  EX8200-PWR-AC3KR
PSU 5          REV 05  740-023002  COUPAEAEAA  EX8200-PWR-AC3KR
Top Fan Tray
Bottom Fan Tray

```

show chassis hardware clei-models (T1600 Router)

```

user@host> show chassis hardware clei-models
Hardware inventory:

```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 03	710-005608		CHAS-BP-T640-S
FPM Display	REV 05	710-002897		CRAFT-T640-S
CIP	REV 06	710-002895		CIP-L-T640-S
PEM 0	Rev 07	740-017906	IPUPAC7KTA	PWR-T1600-3-80-DC-S
PEM 1	Rev 18	740-002595		PWR-T-DC-S
SCG 0	REV 15	710-003423		SCG-T-S
Routing Engine 0	REV 08	740-014082		RE-A-2000-4096-S
Routing Engine 1	REV 07	740-014082		RE-A-2000-4096-S
CB 0	REV 05	710-007655		CB-T-S
CB 1	REV 03	710-017707		CB-T-S
FPC 0	REV 07	710-013558		T640-FPC2-E2
PIC 0	REV 01	750-010618		PB-4GE-SFP
PIC 1	REV 06	750-001900		PB-10C48-SON-SMSR
PIC 2	REV 14	750-001901		PB-40C12-SON-SMIR
PIC 3	REV 07	750-001900		PB-10C48-SON-SMSR
FPC 1	REV 06	710-013553		T640-FPC1-E2
PIC 0	REV 08	750-001072		P-1GE-SX
PIC 1	REV 10	750-012266		PB-4GE-TYPE1-SFP-IQ2
PIC 2	REV 22	750-005634		PB-1CHOC12SMIR-QPP
FPC 2				
PIC 0	REV 16	750-007141		PC-10GE-SFP
PIC 1	REV 06	750-015217		PC-8GE-TYPE3-SFP-IQ2
PIC 2	REV 05	750-004695		PC-TUNNEL
PIC 3	REV 17	750-009553		PC-40C48-SON-SFP
FPC 3	REV 01	710-010154		T640-FPC3-E
PIC 0	REV 07	750-012793		PC-1XGE-TYPE3-XFP-IQ2
PIC 1	REV 25	750-007141		PC-10GE-SFP
PIC 2	REV 17	750-009553		PC-40C48-SON-SFP
PIC 3	REV 32	750-003700		PC-10C192-SON-VSR
FPC 4	REV 16	710-013037		T1600-FPC4-ES
PIC 1	REV 06	750-034781		PD-1CE-CFP
FPC 5	REV 02	710-013037		T1600-FPC4-ES
PIC 0	REV 16	750-012518		PD-40C192-SON-XFP
PIC 1	REV 01	750-010850		PD-10C768-SON-SR
FPC 6	REV 14	710-013037		T1600-FPC4-ES
PIC 0	REV 11	750-017405		PD-4XGE-XFP
PIC 1	REV 13	750-017405		PD-4XGE-XFP
FPC 7	REV 09	710-007529		T640-FPC3
PIC 0	REV 10	750-012793		PC-1XGE-TYPE3-XFP-IQ2
PIC 1	REV 01	750-015217		PC-8GE-TYPE3-SFP-IQ2
PIC 2	REV 01	750-015217		PC-8GE-TYPE3-SFP-IQ2
PIC 3	REV 15	750-009450		PC-10C192-SON-SR2
SIB 0	REV 07	710-013074		SIB-I-T1600-S
SIB 1	REV 07	710-013074		SIB-I-T1600-S
SIB 2	REV 07	710-013074		SIB-I-T1600-S
SIB 3	REV 07	710-013074		SIB-I-T1600-S
SIB 4	REV 07	710-013074		SIB-I-T1600-S
Fan Tray 0				FANTRAY-T-S

```

Fan Tray 1
Fan Tray 2
FANTRAY-T-S
FAN-REAR-TX-T640-S

```

show chassis hardware detail (EX4200 Switch)

```

user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number  Description
Chassis
Routing Engine 0 REV 11    750-021256   BM0208327733  EX4200-24T, 8 POE
Routing Engine 0 REV 11    750-021256   BM0208327733  EX4200-24T, 8 POE
FPC 0
CPU              BUILTIN    BUILTIN      FPC CPU
PIC 0            BUILTIN    BUILTIN      24x 10/100/1000 Base-T
PIC 1            REV 03B    711-021270   AR0208162285  4x GE SFP
BRD              REV 08     711-021264   AK0208328289  EX4200-24T, 8 POE
Power Supply 0   REV 03     740-020957   AT0508346354  PS 320W AC
Fan Tray

```

show chassis hardware models (EX4500 Switch)

```

user@host> show chassis hardware models
Hardware inventory:
Item             Version  Part number  Serial number  FRU model number
Routing Engine 0 REV 01    750-035700   GG0210271867  EX4500-40F-FB-C
FPC 0            REV 01    750-035700   GG0210271867  EX4500-40F-FB-C
PIC 0            BUILTIN    BUILTIN      EX4500-40F-FB-C
Power Supply 1   REV 01    740-029654   H884FS00JC09  EX4500-PWR1-AC-FB

```

show chassis hardware (J6350 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number  Description
Chassis
Midplane         REV 03    710-014593   NP1265        JSR6350
System IO        REV 01    710-016210   NN9950        JX350 System IO
Crypto Module
Routing Engine   REV 08    710-015273   NM6509        RE-J6350-3400
ad0             248 MB   256MB CKS    00102006C24A00000039 Compact
Flash
FPC 0
PIC 0
FPC 1            REV 06    750-010355   AI07030023    FPC
PIC 0            REV 06    750-011148   AJ06520151    FPC
FPC 3            REV 06    750-011148   AJ06520151    FPC
PIC 0            REV 06    750-013492   NC4170        FPC
FPC 6            REV 06    750-013492   NC4170        FPC
PIC 0
Power Supply 0

```

show chassis hardware (J6300 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number  Description
Chassis
Midplane         REV 02.04 710-010001   CORE99570     J6300
System IO        REV 02.00 710-010003   CORE100848    System IO board
Routing Engine   RevX2.6   750-010006   IWGS40735390  RE-J.3
FPC 0
PIC 0

```

FPC 1	RevX2.0	750-011380	N3960005	FPC
PIC 0				1xADSL pic Annex A
FPC 2	RevX2.0	750-011380	N3960002	FPC
PIC 0				1xADSL pic Annex B
FPC 3	REV 03	750-010354	N0780028	FPC
PIC 0				1x T3

show chassis hardware (M7i Router)

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			31959	M7i
Midplane	REV 02	710-008761	CA0209	M7i Midplane
Power Supply 0	Rev 04	740-008537	PD10272	AC Power Supply
Routing Engine	REV 01	740-008846	1000396803	RE-5.0
CFEB	REV 02	750-009492	CA0166	Internet Processor IIv1
FPC 0				E-FPC
PIC 0	REV 04	750-003163	HJ6416	1x G/E, 1000 BASE-SX
PIC 1	REV 04	750-003163	HJ6423	1x G/E, 1000 BASE-SX
PIC 2	REV 04	750-003163	HJ6421	1x G/E, 1000 BASE-SX
PIC 3	REV 02	750-003163	HJ0425	1x G/E, 1000 BASE-SX
FPC 1				E-FPC
PIC 2	REV 01	750-009487	HM2275	ASP - Integrated
PIC 3	REV 01	750-009098	CA0142	2x F/E, 100 BASE-TX

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			B1157	M7i
Midplane	REV 05	710-008761	DM0840	M7i Midplane
Power Supply 0	Rev 08	740-008537	TE53755	AC Power Supply
Routing Engine	REV 07	740-011202	1000736567	RE-850
CFEB	REV 09	750-010463	DK6952	Internet Processor II
FPC 0				E-FPC
PIC 0	REV 12	750-012838	DL7993	4x 1GE(LAN), IQ2
Xcvr 0	REV 01	740-011614	PD94TDJ	SFP-LX10
Xcvr 1	REV 01	740-011615	PAD5EER	UNSUPPORTED
Xcvr 2	REV 01	740-011614	PD94THU	SFP-LX10
Xcvr 3		NON-JNPR	PDC2E7A	SFP-LX10
PIC 1	REV 03	750-023116	JT0203	4x CHSTM1 SDH CE SFP
Xcvr 0	REV 01	740-012434	AGT063832PS	SFP-SR
Xcvr 1	REV 01	740-012434	AGT063832LY	SFP-SR
Xcvr 3	REV 01	740-016064	C06J19018	SFP-LR
PIC 2	REV 15	750-014895	DM5757	MultiServices 100
PIC 3	REV 01	750-025390	JW9448	12x T1/E1 CE
FPC 1				E-FPC
PIC 2		BUILTIN	BUILTIN	1x Tunnel
PIC 3	REV 09	750-009099	DM0899	1x G/E, 1000 BASE
Xcvr 0	REV 01	740-012434	AGT07150HGJ	UNSUPPORTED
Fan Tray				Rear Fan Tray

show chassis hardware (M10 Router)

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			1122	M10
Midplane	REV 1.1	710-001950	S/N AC6626	
Power supply A	Rev 01	740-002497	S/N LC36095	AC
Power supply B	Rev 01	740-002497	S/N LC36100	AC
Display	REV 1.2	710-001995	S/N AC6656	
Host			18000005dfb3fb01	teknon

FEB	REV 01	710-001948	S/N AC6632	Internet Processor II
FPC 0				
PIC 0	REV 08	750-001072	S/N AB2485	1x G/E, 1000 BASE-SX
PIC 1	REV 01	750-000613	S/N AA1048	1x OC-12 SONET, SMIR
FPC 1				
Fan Tray 0				FANTRAY-M10I-S
Fan Tray 1				FANTRAY-M10I-S

show chassis hardware models (M10 Router)

```
user@host> show chassis hardware models
```

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 04	710-008920		CHAS-MP-M10i-S
Power Supply 0	Rev 06	740-008537		PWR-M10i-M7i-AC-S
Power Supply 1	Rev 06	740-008537		PWR-M10i-M7i-AC-S
HCM 0	REV 03	710-010580		HCM-M10i-S
HCM 1	REV 03	710-010580		HCM-M10i-S
Routing Engine 0	REV 09	740-009459		RE-400-256-S
CFEB 0	REV 05	750-010465		FEB-M10i-M7i-S
FPC 0				
PIC 0	REV 10	750-002971		PE-40C3-SON-MM
PIC 1	REV 11	750-002992		PE-4FE-TX
PIC 2	REV 03	750-002977		PE-20C3-ATM-MM
PIC 3	REV 08	750-005724		PE-20C3-ATM2-MM
FPC 1				
PIC 2	REV 12	750-008425		PE-AS
PIC 3	REV 13	750-005636		PE-4CHDS3-QPP
Fan Tray 0				FANTRAY-M10I-S
Fan Tray 1				FANTRAY-M10I-S

show chassis hardware (M20 Router)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			20033	M20
Backplane	REV 07	710-001517	S/N AA7940	
Power supply B	Rev 01	740-001465	S/N 000001	AC
Display	REV 02	710-001519	S/N AA9704	
Host 0			98000004f8f27501	teknor
SSB slot 0	REV 01	710-001951	S/N AD5905	Internet Processor II
SSRAM bank 0	REV 01	710-001385	S00480	2 MB
SSRAM bank 1	REV 01	710-001385	S00490	2 MB
SSRAM bank 2	REV 01	710-001385	S001:?	2 MB
SSRAM bank 3	REV 01	710-001385	S00483	2 MB
SSB slot 1	N/A	N/A	N/A	Backup
FPC 1	REV 01	710-001292	S/N AB7528	
SSRAM	REV 01	710-000077	S/N 304209	1 MB
SDRAM bank 0	REV 01	710-000099	S/N 000603	64 MB
SDRAM bank 1	REV 01	710-000099	S/N 000414	64 MB
PIC 0	REV 03	750-000612	S/N AB8433	2x OC-3 ATM, MM
PIC 1	REV 01	750-000616	S/N AA1168	1x OC-12 ATM, MM
PIC 2	REV 01	750-000613	S/N AA1008	1x OC-12 SONET, SMIR
PIC 3	REV 01	750-002501	S/N AD5810	4x E3
FPC 2	REV 01	710-001292	S/N AC0119	
SSRAM	REV 01	710-000077	S/N 503241	1 MB
SDRAM bank 0	REV 01	710-000099	S/N 306835	64 MB
SDRAM bank 1	REV 01	710-000099	S/N 306832	64 MB
Fan Tray 0				Front Upper Fan Tray
Fan Tray 1				Front Middle Fan Tray

Fan Tray 2
Fan Tray 3

Front Bottom Fan Tray
Rear Fan Tray

show chassis hardware models (M20 Router)

```
user@host> show chassis hardware models
Hardware inventory:
Item          Version  Part number  CLEI code  FRU model number
Backplane     REV 03   710-002334
Power Supply A REV 06   740-001465
Display       REV 04   710-001519
Routing Engine 0 REV 06   740-003239
Routing Engine 1 REV 06   740-003239
SSB 0         REV 02   710-001951
SSB 1         N/A      N/A
FPC 0         REV 03   710-003308
  PIC 0       REV 08   750-002303
  PIC 1       REV 07   750-004745
  PIC 2       REV 03   750-002965
FPC 1         REV 03   710-003308
  PIC 0       REV 03   750-002914
Fan Tray 0
Fan Tray 1
Fan Tray 2
Fan Tray 3
FPC-E
P-4FE-TX
P-2MCD53
PE-4CHDS3
FPC-E
P-20C3-ATM-MM
FANTRAY-F-M20-S
FANTRAY-F-M20-S
FANTRAY-F-M20-S
FANTRAY-R-M20-S
```

show chassis hardware (M40 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Backplane     REV 02   710-000073   S/N AA0053
Power supply A Rev 2     740-000235   S/N 000042    DC
Maxicab       REV X1   710-000229   S/N AA0139
Minicab       REV X1   710-000482   S/N AA0201
Display       REV 06   710-000150   S/N AA0905
Host
SCB
  SSRAM bank 0 REV 02   710-000077   S/N AA2267    1 MB
  SSRAM bank 1 REV 02   710-000077   S/N AA2270    1 MB
  SSRAM bank 2 REV 02   710-000077   S/N AA2269    1 MB
  SSRAM bank 3 REV 02   710-000077   S/N AA2268    1 MB
FPC 0
  SSRAM       REV 01   710-000175   S/N AA0048    1 MB
  SDRAM bank 0 REV 01   710-000099   S/N AA2332    64 MB
  SDRAM bank 1 REV X1   710-000099   S/N AA2337    64 MB
  PIC 0       REV 04   750-000613   S/N aa0343    1x OC-12 SONET, SMIR
  PIC 1       REV 04   750-000613   S/N AA0379    1x OC-12 SONET, SMIR
  PIC 2       REV 04   750-000613   S/N AA0377    1x OC-12 SONET, SMIR
  PIC 3       REV 04   750-000613   S/N AA0378    1x Tunnel
FPC 2
  SSRAM       REV 02   710-000077   S/N AA2288    1 MB
  SDRAM bank 0 REV 01   710-000099   S/N AA2331    64 MB
  SDRAM bank 1 REV 01   710-000099   S/N AA2330    64 MB
  PIC 0       REV X1   750-000603   S/N AA0143    4x OC-3 SONET, SMIR
  PIC 1       REV X1   750-000615   S/N AA0149    4x OC-3 SONET, MM
  PIC 2       REV X1   750-000611   S/N AA0148    4x OC-3 SONET, MM
  PIC 3       REV 04   750-000613   S/N AA0330    1x OC-12 SONET, SMIR
FPC 4
  SSRAM       REV 01   710-000077   S/N AA2327    1 MB
  SDRAM bank 0 REV 01   710-000099   S/N AA2329    64 MB
```

SDRAM bank 1	REV 01	710-000099	S/N AA2328	64 MB
PIC 0	REV 04	750-000613	S/N AA0320	1x OC-12 SONET, SMIR
PIC 2	REV 05	750-000616	S/N AA1341	1x OC-12 ATM, MM
PIC 3	REV 08	750-001072	S/N AB2462	1x G/E, 1000 BASE-SX
FPC 5	REV 10	710-000175	S/N AA7663	
SSRAM	REV 01	710-000077	S/N 501590	1 MB
SDRAM bank 0	REV 01	710-000099	S/N 300949	64 MB
SDRAM bank 1	REV 01	710-000099	S/N 300868	64 MB
PIC 1	REV 01	750-001323	S/N AB1670	1x Tunnel

show chassis hardware (M40e Router)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis				m40e
Midplane	REV 01	710-005071	AX3671	
FPM CMB	REV 03	710-001642	AR9074	
FPM Display	REV 03	710-001647	AR7331	
CIP	REV 04	710-002649	BB4449	
PEM 0	Rev 01	740-003787	MC12364	Power Entry Module
PEM 1	Rev 01	740-003787	MC12383	Power Entry Module
PCG 0	REV 07	710-001568	AG1332	
PCG 1	REV 07	710-001568	AR3789	
Host 0			3e000007c8176601	Present
MCS 0	REV 11	710-001226	AN5813	
SFM 0 SPP	REV 07	710-001228	AG4676	
SFM 0 SPR	REV 05	710-002189	AE4735	Internet Processor II
SFM 1 SPP	REV 07	710-001228	AP1347	
SFM 1 SPR	REV 05	710-002189	BE0063	Internet Processor II
FPC 0	REV 01	710-011725	BE0669	M40e-EP-FPC Type 1
CPU	REV 01	710-004600	BD9504	
PIC 0	REV 03	750-003737	AY3991	4x G/E, 1000 BASE-SX
FPC 1	REV 01	710-005197	BD9842	M40e-FPC Type 2
CPU	REV 01	710-004600	BB4869	
PIC 0	REV 07	750-001900	AR8278	1x OC-48 SONET, SMSR
FPC 2	REV 02	710-005197	BD9824	M40e-FPC Type 2
CPU	REV 01	710-004600	BD9531	
PIC 0	REV 03	750-003737	AY3986	4x G/E, 1000 BASE-SX
FPC 4	REV 02	710-005078	BE0664	M40e-FPC Type 1
CPU	REV 01	710-004600	BD9559	
PIC 0	REV 03	750-001894	AG7963	1x G/E, 1000 BASE-SX
PIC 2	REV 01	750-002575	AF2472	4x OC-3 SONET, SMIR
FPC 6	REV 02	710-005078	BE0652	M40e-FPC Type 1
CPU	REV 01	710-004600	BD9607	
PIC 0	REV 02	750-002911	AN2286	4x F/E, 100 BASE-TX
PIC 2	REV 01	750-002577	AP6345	4x OC-3 SONET, MM

show chassis hardware (M120 Router)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN000054AC	M120
Midplane	REV 01	710-013667	RB4170	M120 Midplane
FPM Board	REV 02	710-011407	CJ9186	M120 FPM Board
FPM Display	REV 02	710-011405	CJ9173	M120 FPM Display
FPM CIP	REV 02	710-011410	CJ9221	M120 FPM CIP
PEM 0	Rev 05	740-011936	RM28320	AC Power Entry Module
PEM 1	Rev 05	740-011936	RM28321	AC Power Entry Module
Routing Engine 0	REV 03	740-014080	1000642883	RE-A-1000

CB 0	REV 03	710-011403	CM8346	M120 Control Board
CB 1	REV 06	710-011403	CP6728	M120 Control Board
FPC 1	REV 02	710-015908	CP6925	M120 CFPC 10GE
PIC 0		BUILTIN	BUILTIN	1x 10GE(LAN/WAN) XFP
Xcvr 0	REV 01	740-014279	62E204N00007	XFP-10G-LR
FPC 3	REV 03	710-011393	CJ9234	M120 FPC Type 2
PIC 0	REV 16	750-008155	NB5229	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	P9F15JB	SFP-SX
Xcvr 1	REV 01	740-007326	P4Q0R9G	SFP-SX
PIC 1	REV 09	750-007745	CG4360	4x OC-3 SONET, SMIR
PIC 2	REV 16	750-008155	ND7787	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	P9F12AS	SFP-SX
Xcvr 1	REV 01	740-011613	P9F1ALU	SFP-SX
PIC 3	REV 07	750-011800	JW1284	8x 1GE(LAN), IQ2
Xcvr 0	REV 01	740-011613	P9F1AM6	SFP-SX
Xcvr 6	REV 01	740-011613	P9F16NN	SFP-SX
Xcvr 7	REV 01	740-011782	P8C29Y7	SFP-SX
Board B	REV 02	710-011395	CN3754	M120 FPC Mezz
FPC 4	REV 02	710-011398	CP6741	M120 FPC Type 3
PIC 0	REV 16	750-007141	NB2855	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011782	P922A1F	SFP-SX
Xcvr 1	REV 01	740-011782	P922A16	SFP-SX
Xcvr 2	REV 01	740-011782	P922A0U	SFP-SX
Xcvr 3	REV 01	740-011782	P9229UZ	SFP-SX
Xcvr 4	REV 01	740-009029	P11JXWP	SFP-LX
Xcvr 6	REV 01	740-011613	P9F1ALW	SFP-SX
FPC 5	REV 01	710-011388	CJ9088	M120 FPC Type 1
PIC 0	*** Hardware Not Supported ***			
PIC 1	REV 05	750-012052	NB0410	1x CHOC3 IQ SONET, SMLR
PIC 2	REV 01	750-013167	CM3824	4x CHDS3 IQ
PIC 3	REV 01	750-010240	CB5366	1x G/E SFP, 1000 BASE
Board B	REV 01	710-011390	CJ9103	M120 FPC Mezz Board
FEB 3	REV 04	710-011663	CP6673	M120 FEB
FEB 4	REV 04	710-011663	CJ9368	M120 FEB
FEB 5	REV 04	710-011663	CJ9386	M120 FEB
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Top Fan Tray
Fan Tray 3				Rear Bottom Fan Tray

show chassis hardware detail (M120 Router)

```

user@host> show chassis hardware detail
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN000054AC     M120
Midplane      REV 01    710-013667   RB4170         M120 Midplane
FPM Board     REV 02    710-011407   CJ9186         M120 FPM Board
FPM Display   REV 02    710-011405   CJ9173         M120 FPM Display
FPM CIP       REV 02    710-011410   CJ9221         M120 FPM CIP
PEM 0         Rev 05    740-011936   RM28320        AC Power Entry Module
PEM 1         Rev 05    740-011936   RM28321        AC Power Entry Module
Routing Engine 0 REV 03    740-014080   1000642883     RE-A-1000
  ad0         248 MB   SILICONSYSTEMS INC 256M 126CT505S0763SC00110 Compact Flash
  ad2         38154 MB HTE541040G9SA00    MPBBTOX2HS2E3M Hard Disk
CB 0          REV 03    710-011403   CM8346         M120 Control Board
CB 1          REV 06    710-011403   CP6728         M120 Control Board
FPC 1         REV 02    710-015908   CP6925         M120 CFPC 10GE
PIC 0         BUILTIN  BUILTIN      BUILTIN        1x 10GE(LAN/WAN) XFP

```

Xcvr 0	REV 01	740-014279	62E204N00007	XFP-10G-LR
FPC 3	REV 03	710-011393	CJ9234	M120 FPC Type 2
PIC 0	REV 16	750-008155	NB5229	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	P9F15JB	SFP-SX
Xcvr 1	REV 01	740-007326	P4Q0R9G	SFP-SX
PIC 1	REV 09	750-007745	CG4360	4x OC-3 SONET, SMIR
PIC 2	REV 16	750-008155	ND7787	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	P9F12AS	SFP-SX
Xcvr 1	REV 01	740-011613	P9F1ALU	SFP-SX
PIC 3	REV 07	750-011800	JW1284	8x 1GE(LAN), IQ2
Xcvr 0	REV 01	740-011613	P9F1AM6	SFP-SX
Xcvr 6	REV 01	740-011613	P9F16NN	SFP-SX
Xcvr 7	REV 01	740-011782	P8C29Y7	SFP-SX
Board B	REV 02	710-011395	CN3754	M120 FPC Mezz
FPC 4	REV 02	710-011398	CP6741	M120 FPC Type 3
PIC 0	REV 16	750-007141	NB2855	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011782	P922A1F	SFP-SX
Xcvr 1	REV 01	740-011782	P922A16	SFP-SX
Xcvr 2	REV 01	740-011782	P922A0U	SFP-SX
Xcvr 3	REV 01	740-011782	P9229UZ	SFP-SX
Xcvr 4	REV 01	740-009029	P11JXWP	SFP-LX
Xcvr 6	REV 01	740-011613	P9F1ALW	SFP-SX
FPC 5	REV 01	710-011388	CJ9088	M120 FPC Type 1
PIC 0	*** Hardware Not Supported ***			
PIC 1	REV 05	750-012052	NB0410	1x CHOC3 IQ SONET, SMLR
PIC 2	REV 01	750-013167	CM3824	4x CHDS3 IQ
PIC 3	REV 01	750-010240	CB5366	1x G/E SFP, 1000 BASE
Board B	REV 01	710-011390	CJ9103	M120 FPC Mezz Board
FEB 3	REV 04	710-011663	CP6673	M120 FEB
FEB 4	REV 04	710-011663	CJ9368	M120 FEB
FEB 5	REV 04	710-011663	CJ9386	M120 FEB
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Top Fan Tray
Fan Tray 3				Rear Bottom Fan Tray

show chassis hardware models (M120 Router)

```
user@host> show chassis hardware models
Hardware inventory:
```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 01	710-013667		
FPM CIP	REV 02	710-011410		CRAFT-M120-S
PEM 0	Rev 05	740-011936		PWR-M120-AC-S
PEM 1	Rev 05	740-011936		PWR-M120-AC-S
Routing Engine 0	REV 03	740-014080		RE-A-1000-2048-S
CB 0	REV 03	710-011403		CB-M120-S
CB 1	REV 06	710-011403		CB-M120-S
FPC 1	REV 02	710-015908		M120-cFPC-1XGE-XFP
FPC 3				
PIC 0	REV 16	750-008155		PB-2GE-SFP-QPP
PIC 1	REV 09	750-007745		PC-40C3-SON-SMIR
PIC 2	REV 16	750-008155		PB-2GE-SFP-QPP
PIC 3	REV 07	750-011800		PB-8GE-TYPE2-SFP-IQ2
FPC 4				
PIC 0	REV 16	750-007141		PC-10GE-SFP
FPC 5				
PIC 1	REV 05	750-012052		PB-1CHOC3-SMIR-QPP
PIC 2	REV 01	750-013167		PE-4CHDS3-QPP

PIC 3	REV 01	750-010240	PB-1GE-SFP
Fan Tray 0			FFANTRAY-M120-S
Fan Tray 1			FFANTRAY-M120-S
Fan Tray 2			RFANTRAY-M120-S
Fan Tray 3			RFANTRAY-M120-S

show chassis hardware (M160 Router)

```
user@host> show chassis hardware
```

Item	Version	Part number	Serial number	Description
Chassis			101	M160
Midplane	REV 02	710-001245	S/N AB4107	
FPM CMB	REV 01	710-001642	S/N AA2911	
FPM Display	REV 01	710-001647	S/N AA2999	
CIP	REV 02	710-001593	S/N AA9563	
PEM 0	Rev 01	740-001243	S/N KJ35769	DC
PEM 1	Rev 01	740-001243	S/N KJ35765	DC
PCG 0	REV 01	710-001568	S/N AA9794	
PCG 1	REV 01	710-001568	S/N AA9804	
Host 1			da000004f8d57001	teknor
MCS 1	REV 03	710-001226	S/N AA9777	
SFM 0 SPP	REV 04	710-001228	S/N AA2975	
SFM 0 SPR	REV 02	710-001224	S/N AA9838	Internet Processor I
SFM 1 SPP	REV 04	710-001228	S/N AA2860	
SFM 1 SPR	REV 01	710-001224	S/N AB0139	Internet Processor I
FPC 0	REV 03	710-001255	S/N AA9806	FPC Type 1
CPU	REV 02	710-001217	S/N AA9590	
PIC 1	REV 05	750-000616	S/N AA1527	1x OC-12 ATM, MM
PIC 2	REV 05	750-000616	S/N AA1535	1x OC-12 ATM, MM
PIC 3	REV 01	750-000616	S/N AA1519	1x OC-12 ATM, MM
FPC 1	REV 02	710-001611	S/N AA9523	FPC Type 2
CPU	REV 02	710-001217	S/N AA9571	
PIC 0	REV 03	750-001900	S/N AA9626	1x STM-16 SDH, SMIR
PIC 1	REV 01	710-002381	S/N AD3633	2x G/E, 1000 BASE-SX
FPC 2				FPC Type OC192
CPU	REV 03	710-001217	S/N AB3329	
PIC 0	REV 01			1x OC-192 SM SR-2
Fan Tray 0				Rear Bottom Blower
Fan Tray 1				Rear Top Blower
Fan Tray 2				Front Top Blower
Fan Tray 3				Front Fan Tray

show chassis hardware models (M160 Router)

```
user@host> show chassis hardware models
```

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 03	710-009120		CHAS-BP-M320-S
FPM Display	REV 02	710-009351		CRAFT-M320-S
CIP	REV 03	710-005926		CIP-M320-S
PEM 2	Rev X4	740-009148		PWR-M-DC-S
PEM 3	Rev X4	740-009148		PWR-M-DC-S
Routing Engine 0	REV 02	740-008883		RE-1600-2048-S
Routing Engine 1	REV 02	740-008883		RE-1600-2048-S
FPC 0	REV 02	710-010419		M320-FPC1
PIC 0	REV 01	750-001323		P-TUNNEL
PIC 1	REV 02	750-002987		PE-10C12-SON-SMIR
PIC 2	REV 04	750-001894		PB-1GE-SX
PIC 3	REV 04	750-001896		PB-10C12-SON-SMIR
FPC 1	REV 02	710-010419		M320-FPC1
PIC 0	REV 04	750-001894		PB-1GE-SX

PIC 1	REV 04	750-001894	PB-1GE-SX
PIC 3	REV 03	750-001894	PB-1GE-SX
FPC 2	REV 02	710-010419	M320-FPC1
PIC 0	REV 10	750-005634	PB-1CHOC12SMIR-QPP
PIC 1	REV 10	750-005634	PB-1CHOC12SMIR-QPP
PIC 2	REV 07	750-005634	PB-1CHOC12SMIR-QPP
PIC 3	REV 07	750-005634	PB-1CHOC12SMIR-QPP
PIC 1	REV 10	750-005634	PB-1CHOC12SMIR-QPP
PIC 2	REV 07	750-005634	PB-1CHOC12SMIR-QPP
PIC 3	REV 07	750-005634	PB-1CHOC12SMIR-QPP
FPC 3			
PIC 0	REV 03	750-001895	PB-10C12-SON-MM
PIC 1	REV 04	750-001894	PB-1GE-SX
PIC 3	REV 04	750-003141	PB-1GE-SX-B
FPC 4	REV 02	710-010419	M320-FPC1
FPC 5	REV 02	710-010419	M320-FPC1
FPC 6	REV 02	710-010419	M320-FPC1
FPC 7			
PIC 0	REV 15	750-001901	PB-40C12-SON-SMIR
PIC 1	REV 06	750-001900	PB-10C48-SON-SMSR
PIC 2	REV 07	750-001900	PB-10C48-SON-SMSR
PIC 3	REV 05	750-003737	PB-4GE-SX
SIB 0	REV 03	710-009184	SIB-M-S
SIB 1	REV 03	710-009184	SIB-M-S
SIB 2	REV 03	710-009184	SIB-M-S
SIB 3	REV 03	710-009184	SIB-M-S
Fan Tray 0			FFANTRAY-M320-S
Fan Tray 1			FFANTRAY-M320-S
Fan Tray 2			RFANTRAY-M320-S

show chassis hardware detail (M160 Router)

```
user@host> show chassis hardware detail
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			101	M160
Midplane	REV 02	710-001245	S/N AB4107	
FPM CMB	REV 01	710-001642	S/N AA2911	
FPM Display	REV 01	710-001647	S/N AA2999	
CIP	REV 02	710-001593	S/N AA9563	
PEM 0	Rev 01	740-001243	S/N KJ35769	DC
PEM 1	Rev 01	740-001243	S/N KJ35765	DC
PCG 0	REV 01	710-001568	S/N AA9794	
PCG 1	REV 01	710-001568	S/N AA9804	
Host 1			da000004f8d57001	teknor
MCS 1	REV 03	710-001226	S/N AA9777	
SFM 0 SPP	REV 04	710-001228	S/N AA2975	
SFM 0 SPR	REV 02	710-001224	S/N AA9838	Internet Processor I
SSRAM bank 0	REV 01	710-000077	S/N 306456	1 MB
SSRAM bank 1	REV 01	710-000077	S/N 306474	1 MB
SSRAM bank 2	REV 01	710-000077	S/N 306388	1 MB
SSRAM bank 3	REV 01	710-000077	S/N 306392	1 MB
SFM 1 SPP	REV 04	710-001228	S/N AA2860	
SFM 1 SPR	REV 01	710-001224	S/N AB0139	Internet Processor I
SSRAM bank 0	REV 01	710-000077	S/N 302917	1 MB
SSRAM bank 1	REV 01	710-000077	S/N 302662	1 MB
SSRAM bank 2	REV 01	710-000077	S/N 302593	1 MB
SSRAM bank 3	REV 01	710-000077	S/N 100160	1 MB
FPC 0	REV 03	710-001255	S/N AA9806	FPC Type 1
CPU	REV 02	710-001217	S/N AA9590	
SSRAM	REV 01	710-000077	S/N 302836	1 MB

SDRAM 0	REV 01	710-001196	S00141	32 MB
SDRAM 1	REV 01	710-001196	S0010;	32 MB
SSRAM	REV 01	710-000077	S/N 302633	1 MB
SDRAM 0	REV 01	710-001196	S00143	32 MB
SDRAM 1	REV 01	710-001196	S00115	32 MB
SSRAM	REV 01	710-000077	S/N 302952	1 MB
SDRAM 0	REV 01	710-001196	S00135	32 MB
SDRAM 1	REV 01	710-001196	S001=3	32 MB
SSRAM	REV 01	710-000077	S/N 302892	1 MB
SDRAM 0	REV 01	710-001196	S000?6	32 MB
SDRAM 1	REV 01	710-001196	S001=5	32 MB
PIC 1	REV 05	750-000616	S/N AA1527	1x OC-12 ATM, MM
PIC 2	REV 05	750-000616	S/N AA1535	1x OC-12 ATM, MM
PIC 3	REV 01	750-000616	S/N AA1519	1x OC-12 ATM, MM
FPC 1	REV 02	710-001611	S/N AA9523	FPC Type 2
CPU	REV 02	710-001217	S/N AA9571	
SSRAM	REV 01	710-000077	S/N 306340	1 MB
SDRAM 0	REV 01	710-001196	S00012	32 MB
SDRAM 1	REV 01	710-001196	S0001?	32 MB
SSRAM	REV 01	710-000077	S/N 306454	1 MB
SDRAM 0	REV 01	710-001196	S00028	32 MB
SDRAM 1	REV 01	710-001196	S0002?	32 MB
SSRAM	REV 01	710-000077	S/N 306492	1 MB
SDRAM 0	REV 01	710-001196	S00015	32 MB
SDRAM 1	REV 01	710-001196	S00031	32 MB
SSRAM	REV 01	710-000077	S/N 306363	1 MB
SDRAM 0	REV 01	710-001196	S00013	32 MB
SDRAM 1	REV 01	710-001196	S00032	32 MB
PIC 0	REV 03	750-001900	S/N AA9626	1x STM-16 SDH, SMIR
PIC 1	REV 01	710-002381	S/N AD3633	2x G/E, 1000 BASE-SX
FPC 2				FPC Type OC192
... SSRAM	REV 01	710-000077	S/N 306466	1 MB

show chassis hardware (M320 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               67245         M320
Midplane      REV 05   710-009120   RB1202        M320 Midplane
FPM GBUS      REV 04   710-005928   HZ5697        M320 Board
FPM Display   REV 05   710-009351   HR1464        M320 FPM Display
CIP           REV 04   710-005926   HT8672        M320 CIP
PEM 0         Rev 05   740-009148   QK34208       DC Power Entry Module
PEM 1         Rev 05   740-009148   QK34262       DC Power Entry Module
PEM 2         Rev 05   740-009148   QF10449       DC Power Entry Module
PEM 3         Rev 05   740-009148   QJ18257       DC Power Entry Module
Routing Engine 0 REV 06   740-008883   P11123901185 RE-4.0
CB 0          REV 07   710-009115   JB2382        M320 Control Board
FPC 0         REV 02   710-005017   CD9926        M320 FPC Type 2
CPU           REV 01   710-011659   CJ6940        M320 PCA SCPU
PIC 0         REV 07   750-001900   AT1594        1x OC-48 SONET, SMSR
PIC 1         REV 03   750-001850   HS2746        1x Tunnel
PIC 2         REV 05   750-010618   JE7117        4x G/E SFP, 1000 BASE
PIC 3         REV 06   750-001900   HE6083        1x OC-48 SONET, SMSR
FPC 2         REV 02   710-005017   CH0319        M320 FPC Type 1
CPU           REV 01   710-011659   CJ6942        M320 PCA SCPU
PIC 0         REV 05   750-003034   BD8705        4x OC-3 SONET, SMIR
FPC 5         REV 02   710-005017   CD9938        M320 FPC Type 2
CPU
FPC 7         REV 02   710-005017   CD9934        M320 FPC Type 2

```

CPU				
SIB 0	REV 09	710-009184	JA6540	M320 SIB
SIB 1	REV 09	710-009184	HV9511	M320 SIB
SIB 2	REV 09	710-009184	HW2057	M320 SIB
SIB 3	REV 09	710-009184	JA6687	M320 SIB
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray

show chassis hardware models (M320 Router)

```

user@host> show chassis hardware models
Hardware inventory:
Item                Version  Part number  CLEI code  FRU model number
Midplane            REV 03    710-009120
FPM Display         REV 02    710-009351
CIP                 REV 03    710-005926
PEM 2               Rev X4    740-009148
PEM 3               Rev X4    740-009148
Routing Engine 0    REV 02    740-008883
Routing Engine 1    REV 02    740-008883
FPC 0               REV 02    710-010419
  PIC 0             REV 01    750-001323
  PIC 1             REV 02    750-002987
  PIC 2             REV 04    750-001894
  PIC 3             REV 04    750-001896
FPC 1               REV 02    710-010419
  PIC 0             REV 04    750-001894
  PIC 1             REV 04    750-001894
  PIC 3             REV 03    750-001894
FPC 2               REV 02    710-010419
  PIC 0             REV 10    750-005634
  PIC 1             REV 10    750-005634
  PIC 2             REV 07    750-005634
  PIC 3             REV 07    750-005634
  PIC 1             REV 10    750-005634
  PIC 2             REV 07    750-005634
  PIC 3             REV 07    750-005634
FPC 3
  PIC 0             REV 03    750-001895
  PIC 1             REV 04    750-001894
  PIC 3             REV 04    750-003141
FPC 4               REV 02    710-010419
FPC 5               REV 02    710-010419
FPC 6               REV 02    710-010419
FPC 7
  PIC 0             REV 15    750-001901
  PIC 1             REV 06    750-001900
  PIC 2             REV 07    750-001900
  PIC 3             REV 05    750-003737
SIB 0               REV 03    710-009184
SIB 1               REV 03    710-009184
SIB 2               REV 03    710-009184
SIB 3               REV 03    710-009184
Fan Tray 0
Fan Tray 1
Fan Tray 2

```

show chassis hardware (MX5 Router)

```

user@host> show chassis hardware

```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			E1368	MX5-T
Midplane	REV 01	711-038215	YF5288	MX5-T
PEM 0	Rev 04	740-028288	VA01215	AC Power Entry Module
PEM 1	Rev 04	740-028288	VA01218	AC Power Entry Module
Routing Engine		BUILTIN	BUILTIN	Routing Engine
TFEB 0		BUILTIN	BUILTIN	Forwarding Engine
Processor				
QXM 0	REV 05	711-028408	ZA9136	MPC QXM
FPC 0		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0		BUILTIN	BUILTIN	4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	4x 10GE XFP
FPC 1		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0	REV 24	750-028392	YX9820	3D 20x 1GE(LAN) SFP
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-031851	AM1045SUAQ3	SFP-SX
Xcvr 1	REV 01	740-031851	AM1045SUAPA	SFP-SX
Xcvr 2	REV 01	740-031851	AM1045SUAN7	SFP-SX
Xcvr 3	REV 01	740-031851	AM1045SU91Q	SFP-SX
Xcvr 4	REV 01	740-031851	AM1045SUDDR	SFP-SX
Xcvr 9	REV 01	740-011613	AM0848SB6A1	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-031851	AM1045SUANO	SFP-SX
Xcvr 1	REV 01	740-011613	AS0812S0719	SFP-SX
Xcvr 2	REV 01	740-011613	AM0821SA121	SFP-SX
Xcvr 3	REV 01	740-011613	PF21K21	SFP-SX
Xcvr 4	REV 01	740-011613	AM0848SB69Z	SFP-SX
Xcvr 5	REV 01	740-011782	P9POXV3	SFP-SX
Xcvr 6	REV 01	740-011613	AM0812S8WJN	SFP-SX
Xcvr 7	REV 01	740-011613	PAM3G9Q	SFP-SX
Xcvr 8	REV 01	740-011613	AM0848SB4A6	SFP-SX
Xcvr 9	REV 01	740-011782	P9MOU37	SFP-SX
MIC 1	REV 20	750-028380	ZG2657	3D 2x 10GE XFP
PIC 2		BUILTIN	BUILTIN	1x 10GE XFP
PIC 3		BUILTIN	BUILTIN	1x 10GE XFP
Fan Tray				Fan Tray

show chassis hardware (MX10 Router)

user@host> show chassis hardware

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			E1372	MX10-T
Midplane	REV 01	711-038211	YF5285	MX10-T
PEM 0	Rev 04	740-028288	VB01678	AC Power Entry Module
Routing Engine		BUILTIN	BUILTIN	Routing Engine
TFEB 0		BUILTIN	BUILTIN	Forwarding Engine
Processor				
QXM 0	REV 05	711-028408	ZA9053	MPC QXM
FPC 0		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0		BUILTIN	BUILTIN	4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	4x 10GE XFP
FPC 1		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0	REV 24	750-028392	YX9436	3D 20x 1GE(LAN) SFP
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-031851	AM1107SUFQW	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Fan Tray				Fan Tray

show chassis hardware (MX40 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number  Description
Chassis                               E1367         MX40-T
Midplane          REV 01   711-038211   YF5284        MX40-T
PEM 0             Rev 04   740-028288   VB01680       AC Power Entry Module
PEM 1             Rev 04   740-028288   VB01700       AC Power Entry Module
Routing Engine    BUILTIN  BUILTIN      BUILTIN       Routing Engine
TFEB 0            BUILTIN  BUILTIN      BUILTIN       Forwarding Engine
Processor
  QXM 0           REV 05   711-028408   ZA9048        MPC QXM
  FPC 0            BUILTIN  BUILTIN      BUILTIN       MPC BUILTIN
    MIC 0           BUILTIN  BUILTIN      BUILTIN       4x 10GE XFP
      PIC 0          BUILTIN  BUILTIN      BUILTIN       4x 10GE XFP
        Xcvr 0       REV 01   740-014279   M7067UPP      XFP-10G-LR
        Xcvr 1       NON-JNPR  K9J02UN      XFP-10G-LR
  FPC 1            BUILTIN  BUILTIN      BUILTIN       MPC BUILTIN
    MIC 0           REV 24   750-028392   YX3504        3D 20x 1GE(LAN) SFP
      PIC 0          BUILTIN  BUILTIN      BUILTIN       10x 1GE(LAN) SFP
        Xcvr 0       REV 01   740-011613   AM0812S8WTE   SFP-SX
        Xcvr 1       REV 01   740-011613   PFA6KV2       SFP-SX
        Xcvr 2       REV 01   740-031851   AM1045SUDDM   SFP-SX
        Xcvr 3       REV 01   740-011613   PD63C7M       SFP-SX
        Xcvr 4       REV 01   740-011613   PD63DJY       SFP-SX
        Xcvr 5       REV 02   740-011613   AA0950STLL9   SFP-SX
        Xcvr 6       REV 01   740-011782   PAR1YHC       SFP-SX
        Xcvr 7       REV 01   740-011782   P9POXXL       SFP-SX
        Xcvr 8       REV 01   740-011613   PD63D95       SFP-SX
        Xcvr 9       REV 01   740-031851   AM1045SU9B8   SFP-SX
      PIC 1          BUILTIN  BUILTIN      BUILTIN       10x 1GE(LAN) SFP
        Xcvr 0       REV 01   740-011613   PF21L3Z       SFP-SX
        Xcvr 1       REV 01   740-031851   AM1045SU7M9   SFP-SX
        Xcvr 2       REV 01   740-031851   AM1045SUAPT   SFP-SX
        Xcvr 3       REV 01   740-011613   PFF2BZH       SFP-SX
        Xcvr 4       REV 01   740-031851   AM1045SUDDN   SFP-SX
        Xcvr 5       REV 01   740-031851   AM1039S00ZR   SFP-SX
        Xcvr 6       REV 01   740-031851   AM1045SUD6Y   SFP-SX
        Xcvr 8       REV 01   740-011613   PFM1QBS       SFP-SX
        Xcvr 9       REV 01   740-011613   PFF2E25       SFP-SX
    MIC 1           REV 01   750-021130   KG4391        3D 2x 10GE XFP
      PIC 2          BUILTIN  BUILTIN      BUILTIN       1x 10GE XFP
        Xcvr 0       REV 01   740-011571   C645XJ04G     XFP-10G-SR
      PIC 3          BUILTIN  BUILTIN      BUILTIN       1x 10GE XFP
        Xcvr 0       NON-JNPR  CA49BK0AE     XFP-10G-SR
Fan Tray

```

show chassis hardware (Fixed MX80 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number  Description
Chassis                               MX80-48T
Midplane          REV 01   711-031603   KF9250        MX80-48T
Routing Engine    BUILTIN  BUILTIN      BUILTIN       Routing Engine
FEB 0            BUILTIN  BUILTIN      BUILTIN       Forwarding Engine Board
FPC 0            BUILTIN  BUILTIN      BUILTIN       MPC BUILTIN
  MIC 0           BUILTIN  BUILTIN      BUILTIN       4x 10GE XFP
    PIC 0          BUILTIN  BUILTIN      BUILTIN       4x 10GE XFP
      Xcvr 0       NON-JNPR  M6439D41     XFP-10G-LR

```


Xcvr 1	REV 01	740-014279	6XE931N00202	XFP-10G-LR
Xcvr 2	REV 01	740-014289	C715XU05F	XFP-10G-SR
Xcvr 3	REV 01	740-014289	C650XU0EP	XFP-10G-SR
FPC 1		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0	REV 01	711-029399	JR6981	12x 1GE(LAN) RJ45
PIC 0		BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
PIC 1		BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
MIC 1	REV 01	BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
PIC 2		BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
PIC 3		BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
Fan Tray				Fan Tray

show chassis hardware (Modular MX80 Router)

```
user@host> show chassis hardware
```

```
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               MX80
Midplane      REV 02    711-031594   JR7084         MX80
PEM 0         Rev 01    740-028288   000018         AC Power Entry Module
Routing Engine                               Routing Engine
FEB 0                               BUILTIN       BUILTIN       Forwarding Engine Board

QXM 0         REV 05    711-028408   JR7041         MPC QXM
FPC 0                               BUILTIN       BUILTIN       MPC BUILTIN
MIC 0                               BUILTIN       BUILTIN       4x 10GE XFP
PIC 0                               BUILTIN       BUILTIN       4x 10GE XFP
FPC 1                               BUILTIN       BUILTIN       MPC BUILTIN
MIC 0         REV 02    750-028380   JR6598         3D 2x 10GE XFP
PIC 0                               BUILTIN       BUILTIN       1x 10GE XFP
Xcvr 0        REV 01    740-014289   T07M86365     XFP-10G-SR
PIC 1                               BUILTIN       BUILTIN       1x 10GE XFP
Xcvr 0        REV 01    740-014289   T07M71094     XFP-10G-SR
MIC 1         REV 02    750-028380   JG8548         3D 2x 10GE XFP
PIC 2                               BUILTIN       BUILTIN       1x 10GE XFP
Xcvr 0        REV 02    740-014289   T08L86302     XFP-10G-SR
PIC 3                               BUILTIN       BUILTIN       1x 10GE XFP
Xcvr 0        REV 02    740-014289   C810XU0BA     XFP-10G-SR
Fan Tray                               Fan Tray
```

show chassis hardware (MX240 Router)

```
user@host> show chassis hardware
```

```
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               MX240
Midplane      REV 01    710-021041   TR1502         MX240 Backplane
FPM Board     REV 01    710-017254   KD4017         Front Panel Display
PEM 0         Rev 02    740-017330   000332         PS 1.2-1.7kW; 100-240V
AC in
PEM 1         Rev 02    740-017330   000226         PS 1.2-1.7kW; 100-240V
AC in
Routing Engine 0 REV 06    740-013063   1000703522     RE-S-2000
Routing Engine 1 REV 06    740-015113   1000687625     RE-S-1300
CB 0          REV 07    710-013385   KC9057         MX SCB
CB 1          REV 05    710-013385   JY4760         MX SCB
FPC 1         REV 01    750-021679   KC7340         DPCE 40x 1GE R
CPU           REV 06    710-013713   KD4078         DPC PMB
PIC 0          BUILTIN  BUILTIN       BUILTIN       10x 1GE(LAN)
Xcvr 0        REV 01    740-011613   P9F18ME        SFP-SX
```

PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN)
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN)
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN)
FPC 2	REV 04	710-016669	JS4529	DPCE 40x 1GE R EQ
CPU	REV 06	710-013713	KB3969	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3Y79	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3XU8	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3YG6	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3XUG	SFP-SX
Xcvr 4	REV 01	740-011613	PBG3XTJ	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3ZUM	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3Y5H	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3UZT	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3US1	SFP-SX
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3YG7	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3XZ9	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3XTY	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3UZG	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3Y8W	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3YVX	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3YB3	SFP-SX
Xcvr 3	REV 01	740-011613	PBG43VQ	SFP-SX
Fan Tray 0	REV 01	710-021113	JS4642	MX240 Fan Tray

show chassis hardware detail (MX240 Router with RE Displaying DIMM information)

```
user@host> show chassis hardware detail
```

Item	Version	Part number	Serial number	Description
Chassis			JN11279B4AFC	MX240 Backplane
Midplane	REV 07	760-021404	TS2474	MX240 Backplane
FPM Board	REV 03	760-021392	XC2643	Front Panel Display
PEM 0	Rev 03	740-017343	QCS0908A068	DC Power Entry Module
Routing Engine 0	REV 01	740-031117	AARCH00	RE-S-1800x4
ad0 3764 MB	STEC M2+	CF 9.0.2	STIM2Q3209239145303	Removable Compact Flash
ad1 28626 MB	WDC SSD-F0030S-5000		C933Z036237215548S00	Compact Flash
usb0 (addr 1)	EHCI root hub 0		Intel	uhub0
usb0 (addr 2)	product 0x0020 32		vendor 0x8087	uhub1
DIMM 0	VL31B5263E-F8S DIE REV-0 PCB REV-0			MFR ID-ce80
DIMM 1	VL31B5263E-F8S DIE REV-0 PCB REV-0			MFR ID-ce80
DIMM 2	VL31B5263E-F8S DIE REV-0 PCB REV-0			MFR ID-ce80
DIMM 3	SL31B5263E-F8S DIE REV-0 PCB REV-0			MFR ID-ce80
CB 0	REV 03	710-021523	XD7225	MX SCB
Fan Tray 0	REV 01	710-021113	WZ4986	MX240 Fan Tray

show chassis hardware (MX240 Router with Enhanced MX SCB)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN10C7F7EAFC	MX240
Midplane	REV 01	710-021041	TR1502	MX240 Backplane
FPM Board	REV 01	710-017254	KD4017	Front Panel Display
PEM 0	Rev 02	740-017330	000332	PS 1.2-1.7kW; 100-240V
AC in				
PEM 1	Rev 02	740-017330	000226	PS 1.2-1.7kW; 100-240V
AC in				

Routing Engine 0	REV 06	740-013063	1000703522	RE-S-2000
Routing Engine 1	REV 06	740-015113	1000687625	RE-S-1300
CB 0	REV 02	710-031391	YE8494	Enhanced MX SCB
CB 1	REV 05	710-031391	YOP5764	Enhanced MX SCB
FPC 1	REV 01	750-021679	KC7340	DPCE 40x 1GE R
CPU	REV 06	710-013713	KD4078	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011613	P9F18ME	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN)
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN)
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN)
FPC 2	REV 04	710-016669	JS4529	DPCE 40x 1GE R EQ
CPU	REV 06	710-013713	KB3969	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3Y79	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3XU8	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3YG6	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3XUG	SFP-SX
Xcvr 4	REV 01	740-011613	PBG3XTJ	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3ZUM	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3Y5H	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3UZT	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3US1	SFP-SX
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3YG7	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3XZ9	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3XTY	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3UZG	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3Y8W	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3YVX	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3YB3	SFP-SX
Xcvr 3	REV 01	740-011613	PBG43VQ	SFP-SX
Fan Tray 0	REV 01	710-021113	JS4642	MX240 Fan Tray

show chassis hardware (MX480 Router)

```
user@host> show chassis hardware
```

Hardware inventory:				
Item	Version	Part number	Serial number	Description
Chassis			JN10C7F7FAFB	MX480
Midplane	REV 04	710-017414	TR2071	MX480 Midplane
FPM Board	REV 02	710-017254	KB8459	Front Panel Display
PEM 0	Rev 02	740-017330	QCS07519029	PS 1.2-1.7kW; 100-240V
AC in				
PEM 1	Rev 02	740-017330	QCS07519041	PS 1.2-1.7kW; 100-240V
AC in				
PEM 2	Rev 02	740-017330	QCS07519097	PS 1.2-1.7kW; 100-240V
AC in				
Routing Engine 0	REV 07	740-013063	1000733381	RE-S-2000
Routing Engine 1	REV 07	740-013063	1000733540	RE-S-2000
CB 0	REV 07	710-013385	KA8022	MX SCB
CB 1	REV 07	710-013385	KA8303	MX SCB
FPC 0	REV 09	750-020452	KA8660	DPCE 40x 1GE X EQ
CPU	REV 06	710-013713	KA8185	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) EQ

PIC 3	BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Fan Tray			Left Fan Tray

show chassis hardware (MX480 Router with Enhanced MX SCB)

```

user@host> show chassis hardware
Hardware inventory:
Item                Version  Part number  Serial number  Description
Chassis              REV 04    710-017414   JN10C7F7FAFB   MX480
Midplane              REV 02    710-017254   TR2071         MX480 Midplane
FPM Board             Rev 02    740-017330   KB8459         Front Panel Display
PEM 0                 Rev 02    740-017330   QCS07519029    PS 1.2-1.7kW; 100-240V
AC in
PEM 1                 Rev 02    740-017330   QCS07519041    PS 1.2-1.7kW; 100-240V
AC in
PEM 2                 Rev 02    740-017330   QCS07519097    PS 1.2-1.7kW; 100-240V
AC in
Routing Engine 0     REV 07    740-013063   1000733381     RE-S-2000
Routing Engine 1     REV 07    740-013063   1000733540     RE-S-2000
CB 0                  REV 07    710-013385   KA8022         Enhanced MX SCB
CB 1                  REV 07    710-013385   KA8303         Enhanced MX SCB
FPC 0                 REV 09    750-020452   KA8660         DPCE 40x 1GE X EQ
CPU                   REV 06    710-013713   KA8185         DPC PMB
PIC 0                 BUILTIN   BUILTIN      10x 1GE(LAN) EQ
PIC 1                 BUILTIN   BUILTIN      10x 1GE(LAN) EQ
PIC 2                 BUILTIN   BUILTIN      10x 1GE(LAN) EQ
PIC 3                 BUILTIN   BUILTIN      10x 1GE(LAN) EQ
Fan Tray              Left Fan Tray

```

show chassis hardware (MX960 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item                Version  Part number  Serial number  Description
Chassis              REV 01    710-013698   AA6082         MX960
Midplane              Rev 01    740-013110   000008         MX960 Midplane
PIM                   Rev 01    740-013682   000038         Power Inlet Module
PEM 2                 Rev 01    740-013682   000038         PS 1.7kW; 200-240VAC in
PEM 3                 REV 00    740-015113   1000617944     RE-S-1300
Routing Engine 0     REV 05    710-013725   JK6947         MX960 Test SCB
CB 0                  REV 01    710-013305   JM7617         MX960 Test DPC
FPC 4                 REV 01    710-013305   JL9634         MX960 Test DPC
CPU
PIC 0                 BUILTIN   BUILTIN      1x 10GE(LAN/WAN)
PIC 1                 BUILTIN   BUILTIN      10x 1GE
FPC 7                 REV 01    710-013305   JL9634         MX960 Test DPC
CPU
PIC 0                 BUILTIN   BUILTIN      1x 10GE(LAN/WAN)
Xcvr 0                NON-JNPR   MYBG65I82C    XFP-10G-SR
PIC 1                 BUILTIN   BUILTIN      10x 1GE
Xcvr 1                REV 01    740-011782   P7N0368        SFP-SX
Xcvr 4                REV 01    740-011782   P8J1W27        SFP-SX
Xcvr 6                REV 01    740-011782   P8J1VSD        SFP-SX
Xcvr 9                REV 01    740-011782   P8J1W25        SFP-SX
Fan Tray 0
Fan Tray 1

```

show chassis hardware (MX960 Router with Bidirectional Optics)

```

user@host> show chassis hardware
Hardware inventory:
Item                Version  Part number  Serial number  Description

```

Chassis			JN10BA5B9AFA	MX960
Midplane	REV 03	710-013698	TR0234	MX960 Backplane
FPM Board	REV 03	710-014974	JA0878	Front Panel Display
PDM	Rev 03	740-013110	QCS11135028	Power Distribution Module
PEM 0	Rev 03	740-013682	QCS11154036	PS 1.7kW; 200-240VAC in
PEM 1	Rev 03	740-013682	QCS11154010	PS 1.7kW; 200-240VAC in
PEM 2	Rev 03	740-013682	QCS11154022	PS 1.7kW; 200-240VAC in
Routing Engine 0	REV 06	740-013063	1000691458	RE-S-2000
CB 0	REV 07	710-013385	KA2190	MX SCB
CB 1	REV 07	710-013385	KA0837	MX SCB
FPC 3	REV 02	750-018122	KB3890	DPCE 40x 1GE R
CPU				
FPC 4	REV 01	750-018122	KB3889	DPCE 40x 1GE R
CPU	REV 06	710-013713	KB3976	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 1	REV 01	740-020426	4910549	SFP-1000BASE-BX40-D
Xcvr 2	REV 01	740-020426	4910551	SFP-1000BASE-BX40-D
Xcvr 5	REV 01	740-021340	77E245N00006	SFP-1000BASE-BX10-U
Xcvr 6	REV 01	740-020425	4882821	SFP-1000BASE-BX40-U
Xcvr 8	REV 01	740-020425	4882820	SFP-1000BASE-BX40-U
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-020465	77E555N00894	SFP-1000BASE-BX10-D
Xcvr 1	REV 01	740-020465	75E467X00818	SFP-1000BASE-BX10-D
Xcvr 2	REV 01	740-020465	75E467X00573	SFP-1000BASE-BX10-D
Xcvr 3	REV 01	740-020465	4888227	SFP-1000BASE-BX10-D
Xcvr 4	REV 01	740-020465	4888241	SFP-1000BASE-BX10-D
Xcvr 5	REV 01	740-021340	77E245N00005	SFP-1000BASE-BX10-U
Xcvr 6	REV 01	740-021340	76E245X00487	SFP-1000BASE-BX10-U
Xcvr 7	REV 01	740-021341	5255889	SFP-1000BASE-BX10-U
Xcvr 8	REV 01	740-021341	5255887	SFP-1000BASE-BX10-U
Xcvr 9	REV 01	740-021340	77E245N00004	SFP-1000BASE-BX10-U
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-020424	5007582	SFP-1000BASE-BX10-D
Xcvr 1	REV 01	740-020424	4888187	SFP-1000BASE-BX10-D
Xcvr 2	REV 01	740-020424	4656500	SFP-1000BASE-BX10-D
Xcvr 5	REV 01	740-021341	5255886	SFP-1000BASE-BX10-U
Xcvr 7	REV 01	740-021340	77E245N00003	SFP-1000BASE-BX10-U
Xcvr 8	REV 01	740-021341	5255888	SFP-1000BASE-BX10-U
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-017726	74S184H30341	SFP-EX
Xcvr 1	REV 01	740-017726	4814061	SFP-EX
Xcvr 5	REV 01	740-017726	6ZS184H31108	SFP-EX
Xcvr 9	REV 01	740-021340	76E245X00486	SFP-1000BASE-BX10-U
Fan Tray 0				
Fan Tray 1	REV 03	740-014971	TP0850	Fan Tray

show chassis hardware (MX960 Router with Enhanced MX SCB)

```
user@host> show chassis hardware
```

Hardware inventory:				
Item	Version	Part number	Serial number	Description
Chassis			JN1096805AFA	MX960
Midplane	REV 03	710-013698	TR0183	MX960 Backplane
Fan Extender	REV 02	710-018051	JY5227	Extended Cable Manager
FPM Board	REV 03	710-014974	JZ6876	Front Panel Display
PDM	Rev 03	740-013110	QCS11035023	Power Distribution Module
PEM 1	Rev 03	740-013682	QCS1109400L	PS 1.7kW; 200-240VAC in
PEM 2	Rev 03	740-013682	QCS11094015	PS 1.7kW; 200-240VAC in
PEM 3	Rev 03	740-013682	QCS11094012	PS 1.7kW; 200-240VAC in
Routing Engine 0	REV 06	740-013063	1000687969	RE-S-2000
Routing Engine 1	REV 06	740-013063	1000687955	RE-S-2000

CB 0	REV 11	750-031391	YZ6072	Enhanced MX SCB
CB 1	REV 11	750-031391	YZ6068	Enhanced MX SCB
CB 2	REV 11	750-031391	YZ6081	Enhanced MX SCB
FPC 0	REV 01	750-018122	KA5576	DPCE 40x 1GE R
CPU	REV 06	710-013713	KB3961	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011613	P9F18GF	SFP-SX
Xcvr 2	REV 01	740-011782	P9M0TL9	SFP-SX
Xcvr 7	REV 01	740-011782	P9P0XXH	SFP-SX
Xcvr 9	REV 01	740-011782	P9M0TN1	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011613	PAJ4UHC	SFP-SX
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011613	PFF2CD0	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3ZUT	SFP-SX
Xcvr 2	REV 01	740-011613	PFF2DDV	SFP-SX
Xcvr 5	REV 01	740-011613	P8E2SST	SFP-SX
Xcvr 9	REV 01	740-011782	PB8329N	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-026192	1U0201084503342	SFP-100BASE-BX10-U
Xcvr 1	REV 01	740-026193	1U1201084503313	SFP-100BASE-BX10-D
Xcvr 2	REV 01	740-011613	PAJ4Y5B	SFP-SX
Xcvr 6	REV 01	740-011782	P9M0U3M	SFP-SX
Xcvr 7	REV 01	740-011782	P9M0TLA	SFP-SX
FPC 1	REV 16	750-031089	YL0719	MPC Type 2 3D
CPU	REV 06	711-030884	YL1463	MPC PMB 2G
MIC 0	REV 07	750-028387	JR6500	3D 4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 01	740-014279	733019A00154	XFP-10G-LR
Xcvr 1	REV 02	740-014289	T09F55034	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 01	740-014279	913019B00791	XFP-10G-LR
Xcvr 1	REV 01	740-014289	98S803A90384	XFP-10G-SR
MIC 1	REV 24	750-028387	YJ3950	3D 4x 10GE XFP
PIC 2		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 02	740-014279	T10B36134	XFP-10G-LR
Xcvr 1	REV 01	740-014289	T07M86354	XFP-10G-SR
PIC 3		BUILTIN	BUILTIN	2x 10GE XFP
FPC 2	REV 08	710-014219	JY9654	DPCE 4x 10GE R
CPU	REV 06	710-013713	JZ6549	DPC PMB
PIC 0		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
PIC 1		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
PIC 2		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
Xcvr 0	REV 03	740-011571	C931BK028	XFP-10G-SR
PIC 3		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
FPC 3	REV 10	750-024199	XJ6692	MX FPC Type 3
CPU	REV 03	710-022351	XF5182	DPC PMB
PIC 0	REV 17	750-009553	RJ2945	4x OC-48 SONET
Xcvr 1	REV 01	740-011785	PCP3YLL	SFP-SR
Xcvr 3	REV 01	740-011785	PDSOMRY	SFP-SR
PIC 1	REV 32	750-003700	DP2113	1x OC-192 12xMM VSR
FPC 5	REV 25	750-028467	YM8256	MPC 3D 16x 10GE
CPU	REV 10	711-029089	YL3029	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 1	REV 01	740-031980	AHNOX1Z	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
FPC 7	REV 02	750-031092	JR6658	MPC Type 1 3D Q
CPU	REV 01	711-030884	JZ9038	MPC PMB 2G
MIC 0	REV 08	750-028392	JZ8737	3D 20x 1GE(LAN) SFP

PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-011782	PBE2C6Y	SFP-SX
Xcvr 2		NON-JNPR	U8105N8	SFP-SX
Xcvr 4	REV 01	740-011613	PFM18EF	SFP-SX
Xcvr 7	REV 01	740-011613	PFF2AM8	SFP-SX
Xcvr 8	REV 01	740-011613	PFF2CT6	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-011782	PB82VHH	SFP-SX
Xcvr 1	REV 01	740-011613	PFF2CSW	SFP-SX
Xcvr 9	REV 01	740-011613	PFF2BY0	SFP-SX
QXM 0	REV 04	711-028408	JR6372	MPC QXM
FPC 8	REV 05	750-024387	JW9754	MX FPC Type 2
CPU	REV 03	710-022351	KF1651	DPC PMB
PIC 0	REV 08	750-014730	DM3664	4x OC-3 1x OC-12 SFP
Xcvr 0	REV 01	740-016065	81S290N00077	SFP-SR
Xcvr 1		NON-JNPR	2191844	SFP-SR
Xcvr 2	REV 01	740-011618	PD81EE5	SFP-IR
PIC 1	REV 08	750-014637	DM3671	4x OC-12-3 SFP
Xcvr 0	REV 01	740-011785	PCK3UNK	SFP-SR
Xcvr 3	REV 01	740-011785	PDSOMPZ	SFP-SR
FPC 10	REV 04	710-013699	JY4654	DPCE 40x 1GE R
CPU	REV 05	710-013713	JS9717	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 5	REV 01	740-011782	PAR1L72	SFP-SX
Xcvr 6	REV 01	740-011782	P8N1YQ4	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN)
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011782	P8Q2AVL	SFP-SX
Xcvr 5	REV 01	740-011782	PAR1L7B	SFP-SX
Xcvr 6	REV 01	740-011782	PAR1L2J	SFP-SX
Xcvr 8	REV 01	740-011782	P8N1YMY	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN)
Fan Tray 0	REV 03	740-014971	TP0567	Fan Tray
Fan Tray 1	REV 03	740-014971	TP0702	Fan Tray

show chassis hardware models (MX960 Router with Enhanced MX SCB)

```
user@host> show chassis hardware models
```

Hardware inventory:				
Item	Version	Part number	Serial number	FRU model number
Midplane	REV 03	710-013698	TR0183	CHAS-BP-MX960-S
Fan Extender	REV 02	710-018051	JY5227	ECM-MX960
FPM Board	REV 03	710-014974	JZ6876	CRAFT-MX960-S
Routing Engine 0	REV 06	740-013063	1000687969	RE-S-2000-4096-S
Routing Engine 1	REV 06	740-013063	1000687955	RE-S-2000-4096-S
CB 0	REV 11	750-031391	YZ6072	SCBE-MX-S
CB 1	REV 11	750-031391	YZ6068	SCBE-MX-S
CB 2	REV 11	750-031391	YZ6081	SCBE-MX-S
FPC 0	REV 01	750-018122	KA5576	DPCE-R-40GE-SFP
FPC 1	REV 16	750-031089	YL0719	MX-MPC2-3D
MIC 0	REV 07	750-028387	JR6500	MIC-3D-4XGE-XFP
MIC 1	REV 24	750-028387	YJ3950	MIC-3D-4XGE-XFP
FPC 2	REV 08	710-014219	JY9654	DPC-R-4XGE-XFP
FPC 3	REV 10	750-024199	XJ6692	MX-FPC3
PIC 0	REV 17	750-009553	RJ2945	PC-40C48-SON-SFP
PIC 1	REV 32	750-003700	DP2113	PC-10C192-SON-VSR
FPC 5	REV 25	750-028467	YM8256	MPC-3D-16XGE-SFPP
FPC 7	REV 02	750-031092	JR6658	MX-MPC1-3D-Q
MIC 0	REV 08	750-028392	JZ8737	MIC-3D-20GE-SFP
FPC 8	REV 05	750-024387	JW9754	MX-FPC2
PIC 0	REV 08	750-014730	DM3664	PB-40C3-10C12-SON2-SFP

PIC 1	REV 08	750-014637	DM3671	PB-40C3-40C12-SON-SFP
FPC 10	REV 04	710-013699	JY4654	DPC-R-40GE-SFP
Fan Tray 0	REV 03	740-014971	TP0567	FFANTRAY-MX960-S
Fan Tray 1	REV 03	740-014971	TP0702	FFANTRAY-MX960-S

show chassis hardware detail (MX960 Router)

```

user@host> show chassis hardware detail
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis
Midplane      REV 01    710-013698   AA6082         MX960 Midplane
PIM           Rev 01    740-013110   000008         Power Inlet Module
PEM 2
PEM 3         Rev 01    740-013682   000038         PS 1.7kW; 200-240VAC in
Routing Engine 0 REV 00    740-015113   1000617944     RE-S-1300
  ad0         245 MB   SanDisk     SDCFB-256      111419E1805T1141 Compact Flash
  ad2         38154 MB FUJITSU     MHT2040BH      NROWT5925N77    Hard Disk
CB 0          REV 05    710-013725   JK6947         MX960 Test SCB
FPC 4         REV 01    710-013305   JM7617         MX960 Test DPC
CPU
PIC 0
PIC 1
FPC 7         REV 01    710-013305   JL9634         MX960 Test DPC
CPU
PIC 0
  Xcvr 0
PIC 1
  Xcvr 1       REV 01    740-011782   P7N0368        SFP-SX
  Xcvr 4       REV 01    740-011782   P8J1W27        SFP-SX
  Xcvr 6       REV 01    740-011782   P8J1VSD        SFP-SX
  Xcvr 9       REV 01    740-011782   P8J1W25        SFP-SX
Fan Tray 0
Fan Tray 1

```

show chassis hardware (MX2010 Router)

```

user@host > show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis
Midplane      REV 01    750-044636   ABAB8506       Lower Backplane
Midplane 1    REV 01    711-044557   ZY8296         Upper Backplane
PMP           REV 03    711-032426   ACAJ1388       Power Midplane
FPM Board     REV 06    711-032349   ZX8744         Front Panel Display
PSM 4         REV 0C    740-033727   VK00254        DC 52V Power Supply
Module
PSM 5         REV 0B    740-033727   VG00015        DC 52V Power Supply
Module
PSM 6         REV 0B    740-033727   VH00097        DC 52V Power Supply
Module
PSM 7         REV 0C    740-033727   VJ00151        DC 52V Power Supply
Module
PSM 8         REV 0C    740-033727   VJ00149        DC 52V Power Supply
Module
PDM 0         REV 0B    740-038109   WA00008        DC Power Dist Module
PDM 1         REV 0B    740-038109   WA00014        DC Power Dist Module
Routing Engine 0 REV 02    740-041821   9009094134     RE-S-1800x4
Routing Engine 1 REV 02    740-041821   9009094141     RE-S-1800x4
CB 0          REV 08    750-040257   CAAB3491       Control Board
CB 1          REV 08    750-040257   CAAB3489       Control Board

```


SPMB 0	REV 02	711-041855	CAAA6135	PMB Board
SPMB 1	REV 02	711-041855	CAAA6137	PMB Board
SFB 0	REV 06	711-032385	ZV1828	Switch Fabric Board
SFB 1	REV 07	711-032385	ZZ2568	Switch Fabric Board
SFB 2	REV 07	711-032385	ZZ2563	Switch Fabric Board
SFB 3	REV 07	711-032385	ZZ2564	Switch Fabric Board
SFB 4	REV 07	711-032385	ZZ2580	Switch Fabric Board
SFB 5	REV 07	711-032385	ZZ2579	Switch Fabric Board
SFB 6	REV 07	711-032385	CAAB4882	Switch Fabric Board
SFB 7	REV 07	711-032385	CAAB4898	Switch Fabric Board
FPC 0	REV 33	750-028467	CAAB1919	MPC 3D 16x 10GE
CPU	REV 11	711-029089	CAAB7174	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMH02RE	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMH038C	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AMH0390	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AMG0SUA	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMH0579	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMG0SGP	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AMH04SV	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AMH04X3	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMH0135	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMH02NC	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AMH02XB	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AMH02PN	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMH057Y	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMG0JHE	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AMH02HT	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AMH04V4	SFP+-10G-SR
FPC 1	REV 21	750-033205	ZG5027	MPCE Type 3D
CPU	REV 04	711-035209	YT4780	HMPC PMB 2G
MIC 0	REV 03	750-033307	ZV6299	10X10GE SFPP
PIC 0		BUILTIN	BUILTIN	10X10GE SFPP
Xcvr 0	REV 01	740-031980	083363A00410	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	083363A00334	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	113363A00125	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	083363A00953	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AHR013D	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJ40JUR	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJ40JKL	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJ30ECK	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	19T511100864	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	19T511100868	SFP+-10G-SR
MIC 1	REV 03	750-033307	ZV6268	10X10GE SFPP
PIC 2		BUILTIN	BUILTIN	10X10GE SFPP
Xcvr 0	REV 01	740-031980	AJCOJML	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ403PC	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJ10N25	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJ40JF4	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AJ40JSJ	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJ403V7	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJ40JN3	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJ40JSU	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	19T511100468	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	19T511101363	SFP+-10G-SR
FPC 8	REV 22	750-031089	ZT9746	MPC Type 2 3D
CPU	REV 06	711-030884	ZS1271	MPC PMB 2G
MIC 0	REV 26	750-028392	ABBS1150	3D 20x 1GE(LAN) SFP

PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-031851	PLG023C	SFP-SX
Xcvr 1	REV 01	740-031851	PLG09C6	SFP-SX
Xcvr 2	REV 02	740-011613	AM0950SF9L7	SFP-SX
Xcvr 3	REV 02	740-011613	AM1001SFN1H	SFP-SX
Xcvr 4	REV 02	740-011613	AM1001SFM9D	SFP-SX
Xcvr 5	REV 02	740-011613	AM1001SFLTJ	SFP-SX
Xcvr 6	REV 01	740-031851	AC1108S03L9	SFP-SX
Xcvr 7	REV 01	740-031851	AC1102S00NC	SFP-SX
Xcvr 8	REV 01	740-031851	AC1102S00MX	SFP-SX
Xcvr 9	REV 01	740-031851	AC1102S0085	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-031851	AC1102S00KU	SFP-SX
Xcvr 1	REV 01	740-031851	AC1102S00NG	SFP-SX
Xcvr 2	REV 01	740-031851	AC1102S00K3	SFP-SX
Xcvr 3	REV 01	740-031851	AC1102S008R	SFP-SX
Xcvr 4	REV 01	740-031851	AM1107SUFVJ	SFP-SX
Xcvr 5	REV 01	740-031851	AC1108S03LG	SFP-SX
MIC 1	REV 26	750-028387	ABBR9582	3D 4x 10GE XFP
PIC 2		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	T10A91703	XFP-10G-SR
Xcvr 1		NON-JNPR	T09L42604	XFP-10G-SR
PIC 3		BUILTIN	BUILTIN	2x 10GE XFP
FPC 9	REV 11	750-036284	ZL3591	MPC 3D 16x 10GE EM
CPU	REV 10	711-029089	ZL0513	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	1YT517101825	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	1YT517101821	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	1YT517101682	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALQ13R6	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	1YT517101828	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	1YT517101716	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	1YT517101732	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALP0TR1	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	1YT517101741	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	1YT517101829	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	1YT517101669	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALQ14E3	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	1YT517101826	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	1YT517101817	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	1YT517101735	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALQ159A	SFP+-10G-SR
ADC 0	REV 05	750-043596	CAAC2073	Adapter Card
ADC 1	REV 01	750-043596	ZV4117	Adapter Card
ADC 8	REV 01	750-043596	ZV4107	Adapter Card
ADC 9	REV 02	750-043596	ZW1555	Adapter Card
Fan Tray 0	REV 2A	760-046960	ACAY0015	172mm FanTray - 6 Fans
Fan Tray 1	REV 2A	760-046960	ACAY0019	172mm FanTray - 6 Fans
Fan Tray 2	REV 2A	760-046960	ACAY0020	172mm FanTray - 6 Fans
Fan Tray 3	REV 2A	760-046960	ACAY0021	172mm FanTray - 6 Fans

show chassis hardware detail (MX2010 Router)

```
user@host > show chassis hardware detail
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN11E233DAFK	MX2010
Midplane	REV 26	750-044636	ABAB9357	Lower Backplane

Midplane 1	REV 01	711-044557	ABAB8643	Upper Backplane
PMP	REV 04	711-032426	ACAJ1677	Power Midplane
FPM Board	REV 08	760-044634	ABBV9726	Front Panel Display
PSM 0	REV 01	740-045050	1E02224000P	DC 52V Power Supply
Module				
PSM 1	REV 01	740-045050	1E02224000M	DC 52V Power Supply
Module				
PSM 2	REV 01	740-045050	1E022240010	DC 52V Power Supply
Module				
PSM 3	REV 01	740-045050	1E02224000G	DC 52V Power Supply
Module				
PSM 4	REV 01	740-045050	1E022240013	DC 52V Power Supply
Module				
PSM 5	REV 01	740-045050	1E022240007	DC 52V Power Supply
Module				
PSM 6	REV 01	740-045050	1E02224001C	DC 52V Power Supply
Module				
PSM 7	REV 01	740-045050	1E02224001D	DC 52V Power Supply
Module				
PSM 8	REV 01	740-045050	1E02224001B	DC 52V Power Supply
Module				
PDM 0	REV 01	740-045234	1E262250067	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009099704	RE-S-1800x4
ad0 3831 MB		UGB30SFA4000T1	SFA4000T1 00000651	Compact Flash
ad1 30533 MB		UGB94BPH32H0S1-KCI	11000019592	Disk 1
usb0 (addr 1)		EHCI root hub 0	Intel	uhub0
usb0 (addr 2)		product 0x0020 32	vendor 0x8087	uhub1
DIMM 0		SGU04G72H1BD2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80
DIMM 1		SGU04G72H1BD2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80
DIMM 2		SGU04G72H1BD2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80
DIMM 3		SGU04G72H1BD2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80
Routing Engine 1	REV 02	740-041821	9009099706	RE-S-1800x4
ad0 3998 MB		Virtium - TuffDrive	VCF P1T0200262860208 114	Compact Flash
ad1 30533 MB		UGB94ARF32H0S3-KC	UNIGEN-499551-000404	Disk 1
CB 0	REV 13	750-040257	CAAF8436	Control Board
CB 1	REV 13	750-040257	CAAF8434	Control Board
SPMB 0	REV 02	711-041855	ABBV3825	PMB Board
SPMB 1	REV 02	711-041855	ABBV3833	PMB Board
SFB 0	REV 05	711-044466	ABBX5682	Switch Fabric Board
SFB 1	REV 05	711-044466	ABBX5676	Switch Fabric Board
SFB 2	REV 05	711-044466	ABBX5665	Switch Fabric Board
SFB 3	REV 05	711-044466	ABBX5699	Switch Fabric Board
SFB 4	REV 05	711-044466	ABBX5603	Switch Fabric Board
SFB 5	REV 05	711-044466	ABBX5587	Switch Fabric Board
SFB 6	REV 05	711-044466	ABBX5607	Switch Fabric Board
SFB 7	REV 05	711-044466	ABBX5669	Switch Fabric Board
FPC 0	REV 09	750-037355	CAAF0924	MPC Type 4-2
CPU	REV 08	711-035209	CAAB9842	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	19T511101656	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AMA04RU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00558	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B10M00202	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	X12J00328	CFP-100G-SR10
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	AMA088W	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B10L04211	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	19T511101602	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B10L04151	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP

Xcvr 0		NON-JNPR	X12J00332	CFP-100G-SR10
FPC 1	REV 18	750-033205	ZE0128	MPCE Type 3D
CPU	REV 06	711-035209	ZG5431	HMPC PMB 2G
MIC 0	REV 15	750-033199	ZP6435	1X100GE CFP
PIC 0		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	J11E46118	CFP-100G-LR4
MIC 1	REV 15	750-033199	ZP6442	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	UMN03T4	CFP-100G-LR4
FPC 2	REV 16	750-037358	CAAL1001	MPC Type 4-1
CPU	REV 08	711-035209	CAAK7927	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	193363A00589	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00028	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00376	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00016	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00499	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	973152A00039	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11E01239	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	973152A00058	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	B10M00075	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00014	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AMA0638	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00063	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AMA0629	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	973152A00053	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00344	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	973152A00046	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA062M	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00080	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00580	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00064	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	093363A01494	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	973152A00020	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	123363A00047	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	973152A00072	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-021308	03DZ06A01033	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00022	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	03DZ06A01026	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00013	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	03DZ06A01028	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	973152A00079	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	03DZ06A01018	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	973152A00025	SFP+-10G-SR
FPC 3	REV 33	750-028467	CAAF5400	MPC 3D 16x 10GE
CPU	REV 11	711-029089	CAAH7626	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00066	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00021	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	973152A00062	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00027	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00065	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00069	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	973152A00026	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00003	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00035	SFP+-10G-SR

Xcvr 1	REV 01	740-021308	973152A00004	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	973152A00049	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00055	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00010	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00001	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	973152A00073	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00012	SFP+-10G-SR
FPC 4	REV 21	750-033205	ZG5028	MPCE Type 3D
CPU	REV 05	711-035209	YX3911	HMPC PMB 2G
MIC 0	REV 03	750-036233	ZL2036	2X40GE QSFP
PIC 0		BUILTIN	BUILTIN	2X40GE QSFP
Xcvr 0	REV 01	740-032986	QB220708	QSFP+-40G-SR4
Xcvr 1	REV 01	740-032986	QB220735	QSFP+-40G-SR4
MIC 1	REV 03	750-036233	ZL2028	2X40GE QSFP
PIC 2		BUILTIN	BUILTIN	2X40GE QSFP
Xcvr 0	REV 01	740-032986	QB220727	QSFP+-40G-SR4
Xcvr 1	REV 01	740-032986	QB220715	QSFP+-40G-SR4
FPC 5	REV 11	750-037358	CAAE2196	MPC Type 4-1
CPU	REV 08	711-035209	CAAD9074	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA062S	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AMA062P	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AMA052R	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA0632	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00564	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	193363A00229	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00363	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00278	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA04CC	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AD0927A001W	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AMA04N2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA062U	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00491	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	183363A01511	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00565	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00405	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA07QX	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AMA06MS	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00318	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	193363A00402	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00174	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	193363A00388	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00377	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00234	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA062T	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	193363A00550	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00364	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA0630	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00509	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	193363A00459	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	113363A00191	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00352	SFP+-10G-SR
FPC 6	REV 33	750-028467	CAAF5552	MPC 3D 16x 10GE
CPU	REV 11	711-029089	CAAH7601	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AD0927A0036	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AD0927A003M	SFP+-10G-SR

Xcvr 2	REV 01	740-021308	AD0927A003G	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AD0927A0031	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	193363A00331	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	193363A00325	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00417	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A02509	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	T09K75140	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11A04356	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01952	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01914	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	T09K75157	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	T09K75194	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01926	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01936	SFP+-10G-SR
FPC 7	REV 16	750-037358	CAAL1012	MPC Type 4-1
CPU	REV 08	711-035209	CAAJ3851	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA04NK	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11F00260	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11E02192	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA04CP	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AJ40JJK	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11F00238	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B10M00275	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00211	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	B11D05577	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11G00586	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AMA08B7	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA04Q0	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11D05840	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11E00467	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11E00029	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	19T511101712	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	193363A00568	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B10M00166	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B10M00212	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11D05823	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	03DZ06A01005	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	03DZ06A01003	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	03DZ06A01009	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	03DZ06A01004	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-021308	03DZ06A01017	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	03DZ06A01016	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	03DZ06A01024	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	03DZ06A01008	SFP+-10G-SR
Xcvr 4	REV 01	740-030658	AD0946A02UH	SFP+-10G-USR
Xcvr 5	REV 01	740-021308	T09J67913	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AD0837ES09G	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	03DZ06A01015	SFP+-10G-SR
FPC 8	REV 03	750-045372	CAAD3111	MPCE Type 3D
CPU	REV 08	711-035209	CAAD8033	HMPC PMB 2G
MIC 0	REV 03	750-036233	ZL2032	2X40GE QSFP
PIC 0		BUILTIN	BUILTIN	2X40GE QSFP
Xcvr 0	REV 01	740-032986	QB230273	QSFP+-40G-SR4
Xcvr 1	REV 01	740-032986	QB230254	QSFP+-40G-SR4

MIC 1	REV 03	750-036233	ZL2021	2X40GE QSFP
PIC 2		BUILTIN	BUILTIN	2X40GE QSFP
Xcvr 0	REV 01	740-032986	QB390962	QSFP+-40G-SR4
Xcvr 1	REV 01	740-032986	QB390960	QSFP+-40G-SR4
FPC 9	REV 09	750-037355	CAAF1531	MPC Type 4-2
CPU	REV 08	711-035209	CAAB9927	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	193363A00525	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	193363A00504	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00368	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJ40JSS	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	123363A00042	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B10M00023	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJ802EM	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11E02348	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
ADC 0	REV 13	750-043596	ABBX5532	Adapter Card
ADC 1	REV 13	750-043596	ABBX5550	Adapter Card
ADC 2	REV 13	750-043596	ABBX5571	Adapter Card
ADC 3	REV 13	750-043596	ABBX5568	Adapter Card
ADC 4	REV 13	750-043596	ABBX5556	Adapter Card
ADC 5	REV 13	750-043596	ABBX5553	Adapter Card
ADC 6	REV 13	750-043596	ABBX5541	Adapter Card
ADC 7	REV 13	750-043596	ABBX5578	Adapter Card
ADC 8	REV 13	750-043596	ABBX5560	Adapter Card
ADC 9	REV 07	750-043596	ABBV7188	Adapter Card
Fan Tray 0	REV 03	760-046960	ACAY0127	172mm FanTray - 6 Fans
Fan Tray 1	REV 2A	760-046960	ACAY0068	172mm FanTray - 6 Fans
Fan Tray 2	REV 2A	760-046960	ACAY0072	172mm FanTray - 6 Fans
Fan Tray 3	REV 2A	760-046960	ACAY0070	172mm FanTray - 6 Fans

show chassis hardware extensive (MX2010 Router)

```

user@host > show chassis hardware extensive
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis
Jedec Code:   0x7fb0          EEPROM Version: 0x02
S/N:          JN11E233DAFK
Assembly ID:  0x0557          Assembly Version: 00.00
Date:         00-00-0000      Assembly Flags:  0x00
ID: MX2010
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 05 57 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: 4a 4e 31 31 45 32 33 33 44 41 46 4b 00 00 00 00
Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane     REV 26   750-044636  ABAB9357      Lower Backplane
Jedec Code:   0x7fb0          EEPROM Version: 0x02
P/N:         750-044636      S/N:          S/N ABAB9357
Assembly ID:  0x0b66          Assembly Version: 01.26
Date:         08-28-2012      Assembly Flags:  0x00
Version:      REV 26         CLEI Code:     PROTOXCLEI

```

```

ID: Lower Backplane          FRU Model Number:  PROTO-ASSEMBLY
Board Information Record:
  Address 0x00: ad 01 08 00 2c 21 72 70 a0 00 ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 66 01 1a 52 45 56 20 32 36 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 36 33 36 00 00
  Address 0x20: 53 2f 4e 20 41 42 41 42 39 33 35 37 00 1c 08 07
  Address 0x30: dc ff ff ff ad 01 08 00 2c 21 72 70 a0 00 ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff
Midplane 1          REV 01    711-044557    ABAB8643          Upper Backplane
Jedec Code: 0x7fb0          EEPROM Version: 0x01
P/N: 711-044557          S/N: S/N ABAB8643
Assembly ID: 0x0b65          Assembly Version: 01.01
Date: 07-27-2012          Assembly Flags: 0x00
Version: REV 01
ID: Upper Backplane
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 01 ff 0b 65 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 34 34 35 35 37 00 00
  Address 0x20: 53 2f 4e 20 41 42 41 42 38 36 34 33 00 1b 07 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
PMP          REV 04    711-032426    ACAJ1677          Power Midplane
Jedec Code: 0x7fb0          EEPROM Version: 0x01
P/N: 711-032426          S/N: S/N ACAJ1677
Assembly ID: 0x045d          Assembly Version: 01.04
Date: 07-20-2012          Assembly Flags: 0x00
Version: REV 04
ID: Power Midplane
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 5d 01 04 52 45 56 20 30 34 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 33 32 34 32 36 00 00
  Address 0x20: 53 2f 4e 20 41 43 41 4a 31 36 37 37 00 14 07 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
  Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM Board          REV 08    760-044634    ABBV9726          Front Panel Display
Jedec Code: 0x7fb0          EEPROM Version: 0x02
P/N: 760-044634          S/N: S/N ABBV9726
Assembly ID: 0x0b64          Assembly Version: 01.08
Date: 09-10-2012          Assembly Flags: 0x00
Version: REV 08          CLEI Code: IPMYA4EJRA
ID: Front Panel Display          FRU Model Number: MX2010-CRAFT-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 64 01 08 52 45 56 20 30 38 00 00
  Address 0x10: 00 00 00 00 37 36 30 2d 30 34 34 36 33 34 00 00
  Address 0x20: 53 2f 4e 20 41 42 42 56 39 37 32 36 00 0a 09 07

```



```

Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 4d 59 41 34 45 4a 52 41 4d
Address 0x50: 58 32 30 31 30 2d 43 52 41 46 54 2d 53 00 00 00
Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 93 ff ff ff ff ff ff ff ff ff ff ff ff
PSM 0          REV 01   740-045050   1E02224000P   DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version: 0x02
P/N:           740-045050      S/N:           1E02224000P
Assembly ID:   0x0478          Assembly Version: 01.01
Date:          12-06-2012      Assembly Flags: 0x00
Version:       REV 01          CLEI Code:     XXXXXXXXXX
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-HC-DC-S-A
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 30 35 30 00 00
Address 0x20: 31 45 30 32 32 32 34 30 30 30 50 00 00 06 0c 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 58 58 58 58 58 58 58 58 58 58 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 48 43 2d 44 43 2d
Address 0x60: 53 2d 41 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 4a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 1          REV 01   740-045050   1E02224000M   DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version: 0x02
P/N:           740-045050      S/N:           1E02224000M
Assembly ID:   0x0478          Assembly Version: 01.01
Date:          12-06-2012      Assembly Flags: 0x00
Version:       REV 01          CLEI Code:     XXXXXXXXXX
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-HC-DC-S-A
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 30 35 30 00 00
Address 0x20: 31 45 30 32 32 32 34 30 30 30 4d 00 00 06 0c 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 58 58 58 58 58 58 58 58 58 58 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 48 43 2d 44 43 2d
Address 0x60: 53 2d 41 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 4a 00 00 00 00 00 00 00 00 00 00 00 00
...
PDM 0          REV 01   740-045234   1E262250067   DC Power Dist Module
Jedec Code:    0x7fb0          EEPROM Version: 0x02
P/N:           740-045234      S/N:           1E262250067
Assembly ID:   0x047b          Assembly Version: 01.01
Date:          06-28-2012      Assembly Flags: 0x00
Version:       REV 01          CLEI Code:     IPUPAJSKAA
ID: DC Power Dist Module      FRU Model Number: MX2000-PDM-DC-S-A
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 7b 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 32 33 34 00 00
Address 0x20: 31 45 32 36 32 32 35 30 30 36 37 00 00 1c 06 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4a 53 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 44 4d 2d 44 43 2d 53 2d 41
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff

```

```

Address 0x70: ff ff ff 89 00 00 00 00 00 00 00 00 00 00 00 00
Routing Engine 0 REV 02 740-041821 9009099704 RE-S-1800x4
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 740-041821 S/N: 9009099704
Assembly ID: 0x09c0 Assembly Version: 01.02
Date: 03-15-2012 Assembly Flags: 0x00
Version: REV 02
ID: RE-S-1800x4 FRU Model Number: RE-S-1800X4-16G-S
Board Information Record:
Address 0x00: 54 32 30 32 37 44 41 2d 34 34 47 42 23 41 23 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 c0 01 02 52 45 56 20 30 32 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 31 38 32 31 00 00
Address 0x20: 39 30 30 39 30 39 39 37 30 34 00 00 00 0f 03 07
Address 0x30: dc ff ff ff 54 32 30 32 37 44 41 2d 34 34 47 42
Address 0x40: 23 41 23 00 01 00 00 00 00 00 00 00 00 00 00 52
Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 8c ff ff ff ff ff ff ff ff ff ff ff ff
ad0 3831 MB UGB30SFA4000T1 SFA4000T1 00000651 Compact Flash
ad1 30533 MB UGB94BPH32H0S1-KCI 11000019592 Disk 1
usb0 (addr 1) EHCI root hub 0 Intel uhub0
usb0 (addr 2) product 0x0020 32 vendor 0x8087 uhub1
DIMM 0 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 1 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 2 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 3 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1 REV 02 740-041821 9009099706 RE-S-1800x4
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 740-041821 S/N: 9009099706
Assembly ID: 0x09c0 Assembly Version: 01.02
Date: 02-23-2012 Assembly Flags: 0x00
Version: REV 02
ID: RE-S-1800x4 FRU Model Number: RE-S-1800X4-16G-S
Board Information Record:
Address 0x00: 54 32 30 32 37 44 41 2d 34 34 47 42 23 41 23 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 c0 01 02 52 45 56 20 30 32 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 31 38 32 31 00 00
Address 0x20: 39 30 30 39 30 39 39 37 30 36 00 00 00 17 02 07
Address 0x30: dc ff ff ff 54 32 30 32 37 44 41 2d 34 34 47 42
Address 0x40: 23 41 23 00 01 00 00 00 00 00 00 00 00 00 00 52
Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 8c ff ff ff ff ff ff ff ff ff ff ff ff
ad0 3998 MB Virtium - TuffDrive VCF P1T0200262860208 114 Compact Flash
ad1 30533 MB UGB94ARF32H0S3-KC UNIGEN-499551-000404 Disk 1
CB 0 REV 13 750-040257 CAAF8436 Control Board
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 750-040257 S/N: S/N CAAF8436
Assembly ID: 0x0b26 Assembly Version: 01.13
Date: 08-29-2012 Assembly Flags: 0x00
Version: REV 13 CLEI Code: PROTOXCLEI
ID: Control Board FRU Model Number: PROTO-ASSEMBLY
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 26 01 0d 52 45 56 20 31 33 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 34 30 32 35 37 00 00
Address 0x20: 53 2f 4e 20 43 41 41 46 38 34 33 36 00 1d 08 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff

```

```

Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff

...
SPMB 0          REV 02    711-041855    ABBV3825          PMB Board
Jedec Code:    0x7fb0          EEPROM Version:    0x01
P/N:          711-041855          S/N:          S/N ABBV3825
Assembly ID:   0x0b29          Assembly Version: 01.02
Date:         08-14-2012        Assembly Flags: 0x00
Version:      REV 02
ID: PMB Board
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 0b 29 01 02 52 45 56 20 30 32 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 34 31 38 35 35 00 00
Address 0x20: 53 2f 4e 20 41 42 42 56 33 38 32 35 00 0e 08 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00

...
SFB 0          REV 05    711-044466    ABBX5682          Switch Fabric Board
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:          711-044466          S/N:          S/N ABBX5682
Assembly ID:   0x0b25          Assembly Version: 01.05
Date:         09-07-2012        Assembly Flags: 0x00
Version:      REV 05          CLEI Code:      PROTOXCLEI
ID: Switch Fabric Board          FRU Model Number: PROTO-ASSEMBLY
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 25 01 05 52 45 56 20 30 35 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 34 34 36 36 00 00
Address 0x20: 53 2f 4e 20 41 42 42 58 35 36 38 32 00 07 09 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c2 00 00 00 01 00 00 00 00 00 00 48 00

...
FPC 0          REV 09    750-037355    CAAF0924          MPC Type 4-2
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:          750-037355          S/N:          S/N CAAF0924
Assembly ID:   0x0b4e          Assembly Version: 01.09
Date:         05-21-2012        Assembly Flags: 0x00
Version:      REV 09          CLEI Code:      PROTOXCLEI
ID: MPC Type 4-2          FRU Model Number: MPC4E-2CGE-8XGE
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 4e 01 09 52 45 56 20 30 39 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 33 35 35 00 00
Address 0x20: 53 2f 4e 20 43 41 41 46 30 39 32 34 00 15 05 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 4d
Address 0x50: 50 43 34 45 2d 32 43 47 45 2d 38 58 47 45 00 00
Address 0x60: 00 00 00 00 00 00 30 39 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c6 ff ff ff ff ff ff ff ff ff ff ff ff

```

```

CPU          REV 08   711-035209   CAAB9842           HMPC PMB 2G
Jedec Code:  0x7fb0           EEPROM Version:  0x01
P/N:         711-035209       S/N:           S/N CAAB9842
Assembly ID: 0x0b04           Assembly Version: 01.08
Date:        05-17-2012       Assembly Flags: 0x00
Version:     REV 08
ID: HMPC PMB 2G
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 0b 04 01 08 52 45 56 20 30 38 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 35 32 30 39 00 00
Address 0x20: 53 2f 4e 20 43 41 41 42 39 38 34 32 00 11 05 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
PIC 0          BUILTIN          BUILTIN          4x10GE SFPP
Jedec Code:  0x0000           EEPROM Version:  0x00
P/N:         BUILTIN          S/N:           BUILTIN
Assembly ID: 0x0a53           Assembly Version: 00.00
Date:        00-00-0000       Assembly Flags: 0x00
ID: 4x10GE SFPP
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 00 00 00 00 0a 53 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 4d 58 43 00
Address 0x20: 42 55 49 4c 54 49 4e 00 4d 58 43 00 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 c0 02 ae 64 00 00 00 00 0a 52 00 00
Xcvr 0      REV 01   740-021308   19T511101656       SFP+-10G-SR
Xcvr 1      REV 01   740-031980   AMA04RU             SFP+-10G-SR
Xcvr 2      REV 01   740-031980   193363A00558       SFP+-10G-SR
Xcvr 3      REV 01   740-031980   B10M00202          SFP+-10G-SR
...
ADC 0      REV 13   750-043596   ABBX5532           Adapter Card
Jedec Code:  0x7fb0           EEPROM Version:  0x02
P/N:         750-043596       S/N:           S/N ABBX5532
Assembly ID: 0x0b3d           Assembly Version: 01.13
Date:        09-12-2012       Assembly Flags: 0x00
Version:     REV 13          CLEI Code:      IPUCBA8CAA
ID: Adapter Card          FRU Model Number: MX2000-LC-ADAPTER
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 3d 01 0d 52 45 56 20 31 33 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 34 33 35 39 36 00 00
Address 0x20: 53 2f 4e 20 41 42 42 58 35 35 33 32 00 0c 09 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 43 42 41 38 43 41 41 4d
Address 0x50: 58 32 30 30 30 2d 4c 43 2d 41 44 41 50 54 45 52
Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 3a 00 00 00 00 00 00 00 00 00 00 00 00
...

```

show chassis hardware models (MX2010 Router)

```
user@host > show chassis hardware models
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	FRU model number
FPM Board	REV 06	711-032349	ZX8744	711-032349
PSM 4	REV 0C	740-033727	VK00254	000000000000000000000000
PSM 5	REV 0B	740-033727	VG00015	000000000000000000000000
PSM 6	REV 0B	740-033727	VH00097	000000000000000000000000
PSM 7	REV 0C	740-033727	VJ00151	000000000000000000000000
PSM 8	REV 0C	740-033727	VJ00149	000000000000000000000000
PDM 0	REV 0B	740-038109	WA00008	
PDM 1	REV 0B	740-038109	WA00014	
Routing Engine 0	REV 02	740-041821	9009094134	RE-S-1800X4-16G-S
Routing Engine 1	REV 02	740-041821	9009094141	RE-S-1800X4-16G-S
CB 0	REV 08	750-040257	CAAB3491	750-040257
CB 1	REV 08	750-040257	CAAB3489	750-040257
SFB 0	REV 06	711-032385	ZV1828	711-032385
SFB 1	REV 07	711-032385	ZZ2568	711-032385
SFB 2	REV 07	711-032385	ZZ2563	711-032385
SFB 3	REV 07	711-032385	ZZ2564	711-032385
SFB 4	REV 07	711-032385	ZZ2580	711-032385
SFB 5	REV 07	711-032385	ZZ2579	711-0323856
SFB 6	REV 07	711-032385	CAAB4882	711-044170
SFB 7	REV 07	711-032385	CAAB4898	711-044170
FPC 0	REV 33	750-028467	CAAB1919	MPC-3D-16XGE-SFPP
FPC 1	REV 21	750-033205	ZG5027	MX-MPC3-3D
MIC 0	REV 03	750-033307	ZV6299	MIC3-3D-10XGE-SFPP
MIC 1	REV 03	750-033307	ZV6268	MIC3-3D-10XGE-SFPP
FPC 8	REV 22	750-031089	ZT9746	MX-MPC2-3D
MIC 0	REV 26	750-028392	ABBS1150	MIC-3D-20GE-SFP
MIC 1	REV 26	750-028387	ABBR9582	MIC-3D-4XGE-XFP
FPC 9	REV 11	750-036284	ZL3591	MPCE-3D-16XGE-SFPP
ADC 0	REV 05	750-043596	CAAC2073	750-043596
ADC 1	REV 01	750-043596	ZV4117	750-043596
ADC 8	REV 01	750-043596	ZV4107	750-043596
ADC 9	REV 02	750-043596	ZW1555	750-043596
Fan Tray 0	REV 2A	760-046960	ACAY0015	
Fan Tray 1	REV 2A	760-046960	ACAY0019	
Fan Tray 2	REV 2A	760-046960	ACAY0020	
Fan Tray 3	REV 2A	760-046960	ACAY0021	

show chassis hardware clei-models (MX2010 Routers)

```
user@host > show chassis hardware clei-models
```

```
Hardware inventory:
```

Item	Version	Part number	CLEI code	FRU model number
FPM Board	REV 06	711-032349	PROTOXCLEI	711-032349
PSM 4	REV 0C	740-033727	0000000000	000000000000000000000000
PSM 5	REV 0B	740-033727	0000000000	000000000000000000000000
PSM 6	REV 0B	740-033727	0000000000	000000000000000000000000
PSM 7	REV 0C	740-033727	0000000000	000000000000000000000000
PSM 8	REV 0C	740-033727	0000000000	000000000000000000000000
PDM 0	REV 0B	740-038109		
PDM 1	REV 0B	740-038109		
Routing Engine 0	REV 02	740-041821		RE-S-1800X4-16G-S
Routing Engine 1	REV 02	740-041821		RE-S-1800X4-16G-S
CB 0	REV 08	750-040257	PROTOXCLEI	750-040257
CB 1	REV 08	750-040257	PROTOXCLEI	750-040257
SFB 0	REV 06	711-032385	PROTOXCLEI	711-032385
SFB 1	REV 07	711-032385	PROTOXCLEI	711-032385

SFB 2	REV 07	711-032385	PROTOXCLEI	711-032385
SFB 3	REV 07	711-032385	PROTOXCLEI	711-032385
SFB 4	REV 07	711-032385	PROTOXCLEI	711-032385
SFB 5	REV 07	711-032385	PROTOXCLEI	711-0323856
SFB 6	REV 07	711-032385	PROTOXCLEI	711-044170
SFB 7	REV 07	711-032385	PROTOXCLEI	711-044170
FPC 0	REV 33	750-028467		MPC-3D-16XGE-SFPP
FPC 1	REV 21	750-033205		MX-MPC3-3D
MIC 0	REV 03	750-033307	PROTOXCLEI	MIC3-3D-10XGE-SFPP
MIC 1	REV 03	750-033307	PROTOXCLEI	MIC3-3D-10XGE-SFPP
FPC 8	REV 22	750-031089	COUIBAYBAA	MX-MPC2-3D
MIC 0	REV 26	750-028392	COUIA15BAA	MIC-3D-20GE-SFP
MIC 1	REV 26	750-028387	COUIA16BAA	MIC-3D-4XGE-XFP
FPC 9	REV 11	750-036284	CMUIACGBAA	MPCE-3D-16XGE-SFPP
ADC 0	REV 05	750-043596	PROTOXCLEI	750-043596
ADC 1	REV 01	750-043596	PROTOXCLEI	750-043596
ADC 8	REV 01	750-043596	PROTOXCLEI	750-043596
ADC 9	REV 02	750-043596	PROTOXCLEI	750-043596
Fan Tray 0	REV 2A	760-046960		
Fan Tray 1	REV 2A	760-046960		
Fan Tray 2	REV 2A	760-046960		
Fan Tray 3	REV 2A	760-046960		

show chassis hardware (MX2020 Router)

```
user@host > show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN11E2227AFJ	MX2020
Midplane	REV 27	750-040240	ABAB9384	Lower Power Midplane
Midplane 1	REV 04	711-032386	ABAB9386	Upper Backplane
PMP 1	REV 05	711-032428	ACAJ1579	Upper Power Midplane
PMP 0	REV 04	711-032426	ACAJ1524	Lower Power Midplane
FPM Board	REV 06	760-040242	ABBT8837	Front Panel Display
PSM 0	REV 01	740-045050	1E022240056	DC 52V Power Supply
Module				
PSM 1	REV 01	740-045050	1E022240054	DC 52V Power Supply
Module				
PSM 2	REV 01	740-045050	1E02224005H	DC 52V Power Supply
Module				
PSM 3	REV 01	740-045050	1E022240053	DC 52V Power Supply
Module				
PSM 4	REV 01	740-045050	1E02224004K	DC 52V Power Supply
Module				
PSM 7	REV 01	740-045050	1E02224006W	DC 52V Power Supply
Module				
PSM 8	REV 01	740-045050	1E022240062	DC 52V Power Supply
Module				
PSM 9	REV 01	740-045050	1E02224005B	DC 52V Power Supply
Module				
PSM 10	REV 01	740-045050	1E02224005A	DC 52V Power Supply
Module				
PSM 11	REV 01	740-045050	1E022240052	DC 52V Power Supply
Module				
PSM 12	REV 01	740-045050	1E022240051	DC 52V Power Supply
Module				
PSM 13	REV 01	740-045050	1E022240058	DC 52V Power Supply
Module				
PSM 14	REV 01	740-045050	1E02224004L	DC 52V Power Supply
Module				
PSM 15	REV 01	740-045050	1E02224005M	DC 52V Power Supply

Module				
PSM 16	REV 01	740-045050	1E02224006S	DC 52V Power Supply
Module				
PSM 17	REV 01	740-045050	1E02224005Z	DC 52V Power Supply
Module				
PDM 0	REV 01	740-045234	1E012150033	DC Power Dist Module
PDM 1	REV 01	740-045234	1E012150027	DC Power Dist Module
PDM 2	REV 01	740-045234	1E012150028	DC Power Dist Module
PDM 3	REV 01	740-045234	1E012150045	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009089704	RE-S-1800x4
Routing Engine 1	REV 02	740-041821	9009094138	RE-S-1800x4
CB 0	REV 14	750-040257	CAAF8430	Control Board
CB 1	REV 08	750-040257	CAAB3482	Control Board
SPMB 0	REV 01	711-041855	ZS2290	PMB Board
SPMB 1	REV 02	711-041855	CAAA6141	PMB Board
SFB 0	REV 03	711-044466	ABBV6789	Switch Fabric Board
SFB 1	REV 05	711-044466	ABBX5666	Switch Fabric Board
SFB 2	REV 05	711-044466	ABBX5678	Switch Fabric Board
SFB 3	REV 05	711-044466	ABBX5687	Switch Fabric Board
SFB 4	REV 05	711-044466	ABBX5609	Switch Fabric Board
SFB 5	REV 05	711-044466	ABBX5675	Switch Fabric Board
SFB 6	REV 03	711-044466	ABBV6805	Switch Fabric Board
SFB 7	REV 05	711-044466	ABBX5701	Switch Fabric Board
FPC 0	REV 30	750-028467	ABBN0284	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0507	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00990	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E04357	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01327	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04375	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02760	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02904	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E03963	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00756	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04418	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01077	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01128	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01253	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E01140	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01626	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01075	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01177	SFP+-10G-USR
FPC 1	REV 30	750-028467	ABBN0208	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ1084	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04745	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01570	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E04388	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01439	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04739	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01869	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01675	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01901	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01346	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01288	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01824	SFP+-10G-USR

Xcvr 3	REV 01	740-030658	B11E04312	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02811	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E03847	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01495	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01265	SFP+-10G-USR
FPC 2	REV 30	750-028467	ZM5111	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ZP6607	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LJA	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MFZ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKL	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KF4	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80FBJ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MM2	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LJV	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NXV	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N1H	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLS	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FL5	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL9	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NG2	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80KDU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80MG1	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80MM0	SFP+-10G-SR
FPC 3	REV 30	750-028467	ABB0302	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABB0495	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01581	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01176	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01251	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02752	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00786	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01020	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01023	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02819	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02812	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11D04437	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01279	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01333	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00978	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01018	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01784	SFP+-10G-USR
Xcvr 3	REV 01	740-031980	AK80NKP	SFP+-10G-SR
FPC 4	REV 30	750-028467	ABB0308	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABB11095	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04305	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01147	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01195	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01743	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01892	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02880	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00725	SFP+-10G-USR

Xcvr 3	REV 01	740-030658	B11E01057	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02816	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11C04501	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E02764	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00789	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01250	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02847	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00787	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E03803	SFP+-10G-USR
FPC 5	REV 30	750-028467	ABBN0316	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABB11082	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00523	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01848	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01865	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00540	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00422	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00428	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K00423	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01855	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K01847	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00526	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K00529	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00525	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00425	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00530	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01851	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00528	SFP+-10G-SR
FPC 6	REV 32	750-028467	ABBN6832	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN6534	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MB4	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FQ6	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N1F	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLQ	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80KDR	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FGJ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N5G	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KD8	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LET	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80N1X	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NRF	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL2	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N3D	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MRB	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LEQ	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LER	SFP+-10G-SR
FPC 7	REV 32	750-028467	ABBN6811	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7288	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NK8	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80LJG	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LBU	SFP+-10G-SR

Xcvr 3	REV 01	740-031980	AK80N21	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEU	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLM	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NL6	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LES	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEN	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80ME0	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LMG	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80MM1	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MG7	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80KF9	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NRQ	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLE	SFP+-10G-SR
FPC 8	REV 23	750-028467	YN2977	MPC 3D 16x 10GE
CPU	REV 10	711-029089	YP1856	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00875	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00851	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00772	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00882	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00735	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00169	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00726	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00077	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00168	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00676	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00732	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00091	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00725	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00642	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00871	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00853	SFP+-10G-SR
FPC 9	REV 32	750-028467	ABB6798	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6556	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	9ZDZ06A00055	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00239	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AD0915E003K	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AD0915E003A	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MRC	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NL5	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKN	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N3U	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N1T	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ808DJ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NG4	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80FND	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80FKQ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLT	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKR	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LKM	SFP+-10G-SR
FPC 10	REV 32	750-028467	ABBN6813	MPC 3D 16x 10GE

CPU	REV 10	711-029089	ABBK6542	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NA3	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLF	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80MRH	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KE4	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00030	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80L9H	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80ME8	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLR	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NG1	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MCA	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LFC	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LEM	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N9X	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80LAC	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LF2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N8T	SFP+-10G-SR
FPC 11	REV 30	750-028467	ABBN0281	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0526	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01326	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E03973	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00950	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00674	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00775	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E04461	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01074	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02821	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04501	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00757	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01623	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01022	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04359	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02751	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E02736	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01178	SFP+-10G-USR
FPC 12	REV 32	750-028467	ABBN6796	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7259	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K01856	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01853	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01863	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02863	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02668	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02881	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01671	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02627	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02725	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02692	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02730	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03081	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+

Xcvr 0	REV 01	740-031980	163363A02736	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02568	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02747	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02579	SFP+-10G-SR
FPC 13	REV 30	750-028467	ABBN0270	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABB0966	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NL1	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NXW	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KD2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80FMD	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NKQ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MGH	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N38	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL7	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEL	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NKD	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KCY	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LHK	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80M5J	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MBE	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NLG	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LFH	SFP+-10G-SR
FPC 14	REV 32	750-028467	ABBN6790	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN6515	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LZM	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MCC	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KCM	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KE0	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021310	C10F99155	SFP+-10G-LRM
Xcvr 1	REV 01	740-021310	C10F99049	SFP+-10G-LRM
Xcvr 2	REV 01	740-021310	C10F99128	SFP+-10G-LRM
Xcvr 3	REV 01	740-021310	C10F99169	SFP+-10G-LRM
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LF3	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02597	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A03060	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03057	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEX	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FEU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FNM	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AJQQQ5G	SFP+-10G-SR
FPC 15	REV 32	750-028467	ABBN6791	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7289	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00424	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01849	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01862	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01852	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00427	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00430	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01854	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00426	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+

Xcvr 0	REV 01	740-031980	B11K00429	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01864	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01850	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00522	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E01144	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00985	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00796	SFP+-10G-USR
Xcvr 3	REV 01	740-031980	B11K01866	SFP+-10G-SR
FPC 16	REV 30	750-028467	ABBM4592	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0465	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01435	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01052	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01328	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01254	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02738	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02881	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01624	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00889	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02883	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00681	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E04306	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02813	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01801	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02753	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01156	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04324	SFP+-10G-USR
FPC 17	REV 32	750-028467	ABBN6810	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7237	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02638	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02082	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01674	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03058	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03048	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02729	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02566	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02567	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02878	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02739	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01959	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02660	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02731	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02588	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02673	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02654	SFP+-10G-SR
FPC 18	REV 30	750-028467	ABBM4739	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0487	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02569	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02886	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A03082	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	133363A00297	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+

Xcvr 0	REV 01	740-031980	163363A02726	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A03050	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02884	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03076	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02581	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02873	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02582	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03083	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031981	UL70BU6	SFP+-10G-LR
Xcvr 1	REV 01	740-031981	UL50QC6	SFP+-10G-LR
Xcvr 2	REV 01	740-031981	UL708N6	SFP+-10G-LR
Xcvr 3	REV 01	740-031981	UL603KK	SFP+-10G-LR
FPC 19	REV 32	750-028467	ABBN6827	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6508	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A01688	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A01724	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01773	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02593	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03061	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A03056	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02669	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03070	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02572	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02697	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02585	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03052	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02591	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02649	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02577	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02698	SFP+-10G-SR
ADC 0	REV 13	750-043596	ABBX5561	Adapter Card
ADC 1	REV 13	750-043596	ABBX5546	Adapter Card
ADC 2	REV 13	750-043596	ABBX5535	Adapter Card
ADC 3	REV 13	750-043596	ABBX5552	Adapter Card
ADC 4	REV 13	750-043596	ABBX5581	Adapter Card
ADC 5	REV 13	750-043596	ABBX5545	Adapter Card
ADC 6	REV 13	750-043596	ABBX5554	Adapter Card
ADC 7	REV 07	750-043596	ABBV7194	Adapter Card
ADC 8	REV 07	750-043596	ABBV7251	Adapter Card
ADC 9	REV 07	750-043596	ABBV7202	Adapter Card
ADC 10	REV 13	750-043596	ABBX5538	Adapter Card
ADC 11	REV 13	750-043596	ABBX5566	Adapter Card
ADC 12	REV 13	750-043596	ABBX5542	Adapter Card
ADC 13	REV 13	750-043596	ABBX5539	Adapter Card
ADC 14	REV 13	750-043596	ABBX5555	Adapter Card
ADC 15	REV 13	750-043596	ABBX5557	Adapter Card
ADC 16	REV 13	750-043596	ABBX5536	Adapter Card
ADC 17	REV 13	750-043596	ABBX5559	Adapter Card
ADC 18	REV 13	750-043596	ABBX5537	Adapter Card
ADC 19	REV 11	750-043596	ABBW5685	Adapter Card
Fan Tray 0	REV 2A	760-046960	ACAY0030	172mm FanTray - 6 Fans
Fan Tray 1	REV 2A	760-046960	ACAY0039	172mm FanTray - 6 Fans
Fan Tray 2	REV 2A	760-046960	ACAY0033	172mm FanTray - 6 Fans
Fan Tray 3	REV 2A	760-046960	ACAY0062	172mm FanTray - 6 Fans

show chassis hardware detail (MX2020 Router)

```

user@host> show chassis hardware detail
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN11E2227AFJ  MX2020
Midplane      REV 27   750-040240   ABAB9384      Lower Power Midplane
Midplane 1    REV 04   711-032386   ABAB9386      Upper Backplane
PMP 1         REV 05   711-032428   ACAJ1821      Upper Power Midplane
PMP 0         REV 04   711-032426   ACAJ1524      Lower Power Midplane
FPM Board     REV 06   760-040242   ABBT8837      Front Panel Display
PSM 0         REV 01   740-045050   1E02224006G   DC 52V Power Supply
Module
PSM 1         REV 01   740-045050   1E022240053   DC 52V Power Supply
Module
PSM 2         REV 01   740-045050   1E02224004K   DC 52V Power Supply
Module
PSM 3         REV 01   740-045050   1E022240056   DC 52V Power Supply
Module
PSM 4         REV 01   740-045050   1E022240054   DC 52V Power Supply
Module
PSM 5         REV 01   740-045050   1E02224005H   DC 52V Power Supply
Module
PSM 6         REV 01   740-045050   1E02224006S   DC 52V Power Supply
Module
PSM 7         REV 01   740-045050   1E02224005M   DC 52V Power Supply
Module
PSM 8         REV 01   740-045050   1E022240062   DC 52V Power Supply
Module
PSM 9         REV 03   740-045050   1EDB2350095   DC 52V Power Supply
Module
PSM 10        REV 03   740-045050   1EDB235009L   DC 52V Power Supply
Module
PSM 11        REV 03   740-045050   1EDB2350092   DC 52V Power Supply
Module
PSM 12        REV 03   740-045050   1EDB23500AT   DC 52V Power Supply
Module
PSM 13        REV 03   740-045050   1EDB2350094   DC 52V Power Supply
Module
PSM 15        REV 03   740-045050   1EDB235008X   DC 52V Power Supply
Module
PDM 0         REV 01   740-045234   1E012150033   DC Power Dist Module
PDM 1         REV 01   740-045234   1E012150027   DC Power Dist Module
PDM 2         REV 01   740-045234   1E262250072   DC Power Dist Module
Routing Engine 0 REV 02   740-041821   9009094138     RE-S-1800x4
  ad0      3998 MB   Virtium - TuffDisk VCF3 20110825A021D00000064 Compact Flash
  ad1      30533 MB UGB94ARF32H0S3-KC  UNIGEN-499551-000347 Disk 1
  usb0 (addr 1) EHCI root hub 0    Intel      uhub0
  usb0 (addr 2) product 0x0020 32 vendor 0x8087 uhub1
  DIMM 0      SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 1      SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 2      SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 3      SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1 REV 02   740-041821   9009089709     RE-S-1800x4
  ad0      3831 MB   UGB30SFA4000T1     SFA4000T1 00000113 Compact Flash
  ad1      30533 MB UGB94ARF32H0S3-KC  UNIGEN-478612-001044 Disk 1
CB 0         REV 08   750-040257   CAAB3482      Control Board
CB 1         REV 04   750-040257   ZT2864        Control Board
SPMB 0       REV 02   711-041855   CAAA6141      PMB Board
SPMB 1       REV 01   711-041855   ZS2275        PMB Board
SFB 0       REV 05   711-044466   ABBT2161      Switch Fabric Board

```

SFB 1	REV 05	711-044466	ABBT2159	Switch Fabric Board
SFB 2	REV 05	711-044466	ABBX3718	Switch Fabric Board
SFB 3	REV 05	711-044466	ABBT2152	Switch Fabric Board
SFB 4	REV 05	711-044466	ABBT2160	Switch Fabric Board
SFB 5	REV 05	711-044466	ABBT2145	Switch Fabric Board
SFB 6	REV 05	711-044466	ABBT2150	Switch Fabric Board
SFB 7	REV 05	711-044466	ABBT2163	Switch Fabric Board
FPC 0	REV 30	750-028467	ABBN0284	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0507	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00990	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E04357	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01327	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04375	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02760	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02904	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E03963	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00756	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04418	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01077	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01128	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01253	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E01140	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01626	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01075	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01177	SFP+-10G-USR
FPC 1	REV 30	750-028467	ABBN0308	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ1095	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04305	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01147	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01195	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01743	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01892	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02880	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00725	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01057	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02816	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11C04501	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E02764	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00789	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01250	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02847	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00787	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E03803	SFP+-10G-USR
FPC 2	REV 30	750-028467	ABBN0316	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ1082	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00523	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01848	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01865	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00540	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00422	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00428	SFP+-10G-SR

Xcvr 2	REV 01	740-031980	B11K00423	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01855	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K01847	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00526	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K00529	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00525	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00425	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00530	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01851	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00528	SFP+-10G-SR
FPC 3	REV 32	750-028467	ABBN6832	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6534	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MB4	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FQ6	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N1F	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLQ	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80KDR	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FGJ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N5G	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KDR	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LET	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80N1X	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NRF	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL2	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N3D	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MRB	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LEQ	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LER	SFP+-10G-SR
FPC 4	REV 32	750-028467	ABBN6811	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7288	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NK8	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80LJG	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LBU	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N21	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEU	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLM	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NL6	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LES	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEN	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80ME0	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LMG	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80MM1	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MG7	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80KF9	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NRQ	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLE	SFP+-10G-SR
FPC 5	REV 32	750-028467	ABBN6791	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7289	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00424	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01849	SFP+-10G-SR

Xcvr 2	REV 01	740-031980	B11K01862	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01852	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP
Xcvr 0	REV 01	740-031980	B11K00427	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00430	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01854	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00426	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00429	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01864	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01850	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00522	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E01144	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00985	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00796	SFP+-10G-USR
Xcvr 3	REV 01	740-031980	B11K01866	SFP+-10G-SR
FPC 6	REV 30	750-028467	ABBM4592	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0465	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01435	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01052	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01328	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01254	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02738	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02881	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01624	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00889	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02883	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00681	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E04306	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02813	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01801	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02753	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01156	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04324	SFP+-10G-USR
FPC 7	REV 32	750-028467	ABBN6810	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7237	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03058	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02082	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01674	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02638	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03048	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02729	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02566	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02567	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02878	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02739	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01959	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02660	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02731	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02588	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02673	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02654	SFP+-10G-SR

FPC 8	REV 30	750-028467	ABBM4739	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0487	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02569	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02886	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A03082	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	133363A00297	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02726	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A03050	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02884	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03076	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02581	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02873	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02582	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03083	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031981	UL70BU6	SFP+-10G-LR
Xcvr 1	REV 01	740-031981	UL50QC6	SFP+-10G-LR
Xcvr 2	REV 01	740-031981	UL708N6	SFP+-10G-LR
Xcvr 3	REV 01	740-031981	UL603KK	SFP+-10G-LR
FPC 9	REV 32	750-028467	ABBN6827	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6508	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A01688	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A01724	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01773	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02593	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03061	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A03056	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02669	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03070	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02572	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02697	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02585	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03052	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02591	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02649	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02577	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02698	SFP+-10G-SR
FPC 10	REV 30	750-028467	ABBN0302	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0495	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01581	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01176	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01251	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02752	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00786	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01020	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01023	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02819	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02812	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11D04437	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01279	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01333	SFP+-10G-USR

PIC 3			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00978	SFP+-10G-USR	
Xcvr 1	REV 01	740-030658	B11E01018	SFP+-10G-USR	
Xcvr 2	REV 01	740-030658	B11F01784	SFP+-10G-USR	
Xcvr 3	REV 01	740-031980	AK80NKP	SFP+-10G-SR	
FPC 11	REV 32	750-028467	ABBN6790	MPC 3D 16x 10GE	
CPU	REV 10	711-029089	ABBK6515	AMPC PMB	
PIC 0			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LZM	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	AK80MCC	SFP+-10G-SR	
Xcvr 2	REV 01	740-031980	AK80KCM	SFP+-10G-SR	
Xcvr 3	REV 01	740-031980	AK80KE0	SFP+-10G-SR	
PIC 1			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021310	C10F99155	SFP+-10G-LRM	
Xcvr 1	REV 01	740-021310	C10F99049	SFP+-10G-LRM	
Xcvr 2	REV 01	740-021310	C10F99128	SFP+-10G-LRM	
Xcvr 3	REV 01	740-021310	C10F99169	SFP+-10G-LRM	
PIC 2			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LF3	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	163363A02597	SFP+-10G-SR	
Xcvr 2	REV 01	740-031980	163363A03060	SFP+-10G-SR	
Xcvr 3	REV 01	740-031980	163363A03057	SFP+-10G-SR	
PIC 3			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEX	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	AK80FEU	SFP+-10G-SR	
Xcvr 2	REV 01	740-031980	AK80FNM	SFP+-10G-SR	
Xcvr 3	REV 01	740-021308	AJQQQ5G	SFP+-10G-SR	
FPC 12	REV 30	750-028467	ZM5111	MPC 3D 16x 10GE	
CPU	REV 10	711-029089	ZP6607	AMPC PMB	
PIC 0			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LJA	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	AK80MFZ	SFP+-10G-SR	
Xcvr 2	REV 01	740-031980	AK80NKL	SFP+-10G-SR	
Xcvr 3	REV 01	740-031980	AK80KF4	SFP+-10G-SR	
PIC 1			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80FBJ	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	AK80MM2	SFP+-10G-SR	
Xcvr 2	REV 01	740-031980	AK80LJV	SFP+-10G-SR	
Xcvr 3	REV 01	740-031980	AK80NXV	SFP+-10G-SR	
PIC 2			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N1H	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	AK80NLS	SFP+-10G-SR	
Xcvr 2	REV 01	740-031980	AK80FL5	SFP+-10G-SR	
Xcvr 3	REV 01	740-031980	AK80NL9	SFP+-10G-SR	
PIC 3			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NG2	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	AK80KDU	SFP+-10G-SR	
Xcvr 2	REV 01	740-031980	AK80MG1	SFP+-10G-SR	
Xcvr 3	REV 01	740-031980	AK80MM0	SFP+-10G-SR	
FPC 13	REV 30	750-028467	ABBN0208	MPC 3D 16x 10GE	
CPU	REV 10	711-029089	ABB11084	AMPC PMB	
PIC 0			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04745	SFP+-10G-USR	
Xcvr 1	REV 01	740-030658	B11F01570	SFP+-10G-USR	
Xcvr 2	REV 01	740-030658	B11E04388	SFP+-10G-USR	
Xcvr 3	REV 01	740-030658	B11F01439	SFP+-10G-USR	
PIC 1			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04739	SFP+-10G-USR	
Xcvr 1	REV 01	740-030658	B11F01869	SFP+-10G-USR	
Xcvr 2	REV 01	740-030658	B11F01675	SFP+-10G-USR	
Xcvr 3	REV 01	740-030658	B11F01901	SFP+-10G-USR	

PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01346	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01288	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01824	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04312	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02811	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E03847	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01495	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01265	SFP+-10G-USR
FPC 14	REV 23	750-028467	YN2977	MPC 3D 16x 10GE
CPU	REV 10	711-029089	YP1856	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00875	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00851	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00772	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00882	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00735	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00169	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00726	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00077	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00168	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00676	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00732	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00091	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00725	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00642	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00871	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00853	SFP+-10G-SR
FPC 15	REV 32	750-028467	ABBN6798	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6556	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	9ZDZ06A00055	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00239	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AD0915E003K	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AD0915E003A	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MRC	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NL5	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKN	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N3U	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N1T	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ808DJ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NG4	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80FND	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80FKQ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLT	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKR	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LKM	SFP+-10G-SR
FPC 16	REV 30	750-028467	ABBN0270	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ0966	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NL1	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NXW	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KD2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80FMD	SFP+-10G-SR

PIC 1			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NKQ		SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MGH		SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N38		SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL7		SFP+-10G-SR
PIC 2			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80M5J		SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NKD		SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KCY		SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LHK		SFP+-10G-SR
PIC 3			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEL		SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MBE		SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NLG		SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LFH		SFP+-10G-SR
FPC 17	REV 32	750-028467	ABBN6796		MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7259		AMPC PMB
PIC 0			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K01856		SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01853		SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01863		SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02863		SFP+-10G-SR
PIC 1			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02668		SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02881		SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01671		SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02627		SFP+-10G-SR
PIC 2			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02725		SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02692		SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02730		SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03081		SFP+-10G-SR
PIC 3			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02736		SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02568		SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02747		SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02579		SFP+-10G-SR
FPC 18	REV 30	750-028467	ABBN0281		MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0526		AMPC PMB
PIC 0			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01326		SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E03973		SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00950		SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00674		SFP+-10G-USR
PIC 1			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00775		SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E04461		SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01074		SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02821		SFP+-10G-USR
PIC 2			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04501		SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00757		SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01623		SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01022		SFP+-10G-USR
PIC 3			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04359		SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02751		SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E02736		SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01178		SFP+-10G-USR
FPC 19	REV 32	750-028467	ABBN6813		MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6542		AMPC PMB

PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NA3	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLF	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80MRH	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KE4	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00030	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80L9H	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80ME8	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLR	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NG1	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MCA	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LFC	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LEM	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N9X	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80LAC	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LF2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N8T	SFP+-10G-SR
ADC 0	REV 13	750-043596	ABBX5561	Adapter Card
ADC 1	REV 13	750-043596	ABBX5546	Adapter Card
ADC 2	REV 13	750-043596	ABBX5535	Adapter Card
ADC 3	REV 13	750-043596	ABBX5552	Adapter Card
ADC 4	REV 13	750-043596	ABBX5581	Adapter Card
ADC 5	REV 13	750-043596	ABBX5545	Adapter Card
ADC 6	REV 13	750-043596	ABBX5554	Adapter Card
ADC 7	REV 07	750-043596	ABBV7194	Adapter Card
ADC 8	REV 07	750-043596	ABBV7251	Adapter Card
ADC 9	REV 07	750-043596	ABBV7202	Adapter Card
ADC 10	REV 13	750-043596	ABBX5579	Adapter Card
ADC 11	REV 13	750-043596	ABBX5548	Adapter Card
ADC 12	REV 13	750-043596	ABBX5575	Adapter Card
ADC 13	REV 13	750-043596	ABBX5539	Adapter Card
ADC 14	REV 13	750-043596	ABBX5555	Adapter Card
ADC 15	REV 13	750-043596	ABBX5557	Adapter Card
ADC 16	REV 13	750-043596	ABBX5536	Adapter Card
ADC 17	REV 13	750-043596	ABBX5559	Adapter Card
ADC 18	REV 13	750-043596	ABBX5537	Adapter Card
ADC 19	REV 11	750-043596	ABBW5685	Adapter Card
Fan Tray 0	REV 04	760-046960	ACAY0090	172mm FanTray - 6 Fans
Fan Tray 1	REV 04	760-046960	ACAY0088	172mm FanTray - 6 Fans
Fan Tray 2	REV 04	760-046960	ACAY0089	172mm FanTray - 6 Fans
Fan Tray 3	REV 04	760-046960	ACAY0108	172mm FanTray - 6 Fans

show chassis hardware models (MX2020 Router)

```
user@host > show chassis hardware models
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 27	750-040240	ABAB9384	750-040240
FPM Board	REV 06	760-040242	ABBT8837	760-040242
PSM 0	REV 01	740-045050	1E02224006G	MX2000-PSM-HC-DC-S-A
PSM 1	REV 01	740-045050	1E022240053	MX2000-PSM-HC-DC-S-A
PSM 2	REV 01	740-045050	1E02224004K	MX2000-PSM-HC-DC-S-A
PSM 3	REV 01	740-045050	1E022240056	MX2000-PSM-HC-DC-S-A
PSM 4	REV 01	740-045050	1E022240054	MX2000-PSM-HC-DC-S-A
PSM 5	REV 01	740-045050	1E02224005H	MX2000-PSM-HC-DC-S-A
PSM 6	REV 01	740-045050	1E02224006S	MX2000-PSM-HC-DC-S-A
PSM 7	REV 01	740-045050	1E02224005M	MX2000-PSM-HC-DC-S-A
PSM 8	REV 01	740-045050	1E022240062	MX2000-PSM-HC-DC-S-A

PSM 9	REV 03	740-045050	1EDB2350095	MX2000-PSM-DC-S-A
PSM 10	REV 03	740-045050	1EDB235009L	MX2000-PSM-DC-S-A
PSM 11	REV 03	740-045050	1EDB2350092	MX2000-PSM-DC-S-A
PSM 12	REV 03	740-045050	1EDB23500AT	MX2000-PSM-DC-S-A
PSM 13	REV 03	740-045050	1EDB2350094	MX2000-PSM-DC-S-A
PSM 15	REV 03	740-045050	1EDB235008X	MX2000-PSM-DC-S-A
PDM 0	REV 01	740-045234	1E012150033	
PDM 1	REV 01	740-045234	1E012150027	
PDM 2	REV 01	740-045234	1E262250072	MX2000-PDM-DC-S-A
Routing Engine 0	REV 02	740-041821	9009094138	RE-S-1800X4-16G-S
Routing Engine 1	REV 02	740-041821	9009089709	RE-S-1800X4-16G-S
CB 0	REV 08	750-040257	CAAB3482	750-040257
CB 1	REV 04	750-040257	ZT2864	750-040257
SFB 0	REV 05	711-044466	ABBT2161	MX2000-SFB-S
SFB 1	REV 05	711-044466	ABBT2159	MX2000-SFB-S
SFB 2	REV 05	711-044466	ABBX3718	MX2000-SFB-S
SFB 4	REV 05	711-044466	ABBT2160	MX2000-SFB-S
SFB 5	REV 05	711-044466	ABBT2145	MX2000-SFB-S
SFB 7	REV 05	711-044466	ABBT2163	MX2000-SFB-S
FPC 0	REV 30	750-028467	ABBN0284	MPC-3D-16XGE-SFPP
FPC 1	REV 30	750-028467	ABBN0308	MPC-3D-16XGE-SFPP
FPC 2	REV 30	750-028467	ABBN0316	MPC-3D-16XGE-SFPP
FPC 3	REV 32	750-028467	ABBN6832	MPC-3D-16XGE-SFPP
FPC 4	REV 32	750-028467	ABBN6811	MPC-3D-16XGE-SFPP
FPC 5	REV 32	750-028467	ABBN6791	MPC-3D-16XGE-SFPP
FPC 6	REV 30	750-028467	ABBM4592	MPC-3D-16XGE-SFPP
FPC 7	REV 32	750-028467	ABBN6810	MPC-3D-16XGE-SFPP
FPC 8	REV 30	750-028467	ABBM4739	MPC-3D-16XGE-SFPP
FPC 9	REV 32	750-028467	ABBN6827	MPC-3D-16XGE-SFPP
FPC 10	REV 30	750-028467	ABBN0302	MPC-3D-16XGE-SFPP
FPC 11	REV 32	750-028467	ABBN6790	MPC-3D-16XGE-SFPP
FPC 12	REV 30	750-028467	ZM5111	MPC-3D-16XGE-SFPP
FPC 13	REV 30	750-028467	ABBN0208	MPC-3D-16XGE-SFPP
FPC 14	REV 23	750-028467	YN2977	MPC-3D-16XGE-SFPP
FPC 15	REV 32	750-028467	ABBN6798	MPC-3D-16XGE-SFPP
FPC 16	REV 30	750-028467	ABBN0270	MPC-3D-16XGE-SFPP
FPC 17	REV 32	750-028467	ABBN6796	MPC-3D-16XGE-SFPP
FPC 18	REV 30	750-028467	ABBN0281	MPC-3D-16XGE-SFPP
FPC 19	REV 32	750-028467	ABBN6813	MPC-3D-16XGE-SFPP
ADC 0	REV 13	750-043596	ABBX5561	PROTO-ASSEMBLY
ADC 1	REV 13	750-043596	ABBX5546	PROTO-ASSEMBLY
ADC 2	REV 13	750-043596	ABBX5535	MX2000-LC-ADAPTER
ADC 3	REV 13	750-043596	ABBX5552	MX2000-LC-ADAPTER
ADC 4	REV 13	750-043596	ABBX5581	MX2000-LC-ADAPTER
ADC 5	REV 13	750-043596	ABBX5545	PROTO-ASSEMBLY
ADC 6	REV 13	750-043596	ABBX5554	PROTO-ASSEMBLY
ADC 7	REV 07	750-043596	ABBV7194	MX2000-LC-ADAPTER
ADC 8	REV 07	750-043596	ABBV7251	MX2000-LC-ADAPTER
ADC 9	REV 07	750-043596	ABBV7202	MX2000-LC-ADAPTER
ADC 10	REV 13	750-043596	ABBX5579	MX2000-LC-ADAPTER
ADC 12	REV 13	750-043596	ABBX5575	MX2000-LC-ADAPTER
ADC 13	REV 13	750-043596	ABBX5539	PROTO-ASSEMBLY
ADC 14	REV 13	750-043596	ABBX5555	PROTO-ASSEMBLY
ADC 15	REV 13	750-043596	ABBX5557	MX2000-LC-ADAPTER
ADC 16	REV 13	750-043596	ABBX5536	PROTO-ASSEMBLY
ADC 17	REV 13	750-043596	ABBX5559	PROTO-ASSEMBLY
ADC 18	REV 13	750-043596	ABBX5537	PROTO-ASSEMBLY
ADC 19	REV 11	750-043596	ABBW5685	PROTO-ASSEMBLY
Fan Tray 0	REV 04	760-046960	ACAY0090	
Fan Tray 1	REV 04	760-046960	ACAY0088	


```

Fan Tray 2      REV 04   760-046960   ACAY0089
Fan Tray 3      REV 04   760-046960   ACAY0108

```

show chassis hardware clei-models (MX2020 Router)

```
user@ host > show chassis hardware clei-models
```

```
Hardware inventory:
```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 27	750-040240	PROTOXCLEI	750-040240
FPM Board	REV 06	760-040242	PROTOXCLEI	760-040242
PSM 0	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 1	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 2	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 3	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 4	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 5	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 6	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 7	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 8	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 9	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 10	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 11	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 12	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 13	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 15	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PDM 0	REV 01	740-045234		
PDM 1	REV 01	740-045234		
PDM 2	REV 01	740-045234	IPUPAJSKAA	MX2000-PDM-DC-S-A
Routing Engine 0	REV 02	740-041821		RE-S-1800X4-16G-S
Routing Engine 1	REV 02	740-041821		RE-S-1800X4-16G-S
CB 0	REV 08	750-040257	PROTOXCLEI	750-040257
CB 1	REV 04	750-040257	PROTOXCLEI	750-040257
SFB 0	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 1	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 2	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 4	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 5	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 7	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
FPC 0	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 1	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 2	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 3	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 4	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 5	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 6	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 7	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 8	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 9	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 10	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 11	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 12	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 13	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 14	REV 23	750-028467		MPC-3D-16XGE-SFPP
FPC 15	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 16	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 17	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 18	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 19	REV 32	750-028467		MPC-3D-16XGE-SFPP
ADC 0	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 1	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 2	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER

ADC 3	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 4	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 5	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 6	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 7	REV 07	750-043596	PROTOXCLEI	MX2000-LC-ADAPTER
ADC 8	REV 07	750-043596	PROTOXCLEI	MX2000-LC-ADAPTER
ADC 9	REV 07	750-043596	PROTOXCLEI	MX2000-LC-ADAPTER
ADC 10	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 12	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 13	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 14	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 15	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 16	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 17	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 18	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 19	REV 11	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
Fan Tray 0	REV 04	760-046960		
Fan Tray 1	REV 04	760-046960		
Fan Tray 2	REV 04	760-046960		
Fan Tray 3	REV 04	760-046960		

show chassis hardware (MX Series Routers with ATM MIC)

Can I retain this output by just deleting "ATM MIC"?

user@host> show chassis hardware

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN115736EAFc	MX240
Midplane	REV 07	760-021404	ABAA5038	MX240 Backplane
FPM Board	REV 03	760-021392	ABBA2758	Front Panel Display
PEM 0	Rev 01	740-022697	QCS0937C07K	PS 1.2-1.7kW; 100-240V
AC in				
PEM 1	Rev 01	740-022697	QCS0939C04X	PS 1.2-1.7kW; 100-240V
AC in				
PEM 2	Rev 01	740-022697	QCS0937C06B	PS 1.2-1.7kW; 100-240V
AC in				
PEM 3	Rev 01	740-022697	QCS0937C07U	PS 1.2-1.7kW; 100-240V
AC in				
Routing Engine 0	REV 12	740-013063	9009042291	RE-S-2000
Routing Engine 1	REV 12	740-013063	9009042266	RE-S-2000
CB 0	REV 06	710-021523	ABBC1435	MX SCB
CB 1	REV 06	710-021523	ABBC1497	MX SCB
FPC 2	REV 14	750-031088	YH8446	MPC Type 2 3D Q
CPU	REV 06	711-030884	YH9612	MPC PMB 2G
MIC 0				
MIC 1	REV 10	750-036132	ZP7062	2xOC12/8xOC3 CC-CE
PIC 2		BUILTIN	BUILTIN	2xOC12/8xOC3 CC-CE
Xcvr 0		NON-JNPR	23393-00492	UNKNOWN
Xcvr 1		NON-JNPR	23393-00500	UNKNOWN
Xcvr 2		NON-JNPR	23393-00912	UNKNOWN
Xcvr 3	REV 01	740-015638	22216-00575	Load SFP
Xcvr 4	REV 01	740-015638	24145-00110	Load SFP
Xcvr 5	REV 01	740-015638	24145-00016	Load SFP
Xcvr 6	REV 01	740-015638	24145-00175	Load SFP
Xcvr 7		NON-JNPR	23393-00627	UNKNOWN
QXM 0	REV 05	711-028408	YF4681	MPC QXM
QXM 1	REV 05	711-028408	YF4817	MPC QXM
Fan Tray 0	REV 01	710-021113	XL3645	MX240 Fan Tray

show chassis hardware (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```

user@host>show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number  Description
Chassis                               JN11D969BAFA  MX960
Midplane          REV 03   710-013698   ACAA2362      MX960 Backplane
FPM Board         REV 03   710-014974   ZR0639        Front Panel Display
PDM              Rev 03   740-013110   QCS152250SX   Power Distribution Module
PEM 0            Rev 10   740-013683   QCS1512718W   DC Power Entry Module
PEM 1            Rev 10   740-013683   QCS1512702Y   DC Power Entry Module
Routing Engine 0 REV 15   740-013063   9012024667    RE-S-2000
Routing Engine 1 REV 15   740-013063   9012024649    RE-S-2000
CB 0             REV 14   750-031391   ZJ7749        Enhanced MX SCB
CB 1             REV 14   750-031391   ZJ7750        Enhanced MX SCB
CB 2             REV 14   750-031391   ZY9233        Enhanced MX SCB
FPC 0            REV 17   750-031089   YR7434        MPC Type 2 3D
  CPU
FPC 1            REV 11   750-037207   ZW9727        AS-MCC
  CPU            REV 04   711-038173   ZW4817        AS-MCC-PMB
    MIC 0        REV 01   750-037214   ZH3764        AS-MSC
      PIC 0              BUILTIN    BUILTIN      AS-MSC
    MIC 1        REV 01   711-028408   JZ9200        AS-MXC
      PIC 2              BUILTIN    BUILTIN      AS-MXC
FPC 4            REV 30   750-028467   ABBN0232      MPC 3D 16x 10GE
  CPU
FPC 5            REV 04   750-037207   ZK9074        AS-MCC
  CPU
Fan Tray 0       REV 05   740-014971   VT5683        Fan Tray
Fan Tray 1       REV 05   740-014971   VT5684        Fan Tray

```

show chassis hardware extensive (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```

user@host> show chassis hardware extensive

ID: AS-MCC                                FRU Model Number: 750-037207
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 37 01 0b 52 45 56 20 31 31 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 30 37 00 00
  Address 0x20: 53 2f 4e 20 5a 57 39 37 32 37 00 00 00 11 02 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 37
  Address 0x50: 35 30 2d 30 33 37 32 30 37 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 31 31 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 5e ff ff ff ff ff ff ff ff ff ff ff ff
  CPU            REV 04   711-038173   ZW4817        AS-MCC-PMB
Jedec Code: 0x7fb0          EEPROM Version: 0x02
P/N: 711-038173            S/N: S/N ZW4817
Assembly ID: 0x0b38        Assembly Version: 01.04
Date: 12-30-2011          Assembly Flags: 0x00
Version: REV 04
ID: AS-MCC-PMB
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 38 01 04 52 45 56 20 30 34 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 33 38 31 37 33 00 00
  Address 0x20: 53 2f 4e 20 5a 57 34 38 31 37 00 00 00 1e 0c 07
  Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

```

```

Address 0x40: ff ff ff ff 00 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 31 31 2d 30 33 38 31 37 33 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 30 34 00 ff ff ff ff ff ff
Address 0x70: ff ff ff 60 00 00 00 00 00 00 00 00 00 00 00
MIC 0          REV 01    750-037214    ZH3764          AS-MS
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           750-037214      S/N:             S/N ZH3764
Assembly ID:   0x0a44          Assembly Version: 01.01
Date:          07-04-2011      Assembly Flags:   0x00
Version:       REV 01
ID: AS-MS
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff I2C Hex Data:
Address 0x00: 7f b0 02 ff 0a 44 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 31 34 00 00
Address 0x20: 53 2f 4e 20 5a 48 33 37 36 34 00 00 00 04 07 07
Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff f6 c0 03 e1 bc 00 00 00 00 00 00 00 00
PIC 0          BUILTIN      BUILTIN          AS-MS
FPC 4          REV 30    750-028467    ABBN0232          MPC 3D 16x 10GE
Jedec Code:    0x7fb0          EEPROM Version:    0x01

```

show chassis hardware (MX480 Router with MPC4E)

```

user@host> show chassis hardware
Hardware inventory:

```

Item	Version	Part number	Serial number	Description
Chassis			JN10FF57BAFB	MX480
Midplane	REV 05	750-047849	Good	MX480 Midplane
FPM Board	REV 02	710-017254	KG2066	Front Panel Display
PEM 0	Rev 03	740-017330	QCS081590BJ	PS 1.2-1.7kW; 100-240V
AC in				
PEM 1	Rev 03	740-017330	QCS0815908Z	PS 1.2-1.7kW; 100-240V
AC in				
PEM 2	Rev 03	740-029970	QCS1001U001	PS 1.4-2.52kW; 90-264V
AC in				
Routing Engine 0	REV 05	740-031116	9009089502	RE-S-1800x4
Routing Engine 1	REV 05	740-031116	9009089624	RE-S-1800x4
CB 0	REV 02	750-031391	YE8506	Enhanced MX SCB
CB 1	REV 14	750-031391	ZK8265	Enhanced MX SCB
FPC 2	REV 05	750-037358	ZT0638	MPC4E 3D 32XGE
CPU	REV 07	711-035209	ZK3187	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	8X10GE SFPP
PIC 1		BUILTIN	BUILTIN	8X10GE SFPP
PIC 2		BUILTIN	BUILTIN	8X10GE SFPP
PIC 3		BUILTIN	BUILTIN	8X10GE SFPP
FPC 3	REV 06	750-037355	CAAB1144	MPC4E 3D 2CGE+8XGE
CPU	REV 08	711-035209	CAAB1278	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	B11E01439	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11D05809	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	D5418	UNKNOWN
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	X12J00362	CFP-100G-SR10
FPC 4	REV 12.3.10	750-033205	YR9445	MPCE Type 3 3D

CPU
Fan Tray

Enhanced Left Fan Tray

show chassis hardware (MX2020 Router with MPC4E)

user@host> show chassis hardware

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11E188CAFJ	MX2020
Midplane	REV 04	711-032387	ABAC7474	Lower Backplane
Midplane 1	REV 04	711-032386	ABAC7408	Upper Backplane
PMP 1	REV 03	711-032428	ACA11137	Upper Power Midplane
PMP 0	REV 03	711-032426	ACA11016	Lower Power Midplane
FPM Board	REV 06	760-040242	ABBT8832	Front Panel Display
PSM 3	REV 0C	740-033727	VK00255	DC 52V Power Supply
Module				
PSM 4	REV 0C	740-033727	VJ00148	DC 52V Power Supply
Module				
PSM 5	REV 0C	740-033727	VK00207	DC 52V Power Supply
Module				
PSM 6	REV 0C	740-033727	VK00319	DC 52V Power Supply
Module				
PSM 7	REV 0C	740-033727	VK00264	DC 52V Power Supply
Module				
PSM 8	REV 0B	740-033727	VG00025	DC 52V Power Supply
Module				
PSM 13	REV 0C	740-033727	VK00274	DC 52V Power Supply
Module				
PSM 14	REV 0C	740-033727	VJ00167	DC 52V Power Supply
Module				
PSM 15	REV 0C	740-033727	VK00299	DC 52V Power Supply
Module				
PSM 16	REV 0C	740-033727	VK00213	DC 52V Power Supply
Module				
PSM 17	REV 0C	740-033727	VK00253	DC 52V Power Supply
Module				
PDM 0	REV 0B	740-038109	VJ00040	DC Power Dist Module
PDM 2	REV 0B	740-038109	VJ00025	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009089735	RE-S-1800x4
Routing Engine 1	REV 02	740-041821	9009089731	RE-S-1800x4
CB 0	REV 04	750-040257	ZT2846	Control Board
CB 1	REV 04	750-040257	ZT2877	Control Board
SPMB 0	REV 01	711-041855	ZS2282	PMB Board
SPMB 1	REV 01	711-041855	ZS2261	PMB Board
SFB 0	REV 07	711-032385	ZZ2582	Switch Fabric Board
SFB 1	REV 04	711-032385	ZV4229	Switch Fabric Board
SFB 2	REV 07	711-032385	CAAB4902	Switch Fabric Board
SFB 3	REV 07	711-032385	CAAB4891	Switch Fabric Board
SFB 4	REV 07	711-032385	CAAB4883	Switch Fabric Board
SFB 5	REV 07	711-032385	CAAB4889	Switch Fabric Board
SFB 6	REV 06	711-032385	ZV1818	Switch Fabric Board
SFB 7	REV 07	711-032385	CAAB4897	Switch Fabric Board
FPC 0	REV 34	750-031090	ZT9799	MPC Type 2 3D EQ
CPU	REV 06	711-030884	ZS1122	MPC PMB 2G
MIC 0	REV 11	750-033535	CAAD7674	MIC-3D-10C192-XFP
PIC 0		BUILTIN	BUILTIN	MIC-3D-10C192-XFP
Xcvr 0	REV 01	740-014279	753019A00404	XFP-0C192-SR
MIC 1	REV 14	750-031967	ZM6103	MIC-3D-80C30C12-40C48
PIC 2		BUILTIN	BUILTIN	MIC-3D-80C30C12-40C48
Xcvr 0	REV 01	740-011615	PEF1AZP	SFP-IR
Xcvr 1	REV 01	740-011615	PEF1AZN	SFP-IR

Xcvr 2	REV 01	740-021308	ANA0N8S	SFP+-10G-SR
QXM 0	REV 06	711-028408	ZT9339	MPC QXM
QXM 1	REV 06	711-028408	ZT9237	MPC QXM
FPC 9	REV 34	750-031090	ZT9770	MPC Type 2 3D EQ
CPU	REV 06	711-030884	ZS1302	MPC PMB 2G
MIC 0	REV 24	750-028387	YJ3950	3D 4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	T09M52516	XFP-10G-SR
Xcvr 1		NON-JNPR	CA49BK095	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 02	740-014289	C834XU01T	XFP-10G-SR
Xcvr 1		NON-JNPR	T09M52515	XFP-10G-SR
MIC 1	REV 11	750-033535	CAAD7681	MIC-3D-10C192-XFP
PIC 2		BUILTIN	BUILTIN	MIC-3D-10C192-XFP
Xcvr 0	REV 01	740-014279	KBQ02BE	XFP-OC192-SR
QXM 0	REV 06	711-028408	ZT9151	MPC QXM
QXM 1	REV 06	711-028408	ZT9116	MPC QXM
FPC 10	REV 27	750-033205	ZL6215	MPCE Type 3 3D
CPU	REV 07	711-035209	ZK9038	HMPC PMB 2G
MIC 0	REV 18	750-028380	YG6885	3D 2x 10GE XFP
PIC 0		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0	REV 01	740-014289	C706XU0AG	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0	REV 02	740-014289	T08L84366	XFP-10G-SR
FPC 14	REV 09	750-037355	CAAF1534	MPC4E 3D 2CGE+8XGE
CPU	REV 08	711-035209	CAAB9879	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	21T511100436	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AHPOGPM	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	123363A00032	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	19T511100477	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	X12J00260	CFP-100G-SR10
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	21T511104086	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	21T511104627	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	21T511104644	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
FPC 19	REV 32	750-028467	ZR2008	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ZT6933	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	19T511100291	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMH02VE	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	23T511102128	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMS15PP	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	123363A00716	SFP+-10G-SR
ADC 0	REV 05	750-043596	CAAC2072	Adapter Card
ADC 9	REV 01	750-043596	ZV4111	Adapter Card
ADC 10	REV 05	750-043596	CAAC2058	Adapter Card
ADC 14	REV 02	750-043596	ZW1561	Adapter Card
ADC 19	REV 01	750-043596	ZV4127	Adapter Card
Fan Tray 0	REV 03	760-046960	ACAY0124	172mm FanTray - 6 Fans
Fan Tray 1	REV 2A	760-046960	ACAY0022	172mm FanTray - 6 Fans
Fan Tray 2	REV 2A	760-046960	ACAY0023	172mm FanTray - 6 Fans
Fan Tray 3	REV 2A	760-046960	ACAY0025	172mm FanTray - 6 Fans

show chassis hardware (T320 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number  Description
Chassis                               19093          T320
Midplane          REV 04   710-004339   BC1436        T320 Backplane
FPM GBUS          REV 03   710-004461   BC1407        T320 FPM Board
FPM Display       REV 04   710-002897   BE0763        FPM Display
CIP               REV 05   710-002895   BB2311        T Series CIP
PEM 0            Rev 01   740-004359   NB12546       Power Entry Module
SCG 0            REV 06   710-004455   AY4522        T320 Sonet
Clock Gen.
Routing Engine 0
CB 0              REV 13   710-002728   BC1577        unknown
Control Board    T Series
CB 1              REV 13   710-002728   BC1595        T Series
Control Board
FPC 1            REV 09   710-007531   HS1572        FPC Type 2
CPU              REV 15   710-001726   HR8763        FPC CPU
PIC 0            REV 01   750-010618   CB5579        4x G/E SFP,
1000 BASE
SFP 0            REV 01   740-007326   P5809Z1       SFP-SX
SFP 1            REV 01   740-007326   P4Q10XU       SFP-SX
SFP 2            NON-JNPR RA45020031    SFP-SX
SFP 3            NON-JNPR RA45020032    SFP-SX
PIC 1            REV 01   750-010618   CD9587        4x G/E SFP,
1000 BASE
SFP 0            NON-JNPR P5A08QZ       SFP-T
SFP 1            REV 01   740-007326   P4Q133K       SFP-SX
SFP 2            REV 01   740-007326   P5809YY       SFP-SX
SFP 3            REV 01   740-007327   4C81704       SFP-LX
MMB 1            REV 03   710-005555   HR9401        MMB-288mbit
PPB 0            REV 04   710-003758   HR2886        PPB Type 2
FPC 2            REV 07   710-005860   HP2392        FPC Type 1
CPU              REV 14   710-001726   HP7797        FPC CPU
PIC 0            REV 02   750-007643   HM0853        1x G/E QPP,
1000 BASE
SFP 0            REV 01   740-007326   P11E9JJ       SFP-SX
MMB 1            REV 02   710-005555   HN2379        MMB-288mbit
PPB 0            REV 04   710-003758   HP8092        PPB Type 2
FPC 3            REV 07   710-005860   HP2393        FPC Type 1
CPU              REV 14   710-001726   HP0968        FPC CPU
PIC 0            REV 01   750-010240   CB5363        1x G/E SFP,
1000 BASE
SFP 0            REV 01   740-007326   P4R0PNH       SFP-SX
PIC 1            REV 03   750-003034   HD2832        4x OC-3 SONET,
SMIR
MMB 1            REV 02   710-005555   HN6307        MMB-288mbit
PPB 0            REV 04   710-003758   HP5051        PPB Type 2
FPC 4            REV 01   710-010845   JD3872        FPC Type 4
CPU              REV 02   710-011481   JB6042        FPC CPU
5                REV 01   710-005802   BC1566        FPC Type 2
CPU              REV 09   710-001726   AY4922        FPC CPU
PIC 0            REV 02   750-008155   BE2114        2x G/E QPP,
1000 BASE
SFP 0            REV 01   740-007326   P4R0PMQ       SFP-SX
SFP 1            REV 01   740-007326   P4R0PN9       SFP-SX
PIC 1            REV 01   750-008155   BE2116        2x G/E QPP,
1000 BASE
SFP 0            REV 01   740-007326   P4R0PNZ       SFP-SX

```

SFP 1		NON-JNPR	2908	SFP-T
MMB 1	REV 01	710-005555	AZ2246	MMB-288mbit
PPB 0	REV 03	710-003758	AY4839	PPB Type 2
FPC 7	REV 01	710-005803	AZ2123	FPC Type 3
...				

show chassis hardware (T640 Router)

```
user@host> show chassis hardware
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			19182	T640
Midplane	REV 04	710-002726	AX5608	T640 Backplane
FPM GBUS	REV 02	710-002901	HE3064	T640 FPM Board
FPM Display	REV 02	710-002897	HE7864	FPM Display
CIP	REV 05	710-002895	HA5024	T Series CIP
PEM 0	Rev 02	740-029522	VH26235	AC PEM 10kW US
PEM 1	Rev 02	740-029522	VH26230	AC PEM 10kW US
SCG 0	REV 03	710-003423	HA4508	T640 Sonet Clock Gen.
Routing Engine 0	REV 02	740-005022	210865700483	RE-3.0 (RE-600)
CB 0	REV 01	710-002728	HD3044	T Series Control Board
FPC 2	REV 04	710-001721	HD5572	FPC Type 3
CPU	REV 06	710-001726	HA4712	FPC CPU
PIC 1	REV 03	750-009567	HV2331	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-009898	USC202R103	XENPAK-SR
PIC 2	REV 03	750-009567	HV2332	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-011268	USC202R112	XENPAK-ZR
PIC 3	REV 03	750-009567	HX4416	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-012056	434TC004	XENPAK-CX4
PIC 4	REV 03	750-009567	HX4420	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-012058	434TC124	XENPAK-LX4
FPC 5	REV 01	710-013553	JE4839	E2-FPC Type 1
CPU	REV 01	710-013569	JW9163	FPC CPU
PIC 0	REV 01	750-009567	HX4419	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-009898	USC202RT05	XENPAK-LR
PIC 1	REV 03	750-009567	HN7426	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-009550	03L90051	XENPAK-ER
PIC 2	REV 03	750-009467	HT7423	1x 10GE(LAN),XENPAK
SFP 0		NON-JNPR		UNKNOWN
PIC 3	REV 04	750-005100	AY4850	1x 10GE(LAN),DWDM
FPC 4	REV 01	710-010845	JD3872	FPC Type 4
CPU	REV 02	710-011481	JB6042	FPC CPU
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray

show chassis hardware models (T640 Router)

```
user@host> show chassis hardware models
Hardware inventory:
```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 04	710-002726		CHAS-BP-T640-S
FPM Display	REV 02	710-002897		CRAFT-T640-S
CIP	REV 05	710-002895		CIP-L-T640-S
PEM 0	Rev 01	740-002595		PWR-T-DC-S
SCG 0	REV 04	710-003423		SCG-T-S
SCG 1	REV 04	710-003423		SCG-T-S
Routing Engine 0	REV 01	740-005022		RE-600-2048-S
Routing Engine 1	REV 07	740-005022		RE-600-2048-S
CB 0	REV 06	710-002726		CHAS-BP-T640-S
CB 1	REV 06	710-002728		CB-L-T-S

FPC 5	REV 05	710-007527	T640-FPC2
PIC 0	REV 05	750-002510	PB-2GE-SX
PIC 1	REV 05	750-001901	PB-40C12-SON-SMIR
FPC 6	REV 03	710-001721	T640-FPC3
PIC 1	REV 01	750-009553	PC-40C48-SON-SFP
SIB 4	REV 02	750-005486	SIB-I-T640-S
Fan Tray 0			FANTRAY-T-S
Fan Tray 1			FANTRAY-T-S
Fan Tray 2			FAN-REAR-TX-T640-S

show chassis hardware extensive (T640 Router)

```

user@host> show chassis hardware extensive
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis
Jedec Code:   0x7fb0          EEPROM Version: 0x01
P/N:          .....          S/N:           .....
Assembly ID:  0x0507          Assembly Version: 00.00
Date:         00-00-0000      Assembly Flags:  0x00
Version:      .....
ID: Gibson LCC Chassis
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 01 ff 05 07 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane      REV 04      710-002726  AX5633
Jedec Code:   0x7fb0          EEPROM Version: 0x01
P/N:          710-002726.      S/N:           S/N AX5633.
Assembly ID:  0x0127          Assembly Version: 01.04
Date:         06-27-2001      Assembly Flags:  0x00
Version:      REV 04.....
ID: Gibson Backplane
Board Information Record:
Address 0x00: ad 01 08 00 00 90 69 0e f8 00 ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 01 27 01 04 52 45 56 20 30 34 00 00
Address 0x10: 00 00 00 00 37 31 30 2d 30 30 32 37 32 36 00 00
Address 0x20: 53 2f 4e 20 41 58 35 36 33 33 00 00 00 1b 06 07
Address 0x30: d1 ff ff ff ad 01 08 00 00 90 69 0e f8 00 ff ff
Address 0x40: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM GBUS      REV 02      710-002901  HE3245
...
FPM Display   REV 02      710-002897  HA4873
...
CIP           REV 05      710-002895  HA4729
...
PEM 1         RevX02     740-002595  MD21815      Power Entry Module
...
SCG 0         REV 04      710-003423  HF6023
...
SCG 1         REV 04      710-003423  HF6061
...
Routing Engine 0 REV 01     740-005022  210865700292 RE-3.0
...
CB 0          REV 06      710-002728  HE3614
...

```

FPC 1	REV 01	710-002385	HE3009	FPC Type 1
...	REV 06	710-001726	HC0010	

show chassis hardware (T4000 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN1172F25AHA  T4000
Midplane      REV 01   710-027486   RC8355        T-series Backplane
FPM GBUS      REV 13   710-002901   BBAE0927      T640 FPM Board
FPM Display   REV 01   710-021387   EF6764        T1600 FPM Display
CIP           REV 06   710-002895   BBAD9210      T-series CIP
PEM 0         REV 01   740-036442   VA00016       Power Entry Module 6x60
SCG 0         REV 18   710-003423   BBAD7248      T640 Sonet Clock Gen.
SCG 1         REV 18   710-003423   BBAE3874      T640 Sonet Clock Gen.
Routing Engine 0 REV 05   740-026941   P737F-002248  RE-DUO-1800
Routing Engine 1 REV 06   740-026941   P737F-002653  RE-DUO-1800
CB 0          REV 09   710-022597   ED0295        LCC Control Board
CB 1          REV 09   710-022597   EA6050        LCC Control Board
FPC 0         REV 26   750-032819   EK1173        FPC Type 5-3D
  CPU         REV 12   711-030686   EJ8584        SNG PMB
  PIC 0       REV 07   750-034624   EF6837        12x10GE (LAN/WAN) SFPP
    Xcvr 0    REV 01   740-031980   123363A01145  SFP+-10G-SR
    Xcvr 1    REV 01   740-031980   123363A01147  SFP+-10G-SR
    Xcvr 2    REV 01   740-031980   AJJ01P3       SFP+-10G-SR
    Xcvr 3    REV 01   740-031980   B10M03256     SFP+-10G-SR
    Xcvr 4    REV 01   740-031980   AJJ01M2       SFP+-10G-SR
    Xcvr 5    REV 01   740-031980   123363A01137  SFP+-10G-SR
    Xcvr 6    REV 01   740-031980   AJJ01PN       SFP+-10G-SR
    Xcvr 7    REV 01   740-031980   AJJ01NW       SFP+-10G-SR
    Xcvr 8    REV 01   740-031980   123363A01139  SFP+-10G-SR
    Xcvr 9    REV 01   740-031980   AJJ01KE       SFP+-10G-SR
    Xcvr 10   REV 01   740-031980   123363A01336  SFP+-10G-SR
    Xcvr 11   REV 01   740-031980   B10M01325     SFP+-10G-SR
  PIC 1       REV 07   750-034624   EF6800        12x10GE (LAN/WAN) SFPP
    Xcvr 0    REV 01   740-031980   AJJ01SA       SFP+-10G-SR
    Xcvr 1    REV 01   740-031980   AJJ01QZ       SFP+-10G-SR
    Xcvr 2    REV 01   740-031980   AJH0217       SFP+-10G-SR
    Xcvr 3    REV 01   740-031980   AJJ01TE       SFP+-10G-SR
    Xcvr 4    REV 01   740-031980   AJJ01KV       SFP+-10G-SR
    Xcvr 5    REV 01   740-031980   AJJ01MU       SFP+-10G-SR
    Xcvr 6    REV 01   740-031980   AJJ01R0       SFP+-10G-SR
    Xcvr 7    REV 01   740-031980   AJJ01TC       SFP+-10G-SR
    Xcvr 8    REV 01   740-031980   AJJ0364       SFP+-10G-SR
    Xcvr 9    REV 01   740-031980   AJD0GV3       SFP+-10G-SR
    Xcvr 10   REV 01   740-031980   B10M03343     SFP+-10G-SR
    Xcvr 11   REV 01   740-031980   AJJ01QJ       SFP+-10G-SR
  LMB 0       REV 05   711-034381   EJ8490        Type-0 LMB
  LMB 1       REV 04   711-035774   EJ8517        Type-1 LMB
  LMB 2       REV 05   711-034381   EJ8489        Type-0 LMB
FPC 3         REV 07   750-032819   EG3637        FPC Type 5-3D
  CPU         REV 09   711-030686   EG0150        SNG PMB
  PIC 0       REV 08   750-035293   EF3657        1x100GE
    Xcvr 0    REV 01   740-032210   C22CQNJ       CFP-100G-LR4
  PIC 1       REV 10   750-034624   BBAN4098      12x10GE (LAN/WAN) SFPP
    Xcvr 0    REV 01   740-031980   B11J04902     SFP+-10G-SR
    Xcvr 1    REV 01   740-031980   B11J04891     SFP+-10G-SR
    Xcvr 2    REV 01   740-031980   AJJ01MX       SFP+-10G-SR
    Xcvr 3    REV 01   740-031980   B11J04183     SFP+-10G-SR
    Xcvr 4    REV 01   740-031980   B11J04894     SFP+-10G-SR

```

Xcvr 5	REV 01	740-031980	B11J04184	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11J04897	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11J04899	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AJJ01TV	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	B11J04057	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AJJ01M4	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B11J04905	SFP+-10G-SR
LMB 0	REV 04	711-034381	EG1524	Type-0 LMB
LMB 1	REV 03	711-035774	EG0345	Type-1 LMB
LMB 2	REV 04	711-034381	EG1522	Type-0 LMB
FPC 5	REV 03	710-033871	BBAJ0768	FPC Type 4-ES
CPU	REV 11	710-016744	BBAH9342	ST-PMB2
PIC 0	REV 09	750-029262	EE6789	100GE
PIC 1	REV 03	750-034781	EE6655	100GE CFP
Xcvr 0	REV 01	740-032210	J11A22334	CFP-100G-LR4
BRIDGE 0	REV 03	711-029995	EE6572	100GE Bridge Board
MMB 0	REV 07	710-025563	BBAJ4657	ST-MMB2
MMB 1	REV 07	710-025563	BBAJ3073	ST-MMB2
FPC 6	REV 05	750-010153	EF4936	FPC Type 5-3D
CPU	REV 06	711-030686	EF4189	SNG PMB
PIC 0	REV 10	750-034624	BBAN4109	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	B11J04895	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11J04898	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11J04021	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11J04903	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11J04311	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J04059	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11J04016	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11J04017	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	B11J04887	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	B11J04297	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B11J04893	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B11J04022	SFP+-10G-SR
PIC 1	REV 02	750-034624	EE3711	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	AJH033X	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJJ01N0	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01SV	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJJ032L	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B10M01593	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJD0FF1	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01NU	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	123363A01305	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	B10M00361	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJJ01M7	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AJJ032X	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AJJ01PG	SFP+-10G-SR
LMB 0	REV 04	711-034381	EF3838	Type-0 LMB
LMB 1	REV 03	711-035774	EF3821	Type-1 LMB
LMB 2	REV 04	711-034381	EF3834	Type-0 LMB
SPMB 0	REV 05	710-023321	ED1990	LCC Switch CPU
SPMB 1	REV 05	710-023321	EA2768	LCC Switch CPU
SIB 0	REV 02	711-036340	EF8802	SIB-HC-3D
SIB 1	REV 07	711-036340	EG2286	SIB-HC-3D
SIB 2	REV 07	711-036340	EG2252	SIB-HC-3D
SIB 3	REV 02	711-036340	EF1358	SIB-HC-3D
SIB 4	REV 02	711-036340	EF8806	SIB-HC-3D
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
-- Rev 2				
Fan Tray 2				Rear Fan Tray -- Rev 3

show chassis hardware (T4000 Router with 16 GB line card chassis (LCC) Routing Engine)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN11BDF2CAHA  T1600
Midplane      REV 01   710-027486   ACAJ0774      T640 Backplane
FPM GBUS      REV 13   710-002901   BBAL6812      T640 FPM Board
FPM Display   REV 04   710-021387   BBAP2679      T1600 FPM Display
CIP           REV 06   710-002895   BBAP4758      T-series CIP
PEM 0         Rev 03   740-026384   XF86421       Power Entry Module 3x80
PEM 1         Rev 03   740-026384   XF86429       Power Entry Module 3x80
SCG 0         REV 18   710-003423   BBAP1896      T640 Sonet Clock Gen.
SCG 1         REV 18   710-003423   BBAN8659      T640 Sonet Clock Gen.
Routing Engine 0 REV 01   740-042243   737F-002238   RE-DUO-1800-16G
Routing Engine 1 REV 01   740-042243   737F-002403   RE-DUO-1800-16G
CB 1          REV 11   710-022597   EK4526        LCC Control Board
CB 1          REV 11   710-022597   EK4527        LCC Control Board
FPC 0         REV 05   710-033871   EK5644        FPC Type 4-ES
CPU           REV 11   710-016744   EK3428        ST-PMB2
PIC 0         REV 20   750-017405   EJ3041        4x 10GE (LAN/WAN) XFP
PIC 1         REV 17   750-026962   EH7536        10x10GE(LAN/WAN) SFPP
MMB 0         REV 07   710-025563   EK6039        ST-MMB2
MMB 1         REV 07   710-025563   EK6086        ST-MMB2
FPC 1         REV 05   710-033871   EK6583        FPC Type 4-ES
CPU           REV 11   710-016744   EK3401        ST-PMB2
PIC 0         REV 17   750-026962   EJ8948        10x10GE(LAN/WAN) SFPP
MMB 0         REV 07   710-025563   EK6202        ST-MMB2
MMB 1         REV 07   710-025563   EK6112        ST-MMB2
SPMB 1        REV 05   710-023321   EK4900        LCC Switch CPU
SIB 0         REV 11   710-013074   EK5958        SIB-I8-SF
SIB 1         REV 11   710-013074   EK4606        SIB-I8-SF
SIB 2         REV 11   710-013074   EK5971        SIB-I8-SF
SIB 3         REV 11   710-013074   EK4609        SIB-I8-SF
SIB 4         REV 11   710-013074   EK4602        SIB-I8-SF
Fan Tray 0    Front Top Fan Tray
Fan Tray 1    Front Bottom Fan Tray
Fan Tray 2    Rear Fan Tray -- Rev 2

```

show chassis hardware clei-models (T4000 Router)

```

user@host> show chassis hardware clei-models
Hardware inventory:
Item          Version  Part number  CLEI code      FRU model number
Midplane      REV 01   710-027486   IPMJ700DRD     CHAS-BP-T1600-S
FPM Display   REV 01   710-021387   IPUPAG6KAA     CRAFT-T1600-S
CIP           REV 06   710-002895   IPUPAG6KAA     CIP-L-T640-S
PEM 0         REV 01   740-036442   IPUPAG6KAA     PWR-T-6-60-DC
SCG 0         REV 18   710-003423   IPUPAG6KAA     SCG-T-S
SCG 1         REV 18   710-003423   IPUPAG6KAA     SCG-T-S
Routing Engine 0 REV 05   740-026941   IPUPAG6KAA     RE-DUO-C1800-8G-S
Routing Engine 1 REV 06   740-026941   IPUPAG6KAA     RE-DUO-C1800-8G-S
CB 0          REV 09   710-022597   IPUPAG6KAA     CB-LCC-S
CB 1          REV 09   710-022597   IPUPAG6KAA     CB-LCC-S
FPC 3
PIC 0         REV 08   750-035293   XXXXXXXXBB     PF-1CGE-CFP
PIC 1         REV 10   750-034624   XXXXXXXXCC     PF-12XGE-SFPP
FPC 5         REV 03   710-033871   IPUCAMBCTD     T1600-FPC4-ES
PIC 1         REV 03   750-034781   IPUIBKLMMA     PD-1CE-CFP-FPC4
FPC 6
PIC 0         REV 10   750-034624   XXXXXXXXCC     PF-12XGE-SFPP

```

Fan Tray 0	FANTRAY-T-S
Fan Tray 1	FANTRAY-T4000-S
Fan Tray 2	FANTRAY-TXP-R-S

show chassis hardware detail (T4000 Router)

```

user@host> show chassis hardware detail
Hardware inventory:
Item                Version  Part number  Serial number  Description
Chassis              REV 01   710-027486   JN1172F25AHA   T4000
Midplane             REV 13   710-002901   RC8355         T-series Backplane
FPM GBUS              REV 01   710-021387   BBAE0927       T640 FPM Board
FPM Display           REV 06   710-002895   BBAD9210       T-series CIP
CIP                   REV 01   740-036442   VA00016        Power Entry Module 6x60
PEM 0                 REV 18   710-003423   BBAD7248       T640 Sonet Clock Gen.
SCG 0                 REV 18   710-003423   BBAE3874       T640 Sonet Clock Gen.
SCG 1                 REV 05   740-026941   P737F-002248   RE-DUO-1800
Routing Engine 0
  ad0 3823 MB SMART CF 2009121602A661576157 Compact Flash
  ad1 59690 MB STEC MACH-8 SSD STM000103FDB Disk 1
Routing Engine 1 REV 06 740-026941 P737F-002653 RE-DUO-1800
  ad0 3823 MB SMART CF 201011150153F52CF52C Compact Flash
  ad1 62720 MB SMART Lite SATA Drive 2010110900150A880A88 Disk 1
CB 0                  REV 09 710-022597 ED0295 LCC Control Board
CB 1                  REV 09 710-022597 EA6050 LCC Control Board
FPC 0                 REV 26 750-032819 EK1173 FPC Type 5-3D
CPU                   REV 12 711-030686 EJ8584 SNG PMB
PIC 0                 REV 07 750-034624 EF6837 12x10GE (LAN/WAN) SFPP
  Xcvr 0              REV 01 740-031980 123363A01145 SFP+-10G-SR
  Xcvr 1              REV 01 740-031980 123363A01147 SFP+-10G-SR
  Xcvr 2              REV 01 740-031980 AJJ01P3 SFP+-10G-SR
  Xcvr 3              REV 01 740-031980 B10M03256 SFP+-10G-SR
  Xcvr 4              REV 01 740-031980 AJJ01M2 SFP+-10G-SR
  Xcvr 5              REV 01 740-031980 123363A01137 SFP+-10G-SR
  Xcvr 6              REV 01 740-031980 AJJ01PN SFP+-10G-SR
  Xcvr 7              REV 01 740-031980 AJJ01NW SFP+-10G-SR
  Xcvr 8              REV 01 740-031980 123363A01139 SFP+-10G-SR
  Xcvr 9              REV 01 740-031980 AJJ01KE SFP+-10G-SR
  Xcvr 10             REV 01 740-031980 123363A01336 SFP+-10G-SR
  Xcvr 11             REV 01 740-031980 B10M01325 SFP+-10G-SR
PIC 1                 REV 07 750-034624 EF6800 12x10GE (LAN/WAN) SFPP
  Xcvr 0              REV 01 740-031980 AJJ01SA SFP+-10G-SR
  Xcvr 1              REV 01 740-031980 AJJ01QZ SFP+-10G-SR
  Xcvr 2              REV 01 740-031980 AJH0217 SFP+-10G-SR
  Xcvr 3              REV 01 740-031980 AJJ01TE SFP+-10G-SR
  Xcvr 4              REV 01 740-031980 AJJ01KV SFP+-10G-SR
  Xcvr 5              REV 01 740-031980 AJJ01MU SFP+-10G-SR
  Xcvr 6              REV 01 740-031980 AJJ01R0 SFP+-10G-SR
  Xcvr 7              REV 01 740-031980 AJJ01TC SFP+-10G-SR
  Xcvr 8              REV 01 740-031980 AJJ0364 SFP+-10G-SR
  Xcvr 9              REV 01 740-031980 AJD0GV3 SFP+-10G-SR
  Xcvr 10             REV 01 740-031980 B10M03343 SFP+-10G-SR
  Xcvr 11             REV 01 740-031980 AJJ01QJ SFP+-10G-SR
LMB 0                 REV 05 711-034381 EJ8490 Type-0 LMB
LMB 1                 REV 04 711-035774 EJ8517 Type-1 LMB
LMB 2                 REV 05 711-034381 EJ8489 Type-0 LMB
FPC 3                 REV 07 750-032819 EG3637 FPC Type 5-3D
CPU                   REV 09 711-030686 EG0150 SNG PMB
PIC 0                 REV 08 750-035293 EF3657 1x100GE
  Xcvr 0              REV 01 740-032210 C22CQNJ CFP-100G-LR4
PIC 1                 REV 10 750-034624 BBAN4098 12x10GE (LAN/WAN) SFPP

```

Xcvr 0	REV 01	740-031980	B11J04902	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11J04891	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01MX	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11J04183	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11J04894	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J04184	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11J04897	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11J04899	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AJJ01TV	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	B11J04057	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AJJ01M4	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B11J04905	SFP+-10G-SR
LMB 0	REV 04	711-034381	EG1524	Type-0 LMB
LMB 1	REV 03	711-035774	EG0345	Type-1 LMB
LMB 2	REV 04	711-034381	EG1522	Type-0 LMB
FPC 5	REV 03	710-033871	BBAJ0768	FPC Type 4-ES
CPU	REV 11	710-016744	BBAH9342	ST-PMB2
PIC 0	REV 09	750-029262	EE6789	100GE
PIC 1	REV 03	750-034781	EE6655	100GE CFP
Xcvr 0	REV 01	740-032210	J11A22334	CFP-100G-LR4
BRIDGE 0	REV 03	711-029995	EE6572	100GE Bridge Board
MMB 0	REV 07	710-025563	BBAJ4657	ST-MMB2
MMB 1	REV 07	710-025563	BBAJ3073	ST-MMB2
FPC 6	REV 05	750-010153	EF4936	FPC Type 5-3D
CPU	REV 06	711-030686	EF4189	SNG PMB
PIC 0	REV 10	750-034624	BBAN4109	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	B11J04895	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11J04898	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11J04021	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11J04903	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11J04311	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J04059	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11J04016	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11J04017	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	B11J04887	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	B11J04297	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B11J04893	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B11J04022	SFP+-10G-SR
PIC 1	REV 02	750-034624	EE3711	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	AJH033X	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJJ01N0	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01SV	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJJ032L	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B10M01593	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJD0FF1	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01NU	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	123363A01305	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	B10M00361	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJJ01M7	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AJJ032X	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AJJ01PG	SFP+-10G-SR
LMB 0	REV 04	711-034381	EF3838	Type-0 LMB
LMB 1	REV 03	711-035774	EF3821	Type-1 LMB
LMB 2	REV 04	711-034381	EF3834	Type-0 LMB
SPMB 0	REV 05	710-023321	ED1990	LCC Switch CPU
SPMB 1	REV 05	710-023321	EA2768	LCC Switch CPU
SIB 0	REV 02	711-036340	EF8802	SIB-HC-3D
SIB 1	REV 07	711-036340	EG2286	SIB-HC-3D
SIB 2	REV 07	711-036340	EG2252	SIB-HC-3D
SIB 3	REV 02	711-036340	EF1358	SIB-HC-3D
SIB 4	REV 02	711-036340	EF8806	SIB-HC-3D

Fan Tray 0
Fan Tray 1
-- Rev 2
Fan Tray 2

Front Top Fan Tray
Front Bottom Fan Tray

Rear Fan Tray -- Rev 3

show chassis hardware models (T4000 Router)

user@host> show chassis hardware models

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 01	710-027486	RC8355	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	EF6764	CRAFT-T1600-S
CIP	REV 06	710-002895	BBAD9210	CIP-L-T640-S
PEM 0	REV 01	740-036442	VA00016	PWR-T-6-60-DC
SCG 0	REV 18	710-003423	BBAD7248	SCG-T-S
SCG 1	REV 18	710-003423	BBAE3874	SCG-T-S
Routing Engine 0	REV 05	740-026941	P737F-002248	RE-DUO-C1800-8G-S
Routing Engine 1	REV 06	740-026941	P737F-002653	RE-DUO-C1800-8G-S
CB 0	REV 09	710-022597	ED0295	CB-LCC-S
CB 1	REV 09	710-022597	EA6050	CB-LCC-S
FPC 3				
PIC 0	REV 08	750-035293	EF3657	PF-1CGE-CFP
PIC 1	REV 10	750-034624	BBAN4098	PF-12XGE-SFPP
FPC 5	REV 03	710-033871	BBAJ0768	T1600-FPC4-ES
PIC 1	REV 03	750-034781	EE6655	PD-1CE-CFP-FPC4
FPC 6				
PIC 0	REV 10	750-034624	BBAN4109	PF-12XGE-SFPP
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T4000-S
Fan Tray 2				FAN-REAR-TXP-LCC

show chassis hardware lcc (TX Matrix Router)

user@host> show chassis hardware lcc 0

lcc0-re0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			65751	T640
Midplane	REV 03	710-005608	RA1408	T640 Backplane
FPM GBUS	REV 09	710-002901	RA2784	T640 FPM Board
FPM Display	REV 05	710-002897	RA2825	FPM Display
CIP	REV 06	710-002895	HT0684	T Series CIP
PEM 0	Rev 11	740-002595	PM18483	Power Entry Module
PEM 1	Rev 11	740-002595	qb13984	Power Entry Module
SCG 0	REV 11	710-003423	HT0022	T640 Sonet Clock Gen.
Routing Engine 0	REV 13	740-005022	210865700363	RE-3.0 (RE-600)
CB 0	REV 03	710-007655	HW1195	Control Board (CB-T)
FPC 1	REV 05	710-007527	HM3245	FPC Type 2
CPU	REV 14	710-001726	HM1084	FPC CPU
PIC 0	REV 02	750-007218	AZ1112	2x OC-12 ATM2 IQ, SMIR
PIC 1	REV 02	750-007745	HG3462	4x OC-3 SONET, SMIR
PIC 2	REV 14	750-001901	BA5390	4x OC-12 SONET, SMIR
PIC 3	REV 09	750-008155	HS3012	2x G/E IQ, 1000 BASE
SFP 0		NON-JNPR	P1186TY	SFP-S
SFP 1	REV 01	740-007326	P11WLTF	SFP-SX
MMB 1	REV 02	710-005555	HL7514	MMB-288mbit
PPB 0	REV 04	710-003758	HM4405	PPB Type 2
PPB 1	REV 04	710-003758	AV1960	PPB Type 2
FPC 2	REV 08	710-010154	HZ3578	E-FPC Type 3

CPU	REV 05	710-010169	HZ3219	FPC CPU-Enhanced
PIC 0	REV 02	750-009567	HX2882	1x 10GE(LAN), XENPAK
SFP 0	REV 01	740-009898	USC202U709	XENPAK-LR
PIC 1	REV 03	750-003336	HJ9954	4x OC-48 SONET, SMSR
PIC 2	REV 01	750-004535	HC0235	1x OC-192 SM SR1
PIC 3	REV 07	750-007141	HX1699	10x 1GE(LAN), 1000 BASE
SFP 0	REV 01	740-007326	2441042	SFP-SX
SFP 1	REV 01	740-007326	2441027	SFP-SX
MMB 0	REV 03	710-010171	HV2365	MMB-5M3-288mbit
MMB 1	REV 03	710-010171	HZ3888	MMB-5M3-288mbit
SPMB 0	REV 09	710-003229	HW5245	T Series Switch CPU
SIB 3	REV 07	710-005781	HR5927	SIB-L8-F16
B Board	REV 06	710-005782	HR5971	SIB-L8-F16 (B)
SIB 4	REV 07	710-005781	HR5903	SIB-L8-F16
B Board	REV 06	710-005782	HZ5275	SIB-L8-F16 (B)

show chassis hardware scc (TX Matrix Router)

```
user@host> show chassis hardware scc
scc-re0:
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis				TX Matrix
Midplane	REV 04	710-004396	RB0014	SCC Midplane
FPM GBUS	REV 04	710-004617	HW9141	SCC FPM Board
FPM Display	REV 04	710-004619	HS5950	SCC FPM
CIP 0	REV 01	710-010218	HV9151	SCC CIP
CIP 1	REV 01	710-010218	HV9152	SCC CIP
PEM 1	Rev 11	740-002595	QB13977	Power Entry Module
Routing Engine 0	REV 05	740-008883	P11123900153	RE-4.0 (RE-1600)
CB 0	REV 01	710-011709	HR5964	Control Board (CB-TX)
SPMB 0	REV 09	710-003229	HW5293	T Series Switch CPU
SIB 3				
SIB 4	REV 01	710-005839	HW1177	SIB-S8-F16
B Board	REV 01	710-005840	HW1202	SIB-S8-F16 (B)

show chassis hardware (T1600 Router)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			B2703	T1600
Midplane	REV 03	710-005608	RC4137	T640 Backplane
FPM GBUS	REV 10	710-002901	DT7062	T640 FPM Board
FPM Display	REV 05	710-002897	DS3067	FPM Display
CIP	REV 06	710-002895	DT3386	T-series CIP
PEM 0	Rev 07	740-017906	UA26344	Power Entry Module 3x80
PEM 1	Rev 18	740-002595	UF38441	Power Entry Module
SCG 0	REV 15	710-003423	DV0941	T640 Sonet Clock Gen.
Routing Engine 0	REV 08	740-014082	9009014502	RE-A-2000
Routing Engine 1	REV 07	740-014082	9009009591	RE-A-2000
CB 0	REV 05	710-007655	JA9360	Control Board (CB-T)
CB 1	REV 03	710-017707	DT3251	Control Board (CB-T)
FPC 0	REV 07	710-013558	DR4253	E2-FPC Type 2
CPU	REV 05	710-013563	DS3902	FPC CPU-Enhanced
PIC 0	REV 01	750-010618	CB5446	4x G/E SFP, 1000 BASE
Xcvr 0	REV 01	740-011613	P9F11CW	SFP-SX
Xcvr 1	REV 01	740-011613	P9F15C2	SFP-SX
Xcvr 2	REV 01	740-011782	PB94K0L	SFP-SX

PIC 1	REV 06	750-001900	HB6399	1x OC-48 SONET, SMSR
PIC 2	REV 14	750-001901	AP1092	4x OC-12 SONET, SMIR
PIC 3	REV 07	750-001900	AR8275	1x OC-48 SONET, SMSR
MMB 1	REV 07	710-010171	DS1524	MMB-5M3-288mbit
FPC 1	REV 06	710-013553	DL9067	E2-FPC Type 1
CPU	REV 04	710-013563	DM1685	FPC CPU-Enhanced
PIC 0	REV 08	750-001072	AB1688	1x G/E, 1000 BASE-SX
PIC 1	REV 10	750-012266	JX5519	4x 1GE(LAN), IQ2
Xcvr 0	REV 01	740-011613	AM0812S8UK6	SFP-SX
Xcvr 2	REV 01	740-011613	AM0812S8UK1	SFP-SX
Xcvr 3	REV 01	740-011782	P8N1YHG	SFP-SX
PIC 2	REV 22	750-005634	DP0083	1x CHOC12 IQ SONET, SMIR
MMB 1	REV 07	710-008923	DN1862	MMB 3M 288-bit
FPC 2	REV 01	710-005548	HJ9899	FPC Type 3
CPU	REV 06	710-001726	HC0586	FPC CPU
PIC 0	REV 16	750-007141	NC9660	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011613	AM0812S8XAR	SFP-SX
Xcvr 1	REV 01	740-011782	P920E7B	SFP-SX
Xcvr 2	REV 01	740-011613	AM0812S8XAU	SFP-SX
Xcvr 4	REV 01	740-011613	AM0812S8XAK	SFP-SX
Xcvr 5	REV 01	740-011613	AM0812S8XAA	SFP-SX
Xcvr 6	REV 01	740-011613	PAJ4NKY	SFP-SX
Xcvr 7	REV 01	740-011613	AM0812S8UJW	SFP-SX
Xcvr 8	REV 01	740-011782	PB81X89	SFP-SX
Xcvr 9	REV 01	740-011613	AM0812S8UJX	SFP-SX
PIC 1	REV 06	750-015217	DK3280	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011782	P8P0A3T	SFP-SX
Xcvr 1	REV 01	740-013111	5090002	SFP-T
Xcvr 2	REV 01	740-011613	AM0814S93BQ	SFP-SX
Xcvr 4		NON-JNPR	PDE0FAN	SFP-SX
Xcvr 5	REV 01	740-011782	P8Q20XY	SFP-SX
Xcvr 6	REV 01	740-011613	AM0812S8UJV	SFP-SX
Xcvr 7	REV 01	740-011613	AM0812S8UP7	SFP-SX
PIC 2	REV 05	750-004695	HT4383	1x Tunnel
PIC 3	REV 17	750-009553	RL0204	4x OC-48 SONET
Xcvr 0	REV 01	740-011785	PDS3T23	SFP-SR
Xcvr 1	REV 01	740-011785	P6Q0F3E	SFP-SR
MMB 0	REV 03	710-004047	HD5843	MMB-288mbit
MMB 1	REV 03	710-004047	HE3208	MMB-288mbit
PPB 0	REV 02	710-002845	HA4524	PPB Type 3
PPB 1	REV 02	710-002845	HA4766	PPB Type 3
FPC 3	REV 01	710-010154	HR0863	E-FPC Type 3
CPU	REV 01	710-010169	HN3422	FPC CPU-Enhanced
PIC 0	REV 07	750-012793	WF5096	1x 10GE(LAN/WAN) IQ2
Xcvr 0		NON-JNPR	M64294TP	XFP-10G-LR
PIC 1	REV 25	750-007141	DV2127	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011613	PFA6LTJ	SFP-SX
Xcvr 1	REV 01	740-011782	P9P0XV4	SFP-SX
Xcvr 2	REV 01	740-011782	P9M0TNX	SFP-SX
Xcvr 4	REV 01	740-011782	P9B0TTP	SFP-SX
Xcvr 5		NON-JNPR	PBS4LED	SFP-SX
PIC 2	REV 17	750-009553	RL0212	4x OC-48 SONET
Xcvr 0	REV 01	740-011785	PDS3T8G	SFP-SR
PIC 3	REV 32	750-003700	DL1279	1x OC-192 12xMM VSR
MMB 0	REV 01	710-010171	HR0821	MMB-288mbit
MMB 1	REV 01	710-010171	HR0818	MMB-288mbit
FPC 4	REV 16	710-013037	EB4919	FPC Type 4-ES
CPU	REV 09	710-016744	BBAA4382	ST-PMB2

PIC 0	REV 03	711-029996	EB1569	100GE
PIC 1	REV 05	711-029999	EB9983	100GE CFP
Xcvr 0	REV 0	740-032210	J10G80746	CFP-100G-LR4
BRIDGE 0	REV 02	711-029995	EB2235	100GE Bridge Board
MMB 0	REV 04	710-025563	BBA7112	ST-MMB2
MMB 1	REV 04	710-025563	BBA7149	ST-MMB2
FPC 5	REV 02	710-013037	DE3407	FPC Type 4-ES
CPU	REV 04	710-016744	DA2124	ST-PMB2
PIC 0	REV 16	750-012518	DF2554	4x OC-192 SONET XFP
Xcvr 0	REV 01	740-014279	AA0745N1FX8	XFP-OC192-SR
Xcvr 1	REV 01	740-014279	AA0748N1HN5	XFP-OC192-SR
Xcvr 2	REV 01	740-014279	AA0748N1HT6	XFP-OC192-SR
Xcvr 3	REV 01	740-014279	AA0744N1EC9	XFP-OC192-SR
PIC 1	REV 01	750-010850	JA0329	1x OC-768 SONET SR
MMB 0	REV 04	710-016036	DE9577	ST-MMB2
MMB 1	REV 04	710-016036	DK4060	ST-MMB2
FPC 6	REV 14	710-013037	DV1431	FPC Type 4-ES
CPU	REV 09	710-016744	DT9020	ST-PMB2
PIC 0	REV 11	750-017405	DM6261	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 01	740-014289	C701XU05Q	XFP-10G-SR
Xcvr 1	REV 01	740-014279	AA0748N1HPT	XFP-10G-LR
Xcvr 2	REV 01	740-014289	T08E19189	XFP-10G-SR
Xcvr 3	REV 01	740-014289	C715XU058	XFP-10G-SR
PIC 1	REV 13	750-017405	DP8772	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 02	740-011571	C850XJ037	XFP-10G-SR
Xcvr 1	REV 02	740-014289	C839XU0L9	XFP-10G-SR
Xcvr 2	REV 02	740-014289	C834XU05A	XFP-10G-SR
Xcvr 3	REV 02	740-014289	C810XU0CE	XFP-10G-SR
MMB 0	REV 01	710-025563	DT8454	ST-MMB2
MMB 1	REV 01	710-025563	DT8366	ST-MMB2
FPC 7	REV 09	710-007529	HZ7624	FPC Type 3
CPU	REV 15	710-001726	HZ1413	FPC CPU
PIC 0	REV 10	750-012793	DM5627	1x 10GE(LAN/WAN) IQ2
Xcvr 0	REV 02	740-011571	C831XJ062	XFP-10G-SR
PIC 1	REV 01	750-015217	JT6762	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011782	P8Q25JU	SFP-SX
Xcvr 1	REV 01	740-011782	P9B0U0K	SFP-SX
PIC 2	REV 01	750-015217	JS4268	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011613	AM0812S8XBZ	SFP-SX
Xcvr 1	REV 01	740-011613	AM0812S8XAP	SFP-SX
Xcvr 2	REV 01	740-011613	AM0812S8XBY	SFP-SX
Xcvr 3	REV 01	740-011613	AM0812S8XBX	SFP-SX
Xcvr 4	REV 01	740-011613	P9F1652	SFP-SX
Xcvr 5	REV 01	740-011782	P8Q21YC	SFP-SX
Xcvr 6	REV 01	740-011782	P8Q27HQ	SFP-SX
Xcvr 7	REV 01	740-011613	P8E2SSU	SFP-SX
PIC 3	REV 15	750-009450	NB6790	1x OC-192 SM SR2
MMB 0	REV 03	710-005555	HZ3450	MMB-288mbit
MMB 1	REV 03	710-005555	HZ3415	MMB-288mbit
PPB 0	REV 04	710-002845	HP0887	PPB Type 3
PPB 1	REV 04	710-002845	HW5255	PPB Type 3
SPMB 0	REV 10	710-003229	HX3699	T-series Switch CPU
SPMB 1	REV 12	710-003229	DT3091	T-series Switch CPU
SIB 0	REV 07	710-013074	DS4747	SIB-I8-SF
SIB 1	REV 07	710-013074	DS4942	SIB-I8-SF
SIB 2	REV 07	710-013074	DS4965	SIB-I8-SF
SIB 3	REV 07	710-013074	DS4990	SIB-I8-SF
SIB 4	REV 07	710-013074	DS4944	SIB-I8-SF
Fan Tray 0				Front Top Fan Tray

Fan Tray 1
Fan Tray 2

Front Bottom Fan Tray
Rear Fan Tray -- Rev 2

show chassis hardware (TX Matrix Plus Router)

```
user@host> show chassis hardware
sfc0-re0:
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN113186EAHB	TXP
Midplane	REV 05	710-022574	TS3822	SFC Midplane
FPM Display	REV 03	710-024027	DW4701	TXP FPM Display
CIP 0	REV 05	710-023792	DW7998	TXP CIP
CIP 1	REV 05	710-023792	DW7999	TXP CIP
PEM 0	Rev 04	740-027463	UM26367	Power Entry Module
PEM 1	Rev 04	740-027463	UM26346	Power Entry Module
Routing Engine 0	REV 06	740-026942	737A-1081	RE-DUO-2600
Routing Engine 1	REV 06	740-026942	737A-1043	RE-DUO-2600
CB 0	REV 05	710-022606	DW4435	SFC Control Board
CB 1	REV 09	710-022606	DW6100	SFC Control Board
SPMB 0		BUILTIN		SFC Switch CPU
SPMB 1		BUILTIN		SFC Switch CPU
SIB F13 0	REV 04	750-024564	DW5764	F13 SIB
B Board	REV 03	710-023431	DW9053	F13 SIB Mezz
SIB F13 3	REV 04	750-024564	DW5785	F13 SIB
B Board	REV 03	710-023431	DW9030	F13 SIB Mezz
SIB F13 6				
SIB F13 8	REV 04	750-024564	DW5752	F13 SIB
B Board	REV 03	710-023431	DW9051	F13 SIB Mezz
SIB F13 11	REV 04	750-024564	DW5782	F13 SIB
B Board	REV 03	710-023431	DW9058	F13 SIB Mezz
SIB F13 12	REV 03	750-024564	DT9466	F13 SIB
B Board	REV 02	710-023431	DT6556	F13 SIB Mezz
SIB F2S 0/0	REV 05	710-022603	DW7898	F2S SIB
B Board	REV 05	710-023787	DW7625	F2S SIB Mezz
SIB F2S 0/2	REV 05	710-022603	DW7811	F2S SIB
B Board	REV 05	710-023787	DW7550	F2S SIB Mezz
SIB F2S 0/4	REV 04	710-022603	DW4873	F2S SIB
B Board	REV 05	710-023787	DW8509	F2S SIB Mezz
SIB F2S 0/6	REV 04	710-022603	DW4867	F2S SIB
B Board	REV 05	710-023787	DW8472	F2S SIB Mezz
SIB F2S 1/0	REV 04	710-022603	DW4871	F2S SIB
B Board	REV 05	710-023787	DW8497	F2S SIB Mezz
SIB F2S 1/2	REV 05	710-022603	DW7868	F2S SIB
B Board	REV 05	710-023787	DW7551	F2S SIB Mezz
SIB F2S 1/4	REV 04	710-022603	DW4854	F2S SIB
B Board	REV 05	710-023787	DW8496	F2S SIB Mezz
SIB F2S 1/6	REV 05	710-022603	DW7889	F2S SIB
B Board	REV 05	710-023787	DW7496	F2S SIB Mezz
SIB F2S 2/0	REV 04	710-022603	DW4852	F2S SIB
B Board	REV 05	710-023787	DW8498	F2S SIB Mezz
SIB F2S 2/2	REV 04	710-022603	DW4845	F2S SIB
B Board	REV 05	710-023787	DW8457	F2S SIB Mezz
SIB F2S 2/4	REV 05	710-022603	DW7802	F2S SIB
B Board	REV 05	710-023787	DW7562	F2S SIB Mezz
SIB F2S 2/6	REV 04	710-022603	DW4822	F2S SIB
B Board	REV 05	710-023787	DW8467	F2S SIB Mezz
SIB F2S 3/0	REV 05	710-022603	DW7815	F2S SIB
B Board	REV 05	710-023787	DW7518	F2S SIB Mezz
SIB F2S 3/2	REV 03	710-022603	DV0068	F2S SIB

B Board	REV 03	710-023787	DT9974	F2S SIB Mezz
SIB F2S 3/4	REV 05	710-022603	DW7874	F2S SIB
B Board	REV 05	710-023787	DW7601	F2S SIB Mezz
SIB F2S 3/6	REV 03	710-022603	DV0033	F2S SIB
B Board	REV 03	710-023787	DT9969	F2S SIB Mezz
SIB F2S 4/0	REV 03	710-022603	DV0043	F2S SIB
B Board	REV 03	710-023787	DT9948	F2S SIB Mezz
SIB F2S 4/2	REV 05	710-022603	DW5446	F2S SIB
B Board	REV 05	710-023787	DW7611	F2S SIB Mezz
SIB F2S 4/4	REV 04	710-022603	DW4826	F2S SIB
B Board	REV 05	710-023787	DW8458	F2S SIB Mezz
SIB F2S 4/6	REV 03	710-022603	DV0026	F2S SIB
B Board	REV 03	710-023787	DT9963	F2S SIB Mezz
Fan Tray 0	REV 02	760-024497	DR8290	Front Fan Tray
Fan Tray 1	REV 02	760-024497	DR8293	Front Fan Tray
Fan Tray 2	REV 05	760-024502	DR8280	Rear Fan Tray
Fan Tray 3				
Fan Tray 4	REV 05	760-024502	DR8276	Rear Fan Tray
Fan Tray 5	REV 02	760-024502	DP5643	Rear Fan Tray

lcc0-re0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11036F8AHA	T1600
Midplane	REV 03	710-017247	RC3799	T-series Backplane
FPM GBUS	REV 10	710-002901	DP7009	T640 FPM Board
FPM Display	REV 01	710-021387	DN7026	T1600 FPM Display
CIP	REV 06	710-002895	DP6024	T-series CIP
PEM 1	Rev 02	740-023211	WA50019	Power Entry Module 4x60A
SCG 0	REV 15	710-003423	DR6757	T640 Sonet Clock Gen.
SCG 1	REV 15	710-003423	DS2225	T640 Sonet Clock Gen.
Routing Engine 0	REV 01	740-026941	737F-1040	RE-DUO-1800
Routing Engine 1	REV 01	740-026941	737F-1016	RE-DUO-1800
CB 0	REV 06	710-022597	DX4011	LCC Control Board
CB 1	REV 06	710-022597	DX4017	LCC Control Board
FPC 1	REV 07	710-013035	DN5847	FPC Type 3-ES
CPU	REV 08	710-016744	DP2570	ST-PMB2
PIC 0	REV 05	750-015217	DB0418	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011782	P8Q27ZG	SFP-SX
Xcvr 1		NON-JNPR	PDA1U0D	SFP-SX
Xcvr 2	REV 01	740-011613	P9F1ALW	SFP-SX
Xcvr 3	REV 01	740-011782	PBA403V	SFP-SX
Xcvr 4		NON-JNPR	PDE09DP	SFP-SX
Xcvr 5	REV 01	740-011782	PCH2P4K	SFP-SX
Xcvr 6	REV 01	740-011782	PB94K0F	SFP-SX
Xcvr 7	REV 01	740-011782	PBA2R2A	SFP-SX
PIC 1	REV 03	750-004424	HJ4020	1x 10GE(LAN),DWDM
PIC 2	REV 01	750-003336	HG6073	4x OC-48 SONET, SMSR
MMB 0	REV 04	710-016036	DP3401	ST-MMB2
FPC 3	REV 12	710-013037	DR1169	FPC Type 4-ES
CPU	REV 08	710-016744	DP9429	ST-PMB2
PIC 0	REV 02	750-010850	JA0332	1x OC-768 SONET SR
MMB 0	REV 04	710-016036	DR0628	ST-MMB2
MMB 1	REV 04	710-016036	DR0592	ST-MMB2
FPC 4	REV 05	710-021534	DR7350	FPC Type 1-ES
CPU	REV 08	710-016744	DP8096	ST-PMB2
PIC 0	REV 04	750-014627	DP9171	4x OC-3 1x OC-12 SFP
Xcvr 0	REV 02	740-011615	PDE2RVR	SFP-SR
PIC 1	REV 22	750-005634	DS5815	1x CHOC12 IQ SONET, SMIR

PIC 2	REV 09	750-002911	CF4539	4x F/E, 100 BASE-TX
PIC 3	REV 08	750-021652	DR2827	1x CHOC12 IQE SONET
Xcvr 0		NON-JNPR	8	UNKNOWN
MMB 0	REV 04	710-016036	DR0809	ST-MMB2
FPC 5	REV 07	710-007529	HS5608	FPC Type 3
CPU	REV 15	710-001726	HX4351	FPC CPU
PIC 0	REV 14	750-009567	WJ8961	1x 10GE(LAN), XENPAK
Xcvr 0	REV 01	740-013170	J05K05961	XENPAK-LR
PIC 1	REV 16	750-007141	JJ8146	10x 1GE(LAN), 1000 BASE
Xcvr 1	REV 01	740-011613	P9F117T	SFP-SX
Xcvr 2	REV 01	740-011782	PBA2VCL	SFP-SX
Xcvr 3	REV 01	740-011782	PB83DRB	SFP-SX
Xcvr 4	REV 01	740-011613	AM0812S8UP8	SFP-SX
PIC 2	REV 12	750-009567	WF3566	1x 10GE(LAN), XENPAK
Xcvr 0	REV 02	740-013170	T07C94489	XENPAK-LR
MMB 0	REV 03	710-005555	HZ1907	MMB-288mbit
MMB 1	REV 03	710-005555	HW5283	MMB-288mbit
PPB 0	REV 04	710-002845	HZ7717	PPB Type 3
PPB 1	REV 04	710-002845	HS0110	PPB Type 3
FPC 6	REV 07	710-013035	DP7486	FPC Type 3-ES
CPU	REV 08	710-016744	DP2545	ST-PMB2
PIC 0	REV 09	750-009567	NE6323	1x 10GE(LAN), XENPAK
Xcvr 0	REV 02	740-013170	T09C71959	XENPAK-LR
PIC 1	REV 06	750-015217	DN4775	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011782	P7E0T6M	SFP-SX
Xcvr 1	REV 01	740-011613	AM0812S8XAY	SFP-SX
Xcvr 2	REV 01	740-011782	P7E0T6J	SFP-SX
Xcvr 3	REV 01	740-011782	PCH2P7D	SFP-SX
Xcvr 4	REV 01	740-011782	P9B0QYT	SFP-SX
Xcvr 5	REV 01	740-011613	AM0812S8WQJ	SFP-SX
Xcvr 6	REV 02	740-013111	9301220	SFP-T
Xcvr 7	REV 01	740-011782	P9B0TZ5	SFP-SX
PIC 2	REV 06	750-015217	DM6747	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011613	PAP0ZB2	SFP-SX
Xcvr 1	REV 01	740-013111	70191002	SFP-T
Xcvr 6	REV 01	740-011782	PBA29H8	SFP-SX
Xcvr 7	REV 01	740-011613	AM0812S8WQG	SFP-SX
MMB 0	REV 04	710-016036	DP3238	ST-MMB2
FPC 7	REV 03	710-021540	DV3154	FPC Type 2-ES
CPU	REV 09	710-016744	DT9053	ST-PMB2
PIC 0	REV 13	750-001901	HB4225	4x OC-12 SONET, SMIR
PIC 1	REV 05	750-001900	AD3644	1x OC-48 SONET, SMSR
PIC 2	REV 10	750-008155	HV0335	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011782	PCH2UKF	SFP-SX
Xcvr 1	REV 01	740-011782	PCH2V19	SFP-SX
PIC 3	REV 03	750-014638	JS9493	1x OC-48-12-3 SFP
Xcvr 0	REV 01	740-011785	P6Q0ENK	SFP-SR
MMB 0	REV 05	710-016036	DP3323	ST-MMB2
SPMB 0	REV 04	710-023321	DX3004	LCC Switch CPU
SPMB 1	REV 04	710-023321	DX3009	LCC Switch CPU
SIB 0	REV 07	710-022594	DW4195	LCC SIB
B Board	REV 07	710-023185	DW3930	LCC SIB Mezz
SIB 1	REV 07	710-022594	DW4179	LCC SIB
B Board	REV 07	710-023185	DW3919	LCC SIB Mezz
SIB 2				
SIB 3	REV 06	710-022594	DT8251	LCC SIB
B Board	REV 06	710-023185	DT5792	LCC SIB Mezz
SIB 4	REV 08	710-022594	DW8014	LCC SIB
B Board	REV 07	710-023185	DW3917	LCC SIB Mezz
Fan Tray 0				Front Top Fan Tray

Fan Tray 1
Fan Tray 2

Front Bottom Fan Tray
Rear Fan Tray -- Rev 3

lcc1-re0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN1102270AHA	T1600
Midplane	REV 04	710-017247	RC5358	T-series Backplane
FPM GBUS	REV 10	710-002901	DS3443	T640 FPM Board
FPM Display	REV 01	710-021387	DS6411	T1600 FPM Display
CIP	REV 06	710-002895	DS4235	T-series CIP
PEM 0	Rev 02	740-023211	VM82438	Power Entry Module 4x60A
SCG 0	REV 15	710-003423	DS6649	T640 Sonet Clock Gen.
SCG 1	REV 15	710-003423	DR6775	T640 Sonet Clock Gen.
Routing Engine 0	REV 01	740-026941	737F-1083	RE-DUO-1800
Routing Engine 1	REV 01	740-026941	737F-1104	RE-DUO-1800
CB 0	REV 06	710-022597	DW8542	LCC Control Board
CB 1	REV 06	710-022597	DW8530	LCC Control Board
FPC 0	REV 02	710-010845	JE2392	FPC Type 4
CPU	REV 02	710-011481	JF6820	FPC CPU-Enhanced
PIC 0	REV 11	750-017405	DP7259	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 01	740-014279	AA0741N1C8T	XFP-10G-LR
Xcvr 1	REV 01	740-014279	AA0746N1GAM	XFP-10G-LR
Xcvr 2	REV 01	740-014279	AA0747N1H0B	XFP-10G-LR
Xcvr 3	REV 01	740-014279	AA0748N1HZ5	XFP-10G-LR
MMB 0	REV 03	710-010842	HY7601	ST-MMB
FPC 1	REV 16	710-013037	BBAA7398	FPC Type 4-ES
CPU	REV 09	710-016744	BBAA2329	ST-PMB2
PIC 0	REV 03	711-029996	EB1575	100GE
PIC 1	REV 06	750-034781	EB9980	100GE CFP
MMB 0	REV 04	710-025563	BBAA5325	ST-MMB2
MMB 1	REV 04	710-025563	BBAA5444	ST-MMB2
FPC 2	REV 16	710-013037	BBAA7185	FPC Type 4-ES
CPU	REV 09	710-016744	BBAA3522	ST-PMB2
PIC 0	REV 03	711-029996	EB1557	100GE
PIC 1	REV 05	750-034781	EB4660	100GE CFP
Xcvr 0	REV 0	740-032210	J10F73666	CFP-100G-LR4
BRIDGE 0	REV 02	711-029995	EB2237	100GE Bridge Board
MMB 0	REV 04	710-025563	BBAA5347	ST-MMB2
MMB 1	REV 04	710-025563	BBAA5401	ST-MMB2
FPC 3	REV 10	710-021534	DZ0941	FPC Type 1-ES
CPU	REV 09	710-016744	DY6364	ST-PMB2
PIC 0	REV 13	750-012266	DK9192	4x 1GE(LAN), IQ2
Xcvr 0	REV 01	740-011613	AM0812S8WVD	SFP-SX
Xcvr 1		NON-JNPR	PDD63Q4	SFP-SX
Xcvr 2		NON-JNPR	PDE4G54	SFP-SX
Xcvr 3		NON-JNPR	PD40MAG	SFP-SX
PIC 1	REV 01	750-007641	HJ2003	1x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	AM0812S8WVG	SFP-SX
PIC 3	REV 17	750-007444	JB6873	1x CHSTM1 IQ SDH, SMIR
MMB 0	REV 04	710-025563	DZ0281	ST-MMB2
FPC 4	REV 06	710-013035	DK0614	FPC Type 3-ES
CPU	REV 07	710-016744	DK1616	ST-PMB2
PIC 0	REV 22	750-007141	DM1870	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011782	PCL3UKW	SFP-SX
Xcvr 1	REV 01	740-011782	P7E0T73	SFP-SX
Xcvr 2	REV 01	740-007326	P4TOWLRL	SFP-SX
Xcvr 3	REV 01	740-011782	PAR1LRL	SFP-SX
Xcvr 4	REV 01	740-011782	P9MOU3Z	SFP-SX

Xcvr 5	REV 01	740-011782	P9M0U0C	SFP-SX
Xcvr 6	REV 01	740-011782	P9M0TLG	SFP-SX
Xcvr 7	REV 01	740-011782	P9M0U0F	SFP-SX
Xcvr 8	REV 01	740-011613	PFA6LAP	SFP-SX
Xcvr 9	REV 01	740-011782	PCH2P0U	SFP-SX
PIC 1	REV 16	750-009450	CV2565	1x OC-192 SM SR2
PIC 2	REV 05	750-004424	HH3057	1x 10GE(LAN),10GBASE-LR
PIC 3	REV 12	750-013423	DP0403	MultiServices 500
MMB 0	REV 04	710-016036	DK1988	ST-MMB2
FPC 5	REV 07	710-013560	DR0004	E2-FPC Type 3
CPU	REV 05	710-013563	DR0089	FPC CPU-Enhanced
PIC 0	REV 11	750-012793	DR6107	1x 10GE(LAN/WAN) IQ2
Xcvr 0	REV 01	740-014289	C743XU074	XFP-10G-SR
PIC 1	REV 01	750-004695	HD5980	1x Tunnel
PIC 2	REV 32	750-003700	DL3770	1x OC-192 12xMM VSR
PIC 3	REV 12	750-009553	WB8901	4x OC-48 SONET
Xcvr 0	REV 01	740-011785	P9D1GTQ	SFP-SR
Xcvr 1	REV 01	740-011785	PDSOMMB	SFP-SR
Xcvr 3	REV 01	740-011785	PDE1KXP	SFP-SR
MMB 0	REV 07	710-010171	DP7374	MMB-5M3-288mbit
MMB 1	REV 07	710-010171	DP7404	MMB-5M3-288mbit
FPC 6	REV 07	710-013035	DM0994	FPC Type 3-ES
CPU	REV 07	710-016744	DM3651	ST-PMB2
PIC 0	REV 07	750-015217	DN4743	8x 1GE(TYPE3), IQ2
Xcvr 3	REV 01	740-011613	AM0812S8XB0	SFP-SX
Xcvr 4	REV 01	740-011782	PB829RB	SFP-SX
Xcvr 5	REV 01	740-011782	P8J1SYX	SFP-SX
PIC 1	REV 03	750-003336	HJ9954	4x OC-48 SONET, SMSR
PIC 3	REV 02	750-012793	JM7665	1x 10GE(LAN/WAN) IQ2
MMB 0	REV 04	710-016036	DN6913	ST-MMB2
FPC 7	REV 08	710-010845	JM3958	FPC Type 4
CPU	REV 04	710-011481	JK3669	FPC CPU-Enhanced
PIC 0	REV 11	750-017405	DP8837	4x 10GE (LAN/WAN) XFP
Xcvr 1	REV 01	740-014279	753019A00277	XFP-10G-LR
Xcvr 2	REV 02	740-011571	C850XJ00P	XFP-10G-SR
Xcvr 3	REV 01	740-014279	AA0813N1RTG	XFP-10G-LR
MMB 0	REV 04	710-010842	JN1971	ST-MMB
SPMB 0	REV 04	710-023321	DW3629	LCC Switch CPU
SPMB 1	REV 04	710-023321	DW3621	LCC Switch CPU
SIB 0	REV 07	710-022594	DW4200	LCC SIB
B Board	REV 07	710-023185	DW3932	LCC SIB Mezz
SIB 1	REV 07	710-022594	DW4193	LCC SIB
B Board	REV 07	710-023185	DW3904	LCC SIB Mezz
SIB 2				
SIB 3	REV 07	710-022594	DW4210	LCC SIB
B Board	REV 06	710-023185	DT5780	LCC SIB Mezz
SIB 4	REV 08	710-022594	DW8019	LCC SIB
B Board	REV 06	710-023185	DT5795	LCC SIB Mezz
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 3

show chassis hardware sfc (TX Matrix Plus Router)

```
user@host> show chassis hardware sfc 0
sfc0-re0:
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN112F007AHB	TXP

Midplane	REV 05	710-022574	TS4027	SFC Midplane
FPM Display	REV 03	710-024027	DX0282	TXP FPM Display
CIP 0	REV 04	710-023792	DW4889	TXP CIP
CIP 1	REV 04	710-023792	DW4887	TXP CIP
PEM 0	Rev 07	740-027463	UM26368	Power Entry Module
Routing Engine 0	REV 01	740-026942	737A-1064	SFC RE
Routing Engine 1	REV 01	740-026942	737A-1082	SFC RE
CB 0	REV 09	710-022606	DW6099	SFC Control Board
CB 1	REV 09	710-022606	DW6096	SFC Control Board
SPMB 0		BUILTIN		SFC Switch CPU
SPMB 1		BUILTIN		SFC Switch CPU
SIB F13 0	REV 04	710-022600	DX0841	F13 SIB
B Board	REV 03	710-023431	DX0966	F13 SIB Mezz
SIB F13 1	REV 04	750-024564	DW5776	F13 SIB
B Board	REV 03	710-023431	DW9028	F13 SIB
SIB F13 3	REV 04	750-024564	DW5762	F13 SIB
B Board	REV 03	710-023431	DW9059	F13 SIB
SIB F13 4	REV 04	750-024564	DW5797	F13 SIB
B Board	REV 03	710-023431	DW9041	F13 SIB
SIB F13 6	REV 04	750-024564	DW5770	F13 SIB
B Board	REV 03	710-023431	DW9079	F13 SIB Mezz
SIB F13 7	REV 04	750-024564	DW5758	F13 SIB
B Board	REV 03	710-023431	DW9047	F13 SIB
SIB F13 8	REV 04	750-024564	DW5761	F13 SIB
B Board	REV 03	710-023431	DW9043	F13 SIB Mezz
SIB F13 9	REV 04	750-024564	DW5754	F13 SIB
B Board	REV 03	710-023431	DW9078	F13 SIB Mezz
SIB F13 11	REV 04	710-022600	DX0826	F13 SIB
B Board	REV 03	710-023431	DX0967	F13 SIB Mezz
SIB F13 12	REV 04	750-024564	DW5794	F13 SIB
B Board	REV 03	710-023431	DW9044	F13 SIB Mezz
SIB F2S 0/0	REV 05	710-022603	DW7897	F2S SIB
B Board	REV 05	710-023787	DW7657	NEO PMB
SIB F2S 0/2	REV 05	710-022603	DW7833	F2S SIB
B Board	REV 05	710-023787	DW7526	NEO PMB
SIB F2S 0/4	REV 05	710-022603	DW7875	F2S SIB
B Board	REV 05	710-023787	DW7588	NEO PMB
SIB F2S 0/6	REV 05	710-022603	DW7860	F2S SIB
B Board	REV 05	710-023787	DW7589	NEO PMB
SIB F2S 1/0	REV 04	710-022603	DW4820	F2S SIB
B Board	REV 05	710-023787	DW8510	NEO PMB
SIB F2S 1/2	REV 05	710-022603	DW7849	F2S SIB
B Board	REV 05	710-023787	DW7525	NEO PMB
SIB F2S 1/4	REV 05	710-022603	DW7927	F2S SIB
B Board	REV 05	710-023787	DW7556	F2S SIB Mezz
SIB F2S 1/6	REV 05	710-022603	DW7866	F2S SIB
B Board	REV 05	710-023787	DW7651	NEO PMB
SIB F2S 2/0	REV 05	710-022603	DW7880	F2S SIB
B Board	REV 05	710-023787	DW7523	NEO PMB
SIB F2S 2/2	REV 05	710-022603	DW7895	F2S SIB
B Board	REV 05	710-023787	DW7591	NEO PMB
SIB F2S 2/4	REV 05	710-022603	DW7907	F2S SIB
B Board	REV 05	710-023787	DW7590	NEO PMB
SIB F2S 2/6	REV 05	710-022603	DW7785	F2S SIB
B Board	REV 05	710-023787	DW7524	NEO PMB
SIB F2S 3/0	REV 05	710-022603	DW7782	F2S SIB
B Board	REV 05	710-023787	DW7634	NEO PMB
SIB F2S 3/2	REV 05	710-022603	DW7793	F2S SIB
B Board	REV 05	710-023787	DW7548	NEO PMB
SIB F2S 3/4	REV 05	710-022603	DW7779	F2S SIB
B Board	REV 05	710-023787	DW7587	NEO PMB

SIB F2S 3/6	REV 05	710-022603	DW7930	F2S SIB
B Board	REV 05	710-023787	DW7505	NEO PMB
SIB F2S 4/0	REV 05	710-022603	DW7867	F2S SIB
B Board	REV 05	710-023787	DW7656	NEO PMB
SIB F2S 4/2	REV 05	710-022603	DW7917	F2S SIB
B Board	REV 05	710-023787	DW7640	NEO PMB
SIB F2S 4/4	REV 05	710-022603	DW7929	F2S SIB
B Board	REV 05	710-023787	DW7643	NEO PMB
SIB F2S 4/6	REV 05	710-022603	DW7870	F2S SIB
B Board	REV 05	710-023787	DW7635	NEO PMB
Fan Tray 0	REV 06	760-024497	DV7831	Front Fan Tray
Fan Tray 1	REV 06	760-024497	DV9614	Front Fan Tray
Fan Tray 2	REV 06	760-024502	DV9618	Rear Fan Tray
Fan Tray 3	REV 06	760-024502	DV9616	Rear Fan Tray
Fan Tray 4	REV 06	760-024502	DV7807	Rear Fan Tray
Fan Tray 5	REV 06	760-024502	DV7828	Rear Fan Tray

show chassis hardware extensive (TX Matrix Plus Router)

```
user@host> show chassis hardware extensive
```

```
sfc0-re0:
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN112F007AHB	TXP

JeDEC Code:	0x7fb0	EEPROM Version:	0x02
		S/N:	JN112F007AHB
Assembly ID:	0x052c	Assembly Version:	00.00
Date:	00-00-0000	Assembly Flags:	0x00
ID:	TXP		

Board Information Record:

Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

I2C Hex Data:

Address 0x00: 7f b0 02 ff 05 2c 00 00 00 00 00 00 00 00 00 00

Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Address 0x20: 4a 4e 31 31 32 46 30 30 37 41 48 42 00 00 00 00

Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Midplane	REV 05	710-022574	TS4027	SFC Midplane
----------	--------	------------	--------	--------------

JeDEC Code:	0x7fb0	EEPROM Version:	0x01
P/N:	710-022574	S/N:	S/N TS4027
Assembly ID:	0x0962	Assembly Version:	01.05
Date:	03-23-2009	Assembly Flags:	0x00
Version:	REV 05		

ID: SFC Midplane

Board Information Record:

Address 0x00: ad 01 ff ff 00 1d b5 14 00 00 ff ff ff ff ff ff

I2C Hex Data:

Address 0x00: 7f b0 01 ff 09 62 01 05 52 45 56 20 30 35 00 00

Address 0x10: 00 00 00 00 37 31 30 2d 30 32 32 35 37 34 00 00

Address 0x20: 53 2f 4e 20 54 53 34 30 32 37 00 00 00 17 03 07

Address 0x30: d9 ff ff ff ad 01 ff ff 00 1d b5 14 00 00 ff ff

Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff

Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

FPM Display	REV 03	710-024027	DX0282	TXP FPM Display
JeDEC Code:	0x7fb0	EEPROM Version:	0x01	

```

P/N:          710-024027      S/N:          S/N DX0282
Assembly ID:  0x096c         Assembly Version: 01.03
Date:         02-10-2009     Assembly Flags:  0x00
Version:      REV 03
ID: TXP FPM Display          FRU Model Number:  CRAFT-TXP
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 09 6c 01 03 52 45 56 20 30 33 00 00
Address 0x10: 00 00 00 00 37 31 30 2d 30 32 34 30 32 37 00 00
Address 0x20: 53 2f 4e 20 44 58 30 32 38 32 00 00 00 0a 02 07
Address 0x30: d9 ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 43
Address 0x50: 52 41 46 54 2d 54 58 50 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
CIP 0          REV 04      710-023792      DW4889          TXP CIP
Jedec Code:    0x7fb0      EEPROM Version:  0x01
P/N:          710-023792      S/N:          S/N DW4889
Assembly ID:   0x0969      Assembly Version: 01.04
Date:         01-26-2009     Assembly Flags:  0x00
Version:      REV 04
ID: TXP CIP          FRU Model Number:  CIP-TXP
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

```

show chassis hardware clei-models (TX Matrix Plus Router)

```

user@host> show chassis hardware clei-models
sfc0-re0:

```

```

-----
Hardware inventory:

```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 05	710-022574		CHAS-BP-TXP-S
FPM Display	REV 03	710-024027		CRAFT-TXP-S
CIP 0	REV 05	710-023792		CIP-TXP-S
CIP 1	REV 05	710-023792		CIP-TXP-S
PEM 0	Rev 04	740-027463	IPUPAFGKTA	PWR-TXP-7-60-DC
PEM 1	Rev 04	740-027463	IPUPAFGKTA	PWR-TXP-7-60-DC
Routing Engine 0	REV 06	740-026942		RE-DUO-C2600-16G-S
Routing Engine 1	REV 06	740-026942		RE-DUO-C2600-16G-S
CB 0	REV 05	710-022606		CB-TXP-S
CB 1	REV 09	710-022606		CB-TXP-S
SIB F13 0	REV 04	750-024564		SIB-TXP-F13
SIB F13 3	REV 04	750-024564		SIB-TXP-F13
SIB F13 8	REV 04	750-024564		SIB-TXP-F13
SIB F13 11	REV 04	750-024564		SIB-TXP-F13
SIB F13 12	REV 03	750-024564		SIB-TXP-F13
SIB F2S 0/0	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 0/2	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 0/4	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 0/6	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 1/0	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 1/2	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 1/4	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 1/6	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 2/0	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 2/2	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 2/4	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 2/6	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 3/0	REV 05	710-022603		SIB-TXP-F2S-S

SIB F2S 3/2	REV 03	710-022603	SIB-TXP-F2S-S
SIB F2S 3/4	REV 05	710-022603	SIB-TXP-F2S-S
SIB F2S 3/6	REV 03	710-022603	SIB-TXP-F2S-S
SIB F2S 4/0	REV 03	710-022603	SIB-TXP-F2S-S
SIB F2S 4/2	REV 05	710-022603	SIB-TXP-F2S-S
SIB F2S 4/4	REV 04	710-022603	SIB-TXP-F2S-S
SIB F2S 4/6	REV 03	710-022603	SIB-TXP-F2S-S
Fan Tray 0	REV 02	760-024497	FANTRAY-TXP-H-S
Fan Tray 1	REV 02	760-024497	FANTRAY-TXP-H-S
Fan Tray 2	REV 05	760-024502	FANTRAY-TXP-V-S
Fan Tray 3			
Fan Tray 4	REV 05	760-024502	FANTRAY-TXP-V-S
Fan Tray 5	REV 02	760-024502	FANTRAY-TXP-V-S

lcc0-re0:

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 03	710-017247		CHAS-BP-T1600-S
FPM Display	REV 01	710-021387		CRAFT-T1600-S
CIP	REV 06	710-002895		CIP-L-T640-S
PEM 1	Rev 02	740-023211	IPUPAC8KTA	PWR-T1600-4-60-DC-S
SCG 0	REV 15	710-003423		SCG-T-S
SCG 1	REV 15	710-003423		SCG-T-S
Routing Engine 0	REV 01	740-026941		RE-DUO-C1800-8G-S
Routing Engine 1	REV 01	740-026941		RE-DUO-C1800-8G-S
CB 0	REV 06	710-022597		CB-LCC-S
CB 1	REV 06	710-022597		CB-LCC-S
FPC 1	REV 07	710-013035		T640-FPC3-ES
PIC 0	REV 05	750-015217		PC-8GE-TYPE3-SFP-IQ2
PIC 1	REV 03	750-004424		PC-1XGE-LR
PIC 2	REV 01	750-003336		PC-40C48-SON-SMSR
FPC 3	REV 12	710-013037		T1600-FPC4-ES
PIC 0	REV 02	750-010850		PD-10C768-SON-SR
FPC 4	REV 05	710-021534		T640-FPC1-ES
PIC 0	REV 04	750-014627		PB-40C3-10C12-SON-SFP
PIC 1	REV 22	750-005634		PB-1CHOC12SMIR-QPP
PIC 2	REV 09	750-002911		PB-4FE-TX
PIC 3	REV 08	750-021652		PB-1CHOC12-STM4-IQE-SFP
FPC 5	REV 07	710-007529		T640-FPC3
PIC 0	REV 14	750-009567		PC-1XGE-XENPAK
PIC 1	REV 16	750-007141		PC-10GE-SFP
PIC 2	REV 12	750-009567		PC-1XGE-XENPAK
FPC 6	REV 07	710-013035		T640-FPC3-ES
PIC 0	REV 09	750-009567		PC-1XGE-XENPAK
PIC 1	REV 06	750-015217		PC-8GE-TYPE3-SFP-IQ2
PIC 2	REV 06	750-015217		PC-8GE-TYPE3-SFP-IQ2
FPC 7	REV 03	710-021540		T640-FPC2-ES
PIC 0	REV 13	750-001901		PB-40C12-SON-SMIR
PIC 1	REV 05	750-001900		PB-10C48-SON-SMSR
PIC 2	REV 10	750-008155		PB-2GE-SFP-QPP
PIC 3	REV 03	750-014638		PB-10C48-SON-B-SFP
SIB 0	REV 07	710-022594		SIB-TXP-T1600-S
SIB 1	REV 07	710-022594		SIB-TXP-T1600-S
SIB 3	REV 06	710-022594		SIB-TXP-T1600-S
SIB 4	REV 08	710-022594		SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

lcc1-re0:

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 04	710-017247		CHAS-BP-T1600-S
FPM Display	REV 01	710-021387		CRAFT-T1600-S
CIP	REV 06	710-002895		CIP-L-T640-S
PEM 0	Rev 02	740-023211	IPUPAC8KTA	PWR-T1600-4-60-DC-S
SCG 0	REV 15	710-003423		SCG-T-S
SCG 1	REV 15	710-003423		SCG-T-S
Routing Engine 0	REV 01	740-026941		RE-DUO-C1800-8G-S
Routing Engine 1	REV 01	740-026941		RE-DUO-C1800-8G-S
CB 0	REV 06	710-022597		CB-LCC-S
CB 1	REV 06	710-022597		CB-LCC-S
FPC 0	REV 02	710-010845		T640-FPC4-ES
PIC 0	REV 11	750-017405		PD-4XGE-XFP
FPC 1	REV 16	710-013037		T1600-FPC4-ES
PIC 1	REV 06	750-034781		PD-1CE-CFP
FPC 2	REV 16	710-013037		T1600-FPC4-ES
PIC 1	REV 05	750-034781		PD-1CE-CFP
FPC 3	REV 10	710-021534		T640-FPC1-ES
PIC 0	REV 13	750-012266		PB-4GE-TYPE1-SFP-IQ2
PIC 1	REV 01	750-007641		PE-1GE-SFP-QPP
PIC 3	REV 17	750-007444		PB-1CHSTM1-SMIR-QPP
FPC 4	REV 06	710-013035		T640-FPC3-ES
PIC 0	REV 22	750-007141		PC-10GE-SFP
PIC 1	REV 16	750-009450		PC-10C192-SON-SR2
PIC 2	REV 05	750-004424		PC-1XGE-LR
PIC 3	REV 12	750-013423		PC-MS-500-3
FPC 5	REV 07	710-013560		T640-FPC3-E2
PIC 0	REV 11	750-012793		PC-1XGE-TYPE3-XFP-IQ2
PIC 1	REV 01	750-004695		PC-TUNNEL
PIC 2	REV 32	750-003700		PC-10C192-SON-VSR
PIC 3	REV 12	750-009553		PC-40C48-SON-SFP
FPC 6	REV 07	710-013035		T640-FPC3-ES
PIC 0	REV 07	750-015217		PC-8GE-TYPE3-SFP-IQ2
PIC 1	REV 03	750-003336		PC-40C48-SON-SMSR
PIC 3	REV 02	750-012793		PC-1XGE-TYPE3-XFP-IQ2
FPC 7	REV 08	710-010845		T640-FPC4-ES
PIC 0	REV 11	750-017405		PD-4XGE-XFP
SIB 0	REV 07	710-022594		SIB-TXP-T1600-S
SIB 1	REV 07	710-022594		SIB-TXP-T1600-S
SIB 3	REV 07	710-022594		SIB-TXP-T1600-S
SIB 4	REV 08	710-022594		SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

show chassis hardware detail (TX Matrix Plus Router)

```
user@host> show chassis hardware detail
sfc0-re0:
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN111B023AHB	TXP
Midplane	REV 01	710-022574	TR7990	SFC Midplane
FPM Display	REV 03	710-024027	DW4699	TXP FPM Display
CIP 0	REV 01	710-023792	DR1437	TXP CIP
CIP 1	REV 02	710-023792	DS4564	TXP CIP
PEM 0	Rev 07	740-027463	UM26360	Power Entry Module
Routing Engine 0	REV 01	740-026942	737A-1024	SFC RE

```

ad0    3887 MB SMART CF                200811050193CEB1CEB1 Compact Flash
ad1    30533 MB SAMSUNG MCBQE32G8MPP-0V SY814A0762    Disk 1
Routing Engine 1 REV 01 740-026942 737A-1024    SFC RE
ad0    3887 MB SMART CF                20081105004C19A019A0 Compact Flash
ad1    30533 MB SAMSUNG MCBQE32G8MPP-0V SY814A0794    Disk 1
CB 0          REV 03 710-022606 DR7134    SFC Control Board
CB 1          REV 01 710-022606 DP8890    SFC Control Board
SPMB 0          BUILTIN                SFC Switch CPU
SPMB 1          BUILTIN                SFC Switch CPU
SIB F13 0      REV 03 750-024564 DT9478    F13 SIB
  B Board      REV 02 710-023431 DT6554    F13 SIB
SIB F13 1      REV 03 750-024564 DT9454    F13 SIB
  B Board      REV 02 710-023431 DT6551    F13 SIB
SIB F2S 0/0    REV 02 710-022603 DT2838    F2S SIB
  B Board      REV 02 710-023787 DT1725    NEO PMB
SIB F2S 0/2    REV 02 710-022603 DT2824    F2S SIB
  B Board      REV 02 710-023787 DT1706    NEO PMB
SIB F2S 0/4    REV 02 710-022603 DT2822    F2S SIB
  B Board      REV 02 710-023787 DT1696    NEO PMB
SIB F2S 0/6    REV 02 710-022603 DT2823    F2S SIB
  B Board      REV 02 710-023787 DT1717    NEO PMB
SIB F2S 1/0    REV 03 710-022603 DV0059    F2S SIB
  B Board      REV 03 710-023787 DT9942    NEO PMB
SIB F2S 1/2    REV 02 710-022603 DT2826    F2S SIB
  B Board      REV 02 710-023787 DT1713    NEO PMB
SIB F2S 1/4    REV 03 710-022603 DV0092    F2S SIB
  B Board      REV 03 710-023787 DV0000    NEO PMB
SIB F2S 1/6    REV 03 710-022603 DV0079    F2S SIB
  B Board      REV 03 710-023787 DT9972    NEO PMB
SIB F2S 2/0    REV 03 710-022603 DV0100    F2S SIB
  B Board      REV 03 710-023787 DT9925    NEO PMB
SIB F2S 2/2    REV 03 710-022603 DV0050    F2S SIB
  B Board      REV 03 710-023787 DV0005    NEO PMB
SIB F2S 2/4    REV 03 710-022603 DV0097    F2S SIB
  B Board      REV 03 710-023787 DT9936    NEO PMB
Fan Tray 0     REV 02 760-024497 DR8286    Front Fan Tray
Fan Tray 1     REV 06 760-024497 DV9624    Front Fan Tray
Fan Tray 2     REV 02 760-024502 DR8259    Rear Fan Tray
Fan Tray 3     REV 02 760-024502 DR8270    Rear Fan Tray
Fan Tray 4     REV 02 760-024502 DR8284    Rear Fan Tray
Fan Tray 5     REV 06 760-024502 DV7813    Rear Fan Tray

```

lcc0-re0:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN1101F27AHA	T1600
Midplane	REV 04	710-017247	RC5317	T Series Backplane
FPM GBUS	REV 10	710-002901	DS8197	T640 FPM Board
FPM Display	REV 01	710-021387	DS6433	T1600 FPM Display
CIP	REV 06	710-002895	DS1493	T Series CIP
PEM 0	Rev 08	740-017906	UD26601	Power Entry Module 3x80
SCG 0	REV 15	710-003423	DP5847	T640 Sonet Clock Gen.
SCG 1	REV 15	710-003423	DR0924	T640 Sonet Clock Gen.
Routing Engine 0	REV 01	740-026942	737F-1024	LCC RE
ad0	3887 MB SMART CF		2008110502B63E513E51	Compact Flash
ad1	30533 MB SAMSUNG MCBQE32G8MPP-0V SY814A1208			Disk 1
Routing Engine 1	REV 01	740-026942	737F-1024	LCC RE
ad0	3887 MB SMART CF		2008110500F9A8A8A8A8	Compact Flash
ad1	30533 MB SAMSUNG MCBQE32G8MPP-0V SY814A1076			Disk 1
CB 0	REV 05	710-022597	DV4264	LCC Control Board

CB 1	REV 03	710-022597	DP8558	LCC Control Board
FPC 0	REV 14	710-013037	DS9967	FPC Type 4-ES
CPU	REV 08	710-016744	DS3989	ST-PMB2
PIC 0	REV 12	750-013198	DL7506	1x Tunnel
PIC 1	REV 12	750-013198	DL7505	1x Tunnel
MMB 0	REV 01	710-025563	DS8524	ST-MMB2
MMB 1	REV 01	710-025563	DS8373	ST-MMB2
FPC 1	REV 14	710-013037	DT0027	FPC Type 4-ES
CPU	REV 09	710-016744	DS7684	ST-PMB2
PIC 0	REV 12	750-013198	DL7512	1x Tunnel
PIC 1	REV 12	750-013198	DL7498	1x Tunnel
MMB 0	REV 01	710-025563	DS8494	ST-MMB2
MMB 1	REV 01	710-025563	DS8436	ST-MMB2
SPMB 0	REV 04	710-023321	DV3867	LCC Switch CPU
SPMB 1	REV 02	710-023321	DP0238	LCC Switch CPU
SIB 0	REV 06	710-022594	DT8268	LCC SIB
B Board	REV 06	710-023185	DT5791	LCC SIB Mezz
SIB 1	REV 06	710-022594	DT8261	LCC SIB
B Board	REV 06	710-023185	DT5769	LCC SIB Mezz
SIB 2	REV 04	710-022594	DS2315	LCC SIB
B Board	REV 06	710-023185	DT5788	LCC SIB Mezz
SIB 3	REV 06	710-022594	DT8253	LCC SIB
B Board	REV 06	710-023185	DT5811	LCC SIB Mezz
SIB 4	REV 06	710-022594	DT8248	LCC SIB
B Board	REV 06	710-023185	DT5812	LCC SIB Mezz
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray

show chassis hardware models (TX Matrix Plus Router)

```
user@host> show chassis hardware models
sfc0-re0:
```

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
FPM Display	REV 03	710-024027	DX0282	CRAFT-TXP
CIP 0	REV 04	710-023792	DW4889	CIP-TXP
CIP 1	REV 04	710-023792	DW4887	CIP-TXP
PEM 0	Rev 07	740-027463	UM26368	yyyyyyyyyyyyyyyyyyyy
Routing Engine 0	REV 01	740-026942	737A-1064	RE-TXP-SFC-DU0-2600-16G
Routing Engine 1	REV 01	740-026942	737A-1082	RE-TXP-SFC-DU0-2600-16G
CB 0	REV 09	710-022606	DW6099	CB-TXP
CB 1	REV 09	710-022606	DW6096	CB-TXP
SIB F13 1	REV 04	750-024564	DW5776	SIB-TXP-F13
SIB F13 3	REV 04	750-024564	DW5762	SIB-TXP-F13
SIB F13 4	REV 04	750-024564	DW5797	SIB-TXP-F13
SIB F13 6	REV 04	750-024564	DW5770	SIB-TXP-F13
SIB F13 7	REV 04	750-024564	DW5758	SIB-TXP-F13
SIB F13 8	REV 04	750-024564	DW5761	SIB-TXP-F13
SIB F13 9	REV 04	750-024564	DW5754	SIB-TXP-F13
SIB F13 12	REV 04	750-024564	DW5794	SIB-TXP-F13
SIB F2S 0/0	REV 05	710-022603	DW7897	
SIB F2S 0/2	REV 05	710-022603	DW7833	
SIB F2S 0/4	REV 05	710-022603	DW7875	
SIB F2S 0/6	REV 05	710-022603	DW7860	
SIB F2S 1/0	REV 04	710-022603	DW4820	
SIB F2S 1/2	REV 05	710-022603	DW7849	
SIB F2S 1/4	REV 05	710-022603	DW7927	SIB-TXP-F2S
SIB F2S 1/6	REV 05	710-022603	DW7866	
SIB F2S 2/0	REV 05	710-022603	DW7880	

SIB F2S 2/2	REV 05	710-022603	DW7895	
SIB F2S 2/4	REV 05	710-022603	DW7907	
SIB F2S 2/6	REV 05	710-022603	DW7785	
SIB F2S 3/0	REV 05	710-022603	DW7782	
SIB F2S 3/2	REV 05	710-022603	DW7793	
SIB F2S 3/4	REV 05	710-022603	DW7779	
SIB F2S 3/6	REV 05	710-022603	DW7930	
SIB F2S 4/0	REV 05	710-022603	DW7867	
SIB F2S 4/2	REV 05	710-022603	DW7917	
SIB F2S 4/4	REV 05	710-022603	DW7929	
SIB F2S 4/6	REV 05	710-022603	DW7870	
Fan Tray 0	REV 06	760-024497	DV7831	FANTRAY-TXP-F
Fan Tray 1	REV 06	760-024497	DV9614	FANTRAY-TXP-F
Fan Tray 2	REV 06	760-024502	DV9618	FANTRAY-TXP-R
Fan Tray 3	REV 06	760-024502	DV9616	FANTRAY-TXP-R
Fan Tray 4	REV 06	760-024502	DV7807	FANTRAY-TXP-R
Fan Tray 5	REV 06	760-024502	DV7828	FANTRAY-TXP-R

lcc0-re0:

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 03	710-017247	RC3765	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	DN5441	CRAFT-T1600-S
CIP	REV 06	710-002895	DP6021	CIP-L-T640-S
PEM 0	Rev 07	740-017906	UA26384	PWR-T1600-3-80-DC-S
PEM 1	Rev 07	740-017906	UA26296	PWR-T1600-3-80-DC-S
SCG 0	REV 15	710-003423	DR0875	SCG-T-S
CB 0	REV 06	710-022597	DW8534	CB-LCC
CB 1	REV 06	710-022597	DW8527	CB-LCC
FPC 4	REV 12	710-013037	DJ8717	T1600-FPC4-ES
PIC 0	REV 11	750-017405	DP8795	PD-4XGE-XFP
PIC 1	REV 11	750-017405	DP8794	PD-4XGE-XFP
FPC 6	REV 14	710-013037	DS5335	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS7634	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS7637	PD-4XGE-XFP
FPC 7	REV 07	710-013035	DM0990	T1600-FPC3-ES
PIC 0	REV 16	750-007141	JJ8067	PC-10GE-SFP
PIC 1	REV 08	750-015749	WE9598	PC-10C192-SON-XFP
PIC 2	REV 10	750-009450	HX6466	PC-10C192-SON-SR2
SIB 0	REV 08	710-022594	DW8033	SIB-TXP-T1600-S
SIB 1	REV 08	710-022594	DW8044	SIB-TXP-T1600-S
SIB 2	REV 08	710-022594	DW8020	SIB-TXP-T1600-S
SIB 3	REV 08	710-022594	DW8063	SIB-TXP-T1600-S
SIB 4	REV 08	710-022594	DW8064	SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

lcc1-re0:

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 04	710-017247	RC5361	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	DS6430	CRAFT-T1600-S
CIP	REV 06	710-002895	DS4239	CIP-L-T640-S
PEM 0	Rev 08	740-017906	UD26649	PWR-T1600-3-80-DC-S
SCG 0	REV 15	710-003423	DP5820	SCG-T-S
CB 0	REV 06	710-022597	DW8523	CB-LCC
CB 1	REV 06	710-022597	DW8528	CB-LCC
FPC 4	REV 12	710-013037	DP8509	T1600-FPC4-ES

PIC 0	REV 11	750-017405	DP8808	PD-4XGE-XFP
PIC 1	REV 11	750-017405	DP7263	PD-4XGE-XFP
FPC 6	REV 14	710-013037	DS9961	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS5532	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS7639	PD-4XGE-XFP
FPC 7	REV 03	710-013035	DF5564	T1600-FPC3-ES
PIC 0	REV 16	750-007141	JJ8063	PC-10GE-SFP
SIB 0	REV 08	710-022594	DW8035	SIB-TXP-T1600-S
SIB 1	REV 10	710-022594	DX7672	SIB-TXP-T1600-S
SIB 2	REV 08	710-022594	DW8060	SIB-TXP-T1600-S
SIB 3	REV 08	710-022594	DW8072	SIB-TXP-T1600-S
SIB 4	REV 08	710-022594	DW8043	SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

lcc2-re0:

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 03	710-017247	RC3956	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	DN7030	CRAFT-T1600-S
CIP	REV 06	710-002895	DM3962	CIP-L-T640-S
PEM 0	Rev 08	740-017906	UD26519	PWR-T1600-3-80-DC-S
PEM 1	Rev 07	740-017906	UC26601	PWR-T1600-3-80-DC-S
SCG 0	REV 15	710-003423	DP0277	SCG-T-S
CB 0	REV 06	710-022597	DW8524	CB-LCC
CB 1	REV 06	710-022597	DW8536	CB-LCC
FPC 4	REV 12	710-013037	DR1194	T1600-FPC4-ES
PIC 0	REV 11	750-017405	DP8811	PD-4XGE-XFP
PIC 1	REV 11	750-017405	DP8823	PD-4XGE-XFP
FPC 5	REV 12	710-013037	DR1184	T1600-FPC4-ES
PIC 1	REV 11	750-017405	DP4744	PD-4XGE-XFP
FPC 6	REV 12	710-013037	DN8622	T1600-FPC4-ES
PIC 0	REV 14	750-012518	JY9924	PD-40C192-SON-XFP
PIC 1	REV 11	750-017405	DP8776	PD-4XGE-XFP
FPC 7	REV 04	710-013560	JR3968	T640-FPC3-E2
PIC 0	REV 16	750-007141	NC9330	PC-10GE-SFP
SIB 0	REV 07	710-022594	DW4217	SIB-TXP-T1600-S
SIB 1	REV 07	710-022594	DW4213	SIB-TXP-T1600-S
SIB 2	REV 07	710-022594	DW4189	SIB-TXP-T1600-S
SIB 3	REV 07	710-022594	DW4173	SIB-TXP-T1600-S
SIB 4	REV 07	710-022594	DW4201	SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

lcc3-re0:

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 04	710-017247	RC5319	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	DS6402	CRAFT-T1600-S
CIP	REV 06	710-002895	DR9973	CIP-L-T640-S
PEM 0	Rev 07	740-017906	UC26496	PWR-T1600-3-80-DC-S
PEM 1	Rev 07	740-017906	UC26599	PWR-T1600-3-80-DC-S
SCG 0	REV 15	710-003423	DP5831	SCG-T-S
CB 0	REV 06	710-022597	DW8533	CB-LCC
CB 1	REV 06	710-022597	DW8538	CB-LCC
FPC 0	REV 14	710-013037	DS5345	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS7641	PD-4XGE-XFP

PIC 1	REV 13	750-017405	DS5479	PD-4XGE-XFP
FPC 1	REV 14	710-013037	DS7338	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS7631	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS7632	PD-4XGE-XFP
FPC 2	REV 14	710-013037	DS9962	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS7581	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS7627	PD-4XGE-XFP
FPC 4	REV 10	710-010845	JZ6573	T640-FPC4-ES
PIC 0	REV 14	750-012518	JT5124	PD-40C192-SON-XFP
FPC 5	REV 14	710-013037	DT0016	T1600-FPC4-ES
PIC 0	REV 14	750-012518	JY9918	PD-40C192-SON-XFP
FPC 7	REV 07	710-013035	DM0967	T1600-FPC3-ES
PIC 0	REV 16	750-007141	JJ8059	PC-10GE-SFP
PIC 1	REV 13	750-004695	DM5712	PC-TUNNEL
SIB 0	REV 07	710-022594	DW4174	SIB-TXP-T1600-S
SIB 1	REV 07	710-022594	DW4207	SIB-TXP-T1600-S
SIB 2	REV 06	710-022594	DT8231	SIB-TXP-T1600-S
SIB 3	REV 07	710-022594	DW4175	SIB-TXP-T1600-S
SIB 4	REV 07	710-022594	DW4209	SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

show chassis hardware (16-Port 10-Gigabit Ethernet MPC with SFP+ Optics [MX Series Routers])

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN112D865AFA	MX960
Midplane			TS3339	MX960 Backplane
FPM Board	REV 03	710-013698	WW6267	Front Panel Display
PDM	Rev 03	740-013110	QCS12485026	Power Distribution
Module				
PEM 0	Rev 04	740-013682	QCS12434086	PS 1.7kW; 200-240VAC
in				
PEM 1	Rev 04	740-013682	QCS1243408Z	PS 1.7kW; 200-240VAC
in				
PEM 2	Rev 04	740-013682	QCS1243407X	PS 1.7kW; 200-240VAC
in				
Routing Engine 0	REV 07	740-015113	9009009677	RE-S-1300
Routing Engine 1	REV 07	740-015113	9009011510	RE-S-1300
CB 0	REV 03	710-021523	XF0394	MX SCB
CB 1	REV 03	710-021523	XF0550	MX SCB
CB 2	REV 03	710-021523	XD7455	MX SCB
FPC 4	REV 02	750-028467	JR6127	MPC M 16x 10GE
CPU	REV 02	711-029089	JX0129	AS PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Fan Tray 0	REV 05	740-014971	TP9990	Fan Tray
Fan Tray 1	REV 05	740-014971	VS1709	Fan Tray

show chassis hardware (MPC3E [MX Series Routers])

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN1101AFEAFB	MX480
Midplane	REV 05	710-017414	TR4444	MX480 Midplane

FPM Board	REV 02	710-017254	KG6056	Front Panel Display
PEM 0	Rev 03	740-017330	QCS082090FC	PS 1.2-1.7kW; 100-240V
PEM 1	Rev 03	740-017330	QCS082090FD	PS 1.2-1.7kW; 100-240V
Routing Engine 0	REV 07	740-013063	9009004124	RE-S-2000
Routing Engine 1	REV 07	740-013063	9009005569	RE-S-2000
CB 0	REV 07	710-021523	XZ3587	MX SCB
CB 1	REV 03	710-021523	KH8306	MX SCB
FPC 1	REV 04.1.07	750-033205	P1240	MPCE Type 3D
CPU	REV 01	711-035209	YL0504	HMPC PMB 2G
MIC 1	REV 10	750-033199	YX4495	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	C22CQNE	CFP-100G-LR4
FPC 2	REV 26	750-016670	KH0045	DPCE 40x 1GE R EQ
CPU	REV 07	710-013713	KF5448	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PF21JHU	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 9	REV 01	740-011613	AM0813S8ZL6	SFP-SX
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 02	740-011613	PGL2KYF	SFP-SX
Xcvr 2	REV 01	740-011613	AM0806S8N4P	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 5	REV 01	740-011613	AM0815S967N	SFP-SX
Xcvr 7	REV 01	740-011613	AM0806S8N1X	SFP-SX
Xcvr 8	REV 01	740-011613	AM0815S967J	SFP-SX
Xcvr 9	REV 01	740-011613	AM0815S967M	SFP-SX
FPC 3	REV 12.2.09	750-033205	YR9443	MPCE Type 3D
CPU	REV 03	711-035209	YL6931	HMPC PMB 2G
MIC 0	REV 05	750-033199	YR3269	1X100GE CFP
PIC 0		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	ULHOKG3	CFP-100G-LR4
MIC 1	REV 02	750-033199	YG3245	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	ULHOKGF	CFP-100G-LR4
FPC 4	REV 12.3.09	750-033205	YR9437	MPCE Type 3D
CPU	REV 03	711-035209	YT5857	HMPC PMB 2G
MIC 0	REV 05	750-033199	YR3295	1X100GE CFP
PIC 0		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	X12000187	CFP-100G-SR10
MIC 1	REV 10	750-033199	YX4518	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	X12J00008	CFP-100G-SR10
FPC 5	REV 06	750-024884	JW9769	MPC Type 2 3D EQ
CPU	REV 02	711-028401	JR6158	MPC PMB 2G Proto
MIC 0	REV 05	750-028387	JR6197	3D 4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 01	740-014289	T07M71112	XFP-10G-SR
Xcvr 1	REV 02	740-014289	T08L85610	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	2x 10GE XFP
MIC 1	REV 22	750-028392	YM0053	3D 20x 1GE(LAN) SFP
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-011613	AM0703S005B	SFP-SX
Xcvr 1	REV 01	740-011613	E07L01352	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 5	REV 01	740-013111	6500217	SFP-T

```

Xcvr 9      REV 02    740-013111    8499527      SFP-T
Fan Tray    Left Fan Tray

```

The PIC number for MIC 1 always starts from 2 (even if the first MIC is a 1X100GE CFP or a legacy MIC).

show chassis hardware (QFX3500 Switches)

```

user@switch> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               QFX3500
Routing Engine 0                       BUILTIN      BUILTIN      QFX Routing Engine
FPC 0          REV 04    750-044071   BBAR3902     QFX3500-48S4Q-AFI
  CPU          BUILTIN      BUILTIN      FPC CPU
  PIC 0        BUILTIN      BUILTIN      48x 10G-SFP+
  PIC 1        BUILTIN      BUILTIN      15x 10G-SFP+
  MGMT BRD     REV 02    750-044063   BBAR0398     QFX3500-MGMT-SFP-AFO
    Xcvr 0     REV 01    740-011614   AC0946S0BD1  SFP-LX10
    Xcvr 1     REV 02    740-013111   A281922      SFP-T
Power Supply 0  Rev 04    740-032091   UI00677      JPSU-650W-AC-AFI
Power Supply 1  REV 00    740-041741   VJ00162      JPSU-650W-AC-AFO
Fan Tray 0                               QFX Fan Tray, Back to
Front Airflow
Fan Tray 1                               QFX Fan Tray, Back to
Front Airflow
Fan Tray 2                               QFX Fan Tray, Back to
Front Airflow

```

show chassis hardware detail (QFX3500 Switches)

```

user@switch> show chassis hardware detail
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               QFX3500
Routing Engine 0                       BUILTIN      BUILTIN      QFX Routing Engine
FPC 0          REV 05    750-036931   EE0823       QFX3500-48S4Q-AFI

CPU          BUILTIN      BUILTIN      FPC CPU
PIC 0        BUILTIN      BUILTIN      48x 10G-SFP+
  Xcvr 0     REV 01    740-030589   S99E270079   SFP+-10G-LPBK
  Xcvr 1     REV 01    740-030589   S9AK450099   SFP+-10G-LPBK
  Xcvr 2     REV 01    740-030589   S99E270078   SFP+-10G-LPBK
  Xcvr 3     REV 01    740-030589   S9AK450098   SFP+-10G-LPBK
  Xcvr 4     REV 01    740-030589   S99E270075   SFP+-10G-LPBK
  Xcvr 5     REV 01    740-030589   S9AK450093   SFP+-10G-LPBK
  Xcvr 6     REV 01    740-030589   S9AK450097   SFP+-10G-LPBK
  Xcvr 7     REV 01    740-030589   S9AK450095   SFP+-10G-LPBK
  Xcvr 8     REV 01    740-030589   S99E270072   SFP+-10G-LPBK
  Xcvr 9     REV 01    740-030589   S99E270073   SFP+-10G-LPBK
  Xcvr 10    REV 01    740-030589   S99E270080   SFP+-10G-LPBK
  Xcvr 11    REV 01    740-030589   S9AK450169   SFP+-10G-LPBK
  Xcvr 12    REV 01    740-030589   S99E270076   SFP+-10G-LPBK
  Xcvr 13    REV 01    740-030589   S9AK450167   SFP+-10G-LPBK
  Xcvr 14    REV 01    740-030589   S9AK450170   SFP+-10G-LPBK
  Xcvr 15    REV 01    740-030589   S9AK450166   SFP+-10G-LPBK
  Xcvr 16    REV 01    740-030589   S9AK450092   SFP+-10G-LPBK
  Xcvr 17    REV 01    740-030589   S9AK450163   SFP+-10G-LPBK
  Xcvr 18    REV 01    740-030589   S9AK450094   SFP+-10G-LPBK
  Xcvr 19    REV 01    740-030589   S9AK450100   SFP+-10G-LPBK

```

Xcvr 20	REV 01	740-030589	S9AK450168	SFP+-10G-LPBK
Xcvr 21	REV 01	740-030589	S9AK450165	SFP+-10G-LPBK
Xcvr 22	REV 01	740-030589	S9AK450073	SFP+-10G-LPBK
Xcvr 23	REV 01	740-030589	S9AK450164	SFP+-10G-LPBK
Xcvr 24	REV 01	740-030589	S9AK450074	SFP+-10G-LPBK
Xcvr 25	REV 01	740-030589	SA62270195	SFP+-10G-LPBK
Xcvr 26	REV 01	740-030589	S9AK450078	SFP+-10G-LPBK
Xcvr 27	REV 01	740-030589	S9AK450024	SFP+-10G-LPBK
Xcvr 28	REV 01	740-030589	S9AK450027	SFP+-10G-LPBK
Xcvr 29	REV 01	740-030589	S9AK450080	SFP+-10G-LPBK
Xcvr 30	REV 01	740-030589	S9AK450030	SFP+-10G-LPBK
Xcvr 31	REV 01	740-030589	S9AK450025	SFP+-10G-LPBK
Xcvr 32	REV 01	740-030589	S9AK450023	SFP+-10G-LPBK
Xcvr 33	REV 01	740-030589	S9AK450075	SFP+-10G-LPBK
Xcvr 34	REV 01	740-030589	S9AK450161	SFP+-10G-LPBK
Xcvr 35	REV 01	740-030589	S9AK450071	SFP+-10G-LPBK
Xcvr 36	REV 01	740-030589	S9AK450072	SFP+-10G-LPBK
Xcvr 37	REV 01	740-030589	S9AK450022	SFP+-10G-LPBK
Xcvr 38	REV 01	740-030589	S9AK450021	SFP+-10G-LPBK
Xcvr 39	REV 01	740-030589	S9AK450175	SFP+-10G-LPBK
Xcvr 40	REV 01	740-030589	S9AK450162	SFP+-10G-LPBK
Xcvr 41	REV 01	740-030589	S99E270074	SFP+-10G-LPBK
Xcvr 42	REV 01	740-030589	S9AK450174	SFP+-10G-LPBK
Xcvr 43	REV 01	740-030589	S9AK450077	SFP+-10G-LPBK
Xcvr 44	REV 01	740-030589	S9AK450076	SFP+-10G-LPBK
Xcvr 45	REV 01	740-030589	S9AK450026	SFP+-10G-LPBK
Xcvr 46	REV 01	740-030589	S9AK450079	SFP+-10G-LPBK
Xcvr 47	REV 01	740-030589	S9AK450029	SFP+-10G-LPBK
PIC 1		BUILTIN	BUILTIN	15x 10G-SFP+
Xcvr 1	REV 01	740-032986	QA170087	QSFP+-40G-SR4
Xcvr 4	REV 01	740-032986	QA360442	QSFP+-40G-SR4
Xcvr 8	REV 01	740-032986	QA170091	QSFP+-40G-SR4
Xcvr 12	REV 01	740-032986	QA170042	QSFP+-40G-SR4
MGMT BRD	REV 08	750-036946	EE0731	QFX3500-MB
Power Supply 0	Rev 04	740-032091	UI00690	QFX PS 650W AC
Power Supply 1	Rev 04	740-032091	UI00679	QFX PS 650W AC
Fan Tray 0				QFX Fan Tray
Fan Tray 1				QFX Fan Tray

show chassis hardware models (QFX3500 Switches)

```

user@switch> show chassis hardware models
Hardware inventory:
Item          Version  Part number  Serial number  FRU model number
Routing Engine 0          BUILTIN    BUILTIN
FPC 0          REV 02    711-032234  EC4074
Power Supply 0  PSMI 2C  11-d65800  --

```

show chassis hardware clei-models (QFX3500 Switches)

```

user@switch> show chassis hardware clei-models
Hardware inventory:
Item          Version  Part number  CLEI code      FRU model number
Routing Engine 0          BUILTIN
FPC 0          REV 02    711-032234
Power Supply 0  PSMI 2C  11-d65800

```

show chassis hardware interconnect-device (QFabric Systems)

```

user@switch> show chassis hardware interconnect-device interconnect1
Hardware inventory:
Item          Version  Part number  Serial number  Description

```

```

Chassis          REV 07
Midplane         REV 07  750-021261  BH0208188289  QFX_olive
CB 0             REV 07  750-021261  BH0208188289  QFX Midplane
                                         QFXIC08-CB4S

```

show chassis hardware node-device (QFabric Systems)

```

user@switch> show chassis hardware node-device node1
Routing Engine 0  BUILTIN      BUILTIN      QFX Routing Engine
node1            REV 05      711-032234    ED3694        QFX3500-48S4Q-AFI

CPU              BUILTIN      BUILTIN
PIC 0            BUILTIN      BUILTIN
Xcvr 8          REV 01      740-030658    AD0946A028B   FPC CPU
                                         48x 10G-SFP+
                                         SFP+-10G-USR
...

```

show chassis hardware (PTX5000 Packet Transport Switch)

```

user@switch> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN11D1FD7AJA  PTX5000
Midplane      REV 03    711-031896    ABAC5589       Midplane-8S
FPM           REV 08    760-030647    EG1679         Front Panel Display
PDU 0         Rev 05    740-032019    ZE00006        DC Power Dist Unit
  PSM 0        Rev 05    740-032022    ZJ00018        DC 12V Power Supply
  PSM 1        Rev 04    740-032022    ZC00052        DC 12V Power Supply
  PSM 2        Rev 04    740-032022    ZD00051        DC 12V Power Supply
  PSM 3        Rev 05    740-032022    ZJ00060        DC 12V Power Supply
CCG 0         REV 04    750-030653    EG3703         Clock Generator
CCG 1         REV 04    750-030653    EG3698         Clock Generator
Routing Engine 0 REV 05    740-026942    P737A-002231   RE-DUO-2600
Routing Engine 1 REV 06    740-026942    P737A-002438   RE-DUO-2600
CB 0          REV 08    750-030625    EG5519         Control Board
CB 1          REV 08    750-030625    EG5516         Control Board
FPC 0         REV 18    750-036844    EJ3080         FPC
  CPU         REV 12    711-030686    EJ3260         SNG PMB
FPC 2         REV 13    750-036844    EG5065         FPC
  CPU         REV 09    711-030686    EG4082         SNG PMB
  PIC 0        REV 14    750-031913    EG5127         24x 10GE(LAN) SFP+
    Xcvr 0      REV 01    740-031980    143363A00240   SFP+-10G-SR
    Xcvr 1      REV 01    740-031981    UK90PZ1        SFP+-10G-LR
    Xcvr 2      REV 01    740-031980    AD1141A04XH    SFP+-10G-SR
    Xcvr 3      REV 01    740-031981    UK90Q46        SFP+-10G-LR
    Xcvr 4      REV 01    740-031980    AD1141A04X4    SFP+-10G-SR
    Xcvr 6      REV 01    740-031980    B11H02560      SFP+-10G-SR
    Xcvr 7      REV 01    740-031980    B11C01589      SFP+-10G-SR
    Xcvr 8      REV 01    740-031980    AD1141A04XF    SFP+-10G-SR
    Xcvr 10     REV 01    740-031980    123363A01094   SFP+-10G-SR
    Xcvr 11     REV 01    740-031980    AK80LKF        SFP+-10G-SR
    Xcvr 12     REV 01    740-031980    183363A01528   SFP+-10G-SR
    Xcvr 14     REV 01    740-031980    193363A01079   SFP+-10G-SR
    Xcvr 15     REV 01    740-031980    AK80MC8        SFP+-10G-SR
    Xcvr 16     REV 01    740-031980    AJC0BHC        SFP+-10G-SR
    Xcvr 19     REV 01    740-021309    J08D26856      SFP+-10G-LR
    Xcvr 21     REV 01    740-031980    AK80KCT        SFP+-10G-SR
    Xcvr 22     REV 01    740-031981    UK90PZL        SFP+-10G-LR
    Xcvr 23     REV 01    740-031980    AK80N1V        SFP+-10G-SR
FPC 3         REV 13    750-036844    EG5074         FPC
  CPU         REV 09    711-030686    EG4064         SNG PMB
  PIC 1        REV 10    750-031903    EG0325         SNG Load

```

FPC 5	REV 06	750-036844	EH3198	FPC
CPU				
PIC 0	REV 14	750-031913	EG5134	24x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LBH	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11B03724	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FMH	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J00818	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00743	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11B06125	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B11H02529	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AK80LFB	SFP+-10G-SR
Xcvr 12	REV 01	740-031980	193363A01061	SFP+-10G-SR
Xcvr 15	REV 01	740-031980	B11J00687	SFP+-10G-SR
Xcvr 16	REV 01	740-031980	193363A00738	SFP+-10G-SR
Xcvr 18	REV 01	740-031980	AK80MQX	SFP+-10G-SR
Xcvr 19	REV 01	740-021309	J08C17257	SFP+-10G-LR
Xcvr 22	REV 01	740-031980	B11J00730	SFP+-10G-SR
Xcvr 23	REV 01	740-031980	AK80KEE	SFP+-10G-SR
PIC 1	REV 08	750-036710	EG3105	2x 40GE CFP
Xcvr 0	REV 01	740-034554	B260HLT	CFP-40G-LR4
Xcvr 1	REV 01	740-034554	B11C02847	CFP-40G-LR4
FPC 6	REV 18	750-036844	EJ4391	FPC
CPU	REV 12	711-030686	EJ3257	SNG PMB
FPC 7	REV 18	750-036844	EJ4382	FPC
CPU	REV 12	711-030686	EJ3238	SNG PMB
SPMB 0	REV 10	711-030686	EG5418	SNG PMB
SPMB 1	REV 09	711-030686	EG5373	SNG PMB
SIB 0	REV 07	750-030631	EG4858	SIB-I-8S
SIB 1	REV 07	750-030631	EG4872	SIB-I-8S
SIB 2	REV 07	750-030631	EG4866	SIB-I-8S
SIB 3	REV 07	750-030631	EG6011	SIB-I-8S
SIB 4	REV 07	750-030631	EG4907	SIB-I-8S
SIB 5	REV 07	750-030631	EG4879	SIB-I-8S
SIB 6	REV 07	750-030631	EG4864	SIB-I-8S
SIB 7	REV 07	750-030631	EG4899	SIB-I-8S
SIB 8	REV 07	750-030631	EG4880	SIB-I-8S
Fan Tray 0	REV 04	760-032784	EG1496	Vertical Fan Tray
Fan Tray 1	REV 04	760-030642	EG1335	Horizontal Fan Tray
Fan Tray 2	REV 02	760-030642	ED4952	Horizontal Fan Tray

show chassis hardware clei-models (PTX5000 Packet Transport Switch)

```
user@switch> show chassis hardware clei-models
```

Hardware inventory:				
Item	Version	Part number	CLEI code	FRU model number
FPM	REV 08	760-030647	PROTOXCLEI	CRAFT-PTX5000-S
PDU 0	Rev 05	740-032019	IPUPAHLKAA	PWR-SAN-PDU-DC
PSM 0	Rev 05	740-032022	IPUPAHNKAA	PSM-PTX-DC-120-S
PSM 1	Rev 04	740-032022	032022XXXX	PWR-SAN-12-DC
PSM 2	Rev 04	740-032022	032022XXXX	PWR-SAN-12-DC
PSM 3	Rev 05	740-032022	IPUPAHNKAA	PSM-PTX-DC-120-S
CCG 0	REV 04	750-030653	PROTOXCLEI	CCG-PTX-S
CCG 1	REV 04	750-030653	PROTOXCLEI	CCG-PTX-S
Routing Engine 0	REV 05	740-026942		RE-DUO-C2600-16G-S
Routing Engine 1	REV 06	740-026942		RE-DUO-C2600-16G-S
CB 0	REV 08	750-030625	PROTOXCLEI	CB-PTX-S
CB 1	REV 08	750-030625	PROTOXCLEI	CB-PTX-S
FPC 0	REV 18	750-036844	PROTOXCLEI	FPC-PTX-P1-A
FPC 2	REV 13	750-036844	PROTOXCLEI	FPC-PTX-P1-A
PIC 0	REV 14	750-031913	PROTOXCLEI	P1-PTX-24-10GE-SFPP
FPC 3	REV 13	750-036844	PROTOXCLEI	FPC-PTX-P1-A

FPC 5				
PIC 0	REV 14	750-031913	PROTOXCLEI	P1-PTX-24-10GE-SFPP
FPC 6	REV 18	750-036844	PROTOXCLEI	FPC-PTX-P1-A
FPC 7	REV 18	750-036844	PROTOXCLEI	FPC-PTX-P1-A
SIB 0	REV 07	750-030631	PROTOXCLEI	SIB-I-PTX5008
SIB 1	REV 07	750-030631	PROTOXCLEI	SIB-I-PTX5008
SIB 2	REV 07	750-030631	PROTOXCLEI	SIB-I-PTX5008
SIB 3	REV 07	750-030631	PROTOXCLEI	SIB-I-PTX5008
SIB 4	REV 07	750-030631	PROTOXCLEI	SIB-I-PTX5008
SIB 5	REV 07	750-030631	PROTOXCLEI	SIB-I-PTX5008
SIB 6	REV 07	750-030631	PROTOXCLEI	SIB-I-PTX5008
SIB 7	REV 07	750-030631	PROTOXCLEI	SIB-I-PTX5008
SIB 8	REV 07	750-030631	PROTOXCLEI	SIB-I-PTX5008
Fan Tray 1	REV 04	760-030642	PROTOXCLEI	FAN-PTX-H-S

show chassis hardware detail (PTX5000 Packet Transport Switch)

```

user@switch> show chassis hardware detail
Hardware inventory:
Item                Version  Part number  Serial number  Description
Chassis              REV 03   711-031896   JN1D1FD7AJA   PTX5000
Midplane             REV 08   760-030647   EG1679        Midplane-8S
FPM                  REV 05   740-032019   ZE00006       Front Panel Display
PDU 0                Rev 05   740-032022   ZJ00018       DC Power Dist Unit
  PSM 0               Rev 04   740-032022   ZC00052       DC 12V Power Supply
  PSM 1               Rev 04   740-032022   ZD00051       DC 12V Power Supply
  PSM 2               Rev 05   740-032022   ZJ00060       DC 12V Power Supply
CCG 0                REV 04   750-030653   EG3703        Clock Generator
CCG 1                REV 04   750-030653   EG3698        Clock Generator
Routing Engine 0     REV 05   740-026942   P737A-002231  RE-DUO-2600
  ad0 3823 MB SMART CF 201006190039C02DC02D Compact Flash
  ad1 62720 MB SMART Lite SATA Drive 2011042300CF4C6B4C6B Disk 1
Routing Engine 1     REV 06   740-026942   P737A-002438  RE-DUO-2600
  ad0 3823 MB SMART CF 20100619053455F055F0 Compact Flash
  ad1 62720 MB SMART Lite SATA Drive 20110423000AE8E7E8E7 Disk 1
CB 0                 REV 08   750-030625   EG5519        Control Board
CB 1                 REV 08   750-030625   EG5516        Control Board
FPC 0                REV 18   750-036844   EJ3080        FPC
  CPU                REV 12   711-030686   EJ3260        SNG PMB
FPC 2                REV 13   750-036844   EG5065        FPC
  CPU                REV 09   711-030686   EG4082        SNG PMB
  PIC 0              REV 14   750-031913   EG5127        24x 10GE(LAN) SFP+
    Xcvr 0            REV 01   740-031980   143363A00240  SFP+-10G-SR
    Xcvr 1            REV 01   740-031981   UK90PZ1       SFP+-10G-LR
    Xcvr 2            REV 01   740-031980   AD1141A04XH   SFP+-10G-SR
    Xcvr 3            REV 01   740-031981   UK90Q46       SFP+-10G-LR
    Xcvr 4            REV 01   740-031980   AD1141A04X4   SFP+-10G-SR
    Xcvr 6            REV 01   740-031980   B11H02560     SFP+-10G-SR
    Xcvr 7            REV 01   740-031980   B11C01589     SFP+-10G-SR
    Xcvr 8            REV 01   740-031980   AD1141A04XF   SFP+-10G-SR
    Xcvr 10           REV 01   740-031980   123363A01094  SFP+-10G-SR
    Xcvr 11           REV 01   740-031980   AK80LKF       SFP+-10G-SR
    Xcvr 12           REV 01   740-031980   183363A01528  SFP+-10G-SR
    Xcvr 14           REV 01   740-031980   193363A01079  SFP+-10G-SR
    Xcvr 15           REV 01   740-031980   AK80MC8       SFP+-10G-SR
    Xcvr 16           REV 01   740-031980   AJC0BHC       SFP+-10G-SR
    Xcvr 19           REV 01   740-021309   J08D26856     SFP+-10G-LR
    Xcvr 21           REV 01   740-031980   AK80KCT       SFP+-10G-SR
    Xcvr 22           REV 01   740-031981   UK90PZL       SFP+-10G-LR
    Xcvr 23           REV 01   740-031980   AK80N1V       SFP+-10G-SR

```

FPC 3	REV 13	750-036844	EG5074	FPC
CPU	REV 09	711-030686	EG4064	SNG PMB
PIC 1	REV 10	750-031903	EG0325	SNG Load
FPC 5	REV 06	750-036844	EH3198	FPC
CPU				
PIC 0	REV 14	750-031913	EG5134	24x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LBH	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11B03724	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FMH	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J00818	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00743	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11B06125	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B11H02529	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AK80LFB	SFP+-10G-SR
Xcvr 12	REV 01	740-031980	193363A01061	SFP+-10G-SR
Xcvr 15	REV 01	740-031980	B11J00687	SFP+-10G-SR
Xcvr 16	REV 01	740-031980	193363A00738	SFP+-10G-SR
Xcvr 18	REV 01	740-031980	AK80MQX	SFP+-10G-SR
Xcvr 19	REV 01	740-021309	J08C17257	SFP+-10G-LR
Xcvr 22	REV 01	740-031980	B11J00730	SFP+-10G-SR
Xcvr 23	REV 01	740-031980	AK80KEE	SFP+-10G-SR
PIC 1	REV 08	750-036710	EG3105	2x 40GE CFP
Xcvr 0	REV 01	740-034554	B260HLT	CFP-40G-LR4
Xcvr 1	REV 01	740-034554	B11C02847	CFP-40G-LR4
FPC 6	REV 18	750-036844	EJ4391	FPC
CPU	REV 12	711-030686	EJ3257	SNG PMB
FPC 7	REV 18	750-036844	EJ4382	FPC
CPU	REV 12	711-030686	EJ3238	SNG PMB
SPMB 0	REV 10	711-030686	EG5418	SNG PMB
SPMB 1	REV 09	711-030686	EG5373	SNG PMB
SIB 0	REV 07	750-030631	EG4858	SIB-I-8S
SIB 1	REV 07	750-030631	EG4872	SIB-I-8S
SIB 2	REV 07	750-030631	EG4866	SIB-I-8S
SIB 3	REV 07	750-030631	EG6011	SIB-I-8S
SIB 4	REV 07	750-030631	EG4907	SIB-I-8S
SIB 5	REV 07	750-030631	EG4879	SIB-I-8S
SIB 6	REV 07	750-030631	EG4864	SIB-I-8S
SIB 7	REV 07	750-030631	EG4899	SIB-I-8S
SIB 8	REV 07	750-030631	EG4880	SIB-I-8S
Fan Tray 0	REV 04	760-032784	EG1496	Vertical Fan Tray
Fan Tray 1	REV 04	760-030642	EG1335	Horizontal Fan Tray
Fan Tray 2	REV 02	760-030642	ED4952	Horizontal Fan Tray

show chassis hardware models (PTX5000 Packet Transport Switch)

user@switch> show chassis hardware models

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
FPM	REV 08	760-030647	EG1679	CRAFT-PTX5000-S
PDU 0	Rev 05	740-032019	ZE00006	PWR-SAN-PDU-DC
PSM 0	Rev 05	740-032022	ZJ00018	PSM-PTX-DC-120-S
PSM 1	Rev 04	740-032022	ZC00052	PWR-SAN-12-DC
PSM 2	Rev 04	740-032022	ZD00051	PWR-SAN-12-DC
PSM 3	Rev 05	740-032022	ZJ00060	PSM-PTX-DC-120-S
CCG 0	REV 04	750-030653	EG3703	CCG-PTX-S
CCG 1	REV 04	750-030653	EG3698	CCG-PTX-S
Routing Engine 0	REV 05	740-026942	P737A-002231	RE-DUO-C2600-16G-S
Routing Engine 1	REV 06	740-026942	P737A-002438	RE-DUO-C2600-16G-S
CB 0	REV 08	750-030625	EG5519	CB-PTX-S
CB 1	REV 08	750-030625	EG5516	CB-PTX-S
FPC 0	REV 18	750-036844	EJ3080	FPC-PTX-P1-A

FPC 2	REV 13	750-036844	EG5065	FPC-PTX-P1-A
PIC 0	REV 14	750-031913	EG5127	P1-PTX-24-10GE-SFPP
FPC 3	REV 13	750-036844	EG5074	FPC-PTX-P1-A
FPC 5				
PIC 0	REV 14	750-031913	EG5134	P1-PTX-24-10GE-SFPP
FPC 6	REV 18	750-036844	EJ4391	FPC-PTX-P1-A
FPC 7	REV 18	750-036844	EJ4382	FPC-PTX-P1-A
SIB 0	REV 07	750-030631	EG4858	SIB-I-PTX5008
SIB 1	REV 07	750-030631	EG4872	SIB-I-PTX5008
SIB 2	REV 07	750-030631	EG4866	SIB-I-PTX5008
SIB 3	REV 07	750-030631	EG6011	SIB-I-PTX5008
SIB 4	REV 07	750-030631	EG4907	SIB-I-PTX5008
SIB 5	REV 07	750-030631	EG4879	SIB-I-PTX5008
SIB 6	REV 07	750-030631	EG4864	SIB-I-PTX5008
SIB 7	REV 07	750-030631	EG4899	SIB-I-PTX5008
SIB 8	REV 07	750-030631	EG4880	SIB-I-PTX5008
Fan Tray 1	REV 04	760-030642	EG1335	FAN-PTX-H-S

show chassis hardware extensive (PTX5000 Packet Transport Switch)

```

user@switch> show chassis hardware extensive
Hardware inventory:
Item          Version  Part number  Serial number  Description
.....
PDU 0         Rev 04    740-032019   UE0003         DC Power Dist Unit
Jedec Code:   0x7fb0          EEPROM Version: 0x02
P/N:          740-032019      S/N:           S/N UE0003
Assembly ID:  0x043d          Assembly Version: 04.00
Date:         11-29-2010     Assembly Flags:  0x00
Version:      Rev 04         CLEI Code:      032022XXXX
ID: DC Power Dist Unit      FRU Model Number: PWR-SAN-PDU-DC
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 3d 04 00 52 65 76 20 30 34 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 33 32 30 31 39 00 00
Address 0x20: 53 2f 4e 20 55 45 30 30 30 33 00 00 00 1d 0b 07
Address 0x30: da ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 30 33 32 30 32 32 58 58 58 58 50
Address 0x50: 57 52 2d 53 41 4e 2d 50 44 55 2d 44 43 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 a3 ff ff ff ff ff ff ff ff ff ff ff ff
PSM 0         Rev 04    740-032022   YG00065        DC 12V Power Supply
Module
Jedec Code:   0x7fb0          EEPROM Version: 0x02
P/N:          740-032022      S/N:           S/N YG00065
Assembly ID:  0x0440          Assembly Version: 04.00
Date:         07-30-2010     Assembly Flags:  0x00
Version:      Rev 04         CLEI Code:      032022XXXX
ID: DC 12V Power Supply Module FRU Model Number: PWR-SAN-12-DC
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 40 04 00 52 65 76 20 30 34 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 33 32 30 32 32 00 00
Address 0x20: 53 2f 4e 20 59 47 30 30 30 36 35 00 00 1e 07 07
Address 0x30: da ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 30 33 32 30 32 32 58 58 58 58 50
Address 0x50: 57 52 2d 53 41 4e 2d 31 32 2d 44 43 20 20 20 20

```

```
Address 0x60: 20 20 20 20 20 20 01 00 ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff 0c ff ff ff ff ff ff ff ff ff ff ff
```

show chassis hardware (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number  Description
Chassis                               JN1100FB1AFB   MX480
Midplane         REV 05   710-017414   TR3310        MX480 Midplane
FPM Board        REV 02   710-017254   KG1872        Front Panel Display
PEM 2            Rev 02   740-017343   QCS0812A00N   DC Power Entry Module
PEM 3            Rev 02   740-017343   QCS0812A00U   DC Power Entry Module
Routing Engine 0 REV 07   740-015113   1000740938    RE-S-1300
CB 0             REV 03   710-021523   KF4630        MX SCB
FPC 1            REV 11   750-037207   ZW9726        AS-MCC
  CPU            REV 04   711-038173   ZW4819        AS-MCC PMB
  MIC 0          REV 06   750-037214   ZW3574        AS-MSC
    PIC 0                BUILTIN      BUILTIN       AS-MSC
  MIC 1          REV 00   750-037211                AS-MXC
    PIC 2                BUILTIN      BUILTIN       AS-MXC
```

show chassis hardware extensive (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis hardware extensive
FPC 1            REV 11   750-037207   ZW9726        AS-MCC
Jedec Code: 0x7fb0          EEPROM Version: 0x02
P/N: 750-037207            S/N: S/N ZW9726
Assembly ID: 0x0b37        Assembly Version: 01.11
Date: 02-17-2012          Assembly Flags: 0x00
Version: REV 11           CLEI Code: PROTOXCLEI
ID: AS-MCC                FRU Model Number: 750-037207
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 37 01 0b 52 45 56 20 31 31 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 30 37 00 00
Address 0x20: 53 2f 4e 20 5a 57 39 37 32 36 00 00 00 11 02 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 35 30 2d 30 33 37 32 30 37 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 31 31 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 5e ff ff ff ff ff ff ff ff ff ff ff ff
CPU            REV 04   711-038173   ZW4819        AS-MCC-PMB
Jedec Code: 0x7fb0          EEPROM Version: 0x02
P/N: 711-038173            S/N: S/N ZW4819
Assembly ID: 0x0b38        Assembly Version: 01.04
Date: 12-30-2011          Assembly Flags: 0x00
Version: REV 04
ID: AS-MCC PMB
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 38 01 04 52 45 56 20 30 34 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 38 31 37 33 00 00
Address 0x20: 53 2f 4e 20 5a 57 34 38 31 39 00 00 00 1e 0c 07
Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 31 31 2d 30 33 38 31 37 33 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 30 34 00 ff ff ff ff ff ff ff
```

```

Address 0x70: ff ff ff 60 00 00 00 00 00 00 00 00 00 00 00
MIC 0          REV 06    750-037214    ZW3574          AS-MS
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           750-037214      S/N:          S/N ZW3574
Assembly ID:   0x0a44          Assembly Version: 01.06
Date:          02-19-2012      Assembly Flags: 0x00
Version:       REV 06          CLEI Code:     PROTOXCLEI
ID: AS-MS      FRU Model Number: 750-037214
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0a 44 01 06 52 45 56 20 30 36 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 31 34 00 00
Address 0x20: 53 2f 4e 20 5a 57 33 35 37 34 00 00 00 13 02 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 35 30 2d 30 33 37 32 31 34 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 30 36 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 60 c0 03 e5 f4 00 00 00 00 00 00 00 00
PIC 0          BUILTIN      BUILTIN          AS-MS
MIC 1          REV 00    750-037211          AS-MXC
Jedec Code:    0x7fb0          EEPROM Version:    0x01
P/N:           750-037211
Assembly ID:   0x0a43          Assembly Version: 01.00
Date:          255-255-65535    Assembly Flags: 0x00
Version:       REV 00
ID: AS-MXC
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 0a 43 01 00 52 45 56 20 30 30 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 31 31 00 00
Address 0x20: 00 00 00 00 00 00 00 00 00 00 00 00 00 ff ff ff
Address 0x30: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff c0 02 e6 6c 7f b0 02 ff 0a 44 01 06
PIC 2          BUILTIN      BUILTIN          AS-MXC

```

show chassis location

Syntax	show chassis location
Syntax (TX Matrix Router)	show chassis location <fpc interface (by-name <i>name</i> by-slot fpc <i>number</i> lcc <i>number</i>) lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show chassis location <fpc interface (by-name <i>name</i> by-slot fpc <i>number</i> lcc <i>number</i>) lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Router)	show chassis location <all-members> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show chassis location <interconnect-device <i>name</i> > <node-device <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the physical location of the chassis. This command can only be used on the master Routing Engine.
Options	<p>none—Display all information about the physical location of the chassis. On a TX Matrix router, display all information about the physical location of the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display all information about the physical location of the TX Matrix Plus router and its attached T1600 routers.</p> <p>all-members—(MX Series routers only) (Optional) Display the physical location of the chassis for all the member routers in the Virtual Chassis configuration.</p> <p>fpc—(TX Matrix and TX Matrix Plus routers only) (Optional) Display the physical location of all Flexible PIC Concentrators (FPCs).</p> <p>interconnect-device <i>name</i>—(QFabric systems only) (Optional) Display the physical location of the Interconnect device.</p> <p>interface by-name <i>name</i>—(TX Matrix and TX Matrix Plus routers only) (Optional) Display the physical location of a specified interface name. On a TX Matrix router, this option displays the FPC number and T640 router (or line-card chassis) number associated with the specified interface. On a TX Matrix Plus router, this option displays the FPC number and T1600 router (or line-card chassis) number associated with the specified interface.</p> <p>interface by-slot fpc <i>number</i> lcc <i>number</i>—(TX Matrix and TX Matrix Plus router only) (Optional) On a TX Matrix router, display the global FPC number of an interface by</p>

specifying its local FPC number and T640 router (or line-card chassis) number. On a TX Matrix Plus router, display the global FPC number of an interface by specifying its local FPC number and T1600 router (or line-card chassis) number.

- The global FPC number is the FPC slot number when all the FPC slots in the routing matrix are considered: **0** through **31**. The local FPC number is the FPC slot number on a particular T640 router.
- For **fpc**, replace **number** with a value from **0** through **7**.
- For **lcc**, replace **number** with a value from **0** through **3**.

lcc number—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display the physical location of a specified T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display the physical location of a specified T1600 router (or line-card chassis) that is connected to a TX Matrix Plus router. Replace **number** with a value from **0** through **3**.

local—(MX Series routers only) (Optional) Display the physical location of the chassis for the local Virtual Chassis member.

member member-id—(MX Series routers only) (Optional) Display the physical location of the chassis for the specified member of the Virtual Chassis configuration. Replace **member-id** with a value of **0** or **1**.

node-device name—(QFabric systems only) (Optional) Display the physical location of the Node device.

scc—(TX Matrix routers only) (Optional) Display the physical location of the TX Matrix router (or switch-card chassis).

sfc—(TX Matrix Plus routers only) (Optional) Display the physical location of the TX Matrix Plus router (or switch-fabric chassis).

Required Privilege Level view

List of Sample Output [show chassis location on page 752](#)
[show chassis location fpc \(TX Matrix Router\) on page 752](#)
[show chassis location interface by-slot \(TX Matrix Router\) on page 752](#)
[show chassis location fpc \(TX Matrix Plus Router\) on page 752](#)
[show chassis location interface by-slot \(TX Matrix Plus Router\) on page 752](#)
[show chassis location \(QFX3500 Switches\) on page 752](#)
[show chassis location \(QFabric Systems\) on page 753](#)

Output Fields [Table 61 on page 751](#) lists the output fields for the **show chassis location** command. Output fields are listed in the approximate order in which they appear.

Table 61: show chassis location Output Fields

Field Name	Field Description
country-code	Country code information.

Table 61: show chassis location Output Fields (*continued*)

Field Name	Field Description
postal-code	Postal code information.
Building	Building information.
Floor	Floor information.
Global FPC	Global FPC number. The FPC slot number, when all FPC slots in the routing matrix are considered. The range of values is 0 through 31 .
LCC	Line-card chassis number. On a TX Matrix router, the number of a particular T640 router connected to the TX Matrix router. On a TX Matrix Plus router, the number of a particular T1600 router connected to the TX Matrix Plus router.
Local FPC	Local FPC number. On a TX Matrix router, the FPC slot number on a particular T640 router. On a TX Matrix Plus router, the FPC slot number on a particular T1600 router.

Sample Output

show chassis location

```
user@host> show chassis location
country-code: US
postal-code: 94404
Building: Building 2, Floor: 2
```

show chassis location fpc (TX Matrix Router)

```
user@host> show chassis location fpc
Global FPC    LCC    Local FPC
    17         2         1
    21         2         5
```

show chassis location interface by-slot (TX Matrix Router)

```
user@host> show chassis location interface by-slot fpc 1 lcc 1
Global FPC: 9
```

show chassis location fpc (TX Matrix Plus Router)

```
user@host> show chassis location fpc
Global FPC    LCC    Local FPC
    0         0         0
    1         0         1
```

show chassis location interface by-slot (TX Matrix Plus Router)

```
user@host> show chassis location interface by-slot fpc 2 lcc 1
Global FPC: 10
```

show chassis location (QFX3500 Switches)

```
user@switch> show chassis location
```

```
country-code: US  
postal-code: 94404  
Building: Building 2, Floor: 2
```

show chassis location (QFabric Systems)

```
user@switch> show chassis location interconnect-device interconnect1  
country-code: US  
postal-code: 94404  
Building: Building 2, Floor: 2
```

show chassis pic

Syntax	show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i>
Syntax (TX Matrix and TX Matrix Plus Routers)	show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i> <lcc <i>number</i> >
Syntax (MX Series Routers and EX Series Switches)	show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i> <all-members> <local> <member <i>member-id</i> >
Syntax (MX2010 and MX2010 3D Universal Edge Routers)	show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i>
Syntax (QFX Series)	show chassis pic <interconnect-device <i>name</i> (fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i>)> <node-device <i>name</i> pic-slot <i>slot-number</i> >
Syntax (ACX Series Universal Access Routers)	show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for QFX Series. Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.
Description	Display status information about the PIC installed in the specified Flexible PIC Concentrator (FPC) and PIC slot.
Options	<p>fpc-slot <i>slot-number</i>—Display information about the PIC in this particular FPC slot:</p> <ul style="list-style-type: none"> On a TX Matrix router, if you specify the number of the T640 router by using the lcc <i>number</i> option (the recommended method), replace <i>slot-number</i> with a value from 0 through 7. Otherwise, replace <i>slot-number</i> with a value from 0 through 31. <p>Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 router by using the lcc <i>number</i> option (the recommended method), replace <i>slot-number</i> with a value from 0 through 7. Otherwise, replace <i>slot-number</i> with a value from 0 through 31. For example, the following commands have the same result:</p> <pre> user@host> show chassis pic fpc-slot 1 lcc 1 pic-slot 1 user@host> show chassis pic fpc-slot 9 pic-slot 1 </pre> <ul style="list-style-type: none"> M120 routers only—Replace <i>slot-number</i> with a value from 0 through 5. MX80 routers only—Replace <i>slot-number</i> with a value from 0 through 1. MX240 routers only—Replace <i>slot-number</i> with a value from 0 through 2.

- MX480 routers only—Replace **slot-number** with a value from 0 through 5.
- MX960 routers only—Replace **slot-number** with a value from 0 through 11.
- MX2020 routers only—Replace **slot-number** with a value from 0 through 19.
- Other routers—Replace **slot-number** with a value from 0 through 7.
- EX Series switches:
 - EX3200 switches and EX4200 standalone switches—Replace **slot-number** with 0.
 - EX4200 switches in a Virtual Chassis configuration—Replace **slot-number** with a value from 0 through 9 (switch's member ID).
 - EX8208 switches—Replace **slot-number** with a value from 0 through 7 (line card).
 - EX8216 switches—Replace **slot-number** with a value from 0 through 15 (line card).
- QFX Series:
 - QFX3500 switches—Replace **slot-number** with 0. In the command output, FPC refers to a line card. The FPC number equals the slot number for the line card.
 - QFabric systems—Replace **slot-number** with any number between 0 and 15. In the command output, FPC refers to a line card. The FPC number equals the slot number for the line card.

all-members—(MX Series routers and EX Series switches only) (Optional) Display PIC information for all member routers in the Virtual Chassis configuration.

interconnect-device name—(QFabric systems only) (Optional) Display PIC information for a specified Interconnect device.

lcc number—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display PIC information for a specified T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display PIC information for a specified T1600 router (or line-card chassis) that is connected to the TX Matrix Plus router. Replace **number** with a value from 0 through 3.

local—(MX Series routers and EX Series switches only) (Optional) Display PIC information for the local Virtual Chassis member.

member member-id—(MX Series routers and EX Series switches only) (Optional) Display PIC information for the specified member of the Virtual Chassis configuration. Replace **member-id** with a value of 0 or 1.

node-device name—(QFabric systems only) (Optional) Display PIC information for a specified Node device.

pic-slot slot-number—Display information about the PIC in this particular PIC slot. For routers, replace **slot-number** with a value from 0 through 3. For EX3200 and EX4200

switches, replace *slot-number* with 0 for built-in network interfaces and 1 for interfaces on uplink modules. For EX8208 and EX8216 switches, replace *slot-number* with 0. For the QFX3500 standalone switch and the QFabric system, replace *slot-number* with 0 or 1.

Required Privilege Level view

Related Documentation

- *request chassis pic*
- [show chassis hardware on page 644](#)
- *Configuring the PIC Type*
- *100-Gigabit Ethernet PIC Overview*

List of Sample Output

- [show chassis pic fpc-slot pic-slot on page 758](#)
- [show chassis pic fpc-slot pic-slot \(PIC Offline\) on page 759](#)
- [show chassis pic fpc-slot pic-slot \(FPC Offline\) on page 759](#)
- [show chassis pic fpc-slot pic-slot \(FPC Not Present\) on page 759](#)
- [show chassis pic fpc-slot pic-slot \(PIC Not Present\) on page 759](#)
- [show chassis pic fpc-slot 3 pic-slot 0 \(M120 Router\) on page 759](#)
- [show chassis pic fpc-slot pic-slot \(MX960 Router with Bidirectional Optics\) on page 759](#)
- [show chassis pic fpc-slot pic-slot \(MX480 Router with 100-Gigabit Ethernet MIC\) on page 760](#)
- [show chassis pic fpc-slot pic-slot \(MX240, MX480, MX960 Routers with Application Services Modular Line Card\) on page 760](#)
- [show chassis pic fpc-slot pic-slot \(MX480 Router with MPC4E\) on page 760](#)
- [show chassis pic fpc-slot pic-slot \(MX2010 Router\) on page 760](#)
- [show chassis pic fpc-slot pic-slot \(MX2020 Router\) on page 760](#)
- [show chassis pic fpc-slot pic-slot \(MX2020 Router with MPC4E\) on page 761](#)
- [show chassis pic fpc-slot pic-slot \(T1600 Router with 100-Gigabit Ethernet PIC\) on page 761](#)
- [show chassis pic fpc-slot pic-slot lcc \(TX Matrix Router\) on page 761](#)
- [show chassis pic fpc-slot pic-slot lcc \(TX Matrix Plus Router\) on page 762](#)
- [show chassis pic fpc-slot pic-slot \(Next-Generation SONET/SDH SFP\) on page 762](#)
- [show chassis pic fpc-slot pic-slot \(12-Port T1/E1\) on page 762](#)
- [show chassis pic fpc-slot 0 pic-slot 1 \(4x CHOC3 SONET CE SFP\) on page 762](#)
- [show chassis pic fpc-slot 0 pic-slot 0 \(SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 763](#)
- [show chassis pic fpc-slot 3 pic-slot 0 \(8-port Channelized SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 763](#)
- [show chassis pic fpc-slot 5 pic-slot 0 \(4-port Channelized SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 763](#)
- [show chassis pic fpc-slot 1 pic-slot 0 \(1-port OC192/STM64 MIC with XFP\) on page 764](#)
- [show chassis pic fpc-slot 1 pic-slot 2 \(8-port DS3/E3 MIC\) on page 764](#)
- [show chassis pic fpc-slot pic-slot \(OTN\) on page 764](#)
- [show chassis pic fpc-slot pic-slot \(QFX3500 Switch\) on page 764](#)
- [show chassis pic interconnect-device fpc-slot pic-slot \(QFabric Systems\) on page 764](#)
- [show chassis pic node-device fpc-slot pic-slot \(QFabric System\) on page 764](#)

[show chassis pic fpc-slot 0 pic-slot 1 \(ACX2000 Universal Access Router\) on page 765](#)
[show chassis pic FPC-slot 1 PIC-slot 0 \(MX Routers with Media Services Blade \[MSB\]\) on page 766](#)
[show chassis pic FPC slot 1, PIC slot 2 \(MX Routers with Media Services Blade \[MSB\]\) on page 766](#)

Output Fields Table 62 on page 757 lists the output fields for the **show chassis pic** command. Output fields are listed in the approximate order in which they appear.

Table 62: show chassis pic Output Fields

Field Name	Field Description
Type	<p>PIC type.</p> <p>NOTE: On the 1-port OC192/STM64 MICs with the SDH framing mode, the type is displayed as MIC-3D-1STM64-XFP and with the SONET framing mode, the type is displayed as MIC-3D-1OC192-XFP. By default, the 1-port OC192/STM64 MICs displays the type as MIC-3D-1OC192-XFP.</p>
ASIC type	Type of ASIC on the PIC.
State	<p>Status of the PIC. State is displayed only when a PIC is in the slot.</p> <ul style="list-style-type: none"> • Online— PIC is online and running. • Offline—PIC is powered down.
PIC version	PIC hardware version.
Uptime	How long the PIC has been online.
Package	(Multiservices PICs only) Services package supported: Layer-2 or Layer-3 .
Port Number	Port number for the PIC.
Cable Type	Type of cable connected to the port: LH , LX , or SX .
PIC Port Information (MX480 Router 100-Gigabit Ethernet CFP)	<p>Port-level information for the PIC.</p> <ul style="list-style-type: none"> • Port—Port number • Cable type—Type of optical transceiver installed. • Fiber type—Type of fiber. SM is single-mode. • Xcvr vendor—Transceiver vendor name. • Xcvr vendor part number—Transceiver vendor part number. • Wavelength—Wavelength of the transmitted signal. Uplinks and downlinks are always 1550 nm. There is a separate fiber for each direction • Xcvr Firmware—Transceiver firmware version.

Table 62: show chassis pic Output Fields (*continued*)

Field Name	Field Description
PIC Port Information (MX960 Router Bidirectional Optics)	Port-level information for the PIC. <ul style="list-style-type: none"> • Port—Port number • Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed. Uplink interfaces display -U. Down link interfaces display -D. • Fiber type—Type of fiber. SM is single-mode. • Xcvr vendor—Transceiver vendor name. • Xcvr vendor part number—Transceiver vendor part number. <ul style="list-style-type: none"> • BX10-10-km bidirectional optics. • BX40-40-km bidirectional optics. • SFP-LX-40-km SFP optics. • Wavelength—Wavelength of the transmitted signal. Uplinks are always 1310 nm. Downlinks are either 1490 nm or 1550 nm.
PIC Port Information (Next-Generation SONET/SDH SFP)	Port-level information for the next-generation SONET/SDH SFP PIC. <ul style="list-style-type: none"> • Port—Port number. • Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed. • Fiber type—Type of fiber: SM (single-mode) or MM (multimode). • Xcvr vendor—Transceiver vendor name. • Xcvr vendor part number—Transceiver vendor part number. • Wavelength—Wavelength of the transmitted signal. Next-generation SONET/SDH SFPs use 1310 nm.
Multirate Mode	Rate-selectability status for the MIC: Enabled or Disabled .
Channelization	Indicates whether channelization is enabled or disabled on the DS3/E3 MIC.

Sample Output

show chassis pic fpc-slot pic-slot

```

user@host> show chassis pic fpc-slot 2 pic-slot 0
PIC fpc slot 2 pic slot 0 information:
  Type                10x 1GE(LAN), 1000 BASE
  ASIC type           H chip
  State               Online
  PIC version         1.1
  Uptime              1 day, 50 minutes, 58 seconds
PIC Port Information:
  Port      Cable      Xcvr      Xcvr Vendor
  Number    Type        Vendor Name  Part Number
  0         GIGE 1000EX  FINISAR CORP.  FTRJ8519P1BNL-J3
  1         GIGE 1000EX  FINISAR CORP.  FTRJ-8519-7D-JUN

```

show chassis pic fpc-slot pic-slot (PIC Offline)

```

user@host> show chassis pic fpc-slot 1 pic-slot 0
PIC fpc slot 1 pic slot 0 information:
  State                               Offline

```

show chassis pic fpc-slot pic-slot (FPC Offline)

```

user@host> show chassis pic fpc-slot 1 pic-slot 0
FPC 1 is not online

```

show chassis pic fpc-slot pic-slot (FPC Not Present)

```

user@host> show chassis pic fpc-slot 4 pic-slot 0
FPC slot 4 is empty

```

show chassis pic fpc-slot pic-slot (PIC Not Present)

```

user@host> show chassis pic fpc-slot 5 pic-slot 2
FPC 5, PIC 2 is empty

```

show chassis pic fpc-slot 3 pic-slot 0 (M120 Router)

```

user@host> show chassis pic fpc-slot 3 pic-slot 0
PC slot 3, PIC slot 0 information:
  Type                2x G/E IQ, 1000 BASE
  ASIC type           IQ GE 2 VLAN-TAG FPGA
  State               Online
  PIC version         1.16
  Uptime              3 hours, 3 minutes

PIC Port Information:
  Port      Cable      Xcvr      Xcvr Vendor
  Number    Type        Vendor Name Part Number
  0         GIGE 1000SX  FINISAR CORP.  FTRJ8519P1BNL-J3
  1         GIGE 1000SX  FINISAR CORP.  FTRJ-8519-7D-JUN

```

show chassis pic fpc-slot pic-slot (MX960 Router with Bidirectional Optics)

```

user@host> show chassis pic fpc-slot 4 pic-slot 1
FPC slot 4, PIC slot 1 information:
  Type                10x 1GE(LAN)
  State               Online
  PIC version         0.0
  Uptime              18 days, 5 hours, 41 minutes, 54 seconds

PIC port information:

```

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	SFP-1000BASE-BX10-D	SM	SumitomoElectric	SBP6H44-J3-BW-49	1490 nm
1	SFP-1000BASE-BX10-D	SM	SumitomoElectric	SBP6H44-J3-BW-49	1490 nm
2	SFP-1000BASE-BX10-D	SM	SumitomoElectric	SBP6H44-J3-BW-49	1490 nm
3	SFP-1000BASE-BX10-D	SM	OCP	TRXBG1LXDBVM2-JW	1490 nm
4	SFP-1000BASE-BX10-D	SM	OCP	TRXBG1LXDBVM2-JW	1490 nm
5	SFP-1000BASE-BX10-U	SM	SumitomoElectric	SBP6H44-J3-BW-31	1310 nm
6	SFP-1000BASE-BX10-U	SM	SumitomoElectric	SBP6H44-J3-BW-31	1310 nm
7	SFP-1000BASE-BX10-U	SM	OCP	TRXBG1LXDBBMH-J1	1310 nm
8	SFP-1000BASE-BX10-U	SM	OCP	TRXBG1LXDBBMH-J1	1310 nm
9	SFP-1000BASE-BX10-U	SM	SumitomoElectric	SBP6H44-J3-BW-31	1310 nm

show chassis pic fpc-slot pic-slot (MX480 Router with 100-Gigabit Ethernet MIC)

```
user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
  Type                1X100GE CFP
  State                Online
  PIC version          2.10
  Uptime               4 minutes, 48 seconds

PIC port information:
  Fiber
  Port Cable type      type  Xcvr vendor      part number      Wavelength
  0    100GBASE LR4    SM   FINISAR CORP.  FTLC1181RDN5-J3  1310 nm

  Xcvr vendor
  firmware version
  1.8
```

show chassis pic fpc-slot pic-slot (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```
user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
  Type                AS-MXC
  State                Online
  PIC version          1.0
  Uptime               11 hours, 18 minutes, 3 seconds
```

show chassis pic fpc-slot pic-slot (MX480 Router with MPC4E)

```
user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 3, PIC slot 0 information:
  Type                4x10GE SFPP
  State                Online
  PIC version          0.0
  Uptime               41 seconds

PIC port information:
  Fiber
  Port Cable type      type  Xcvr vendor      part number      Wave-    Xcvr
  Firmware              length
  0    10GBASE SR      MM   OPNEXT, INC.    TRS2001EM-0014  850 nm  0.0
  1    10GBASE SR      MM   OPNEXT, INC.    TRS2001EM-0014  850 nm  0.0
```

show chassis pic fpc-slot pic-slot (MX2010 Router)

```
user@host> show chassis pic fpc-slot 9 pic-slot 3
FPC slot 9, PIC slot 3 information:
  Type                1X100GE CFP
  State                Online
  PIC version          0.0
  Uptime               14 hours, 51 seconds
```

show chassis pic fpc-slot pic-slot (MX2020 Router)

```
user@host> show chassis pic fpc-slot 19 pic-slot 3
FPC slot 19, PIC slot 3 information:
  Type                4x 10GE(LAN) SFP+
```

```

State                               Online
PIC version                         0.0
Uptime                             1 day, 11 hours, 26 minutes, 36 seconds

PIC port information:

```

Port	Cable type	Fiber type	Xcvr vendor	part number	Wave-length	Xcvr
0	10GBASE SR	MM	SumitomoElectric	SPP5200SR-J6-M	850 nm	0.0
1	10GBASE SR	MM	SumitomoElectric	SPP5200SR-J6-M	850 nm	0.0
2	10GBASE SR	MM	SumitomoElectric	SPP5200SR-J6-M	850 nm	0.0
3	10GBASE SR	MM	SumitomoElectric	SPP5200SR-J6-M	850 nm	0.0

show chassis pic fpc-slot pic-slot (MX2020 Router with MPC4E)

```

user@host> show chassis pic fpc-slot 14 pic-slot 0
FPC slot 14, PIC slot 1 information:
Type                               1X100GE CFP
State                              Online
PIC version                         0.0
Uptime                             1 day, 2 hours, 19 minutes, 18 seconds

PIC port information:

```

Port	Cable type	Fiber type	Xcvr vendor	part number	Wave-length	Xcvr
0	100GBASE SR10	MM	Reflex Photonics	CF-X12-C11801-50	860 nm	4.7

show chassis pic fpc-slot pic-slot (T1600 Router with 100-Gigabit Ethernet PIC)

```

user@host> run show chassis pic fpc-slot 3 pic-slot 1
FPC slot 3, PIC slot 1 information:
Type                               100GE SLOT1
ASIC type                          Brooklyn 100GE FPGA
State                              Online
PIC version                         1.3
Uptime                             10 minutes, 44 seconds

PIC port information:

```

Port	Cable type	Fiber type	Xcvr vendor	part number	Wavelength
0	100GBASE LR4	SM	Opnext Inc.	TRC5E20ENFSF000F	1310 nm

show chassis pic fpc-slot pic-slot lcc (TX Matrix Router)

```

user@host> show chassis pic fpc-slot 1 pic-slot 1 lcc 0
lcc0-re0:
-----
PIC fpc slot 1 pic slot 1 information:
Type                               4x OC-3 SONET, SMIR
ASIC type                          D chip
State                              Online

```

```

PIC version          1.2
Uptime               5 days, 2 hours, 12 minutes, 8 seconds

```

show chassis pic fpc-slot pic-slot lcc (TX Matrix Plus Router)

```

user@host> show chassis pic pic-slot 0 fpc-slot 8
lcc0-re0:
-----
FPC slot 8, PIC slot 0 information:
Type                1x 10GE(LAN/WAN)
State               Online
Uptime              2 hours, 46 minutes, 23 seconds

PIC port information:

Port  Cable type      Fiber
                                type  Xcvr vendor      part number      Wavelength
0      10GBASE ZR        SM    Opnext Inc.      TRF7061BN-LF150  1550 nm
0      10GBASE ZR        SM    FINISAR CORP.    FTRX-1811-3-J2   1550 nm

```

show chassis pic fpc-slot pic-slot (Next-Generation SONET/SDH SFP)

```

user@host> show chassis pic fpc-slot 4 pic-slot 0
FPC slot 4, PIC slot 0 information:
Type                4x OC-3 1x OC-12 SFP
ASIC type            D FPGA
State               Online
PIC version          1.3
Uptime              1 day, 50 minutes, 4 seconds

PIC port information:

Port  Cable type      Fiber
                                type  Xcvr vendor      part number      Wavelength
0      OC48 short reach SM    FINISAR CORP.    FTRJ1321P1BTL-J2 1310 nm
1      OC3 short reach  MM    OCP              TRPA03MM3BAS-JE  1310 nm
2      OC3 short reach  MM    OCP              TRXA03MM3BAS-JW  1310 nm
3      OC12 inter reach SM    FINISAR CORP.    FTLF1322P1BTR    1310 nm

```

show chassis pic fpc-slot pic-slot (12-Port T1/E1)

```

user@host> show chassis pic fpc-slot 0 pic-slot 3
FPC slot 0, PIC slot 3 information:
Type                12x T1/E1 CE
State               Online
PIC version          1.1
CPU load average     1 percent
Interrupt load average 0 percent
Total DRAM size      128 MB
Memory buffer utilization 100 percent
Memory heap utilization 4 percent
Uptime              1 day, 22 hours, 28 minutes, 12 seconds
Internal Clock Synchronization Normal

```

show chassis pic fpc-slot 0 pic-slot 1 (4x CHOC3 SONET CE SFP)

```

user@host> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
Type                4x CHOC3 SONET CE SFP
State               Online
PIC version          1.3
CPU load average     1 percent

```



```

Interrupt load average      0 percent
Total DRAM size            128 MB
Memory buffer utilization   99 percent
Memory heap utilization     4 percent
Uptime                     1 day, 22 hours, 55 minutes, 37 seconds
Internal Clock Synchronization Normal

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	OC3 short reach	MM	AVAGO	HFBR-57E0P-JU2	n/a
1	OC3 short reach	MM	AVAGO	HFBR-57E0P-JU2	n/a
3	OC3 long reach	SM	OPNEXT INC	TRF5456AVLB314	1310 nm

show chassis pic fpc-slot 0 pic-slot 0 (SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```
user@host> show chassis pic fpc-slot 0 pic-slot 0
```

FPC slot 0, PIC slot 0 information:

```

Type                MIC-3D-80C30C12-40C48
State                Online
PIC version          1.8
Uptime               3 days, 22 hours, 3 minutes, 50 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
1	OC12 inter reach	SM	FINISAR CORP	FTRJ1322P1BTR-J3	1310 nm
7	OC12 inter reach	SM	FINISAR CORP	FTRJ1322P1BTR-J3	1310 nm

Multirate Mode Enabled

show chassis pic fpc-slot 3 pic-slot 0 (8-port Channelized SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```
user@host> show chassis pic fpc-slot 3 pic-slot 0
```

FPC slot 3, PIC slot 0 information:

```

Type                MIC-3D-8CHOC3-4CHOC12
State                Online
PIC version          1.9
Uptime               1 hour, 21 minutes, 24 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
1	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
2	OC12 inter reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J2	1310 nm
4	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
5	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
6	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
7	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm

show chassis pic fpc-slot 5 pic-slot 0 (4-port Channelized SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```
user@host> show chassis pic fpc-slot 5 pic-slot 0
```

FPC slot 5, PIC slot 0 information:

```

Type                MIC-3D-4CHOC3-2CHOC12
State                Online
PIC version          1.9
Uptime               1 hour, 21 minutes

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
------	------------	------------	-------------	-------------------------	------------

1	OC12 inter reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
2	OC12 inter reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
3	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm

show chassis pic fpc-slot 1 pic-slot 0 (1-port OC192/STM64 MIC with XFP)

```

user@host> show chassis pic fpc-slot 1 pic-slot 0
FPC slot 1, PIC slot 0 information:
  Type          MIC-3D-10C192-XFP
  State         Online
  PIC version    1.2
  Uptime        1 day, 11 hours, 4 minutes, 6 seconds

PIC port information:
  Port  Cable type      Fiber type  Xcvr vendor  Xcvr vendor  part number  Wavelength
  0     OC192 short reach n/a        FINISAR CORP. FTLX1412M3BCL-J3 1310 nm

```

show chassis pic fpc-slot 1 pic-slot 2 (8-port DS3/E3 MIC)

```

user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
  Type          MIC-3D-8DS3-E3
  State         Online
  PIC version    1.10
  Uptime        4 days, 1 hour, 29 minutes, 19 seconds
  Channelization Mode Disabled

```

show chassis pic fpc-slot pic-slot (OTN)

```

user@host> show chassis pic fpc-slot 5 pic-slot 0
PIC fpc slot 5 pic slot 0 information:
  Type          1x10GE(LAN),OTN
  ASIC type      H chip
  State         Online
  PIC version    1.0
  Uptime        5 minutes, 50 seconds

```

show chassis pic fpc-slot pic-slot (QFX3500 Switch)

```

user@switch> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
  Type 48x 10G-SFP+ Builtin
  State Online
  Uptime 3 days, 3 hours, 5 minutes, 20 seconds

```

show chassis pic interconnect-device fpc-slot pic-slot (QFabric Systems)

```

user@switch> show chassis pic interconnect-device interconnect1 fpc-slot 9 pic-slot 0
FPC slot 9, PIC slot 0 information:
  Type          16x 40G-GE Builtin
  State         Online
  Uptime        2 hours, 47 minutes, 40 seconds

```

show chassis pic node-device fpc-slot pic-slot (QFabric System)

```

user@switch> show chassis pic node-device node1 pic-slot 0
FPC slot node1, PIC slot 0 information:
  Type          48x 10G-SFP+ Builtin
  State         Online
  Uptime        2 hours, 52 minutes, 37 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
1	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
2	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
3	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
4	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
5	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
6	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
7	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
8	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
9	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
10	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
11	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
12	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
13	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
14	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
15	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
16	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
17	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
18	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
19	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
20	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
21	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
22	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
23	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
24	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
25	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
26	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
27	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
28	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
29	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
30	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
31	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
32	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
33	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
34	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
35	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
36	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
37	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
38	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
39	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
40	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
41	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
42	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
43	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
44	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
45	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
46	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
47	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm

show chassis pic fpc-slot 0 pic-slot 1 (ACX2000 Universal Access Router)

```

user@host> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
  Type                8x 1GE(LAN) RJ45 Builtin
  State               Online
  Uptime              6 days, 2 hours, 51 minutes, 11 seconds

```

show chassis pic FPC-slot 1 PIC-slot 0 (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis pic fpc-slot 1 pic-slot 0
FPC slot 1, PIC slot 0 information:
  Type           AS-MSC
  State          Online
  PIC version    1.6
  Uptime         11 hours, 17 minutes, 56 seconds
```

show chassis pic FPC slot 1, PIC slot 2 (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis pic fpc-slot 1 pic-slot 2
  Type           AS-MXC
  State          Online
  PIC version    1.0
  Uptime         11 hours, 18 minutes, 3 seconds
```

show chassis routing-engine

Syntax	show chassis routing-engine <bios <i>slot</i> >
Syntax (EX Series Switches)	show chassis routing-engine < <i>slot</i> >
Syntax (T Series routers)	show chassis routing-engine <bios <i>slot</i> >
Syntax (TX Matrix Routers)	show chassis routing-engine <bios <i>slot</i> > <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Routers)	show chassis routing-engine <bios <i>slot</i> > <lcc <i>number</i> sfc <i>number</i> >
Syntax (QFX Series)	show chassis routing-engine <interconnect-device <i>name</i> > <node-device <i>name</i> >
Syntax (MX Series Routers)	show chassis routing-engine <bios <i>slot</i> > <all-members> <local> <member <i>member-id</i> >
Syntax (MX2010 3D Universal Edge Routers)	show chassis routing-engine <bios <i>slot</i> >
Syntax (MX2020 3D Universal Edge Routers)	show chassis routing-engine <bios <i>slot</i> >
Syntax (ACX Series Universal Access Routers)	show chassis routing-engine
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc option introduced for the TX Matrix Plus router in Junos OS Release in 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>
Description	Display the status of the Routing Engine.

Options	<p>none—Display information about one or more Routing Engines. On a TX Matrix router, display information about all Routing Engines on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display information about all Routing Engines on the TX Matrix Plus router and its attached T1600 routers.</p> <p>all-members—(MX Series routers only) (Optional) Display Routing Engine information for all members of the Virtual Chassis configuration.</p> <p>bios—(Optional) Display the (BIOS) firmware version.</p> <p>interconnect-device <i>number</i>—(QFabric systems only) (Optional) Display Routing Engine information for a specified Interconnect device.</p> <p>lcc <i>number</i>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Routing Engine information for a specified T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display Routing Engine information for a specified T1600 router (or line-card chassis) that is connected to the TX Matrix Plus router. Replace <i>number</i> with a value from 0 through 3.</p> <p>local—(MX Series routers only) (Optional) Display Routing Engine information for the local Virtual Chassis member.</p> <p>member <i>member-id</i>—(MX Series routers only) (Optional) Display Routing Engine information for the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace <i>member-id</i> with a value of 0 or 1.</p> <p>node-device <i>number</i>—(QFabric systems only) (Optional) Display Routing Engine information for a specified Node device.</p> <p>scc—(TX Matrix routers only) (Optional) Display Routing Engine information for the TX Matrix router (or switch-card chassis).</p> <p>sfc <i>number</i>—(TX Matrix Plus routers only) (Optional) Display Routing Engine information for the TX Matrix Plus router (or switch-fabric chassis). Replace <i>number</i> with 0.</p> <p>slot—(Systems with multiple Routing Engines) (Optional) Display information for an individual Routing Engine. Replace <i>slot</i> with 0 or 1. For QFX3500 switches, there is only one Routing Engine, so you do not need to specify the slot number.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• <i>request chassis routing-engine master</i>• <i>Configuring Routing Engine Redundancy</i>• <i>Switching the Global Master and Backup Roles in a Virtual Chassis Configuration</i>
List of Sample Output	<p>show chassis routing-engine (M5 Router) on page 771</p> <p>show chassis routing-engine (M10 Router) on page 772</p> <p>show chassis routing-engine (M20 Router) on page 772</p>

[show chassis routing-engine \(M40 Router\) on page 773](#)
[show chassis routing-engine \(M120 Router\) on page 773](#)
[show chassis routing-engine \(M160 Router\) on page 774](#)
[show chassis routing-engine \(MX240 Router\) on page 774](#)
[show chassis routing-engine \(MX480 Router\) on page 775](#)
[show chassis routing-engine \(MX960 Router\) on page 775](#)
[show chassis routing-engine \(MX2010 Router\) on page 776](#)
[show chassis routing-engine \(MX2020 Router\) on page 776](#)
[show chassis routing-engine \(T320 router\) on page 777](#)
[show chassis routing-engine \(T640 router\) on page 778](#)
[show chassis routing-engine \(T1600 router\) on page 779](#)
[show chassis routing-engine \(T4000 router\) on page 779](#)
[show chassis routing-engine \(TX Matrix Router\) on page 780](#)
[show chassis routing-engine lcc \(TX Matrix Router\) on page 781](#)
[show chassis routing-engine bios \(TX Matrix Router\) on page 782](#)
[show chassis routing-engine \(TX Matrix Plus Router\) on page 782](#)
[show chassis routing-engine lcc \(TX Matrix Plus Router\) on page 783](#)
[show chassis routing-engine bios \(TX Matrix Plus Router\) on page 784](#)
[show chassis routing-engine \(QFX Series\) on page 784](#)
[show chassis routing-engine \(PTX Series Packet Transport Switch\) on page 785](#)
[show chassis routing-engine \(ACX2000 Universal Access Router\) on page 785](#)
[show chassis routing-engine \(ACX1000 Universal Access Router\) on page 786](#)

Output Fields [Table 63 on page 769](#) lists the output fields for the **show chassis routing-engine** command. Output fields are listed in the approximate order in which they appear.

Table 63: show chassis routing-engine Output Fields

Field Name	Field Description
Slot	(Systems with single and multiple Routing Engines) Slot number.
Current state	(Systems with multiple Routing Engines) Current state of the Routing Engine: Master , Backup , or Disabled .
Election priority	(Systems with multiple Routing Engines) Election priority for the Routing Engine: Master or Backup .
Temperature	Temperature of the air flowing past the Routing Engine.
CPU Temperature	Temperature of the CPU.
DRAM	Total DRAM available to the Routing Engine's processor.
Memory utilization	Percentage of Routing Engine memory being used.

Table 63: show chassis routing-engine Output Fields (*continued*)

Field Name	Field Description
CPU utilization	Information about the Routing Engine's CPU utilization: <ul style="list-style-type: none">• User—Percentage of CPU time being used by user processes.• Background—Percentage of CPU time being used by background processes.• Kernel—Percentage of CPU time being used by kernel processes.• Interrupt—Percentage of CPU time being used by interrupts.• Idle—Percentage of CPU time that is idle.
Model	Routing Engine model number.
Serial ID	(Systems with multiple Routing Engines) Identification number of the Routing Engine in this slot.
Start time	Time at which the Routing Engine started running.
Uptime	How long the Routing Engine has been running.
Routing Engine BIOS Version	BIOS version being run by the Routing Engine.

Table 63: show chassis routing-engine Output Fields (*continued*)

Field Name	Field Description
Last reboot reason	<p>Reason for last reboot, including:</p> <ul style="list-style-type: none"> • power cycle/failure—Halt of the Routing Engine using the halt command, powering down using the power button on the chassis or any other method (such as removal of the control board or Routing Engine), and then powering back the Routing Engine. A halt of the operating system also occurs if you enter the request system halt command. You can enter this command to halt the system operations on the chassis or specific Routing Engines. To restart the software, press any key on the keyboard. • watchdog—Reboot due to a hardware watchdog. A watchdog is a hardware monitoring process that examines the health and performance of the router to enable the device to recover from failures. A watchdog checks for problems at certain intervals, and reboots the routing engine if a problem is encountered. • reset-button reset—(Not available on the J Series router or EX Series switch) Reboot due to pressing of the reset button on the Routing Engine. • power-button hard power off—Reboot due to pressing of the power button on the chassis. A powering down of the software also occurs if you enter the request system power-off command. You can enter this command to power down the chassis or specific Routing Engines; you can then restart the software. • misc hardware reason—Reboot due to miscellaneous hardware reasons. • thermal shutdown—Reboot due to the router or switch reaching a critical temperature at which point it is unsafe to continue operations. • hard disk failure—Reboot due to a hard disk or solid-state drive (SSD) failure. • reset from debugger—Reboot due to reset from the debugger. • chassis control reset—Restart the chassis process that manages PICs, FPCs, and other hardware components. The chassis control module that runs the Routing Engine performs management and monitoring functions, and it provides a single access point for operational and maintenance functions. A reset of the chassis management process occurs when you enter the restart chassis-control command. • bios auto recovery reset—Reboot due to a BIOS auto-recovery reset. • could not be determined—Reboot due to an undetermined reason. • Router rebooted after a normal shutdown—Reboot due to a normal shutdown. This reason is displayed if the Routing Engine is powered down by pushing and holding the online/offline button on the Routing Engine faceplate for 30 seconds, and then powered back. A reboot of the software also occurs if you enter the request system reboot command. You can enter this command to reboot the chassis or specific Routing Engines.
Load averages	Routing Engine load averages for the last 1, 5, and 15 minutes.

Sample Output

show chassis routing-engine (M5 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature                25 degrees C / 77 degrees F
  DRAM                       768 MB
  Memory utilization         21 percent
  CPU utilization:
    User                      0 percent
    Background                0 percent
    Kernel                    0 percent
    Interrupt                  0 percent

```

Idle	100 percent
Model	RE-2.0
Serial ID	31000007349bf701
Start time	2003-12-04 09:42:17 PST
Uptime	26 days, 1 hour, 12 minutes, 27 seconds
Last reboot reason	Router rebooted after a normal shutdown
Load averages:	1 minute 5 minute 15 minute
	0.00 0.01 0.00

show chassis routing-engine (M10 Router)

```
user@host> show chassis routing-engine
Routing Engine status:
  Temperature      25 degrees C / 77 degrees F
  DRAM             768 MB
  Memory utilization 21 percent
  CPU utilization:
    User           0 percent
    Background     0 percent
    Kernel         0 percent
    Interrupt      0 percent
    Idle           100 percent
  Model            RE-2.0
  Serial ID        31000007349bf701
  Start time       2003-12-04 09:42:17 PST
  Uptime           26 days, 1 hour, 12 minutes, 27 seconds
  Last reboot reason Router rebooted after a normal shutdown
  Load averages:   1 minute 5 minute 15 minute
                   0.00 0.01 0.00
```

show chassis routing-engine (M20 Router)

```
user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state      Master
  Election priority  Master (default)
  Temperature        29 degrees C / 84 degrees F
  DRAM               768 MB
  Memory utilization 20 percent
  CPU utilization:
    User             1 percent
    Background       0 percent
    Kernel           2 percent
    Interrupt        0 percent
    Idle             97 percent
  Model              RE-2.0
  Serial ID          58000007348d9a01
  Start time         2003-12-30 07:05:47 PST
  Uptime              3 hours, 41 minutes, 14 seconds
  Last reboot reason Router rebooted after a normal shutdown
  Load averages:     1 minute 5 minute 15 minute
                     0.00 0.02 0.00

Routing Engine status:
Slot 1:
  Current state      Backup
  Election priority  Backup (default)
  Temperature        29 degrees C / 84 degrees F
  DRAM               768 MB
  Memory utilization 0 percent
  CPU utilization:
```

```

User                0 percent
Background          0 percent
Kernel              1 percent
Interrupt           0 percent
Idle                99 percent
Model               RE-2.0
Serial ID            d800000734745701
Start time           2003-06-17 16:37:33 PDT
Uptime              195 days, 18 hours, 47 minutes, 9 seconds
Last reboot reason   Router rebooted after a normal shutdown

```

show chassis routing-engine (M40 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature        25 degrees C / 77 degrees F
  DRAM                768 MB
  Memory utilization  21 percent
  CPU utilization:
    User              0 percent
    Background        0 percent
    Kernel             0 percent
    Interrupt          0 percent
    Idle              100 percent
  Model              RE-2.0
  Serial ID           31000007349bf701
  Start time           2003-12-04 09:42:17 PST
  Uptime              26 days, 1 hour, 12 minutes, 27 seconds
  Last reboot reason   Router rebooted after a normal shutdown
  Load averages:      1 minute   5 minute  15 minute
                      0.00        0.01    0.00

```

show chassis routing-engine (M120 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state        Master
  Election priority     Master (default)
  Temperature          46 degrees C / 114 degrees F
  CPU temperature       44 degrees C / 111 degrees F
  DRAM                  2048 MB
  Memory utilization    18 percent
  CPU utilization:
    User                0 percent
    Background          0 percent
    Kernel               5 percent
    Interrupt            0 percent
    Idle                 95 percent
  Model                RE-A-1000
  Serial ID             1000621154
  Start time            2006-10-31 17:10:05 PST
  Uptime                14 minutes, 31 seconds
  Last reboot reason     Router rebooted after a normal shutdown
  Load averages:        1 minute   5 minute  15 minute
                      0.02        0.07    0.07

Routing Engine status:
Slot 1:
  Current state        Backup
  Election priority     Backup (default)
  Temperature          45 degrees C / 113 degrees F

```

```
CPU temperature      42 degrees C / 107 degrees F
DRAM                2048 MB
Memory utilization   15 percent
CPU utilization:
  User               0 percent
  Background         0 percent
  Kernel             0 percent
  Interrupt          0 percent
  Idle               100 percent
Model               RE-A-1000
Serial ID            1000621151
Start time           2006-10-31 17:10:04 PST
Uptime               14 minutes, 30 seconds
Last reboot reason   Router rebooted after a normal shutdown
```

show chassis routing-engine (M160 Router)

```
user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state      Master
  Election priority   Master (default)
  Temperature        43 degrees C / 109 degrees F
  DRAM               2048 MB
  Memory utilization  11 percent
  CPU utilization:
    User             1 percent
    Background       0 percent
    Kernel           2 percent
    Interrupt        0 percent
    Idle             97 percent
  Model              RE-3.0
  Serial ID          210865700403
  Start time         2003-12-23 12:25:55 PST
  Uptime             6 days, 22 hours, 33 minutes, 24 seconds
  Last reboot reason Router rebooted after a normal shutdown
  Load averages:    1 minute   5 minute   15 minute
                   0.24       0.13       0.04

Routing Engine status:
Slot 1:
  Current state      Backup
  Election priority   Backup (default)
  Temperature        40 degrees C / 104 degrees F
  DRAM               2048 MB
  Memory utilization  9 percent
  CPU utilization:
    User             0 percent
    Background       0 percent
    Kernel           0 percent
    Interrupt        0 percent
    Idle             100 percent
  Model              RE-3.0
  Serial ID          210865700332
  Start time         2003-12-23 12:25:55 PST
  Uptime             6 days, 22 hours, 33 minutes, 21 seconds
  Last reboot reason Router rebooted after a normal shutdown
```

show chassis routing-engine (MX240 Router)

```
user@host> show chassis routing-engine
```

```

Routing Engine status:
Slot 0:
  Current state           Backup
  Election priority       Master (default)
  Temperature             40 degrees C / 104 degrees F
  CPU temperature         47 degrees C / 116 degrees F
  DRAM                   3584 MB
  Memory utilization      7 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                0 percent
    Interrupt             0 percent
    Idle                  100 percent
  Model                  RE-S-2000
  Serial ID              1000703522
  Start time             2007-12-19 10:35:40 PST
  Uptime                 16 days, 3 hours, 15 minutes, 23 seconds
  Last reboot reason     Router rebooted after a normal shutdown

```

show chassis routing-engine (MX480 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             41 degrees C / 105 degrees F
  CPU temperature         38 degrees C / 100 degrees F
  DRAM                   2048 MB
  Memory utilization      13 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                2 percent
    Interrupt             0 percent
    Idle                  98 percent
  Model                  RE-S-1300
  Serial ID              1000697044
  Start time             2008-01-04 06:46:08 PST
  Uptime                 8 hours, 17 minutes, 16 seconds
  Last reboot reason     Router rebooted after a normal shutdown

```

show chassis routing-engine (MX960 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             37 degrees C / 98 degrees F
  CPU temperature         37 degrees C / 98 degrees F
  DRAM                   2048 MB
  Memory utilization      18 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                4 percent
    Interrupt             0 percent
    Idle                  96 percent
  Model                  RE-S-1300

```

```

Serial ID          1000617944
Start time         2006-10-26 12:37:13 PDT
Uptime            6 days, 4 hours, 59 minutes, 40 seconds
Last reboot reason Router rebooted after a normal shutdown
Load averages:    1 minute   5 minute   15 minute
                  0.16      0.08      0.02

```

show chassis routing-engine (MX2010 Router)

```
user@host> show chassis routing-engine
```

Routing Engine status:

Slot 0:

```

Current state      Master
Election priority  Master (default)
Temperature        3 degrees C / 37 degrees F
CPU temperature    3 degrees C / 37 degrees F
DRAM              17152 MB
Memory utilization 13 percent
CPU utilization:
  User            0 percent
  Background      0 percent
  Kernel          4 percent
  Interrupt       2 percent
  Idle            95 percent
Model            RE-S-1800x4
Serial ID         9009099704
Start time        2012-10-02 14:33:32 PDT
Uptime           14 hours, 39 minutes, 39 seconds
Last reboot reason Router rebooted after a normal shutdown.
Load averages:    1 minute   5 minute   15 minute
                  0.06      0.05      0.01

```

Routing Engine status:

Slot 1:

```

Current state      Backup
Election priority  Backup (default)
Temperature        1 degrees C / 33 degrees F
CPU temperature    2 degrees C / 35 degrees F
DRAM              17152 MB
Memory utilization 11 percent
CPU utilization:
  User            0 percent
  Background      0 percent
  Kernel          0 percent
  Interrupt       0 percent
  Idle            100 percent
Model            RE-S-1800x4
Serial ID         9009099706
Start time        2012-10-02 10:36:06 PDT
Uptime           18 hours, 36 minutes, 57 seconds
Last reboot reason Router rebooted after a normal shutdown.
Load averages:    1 minute   5 minute   15 minute
                  0.01      0.00      0.00

```

show chassis routing-engine (MX2020 Router)

```
user@host> show chassis routing-engine
```

Routing Engine status:

Slot 0:

```

Current state      Master
Election priority  Master (default)

```

```

Temperature          6 degrees C / 42 degrees F
CPU temperature       6 degrees C / 42 degrees F
DRAM                 17152 MB
Memory utilization    14 percent
CPU utilization:
  User                1 percent
  Background          0 percent
  Kernel              7 percent
  Interrupt            2 percent
  Idle                91 percent
Model                RE-S-1800x4
Serial ID             9009089704
Start time            2012-10-02 11:05:24 PDT
Uptime                2 days, 15 hours, 49 minutes, 13 seconds
Last reboot reason    Router rebooted after a normal shutdown.
Load averages:        1 minute   5 minute   15 minute
                      0.10       0.05       0.01

Routing Engine status:
Slot 1:
  Current state        Backup
  Election priority    Backup (default)
  Temperature          7 degrees C / 44 degrees F
  CPU temperature      5 degrees C / 41 degrees F
  DRAM                 17152 MB
  Memory utilization    12 percent
  CPU utilization:
    User                0 percent
    Background          0 percent
    Kernel              0 percent
    Interrupt            0 percent
    Idle                99 percent
  Model                RE-S-1800x4
  Serial ID             9009094138
  Start time            2012-10-02 11:09:57 PDT
  Uptime                2 days, 15 hours, 44 minutes, 27 seconds
  Last reboot reason    Router rebooted after a normal shutdown.
  Load averages:        1 minute   5 minute   15 minute
                      0.00       0.00       0.00

```

show chassis routing-engine (T320 router)

```

user@host> show chassis routing-engine
Slot 0:
  Current state        Master
  Election priority    Master (default)
  Temperature          51 degrees C / 123 degrees F
  CPU temperature      55 degrees C / 131 degrees F
  DRAM                 3584 MB
  Memory utilization    11 percent
  CPU utilization:
    User                0 percent
    Background          0 percent
    Kernel              2 percent
    Interrupt            0 percent
    Idle                97 percent
  Model                RE-A-2000
  Serial ID             9009010618
  Start time            2012-10-10 01:24:05 PDT
  Uptime                5 days, 10 hours, 49 minutes, 23 seconds
  Last reboot reason    0x1:power cycle/failure
  Load averages:        1 minute   5 minute   15 minute

```

```

                                0.00      0.05      0.04
Routing Engine status:
Slot 1:
  Current state                Backup
  Election priority            Backup (default)
  Temperature                   45 degrees C / 113 degrees F
  CPU temperature               48 degrees C / 118 degrees F
  DRAM                          3584 MB
  Memory utilization            9 percent
  CPU utilization:
    User                        0 percent
    Background                  0 percent
    Kernel                      0 percent
    Interrupt                   0 percent
    Idle                        100 percent
  Model                         RE-A-2000
  Serial ID                     9009003642
  Start time                    2012-10-10 01:24:04 PDT
  Uptime                        5 days, 10 hours, 49 minutes, 28 seconds
  Last reboot reason            0x1:power cycle/failure

```

show chassis routing-engine (T640 router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state                Master
  Election priority            Master (default)
  Temperature                   50 degrees C / 122 degrees F
  CPU temperature               58 degrees C / 136 degrees F
  DRAM                          3584 MB
  Memory utilization            14 percent
  CPU utilization:
    User                        1 percent
    Background                  0 percent
    Kernel                      4 percent
    Interrupt                   1 percent
    Idle                        95 percent
  Model                         RE-A-2000
  Serial ID                     1000686556
  Start time                    2012-10-10 01:24:02 PDT
  Uptime                        5 days, 10 hours, 50 minutes, 27 seconds
  Last reboot reason            0x1:power cycle/failure
  Load averages:               1 minute   5 minute   15 minute
                                1.24      0.33      0.12

Routing Engine status:
Slot 1:
  Current state                Backup
  Election priority            Backup (default)
  Temperature                   44 degrees C / 111 degrees F
  CPU temperature               49 degrees C / 120 degrees F
  DRAM                          3584 MB
  Memory utilization            12 percent
  CPU utilization:
    User                        0 percent
    Background                  0 percent
    Kernel                      0 percent
    Interrupt                   1 percent
    Idle                        99 percent
  Model                         RE-A-2000
  Serial ID                     1000702739

```



```

Start time          2012-10-10 01:24:02 PDT
Uptime              5 days, 10 hours, 50 minutes, 26 seconds
Last reboot reason  0x1:power cycle/failure

```

show chassis routing-engine (T1600 router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state          Master
  Election priority      Master (default)
  Temperature            48 degrees C / 118 degrees F
  CPU temperature        58 degrees C / 136 degrees F
  DRAM                   3584 MB
  Memory utilization     13 percent
  CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel               3 percent
    Interrupt            1 percent
    Idle                 96 percent
  Model                  RE-A-2000
  Serial ID              1000704521
  Start time             2012-10-10 01:23:41 PDT
  Uptime                 5 days, 10 hours, 46 minutes, 56 seconds
  Last reboot reason     0x1:power cycle/failure
  Load averages:        1 minute   5 minute   15 minute
                        0.05         0.03         0.01

Routing Engine status:
Slot 1:
  Current state          Backup
  Election priority      Backup (default)
  Temperature            44 degrees C / 111 degrees F
  CPU temperature        48 degrees C / 118 degrees F
  DRAM                   3584 MB
  Memory utilization     12 percent
  CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel               0 percent
    Interrupt            0 percent
    Idle                 100 percent
  Model                  RE-A-2000
  Serial ID              9009006579
  Start time             2012-10-10 01:23:42 PDT
  Uptime                 5 days, 10 hours, 46 minutes, 54 seconds
  Last reboot reason     0x1:power cycle/failure

```

show chassis routing-engine (T4000 router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state          Master
  Election priority      Master (default)
  Temperature            33 degrees C / 91 degrees F
  CPU temperature        50 degrees C / 122 degrees F
  DRAM                   8960 MB
  Memory utilization     18 percent
  CPU utilization:
    User                 0 percent

```

```

Background          0 percent
Kernel              4 percent
Interrupt            1 percent
Idle                 95 percent
Model                RE-DUO-1800
Serial ID            P737F-002248
Start time           2012-02-09 22:49:53 PST
Uptime               2 hours, 21 minutes, 35 seconds
Last reboot reason   Router rebooted after a normal shutdown.
Load averages:       1 minute   5 minute   15 minute
                      0.00        0.04        0.00

Routing Engine status:
Slot 1:
  Current state      Backup
  Election priority  Backup (default)
  Temperature        32 degrees C / 89 degrees F
  CPU temperature    46 degrees C / 114 degrees F
  DRAM               8960 MB
  Memory utilization 24 percent
  CPU utilization:
    User             0 percent
    Background       0 percent
    Kernel           0 percent
    Interrupt        0 percent
    Idle             99 percent
  Model              RE-DUO-1800
  Serial ID          P737F-002653
  Start time         2012-02-08 20:12:51 PST
  Uptime             1 day, 4 hours, 58 minutes, 28 seconds
  Last reboot reason Router rebooted after a normal shutdown.

```

show chassis routing-engine (TX Matrix Router)

```

user@host> show chassis routing-engine
scc-re0:
-----
Routing Engine status:
Slot 0:
  Current state      Master
  Election priority  Master (default)
  Temperature        34 degrees C / 93 degrees F
  CPU temperature    33 degrees C / 91 degrees F
  DRAM               2048 MB
  Memory utilization 12 percent
  CPU utilization:
    User             0 percent
    Background       0 percent
    Kernel           2 percent
    Interrupt        0 percent
    Idle             98 percent
  Model              RE-4.0
  Serial ID          P11123900153
  Start time         2004-08-05 18:42:05 PDT
  Uptime             9 days, 22 hours, 49 minutes, 50 seconds
  Last reboot reason Router rebooted after a normal shutdown
  Load averages:    1 minute   5 minute   15 minute
                      0.00        0.08        0.07

lcc0-re0:
-----
Routing Engine status:

```

```

Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             33 degrees C / 91 degrees F
  CPU temperature         30 degrees C / 86 degrees F
  DRAM                   2048 MB
  Memory utilization      12 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                1 percent
    Interrupt             0 percent
    Idle                  98 percent
  Model                  RE-3.0
  Serial ID               210865700363
  Start time             2004-08-05 18:42:05 PDT
  Uptime                  9 days, 22 hours, 48 minutes, 20 seconds
  Last reboot reason      Router rebooted after a normal shutdown
  Load averages:         1 minute  5 minute 15 minute
                          0.00      0.02   0.00

```

```
lcc2-re0:
```

```
-----
Routing Engine status:
```

```

Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             34 degrees C / 93 degrees F
  CPU temperature         35 degrees C / 95 degrees F
  DRAM                   2048 MB
  Memory utilization      12 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                2 percent
    Interrupt             0 percent
    Idle                  98 percent
  Model                  RE-4.0
  Serial ID               P11123900126
  Start time             2004-08-05 18:42:05 PDT
  Uptime                  9 days, 22 hours, 49 minutes, 4 seconds
  Last reboot reason      Router rebooted after a normal shutdown
  Load averages:         1 minute  5 minute 15 minute
                          0.01      0.01   0.0

```

show chassis routing-engine lcc (TX Matrix Router)

```
user@host> show chassis routing-engine 0 lcc 0
```

```
lcc0-re0:
```

```
-----
Routing Engine status:
```

```

Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             33 degrees C / 91 degrees F
  CPU temperature         30 degrees C / 86 degrees F
  DRAM                   2048 MB
  Memory utilization      12 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent

```

```

Kernel                1 percent
Interrupt              0 percent
Idle                  98 percent
Model                 RE-3.0
Serial ID              210865700363
Start time             2004-08-05 18:42:05 PDT
Uptime                7 days, 22 hours, 49 minutes, 6 seconds
Last reboot reason     Router rebooted after a normal shutdown
Load averages:         1 minute   5 minute   15 minute
                       0.00       0.00       0.00

```

show chassis routing-engine bios (TX Matrix Router)

```

user@host> show chassis routing-engine bios
scc-re0:

```

```

-----
Routing Engine BIOS Version: V1.0.0
lcc0-re0:

```

```

-----
Routing Engine BIOS Version: V1.0.17
lcc2-re0:

```

```

-----
Routing Engine BIOS Version: V1.0.0

```

show chassis routing-engine (TX Matrix Plus Router)

```

user@host> show chassis routing-engine
sfc0-re0:

```

```

-----
Routing Engine status:

```

Slot 0:

```

Current state           Master
Election priority       Master (default)
Temperature             27 degrees C / 80 degrees F
CPU temperature         42 degrees C / 107 degrees F
DRAM                   3327 MB
Memory utilization      12 percent
CPU utilization:
  User                  0 percent
  Background            0 percent
  Kernel                2 percent
  Interrupt              0 percent
  Idle                  98 percent
Model                  RE-TXP-SFC
Serial ID              737A-1024
Start time             2009-05-11 17:39:49 PDT
Uptime                 3 hours, 45 minutes, 25 seconds
Last reboot reason     Router rebooted after a normal shutdown.
Load averages:         1 minute   5 minute   15 minute
                       0.00       0.00       0.00

```

```

Routing Engine status:

```

Slot 1:

```

Current state           Backup
Election priority       Backup (default)
Temperature             29 degrees C / 84 degrees F
CPU temperature         43 degrees C / 109 degrees F
DRAM                   3327 MB
Memory utilization      11 percent
CPU utilization:
  User                  0 percent
  Background            0 percent

```

```

Kernel                0 percent
Interrupt             0 percent
Idle                 100 percent
Model                RE-TXP-SFC
Serial ID            737A-1024
Start time           2009-05-11 17:08:54 PDT
Uptime               4 hours, 16 minutes, 52 seconds
Last reboot reason   0x1:power cycle/failure

lcc0-re0:
-----
Routing Engine status:
Slot 0:
  Current state       Master
  Election priority   Master (default)
  Temperature         30 degrees C / 86 degrees F
  CPU temperature     43 degrees C / 109 degrees F
  DRAM                3327 MB
  Memory utilization  9 percent
  CPU utilization:
    User              0 percent
    Background        0 percent
    Kernel            2 percent
    Interrupt         0 percent
    Idle              98 percent
  Model              RE-TXP-LCC
  Serial ID          737F-1024
  Start time         2009-05-11 17:40:32 PDT
  Uptime             3 hours, 44 minutes, 51 seconds
  Last reboot reason Router rebooted after a normal shutdown.
  Load averages:    1 minute  5 minute 15 minute
                    0.00      0.00    0.00

Routing Engine status:
Slot 1:
  Current state       Backup
  Election priority   Backup (default)
  Temperature         30 degrees C / 86 degrees F
  CPU temperature     43 degrees C / 109 degrees F
  DRAM                3327 MB
  Memory utilization  9 percent
  CPU utilization:
    User              0 percent
    Background        0 percent
    Kernel            0 percent
    Interrupt         0 percent
    Idle              100 percent
  Model              RE-TXP-LCC
  Serial ID          737F-1024
  Start time         2009-05-06 17:31:32 PDT
  Uptime             5 days, 3 hours, 54 minutes, 19 seconds
  Last reboot reason Router rebooted after a normal shutdown.

```

show chassis routing-engine lcc (TX Matrix Plus Router)

```

user@host> show chassis routing-engine 0 lcc 0
lcc0-re0:
-----
Routing Engine status:
Slot 0:
  Current state       Master
  Election priority   Master (default)

```

```

Temperature          30 degrees C / 86 degrees F
CPU temperature       43 degrees C / 109 degrees F
DRAM                 3327 MB
Memory utilization    9 percent
CPU utilization:
  User                0 percent
  Background          0 percent
  Kernel              2 percent
  Interrupt            0 percent
  Idle                98 percent
Model                RE-TXP-LCC
Serial ID             737F-1024
Start time            2009-05-11 17:40:32 PDT
Uptime                3 hours, 45 minutes, 26 seconds
Last reboot reason    Router rebooted after a normal shutdown.
Load averages:        1 minute   5 minute   15 minute
                      0.00       0.00       0.00

Routing Engine status:
Slot 1:
  Current state        Backup
  Election priority    Backup (default)
  Temperature          30 degrees C / 86 degrees F
  CPU temperature       43 degrees C / 109 degrees F
  DRAM                 3327 MB
  Memory utilization    9 percent
  CPU utilization:
    User                0 percent
    Background          0 percent
    Kernel              0 percent
    Interrupt            0 percent
    Idle                100 percent
  Model                RE-TXP-LCC
  Serial ID             737F-1024
  Start time            2009-05-06 17:31:32 PDT
  Uptime                5 days, 3 hours, 54 minutes, 59 seconds
  Last reboot reason    Router rebooted after a normal shutdown.

```

show chassis routing-engine bios (TX Matrix Plus Router)

```

user@host> show chassis routing-engine bios
sfc0-re0:

```

```

-----
Routing Engine BIOS Version: V0.0.Z

```

```

lcc0-re0:

```

```

-----
Routing Engine BIOS Version: V0.0.N

```

show chassis routing-engine (QFX Series)

```

user@switch> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state Master
  Election priority Master (default)
  DRAM 2820 MB
  Memory utilization 49 percent
  CPU utilization:
    User 1 percent
    Background 0 percent
    Kernel 1 percent

```

```

Interrupt 0 percent
Idle 97 percent
Model QFX3500-48S4Q
Serial ID S/N ED3709
Uptime 3 days, 4 hours, 29 minutes, 42 seconds
Last reboot reason 0x200:chassis control reset
Load averages: 1 minute 5 minute 15 minute
0.37 0.26 0.19

```

show chassis routing-engine (PTX Series Packet Transport Switch)

```

user@switch> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             60 degrees C / 140 degrees F
  CPU temperature         76 degrees C / 168 degrees F
  DRAM                   17152 MB
  Memory utilization      11 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                4 percent
    Interrupt             0 percent
    Idle                  95 percent
  Model                   RE-DUO-2600
  Serial ID               P737A-002231
  Start time              2011-12-21 16:54:37 PST
  Uptime                  25 minutes, 44 seconds
  Last reboot reason      Router rebooted after a normal shutdown.
  Load averages:        1 minute   5 minute   15 minute
                        0.01       0.02       0.06

Routing Engine status:
Slot 1:
  Current state           Backup
  Election priority       Backup (default)
  Temperature             50 degrees C / 122 degrees F
  CPU temperature         64 degrees C / 147 degrees F
  DRAM                   17152 MB
  Memory utilization      10 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                0 percent
    Interrupt             0 percent
    Idle                  99 percent
  Model                   RE-DUO-2600
  Serial ID               P737A-002438
  Start time              2011-12-21 16:52:26 PST
  Uptime                  27 minutes, 49 seconds
  Last reboot reason      Router rebooted after a normal shutdown.

```

show chassis routing-engine (ACX2000 Universal Access Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature             53 degrees C / 127 degrees F
  DRAM                   1536 MB
  Memory utilization      25 percent
  CPU utilization:

```

```
User                0 percent
Background          0 percent
Kernel              0 percent
Interrupt            1 percent
Idle                99 percent
Model               RE-ACX-2000
Start time          2012-05-09 00:57:07 PDT
Uptime              5 days, 3 hours, 16 minutes, 15 seconds
Last reboot reason  Router rebooted after a normal shutdown.
Load averages:      1 minute   5 minute   15 minute
                   0.00        0.03        0.05
```

show chassis routing-engine (ACX1000 Universal Access Router)

```
user@host> show chassis routing-engine
Routing Engine status:
  Temperature        36 degrees C / 96 degrees F
  DRAM               768 MB
  Memory utilization  50 percent
  CPU utilization:
    User              3 percent
    Background        0 percent
    Kernel            6 percent
    Interrupt         0 percent
    Idle              91 percent
  Model              RE-ACX-1000
  Start time         2012-05-10 07:12:23 PDT
  Uptime             4 days, 10 hours, 46 minutes, 53 seconds
  Last reboot reason Router rebooted after a normal shutdown.
  Load averages:    1 minute   5 minute   15 minute
                   0.00        0.00        0.00
```


show chassis temperature-thresholds

Syntax	show chassis temperature-thresholds
Syntax (TX Matrix Routers)	show chassis temperature-thresholds <lcc <i>number</i> scc>
Syntax (TX Matrix Plus Routers)	show chassis temperature-thresholds <lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Routers)	show chassis temperature-thresholds <all-members> <local> <member <i>member-id</i> >
Syntax (MX2010 and MX2020 3D Universal Edge Routers)	show chassis temperature-thresholds
Syntax (QFX Series)	show chassis temperature-thresholds <interconnect-device <i>name</i> > <node-device <i>name</i> >
Release Information	<p>Command introduced in Junos OS Release 8.0.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>sfc command introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.1 for T4000 Core Routers.</p> <p>Command introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>
Description	Display chassis temperature threshold settings, in degrees Celsius.
Options	<p>none—Display the temperature threshold details.</p> <p>all-members—(MX Series routers only) (Optional) Display the chassis temperature threshold settings of all member routers in the Virtual Chassis configuration.</p> <p>interconnect-device <i>name</i>—(QFabric systems only) (Optional) Display the chassis temperature threshold settings of the Interconnect device.</p> <p>lcc <i>number</i>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display the temperature threshold details of a specified T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display the temperature threshold details of a specified T1600 router (or line-card chassis) that is connected to a TX Matrix Plus router. Replace <i>number</i> with a value from 0 through 3.</p> <p>local—(MX Series routers only) (Optional) Display the chassis temperature threshold settings of the local Virtual Chassis member.</p>

member *member-id*—(MX Series routers only) (Optional) Display the chassis temperature threshold settings of the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

node-device *name*—(QFabric systems only) (Optional) Display the chassis temperature threshold settings of the Node device.

scc—(TX Matrix routers only) (Optional) Display the temperature threshold details of the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Display the temperature threshold details of the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

Required Privilege Level view

Related Documentation • *Defining Alarm Thresholds for System Temperature Sensors*

List of Sample Output [show chassis temperature-thresholds on page 789](#)
[show chassis temperature-thresholds \(MX240, MX480, MX960 Routers with Application Services Modular Line Card\) on page 790](#)
[show chassis temperature-thresholds \(MX480 with MPC4E\) on page 790](#)
[show chassis temperature-thresholds \(MX2010 Router\) on page 790](#)
[show chassis temperature-thresholds \(MX2020 Router\) on page 793](#)
[show chassis temperature-thresholds \(MX2020 with MPC4E\) on page 796](#)
[show chassis temperature-thresholds \(T4000 Core Routers\) on page 799](#)
[show chassis temperature-thresholds \(TX Matrix Plus Router\) on page 800](#)
[show chassis temperature-thresholds lcc \(TX Matrix Plus Router\) on page 801](#)
[show chassis temperature-thresholds sfc \(TX Matrix Plus Router\) on page 801](#)
[show chassis temperature-thresholds \(QFX3500 Switch and QFX3600\) on page 802](#)
[show chassis temperature-thresholds interconnect-device \(QFabric System\) on page 802](#)
[show chassis temperature-thresholds \(PTX5000 Packet Transport Switch\) on page 802](#)
[show chassis temperature-thresholds \(MX Routers with Media Services Blade \[MSB\]\) on page 803](#)

Output Fields [Table 64 on page 788](#) lists the output fields for the **show chassis temperature-thresholds** command. Output fields are listed in the approximate order in which they appear.

Table 64: show chassis temperature-thresholds Output Fields

Field name	Field Description
Item	Chassis component. If per FRU per slot thresholds are configured, the components about which information is displayed include the chassis, the Routing Engines, FPCs, and FEBs. If per FRU per slot thresholds are not configured, the components about which information is displayed include the chassis and the Routing Engines.

Table 64: show chassis temperature-thresholds Output Fields (*continued*)

Field name	Field Description
Fan speed	<p>NOTE: On the QFX3500 switch and QFX3600 switch, there are four fan speeds: low, medium-low, medium-high, and high. The fan speed changes at the threshold when going from a low speed to a higher speed. When the fan speed changes from a higher speed to a lower speed, the temperature changes two degrees below the threshold.</p> <p>Temperature threshold settings, in degrees Celsius, for the fans to operate at normal and high speeds.</p> <ul style="list-style-type: none"> • Normal—The fans operate at normal speed if the component is at or below this temperature and all the fans are present and functioning normally. • High—The fans operate at high speed if the component has exceeded this temperature or a fan has failed or is missing. <p>NOTE: For MX480 Routers, there are three fan speeds: Low, Medium, and High.</p> <p>An alarm is not triggered until the temperature exceeds the threshold settings for a yellow alarm or a red alarm.</p>
Yellow alarm	<p>Temperature threshold settings, in degrees Celsius, that trigger a yellow alarm.</p> <ul style="list-style-type: none"> • Normal—The temperature that must be exceeded on the component to trigger a yellow alarm when the fans are running at full speed. • Bad fan—The temperature that must be exceeded on the component to trigger a yellow alarm when one or more fans have failed or are missing.
Red alarm	<p>Temperature threshold settings, in degrees Celsius, that trigger a red alarm.</p> <ul style="list-style-type: none"> • Normal—The temperature that must be exceeded on the component to trigger a red alarm when the fans are running at full speed. • Bad fan—The temperature that must be exceeded on the component to trigger a red alarm when one or more fans have failed or are missing.
Fire Shutdown	<p>(T4000 routers and PTX Series Packet Transport Switches only)—Temperature threshold settings, in degrees Celsius, for the network device to shut down.</p>

Sample Output

show chassis temperature-thresholds

```
user@host> show chassis temperature-thresholds
```

Item	Fan speed (degrees C)		Yellow alarm (degrees C)		Red alarm (degrees C)	
	Normal	High	Normal	Bad fan	Normal	Bad fan
Chassis default	48	54	65	55	75	65
Routing Engine 0	70	80	95	95	110	110
Routing Engine 1	70	80	95	95	110	110
FPC 0	55	60	75	65	90	80
FPC 1	55	60	75	65	90	80
FPC 2	55	60	75	65	90	80
FPC 3	55	60	75	65	90	80
FPC 4	55	60	75	65	90	80
FPC 5	55	60	75	65	90	80
FPC 6	55	60	75	65	90	80
FPC 7	55	60	75	65	90	80

FPC 8	55	60	75	65	90	80
FPC 9	55	60	75	65	90	80
FPC 10	55	60	75	65	90	80
FPC 11	55	60	75	65	90	80

show chassis temperature-thresholds (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```

user@host>show chassis temperature-thresholds
Fan speed      Yellow alarm      Red alarm      Fire Shutdown
(degrees C)      (degrees C)      (degrees C)      (degrees C)
Item
Normal
Chassis default      48    54    65    55    75    65
100
Routing Engine 0      70    80    95    95    110   110
112
Routing Engine 1      70    80    95    95    110   110
112
FPC 0                55    60    75    65    90    80
95
FPC 1                55    60    75    65    90    80
95
FPC 2                55    60    75    65    90    80
95
FPC 4                55    60    75    65    90    80
95
FPC 5                55    60    75    65    90    80
95

```

show chassis temperature-thresholds (MX480 with MPC4E)

```

user@host > show chassis temperature-thresholds
Fan speed      Yellow alarm      Red alarm      Fire Shutdown
(degrees C)      (degrees C)      (degrees C)      (degrees C)
Item
Normal
Chassis default      48    54    65    55    75    65
100
Routing Engine 0      70    80    95    95    110   110
112
Routing Engine 1      70    80    95    95    110   110
112
FPC 2                55    60    75    65    95    80
100
FPC 3                55    60    75    65    95    80
100
FPC 4                55    60    75    65    90    80
95

```

show chassis temperature-thresholds (MX2010 Router)

```

user@host> show chassis temperature-thresholds
Fan speed      Yellow alarm      Red alarm      Fire Shutdown
(degrees C)      (degrees C)      (degrees C)      (degrees C)
Item
Routing Engine 0      70    80    95    95    110   110   112
Routing Engine 1      70    80    95    95    110   110   112
CB 0 IntakeA-Zone0    60    65    78    75    85    80    95
CB 0 IntakeB-Zone1    60    65    78    75    85    80    95

```

CB 0 IntakeC-Zone0	60	65	78	75	85	80	95
CB 0 ExhaustA-Zone0	60	65	78	75	85	80	95
CB 0 ExhaustB-Zone1	60	65	78	75	85	80	95
CB 0 TCBC-Zone0	60	65	78	75	85	80	95
CB 1 IntakeA-Zone0	60	65	78	75	85	80	95
CB 1 IntakeB-Zone1	60	65	78	75	85	80	95
CB 1 IntakeC-Zone0	60	65	78	75	85	80	95
CB 1 ExhaustA-Zone0	60	65	78	75	85	80	95
CB 1 ExhaustB-Zone1	60	65	78	75	85	80	95
CB 1 TCBC-Zone0	60	65	78	75	85	80	95
SPMB 0 Intake	56	62	75	63	83	76	95
SPMB 1 Intake	56	62	75	63	83	76	95
SFB 0 Intake-Zone0	56	62	75	63	82	70	87
SFB 0 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 0 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 0 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 0 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 0 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 0 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 0 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 1 Intake-Zone0	56	62	75	63	82	70	87
SFB 1 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 1 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 1 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 1 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 1 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 1 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 1 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 2 Intake-Zone0	56	62	75	63	82	70	87
SFB 2 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 2 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 2 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 2 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 2 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 2 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 2 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 3 Intake-Zone0	56	62	75	63	82	70	87
SFB 3 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 3 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 3 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 3 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 3 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 3 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 3 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 4 Intake-Zone0	56	62	75	63	82	70	87
SFB 4 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 4 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 4 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 4 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 4 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 4 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 4 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 5 Intake-Zone0	56	62	75	63	82	70	87
SFB 5 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 5 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 5 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 5 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 5 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 5 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 5 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 6 Intake-Zone0	56	62	75	63	82	70	87

SFB 6 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 6 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 6 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 6 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 6 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 6 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 6 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 7 Intake-Zone0	56	62	75	63	82	70	87
SFB 7 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 7 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 7 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 7 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 7 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 7 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 7 SFB-XF0-Zone0	70	80	90	90	107	107	115
FPC 0	55	60	75	65	95	80	100
FPC 1	55	60	75	65	90	80	95
FPC 2	55	60	75	65	95	80	100
FPC 3	55	60	75	65	90	80	95
FPC 4	55	60	75	65	90	80	95
FPC 5	55	60	75	65	95	80	100
FPC 6	55	60	75	65	90	80	95
FPC 7	55	60	75	65	95	80	100
FPC 8	55	60	75	65	90	80	95
FPC 9	55	60	75	65	95	80	100
ADC 0 Intake	56	62	75	63	83	76	95
ADC 0 Exhaust	56	62	75	63	83	76	95
ADC 0 ADC-XF1	70	80	90	90	107	107	115
ADC 0 ADC-XF0	70	80	90	90	107	107	115
ADC 1 Intake	56	62	75	63	83	76	95
ADC 1 Exhaust	56	62	75	63	83	76	95
ADC 1 ADC-XF1	70	80	90	90	107	107	115
ADC 1 ADC-XF0	70	80	90	90	107	107	115
ADC 2 Intake	56	62	75	63	83	76	95
ADC 2 Exhaust	56	62	75	63	83	76	95
ADC 2 ADC-XF1	70	80	90	90	107	107	115
ADC 2 ADC-XF0	70	80	90	90	107	107	115
ADC 3 Intake	56	62	75	63	83	76	95
ADC 3 Exhaust	56	62	75	63	83	76	95
ADC 3 ADC-XF1	70	80	90	90	107	107	115
ADC 3 ADC-XF0	70	80	90	90	107	107	115
ADC 4 Intake	56	62	75	63	83	76	95
ADC 4 Exhaust	56	62	75	63	83	76	95
ADC 4 ADC-XF1	70	80	90	90	107	107	115
ADC 4 ADC-XF0	70	80	90	90	107	107	115
ADC 5 Intake	56	62	75	63	83	76	95
ADC 5 Exhaust	56	62	75	63	83	76	95
ADC 5 ADC-XF1	70	80	90	90	107	107	115
ADC 5 ADC-XF0	70	80	90	90	107	107	115
ADC 6 Intake	56	62	75	63	83	76	95
ADC 6 Exhaust	56	62	75	63	83	76	95
ADC 6 ADC-XF1	70	80	90	90	107	107	115
ADC 6 ADC-XF0	70	80	90	90	107	107	115
ADC 7 Intake	56	62	75	63	83	76	95
ADC 7 Exhaust	56	62	75	63	83	76	95
ADC 7 ADC-XF1	70	80	90	90	107	107	115
ADC 7 ADC-XF0	70	80	90	90	107	107	115
ADC 8 Intake	56	62	75	63	83	76	95
ADC 8 Exhaust	56	62	75	63	83	76	95
ADC 8 ADC-XF1	70	80	90	90	107	107	115
ADC 8 ADC-XF0	70	80	90	90	107	107	115

ADC 9 Intake	56	62	75	63	83	76	95
ADC 9 Exhaust	56	62	75	63	83	76	95
ADC 9 ADC-XF1	70	80	90	90	107	107	115
ADC 9 ADC-XF0	70	80	90	90	107	107	115

show chassis temperature-thresholds (MX2020 Router)

```
user@host> show chassis temperature-thresholds
```

	Fan speed		Yellow alarm		Red alarm		Fire Shutdown
	(degrees C)		(degrees C)		(degrees C)		(degrees C)
Item	Normal	High	Normal	Bad fan	Normal	Bad fan	Normal
Routing Engine 0	70	80	95	95	110	110	112
Routing Engine 1	70	80	95	95	110	110	112
CB 0 IntakeA-Zone0	60	65	78	75	85	80	95
CB 0 IntakeB-Zone1	60	65	78	75	85	80	95
CB 0 IntakeC-Zone0	60	65	78	75	85	80	95
CB 0 ExhaustA-Zone0	60	65	78	75	85	80	95
CB 0 ExhaustB-Zone1	60	65	78	75	85	80	95
CB 0 TCBC-Zone0	60	65	78	75	85	80	95
CB 1 IntakeA-Zone0	60	65	78	75	85	80	95
CB 1 IntakeB-Zone1	60	65	78	75	85	80	95
CB 1 IntakeC-Zone0	60	65	78	75	85	80	95
CB 1 ExhaustA-Zone0	60	65	78	75	85	80	95
CB 1 ExhaustB-Zone1	60	65	78	75	85	80	95
CB 1 TCBC-Zone0	60	65	78	75	85	80	95
SPMB 0 Intake	56	62	75	63	83	76	95
SPMB 1 Intake	56	62	75	63	83	76	95
SFB 0 Intake-Zone0	56	62	75	63	82	70	87
SFB 0 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 0 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 0 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 0 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 0 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 0 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 0 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 1 Intake-Zone0	56	62	75	63	82	70	87
SFB 1 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 1 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 1 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 1 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 1 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 1 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 1 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 2 Intake-Zone0	56	62	75	63	82	70	87
SFB 2 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 2 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 2 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 2 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 2 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 2 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 2 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 3 Intake-Zone0	56	62	75	63	82	70	87
SFB 3 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 3 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 3 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 3 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 3 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 3 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 3 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 4 Intake-Zone0	56	62	75	63	82	70	87

SFB 4 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 4 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 4 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 4 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 4 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 4 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 4 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 5 Intake-Zone0	56	62	75	63	82	70	87
SFB 5 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 5 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 5 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 5 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 5 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 5 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 5 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 6 Intake-Zone0	56	62	75	63	82	70	87
SFB 6 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 6 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 6 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 6 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 6 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 6 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 6 SFB-XF0-Zone0	70	80	90	90	107	107	115
SFB 7 Intake-Zone0	56	62	75	63	82	70	87
SFB 7 Exhaust-Zone1	56	62	75	63	82	70	87
SFB 7 IntakeA-Zone0	56	62	75	63	82	70	87
SFB 7 IntakeB-Zone1	56	62	75	63	82	70	87
SFB 7 Exhaust-Zone0	56	62	75	63	82	70	87
SFB 7 SFB-XF2-Zone1	70	80	90	90	107	107	115
SFB 7 SFB-XF1-Zone0	70	80	90	90	107	107	115
SFB 7 SFB-XF0-Zone0	70	80	90	90	107	107	115
FPC 0	55	60	75	65	90	80	95
FPC 1	55	60	75	65	90	80	95
FPC 2	55	60	75	65	90	80	95
FPC 3	55	60	75	65	90	80	95
FPC 4	55	60	75	65	90	80	95
FPC 5	55	60	75	65	90	80	95
FPC 6	55	60	75	65	90	80	95
FPC 7	55	60	75	65	90	80	95
FPC 8	55	60	75	65	90	80	95
FPC 9	55	60	75	65	90	80	95
FPC 10	55	60	75	65	90	80	95
FPC 11	55	60	75	65	90	80	95
FPC 12	55	60	75	65	90	80	95
FPC 13	55	60	75	65	90	80	95
FPC 14	55	60	75	65	90	80	95
FPC 15	55	60	75	65	90	80	95
FPC 16	55	60	75	65	90	80	95
FPC 17	55	60	75	65	90	80	95
FPC 18	55	60	75	65	90	80	95
FPC 19	55	60	75	65	90	80	95
ADC 0 Intake	56	62	75	63	83	76	95
ADC 0 Exhaust	56	62	75	63	83	76	95
ADC 0 ADC-XF1	70	80	90	90	107	107	115
ADC 0 ADC-XF0	70	80	90	90	107	107	115
ADC 1 Intake	56	62	75	63	83	76	95
ADC 1 Exhaust	56	62	75	63	83	76	95
ADC 1 ADC-XF1	70	80	90	90	107	107	115
ADC 1 ADC-XF0	70	80	90	90	107	107	115
ADC 2 Intake	56	62	75	63	83	76	95
ADC 2 Exhaust	56	62	75	63	83	76	95

ADC 2 ADC-XF1	70	80	90	90	107	107	115
ADC 2 ADC-XF0	70	80	90	90	107	107	115
ADC 3 Intake	56	62	75	63	83	76	95
ADC 3 Exhaust	56	62	75	63	83	76	95
ADC 3 ADC-XF1	70	80	90	90	107	107	115
ADC 3 ADC-XF0	70	80	90	90	107	107	115
ADC 4 Intake	56	62	75	63	83	76	95
ADC 4 Exhaust	56	62	75	63	83	76	95
ADC 4 ADC-XF1	70	80	90	90	107	107	115
ADC 4 ADC-XF0	70	80	90	90	107	107	115
ADC 5 Intake	56	62	75	63	83	76	95
ADC 5 Exhaust	56	62	75	63	83	76	95
ADC 5 ADC-XF1	70	80	90	90	107	107	115
ADC 5 ADC-XF0	70	80	90	90	107	107	115
ADC 6 Intake	56	62	75	63	83	76	95
ADC 6 Exhaust	56	62	75	63	83	76	95
ADC 6 ADC-XF1	70	80	90	90	107	107	115
ADC 6 ADC-XF0	70	80	90	90	107	107	115
ADC 7 Intake	56	62	75	63	83	76	95
ADC 7 Exhaust	56	62	75	63	83	76	95
ADC 7 ADC-XF1	70	80	90	90	107	107	115
ADC 7 ADC-XF0	70	80	90	90	107	107	115
ADC 8 Intake	56	62	75	63	83	76	95
ADC 8 Exhaust	56	62	75	63	83	76	95
ADC 8 ADC-XF1	70	80	90	90	107	107	115
ADC 8 ADC-XF0	70	80	90	90	107	107	115
ADC 9 Intake	56	62	75	63	83	76	95
ADC 9 Exhaust	56	62	75	63	83	76	95
ADC 9 ADC-XF1	70	80	90	90	107	107	115
ADC 9 ADC-XF0	70	80	90	90	107	107	115
ADC 10 Intake	56	62	75	63	83	76	95
ADC 10 Exhaust	56	62	75	63	83	76	95
ADC 10 ADC-XF1	70	80	90	90	107	107	115
ADC 10 ADC-XF0	70	80	90	90	107	107	115
ADC 11 Intake	56	62	75	63	83	76	95
ADC 11 Exhaust	56	62	75	63	83	76	95
ADC 11 ADC-XF1	70	80	90	90	107	107	115
ADC 11 ADC-XF0	70	80	90	90	107	107	115
ADC 12 Intake	56	62	75	63	83	76	95
ADC 12 Exhaust	56	62	75	63	83	76	95
ADC 12 ADC-XF1	70	80	90	90	107	107	115
ADC 12 ADC-XF0	70	80	90	90	107	107	115
ADC 13 Intake	56	62	75	63	83	76	95
ADC 13 Exhaust	56	62	75	63	83	76	95
ADC 13 ADC-XF1	70	80	90	90	107	107	115
ADC 13 ADC-XF0	70	80	90	90	107	107	115
ADC 14 Intake	56	62	75	63	83	76	95
ADC 14 Exhaust	56	62	75	63	83	76	95
ADC 14 ADC-XF1	70	80	90	90	107	107	115
ADC 14 ADC-XF0	70	80	90	90	107	107	115
ADC 15 Intake	56	62	75	63	83	76	95
ADC 15 Exhaust	56	62	75	63	83	76	95
ADC 15 ADC-XF1	70	80	90	90	107	107	115
ADC 15 ADC-XF0	70	80	90	90	107	107	115
ADC 16 Intake	56	62	75	63	83	76	95
ADC 16 Exhaust	56	62	75	63	83	76	95
ADC 16 ADC-XF1	70	80	90	90	107	107	115
ADC 16 ADC-XF0	70	80	90	90	107	107	115
ADC 17 Intake	56	62	75	63	83	76	95
ADC 17 Exhaust	56	62	75	63	83	76	95
ADC 17 ADC-XF1	70	80	90	90	107	107	115

ADC 17 ADC-XF0	70	80	90	90	107	107	115
ADC 18 Intake	56	62	75	63	83	76	95
ADC 18 Exhaust	56	62	75	63	83	76	95
ADC 18 ADC-XF1	70	80	90	90	107	107	115
ADC 18 ADC-XF0	70	80	90	90	107	107	115
ADC 19 Intake	56	62	75	63	83	76	95
ADC 19 Exhaust	56	62	75	63	83	76	95
ADC 19 ADC-XF1	70	80	90	90	107	107	115
ADC 19 ADC-XF0	70	80	90	90	107	107	115

show chassis temperature-thresholds (MX2020 with MPC4E)

```
user@host> show chassis temperature-thresholds
```

Fan speed	Yellow alarm (degrees C)		Red alarm (degrees C)		Fire Shutdown (degrees C)	
(degrees C)	Normal	High	Normal	Bad fan	Normal	Bad fan
Item						
Normal						
Routing Engine 0	70	80	95	95	110	110
112						
Routing Engine 1	70	80	95	95	110	110
112						
CB 0 IntakeA-Zone0	60	65	78	75	85	80
95						
CB 0 IntakeB-Zone1	60	65	78	75	85	80
95						
CB 0 IntakeC-Zone0	60	65	78	75	85	80
95						
CB 0 ExhaustA-Zone0	60	65	78	75	85	80
95						
CB 0 ExhaustB-Zone1	60	65	78	75	85	80
95						
CB 0 TCBC-Zone0	60	65	78	75	85	80
95						
CB 1 IntakeA-Zone0	60	65	78	75	85	80
95						
CB 1 IntakeB-Zone1	60	65	78	75	85	80
95						
CB 1 IntakeC-Zone0	60	65	78	75	85	80
95						
CB 1 ExhaustA-Zone0	60	65	78	75	85	80
95						
CB 1 ExhaustB-Zone1	60	65	78	75	85	80
95						
CB 1 TCBC-Zone0	60	65	78	75	85	80
95						
SPMB 0 Intake	56	62	75	63	83	76
95						
SPMB 1 Intake	56	62	75	63	83	76
95						
SFB 0 Intake-Zone0	56	62	70	70	85	85
89						
SFB 0 Exhaust-Zone1	56	62	70	70	85	85
89						
SFB 0 IntakeA-Zone0	56	62	70	70	85	85
89						
SFB 0 IntakeB-Zone1	56	62	70	70	85	85
89						
SFB 0 Exhaust-Zone0	56	62	70	70	85	85
89						
SFB 0 SFB-XF2-Zone1	70	75	90	85	95	90

100						
SFB 0	SFB-XF1-Zone0	70	75	90	85	95
100						90
SFB 0	SFB-XF0-Zone0	70	75	90	85	95
100						90
SFB 1	Intake-Zone0	56	62	70	70	85
89						85
SFB 1	Exhaust-Zone1	56	62	70	70	85
89						85
SFB 1	IntakeA-Zone0	56	62	70	70	85
89						85
SFB 1	IntakeB-Zone1	56	62	70	70	85
89						85
SFB 1	Exhaust-Zone0	56	62	70	70	85
89						85
SFB 1	SFB-XF2-Zone1	70	75	90	85	95
100						90
SFB 1	SFB-XF1-Zone0	70	75	90	85	95
100						90
SFB 1	SFB-XF0-Zone0	70	75	90	85	95
100						90
SFB 2	Intake-Zone0	56	62	70	70	85
89						85
SFB 2	Exhaust-Zone1	56	62	70	70	85
89						85
SFB 2	IntakeA-Zone0	56	62	70	70	85
89						85
SFB 2	IntakeB-Zone1	56	62	70	70	85
89						85
SFB 2	Exhaust-Zone0	56	62	70	70	85
89						85
SFB 2	SFB-XF2-Zone1	70	75	90	85	95
100						90
SFB 2	SFB-XF1-Zone0	70	75	90	85	95
100						90
SFB 2	SFB-XF0-Zone0	70	75	90	85	95
100						90
SFB 3	Intake-Zone0	56	62	70	70	85
89						85
SFB 3	Exhaust-Zone1	56	62	70	70	85
89						85
SFB 3	IntakeA-Zone0	56	62	70	70	85
89						85
SFB 3	IntakeB-Zone1	56	62	70	70	85
89						85
SFB 3	Exhaust-Zone0	56	62	70	70	85
89						85
SFB 3	SFB-XF2-Zone1	70	75	90	85	95
100						90
SFB 3	SFB-XF1-Zone0	70	75	90	85	95
100						90
SFB 3	SFB-XF0-Zone0	70	75	90	85	95
100						90
SFB 4	Intake-Zone0	56	62	70	70	85
89						85
SFB 4	Exhaust-Zone1	56	62	70	70	85
89						85
SFB 4	IntakeA-Zone0	56	62	70	70	85
89						85
SFB 4	IntakeB-Zone1	56	62	70	70	85
89						85

SFB 4 Exhaust-Zone0 89	56	62	70	70	85	85
SFB 4 SFB-XF2-Zone1 100	70	75	90	85	95	90
SFB 4 SFB-XF1-Zone0 100	70	75	90	85	95	90
SFB 4 SFB-XF0-Zone0 100	70	75	90	85	95	90
SFB 5 Intake-Zone0 89	56	62	70	70	85	85
SFB 5 Exhaust-Zone1 89	56	62	70	70	85	85
SFB 5 IntakeA-Zone0 89	56	62	70	70	85	85
SFB 5 IntakeB-Zone1 89	56	62	70	70	85	85
SFB 5 Exhaust-Zone0 89	56	62	70	70	85	85
SFB 5 SFB-XF2-Zone1 100	70	75	90	85	95	90
SFB 5 SFB-XF1-Zone0 100	70	75	90	85	95	90
SFB 5 SFB-XF0-Zone0 100	70	75	90	85	95	90
SFB 6 Intake-Zone0 89	56	62	70	70	85	85
SFB 6 Exhaust-Zone1 89	56	62	70	70	85	85
SFB 6 IntakeA-Zone0 89	56	62	70	70	85	85
SFB 6 IntakeB-Zone1 89	56	62	70	70	85	85
SFB 6 Exhaust-Zone0 89	56	62	70	70	85	85
SFB 6 SFB-XF2-Zone1 100	70	75	90	85	95	90
SFB 6 SFB-XF1-Zone0 100	70	75	90	85	95	90
SFB 6 SFB-XF0-Zone0 100	70	75	90	85	95	90
SFB 7 Intake-Zone0 89	56	62	70	70	85	85
SFB 7 Exhaust-Zone1 89	56	62	70	70	85	85
SFB 7 IntakeA-Zone0 89	56	62	70	70	85	85
SFB 7 IntakeB-Zone1 89	56	62	70	70	85	85
SFB 7 Exhaust-Zone0 89	56	62	70	70	85	85
SFB 7 SFB-XF2-Zone1 100	70	75	90	85	95	90
SFB 7 SFB-XF1-Zone0 100	70	75	90	85	95	90
SFB 7 SFB-XF0-Zone0 100	70	75	90	85	95	90
FPC 0 95	55	60	75	65	90	80
FPC 9 95	55	60	75	65	90	80
FPC 10	55	60	75	65	90	80

95						
FPC 14	55	60	75	65	95	80
100						
FPC 19	55	60	75	65	90	80
95						
ADC 0 Intake	50	55	60	60	65	65
80						
ADC 0 Exhaust	50	55	60	60	65	65
80						
ADC 0 ADC-XF1	70	75	90	85	95	90
100						
ADC 0 ADC-XF0	70	75	90	85	95	90
100						
ADC 9 Intake	50	55	60	60	65	65
80						
ADC 9 Exhaust	50	55	60	60	65	65
80						
ADC 9 ADC-XF1	70	75	90	85	95	90
100						
ADC 9 ADC-XF0	70	75	90	85	95	90
100						
ADC 10 Intake	50	55	60	60	65	65
80						
ADC 10 Exhaust	50	55	60	60	65	65
80						
ADC 10 ADC-XF1	70	75	90	85	95	90
100						
ADC 10 ADC-XF0	70	75	90	85	95	90
100						
ADC 14 Intake	50	55	60	60	65	65
80						
ADC 14 Exhaust	50	55	60	60	65	65
80						
ADC 14 ADC-XF1	70	75	90	85	95	90
100						
ADC 14 ADC-XF0	70	75	90	85	95	90
100						
ADC 19 Intake	50	55	60	60	65	65
80						
ADC 19 Exhaust	50	55	60	60	65	65
80						
ADC 19 ADC-XF1	70	75	90	85	95	90
100						
ADC 19 ADC-XF0	70	75	90	85	95	90
100						

show chassis temperature-thresholds (T4000 Core Routers)

```
user@host> show chassis temperature-thresholds
```

Item	Fan speed (degrees C)		Yellow alarm (degrees C)		Red alarm (degrees C)		Fire Shutdown (degrees C)
	Normal	High	Normal	Bad fan	Normal	Bad fan	Normal
Chassis default	48	54	65	55	75	65	100
Routing Engine 0	55	65	85	85	100	100	102
Routing Engine 1	55	65	85	85	100	100	102
FPC 0	63	68	75	70	90	83	95
FPC 3	63	68	75	70	90	83	95
FPC 5	56	62	75	63	83	76	95
FPC 6	63	68	75	70	90	83	95
SIB 0	64	70	76	72	87	84	95

SIB 1	64	70	76	72	87	84	95
SIB 2	64	70	76	72	87	84	95
SIB 3	64	70	76	72	87	84	95
SIB 4	64	70	76	72	87	84	95

show chassis temperature-thresholds (TX Matrix Plus Router)

```
user@host> show chassis temperature-thresholds
sfc0-re0:
```

Item	Fan speed (degrees C)		Yellow alarm (degrees C)		Red alarm (degrees C)	
	Normal	High	Normal	Bad fan	Normal	Bad fan
Chassis default	48	54	65	55	75	65
Routing Engine 0	55	65	85	85	100	100
Routing Engine 1	55	65	85	85	100	100
SIB F13 0	64	70	76	72	90	84
SIB F13 3	64	70	76	72	90	84
SIB F13 6	64	70	76	72	90	84
SIB F13 8	64	70	76	72	90	84
SIB F13 11	64	70	76	72	90	84
SIB F13 12	64	70	76	72	90	84
SIB F2S 16	64	70	76	72	90	84
SIB F2S 17	64	70	76	72	90	84
SIB F2S 18	64	70	76	72	90	84
SIB F2S 19	64	70	76	72	90	84
SIB F2S 20	64	70	76	72	90	84
SIB F2S 21	64	70	76	72	90	84
SIB F2S 22	64	70	76	72	90	84
SIB F2S 23	64	70	76	72	90	84
SIB F2S 24	64	70	76	72	90	84
SIB F2S 25	64	70	76	72	90	84
SIB F2S 26	64	70	76	72	90	84
SIB F2S 27	64	70	76	72	90	84
SIB F2S 28	64	70	76	72	90	84
SIB F2S 29	64	70	76	72	90	84
SIB F2S 30	64	70	76	72	90	84
SIB F2S 31	64	70	76	72	90	84
SIB F2S 32	64	70	76	72	90	84
SIB F2S 33	64	70	76	72	90	84
SIB F2S 34	64	70	76	72	90	84
SIB F2S 35	64	70	76	72	90	84

```
lcc0-re0:
```

Item	Fan speed (degrees C)		Yellow alarm (degrees C)		Red alarm (degrees C)	
	Normal	High	Normal	Bad fan	Normal	Bad fan
Chassis default	48	54	65	55	75	65
Routing Engine 0	55	65	85	85	100	100
Routing Engine 1	55	65	85	85	100	100
FPC 1	56	62	75	63	83	76
FPC 3	56	62	75	63	83	76
FPC 4	56	62	75	63	83	76
FPC 6	56	62	75	63	83	76
FPC 7	56	62	75	63	83	76
SIB 0	48	54	65	60	80	75
SIB 1	48	54	65	60	80	75
SIB 2	48	54	65	60	80	75
SIB 3	48	54	65	60	80	75
SIB 4	48	54	65	60	80	75

```
lcc1-re0:
```

Item	Fan speed (degrees C)		Yellow alarm (degrees C)		Red alarm (degrees C)	
	Normal	High	Normal	Bad fan	Normal	Bad fan
Chassis default	48	54	65	55	75	65
Routing Engine 0	55	65	85	85	100	100
Routing Engine 1	55	65	85	85	100	100
FPC 1	56	62	75	63	83	76
FPC 3	56	62	75	63	83	76
FPC 4	56	62	75	63	83	76
FPC 6	56	62	75	63	83	76
...						

show chassis temperature-thresholds lcc (TX Matrix Plus Router)

```
user@host> show chassis temperature-thresholds lcc 1
lcc1-re0:
```

Item	Fan speed (degrees C)		Yellow alarm (degrees C)		Red alarm (degrees C)	
	Normal	High	Normal	Bad fan	Normal	Bad fan
Chassis default	48	54	65	55	75	65
Routing Engine 0	55	65	85	85	100	100
Routing Engine 1	55	65	85	85	100	100
FPC 1	56	62	75	63	83	76
FPC 3	56	62	75	63	83	76
FPC 4	56	62	75	63	83	76
FPC 6	56	62	75	63	83	76
SIB 0	48	54	65	60	80	75
SIB 1	48	54	65	60	80	75
SIB 2	48	54	65	60	80	75
SIB 3	48	54	65	60	80	75
SIB 4	48	54	65	60	80	75

show chassis temperature-thresholds sfc (TX Matrix Plus Router)

```
user@host> show chassis temperature-thresholds sfc 0
sfc0-re0:
```

Item	Fan speed (degrees C)		Yellow alarm (degrees C)		Red alarm (degrees C)	
	Normal	High	Normal	Bad fan	Normal	Bad fan
Chassis default	48	54	65	55	75	65
Routing Engine 0	55	65	85	85	100	100
Routing Engine 1	55	65	85	85	100	100
SIB F13 0	64	70	76	72	90	84
SIB F13 3	64	70	76	72	90	84
SIB F13 6	64	70	76	72	90	84
SIB F13 8	64	70	76	72	90	84
SIB F13 11	64	70	76	72	90	84
SIB F13 12	64	70	76	72	90	84
SIB F2S 16	64	70	76	72	90	84
SIB F2S 17	64	70	76	72	90	84
SIB F2S 18	64	70	76	72	90	84
SIB F2S 19	64	70	76	72	90	84
SIB F2S 20	64	70	76	72	90	84
SIB F2S 21	64	70	76	72	90	84
SIB F2S 22	64	70	76	72	90	84
SIB F2S 23	64	70	76	72	90	84

SIB F2S 24	64	70	76	72	90	84
SIB F2S 25	64	70	76	72	90	84
SIB F2S 26	64	70	76	72	90	84
SIB F2S 27	64	70	76	72	90	84
SIB F2S 28	64	70	76	72	90	84
SIB F2S 29	64	70	76	72	90	84
SIB F2S 30	64	70	76	72	90	84
SIB F2S 31	64	70	76	72	90	84
SIB F2S 32	64	70	76	72	90	84
SIB F2S 33	64	70	76	72	90	84
SIB F2S 34	64	70	76	72	90	84
SIB F2S 35	64	70	76	72	90	84

show chassis temperature-thresholds (QFX3500 Switch and QFX3600)

```
user@switch> show chassis temperature-thresholds
```

Item	Fan speed (degrees C)		Yellow alarm (degrees C)		Red alarm (degrees C)	
	Normal	High	Normal	Bad fan	Normal	Bad fan
FPC Sensor TopLeft I	48	56	53	43	56	46
FPC Sensor TopRight I	46	54	51	41	54	44
FPC Sensor TopLeft E	58	65	62	52	65	55
FPC Sensor TopRight E	56	64	61	51	64	54
FPC Sensor TopMiddle I	58	64	61	51	64	54
FPC Sensor TopMiddle E	67	74	71	61	74	64
FPC Sensor Bottom I	59	67	64	54	67	57
FPC Sensor Bottom E	66	73	70	60	73	63
FPC Sensor Die Temp	69	75	72	62	75	65
FPC Sensor Mgmt Brd I	46	54	51	41	54	44
FPC Sensor Switch I	56	63	60	50	63	53

show chassis temperature-thresholds interconnect-device (QFabric System)

```
user@switch> show chassis temperature-thresholds interconnect-device interconnect1
temperature-thresholds interconnect-device interconnect1
```

Item	Fan speed		Yellow alarm		Red alarm	
	Normal	High	Normal	Bad fan	Normal	Bad fan
Chassis default	48	54	65	55	75	65

show chassis temperature-thresholds (PTX5000 Packet Transport Switch)

```
user@switch> show chassis temperature-thresholds
user@switch> show chassis temperature-thresholds
```

Item	Fan speed (degrees C)		Yellow alarm (degrees C)		Red alarm (degrees C)		Fire Shutdown (degrees C)
	Normal	High	Normal	Bad fan	Normal	Bad fan	Normal
Routing Engine 0	70	75	90	87	102	97	115
Routing Engine 1	70	75	90	87	102	97	115
CB 0 Exhaust A	60	65	78	75	85	80	95
CB 0 Exhaust B	60	65	78	75	85	80	95
CB 1 Exhaust A	60	65	78	75	85	80	95
CB 1 Exhaust B	20	25	65	60	80	75	100
FPC 1 Exhaust A	60	65	78	75	85	80	95
FPC 1 Exhaust B	60	65	78	75	85	80	95
FPC 1 TL0	70	75	90	87	102	97	115
FPC 1 TQ0	70	75	90	87	102	97	115
FPC 1 TL1	70	75	90	87	102	97	115
FPC 1 TQ1	70	75	90	87	102	97	115
FPC 1 TL2	70	75	90	87	102	97	115
FPC 1 TQ2	70	75	90	87	102	97	115

FPC 1 TL3	70	75	90	87	102	97	115
FPC 1 TQ3	70	75	90	87	102	97	115
FPC 2 Exhaust A	60	65	78	75	85	80	95
FPC 2 Exhaust B	60	65	78	75	85	80	95
FPC 2 TL0	70	75	90	87	102	97	115
FPC 2 TQ0	70	75	90	87	102	97	115
FPC 2 TL1	70	75	90	87	102	97	115
FPC 2 TQ1	70	75	90	87	102	97	115
FPC 2 TL2	70	75	90	87	102	97	115
FPC 2 TQ2	70	75	90	87	102	97	115
FPC 2 TL3	70	75	90	87	102	97	115
FPC 2 TQ3	70	75	90	87	102	97	115
PIC 2/0 Ambient	60	65	78	75	85	80	95
PIC 2/0 cfp-2/0/1	60	65	70	67	75	72	85
PIC 2/1 Ambient	60	65	78	75	85	80	95
SIB 0 Exhaust	60	65	78	75	85	80	95
SIB 0 Junction	70	75	90	87	102	97	115
SIB 1 Exhaust	60	65	78	75	85	80	95
SIB 1 Junction	70	75	90	87	102	97	115
SIB 2 Exhaust	60	65	78	75	85	80	95
SIB 2 Junction	70	75	90	87	102	97	115
SIB 3 Exhaust	60	65	78	75	85	80	95
SIB 3 Junction	70	75	90	87	102	97	115
SIB 4 Exhaust	60	65	78	75	85	80	95
SIB 4 Junction	70	75	90	87	102	97	115
SIB 5 Exhaust	60	65	78	75	85	80	95
SIB 5 Junction	70	75	90	87	102	97	115
SIB 6 Exhaust	60	65	78	75	85	80	95
SIB 6 Junction	70	75	90	87	102	97	115
SIB 7 Exhaust	60	65	78	75	85	80	95
SIB 7 Junction	70	75	90	87	102	97	115
SIB 8 Exhaust	60	65	78	75	85	80	95
SIB 8 Junction	70	75	90	87	102	97	115

show chassis temperature-thresholds (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis temperature-thresholds
```

Fan speed	Yellow alarm		Red alarm		Fire Shutdown	
(degrees C)	(degrees C)		(degrees C)		(degrees C)	
Item	Normal	High	Normal	Bad fan	Normal	Bad fan
Normal						
Chassis default	48	54	65	55	75	65
100						
Routing Engine 0	70	80	95	95	110	110
112						
Routing Engine 1	70	80	95	95	110	110
112						
FPC 0	55	60	75	65	90	80
95						
FPC 1	55	60	75	65	90	80
95						
FPC 2	55	60	75	65	90	80
95						
FPC 4	55	60	75	65	90	80
95						
FPC 5	55	60	75	65	90	80
95						

show log

Syntax	show log <filename user <username>>
Syntax (QFabric System)	show log <filename>
Syntax (TX Matrix Routers)	show log <all-lcc lcc <i>number</i> scc> <filename user <username>>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	List log files, display log file contents, or display information about users who have logged in to the router or switch.
Options	<p>none—List all log files.</p> <p><all-lcc lcc <i>number</i> scc>—(TX Matrix routers only) (Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace <i>number</i> with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).</p> <p>filename—(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.</p> <p>user <username>—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include <i>username</i>, display logging information about the specified user.</p>
Required Privilege Level	trace
List of Sample Output	show log on page 804 show log filename on page 805 show log filename (QFabric System) on page 805 show log user on page 806

Sample Output

show log

```

user@host> show log
total 57518
-rw-r--r--  1 root  bin      211663 Oct  1 19:44 dcd
-rw-r--r--  1 root  bin      999947 Oct  1 19:41 dcd.0
-rw-r--r--  1 root  bin      999994 Oct  1 17:48 dcd.1
-rw-r--r--  1 root  bin      238815 Oct  1 19:44 rpd
-rw-r--r--  1 root  bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r--  1 root  bin     1061095 Oct  1 12:13 rpd.1
-rw-r--r--  1 root  bin     1052026 Oct  1 06:08 rpd.2

```

```

-rw-r--r-- 1 root bin      1056309 Sep 30 18:21 rpd.3
-rw-r--r-- 1 root bin      1056371 Sep 30 14:36 rpd.4
-rw-r--r-- 1 root bin      1056301 Sep 30 10:50 rpd.5
-rw-r--r-- 1 root bin      1056350 Sep 30 07:04 rpd.6
-rw-r--r-- 1 root bin      1048876 Sep 30 03:21 rpd.7
-rw-rw-r-- 1 root bin          19656 Oct  1 19:37 wtmp

```

show log filename

```

user@host> show log rpd
Oct  1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT recv len 56 V9 seq 148 op add Type route/if af 2 addr
13.13.13.21 nhop type local nhop 13.13.13.21
Oct  1 18:00:19 KRT recv len 56 V9 seq 149 op add Type route/if af 2 addr
13.13.13.22 nhop type unicast nhop 13.13.13.22
Oct  1 18:00:19 KRT recv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct  1 18:00:19 KRT recv len 144 V9 seq 151 op chnge Type ifdev devindex 44
Oct  1 18:00:19 KRT recv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT recv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct  1 18:00:19 KRT recv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...

```

show log filename (QFabric System)

```

user@qfabric> show log messages
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_REO_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_REO_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492

```

```
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)
```

show log user

```
user@host> show log user
darius  mg2546                Thu Oct  1 19:37  still logged in
darius  mg2529                Thu Oct  1 19:08 - 19:36 (00:28)
darius  mg2518                Thu Oct  1 18:53 - 18:58 (00:04)
root    mg1575                Wed Sep 30 18:39 - 18:41 (00:02)
root    ttyp2      jun.site.per Wed Sep 30 18:39 - 18:41 (00:02)
alex    ttyp1      192.168.1.2  Wed Sep 30 01:03 - 01:22 (00:19)
```

show system alarms

Syntax	show system alarms
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display active system alarms.
Options	This command has no options.
Additional Information	System alarms are preset. They include a configuration alarm that appears when no rescue configuration alarm is set and a license alarm that appears when a software feature is configured and no valid license is configured for the feature. For more information about system alarms, see the <i>Junos OS System Basics Configuration Guide</i> . In Junos OS release 11.1 and later, alarms for fans also show the slot number of the fans in the CLI output.
Required Privilege Level	admin
List of Sample Output	show system alarms on page 807 show system alarms (Fan Tray) on page 807 show system alarms (QFX Series) on page 807

Sample Output

show system alarms

```
user@host> show system alarms
2 alarms currently active
Alarm time          Class    Description
2005-02-24 17:29:34 UTC  Minor    IPsec VPN tunneling usage requires a
license
2005-02-24 17:29:34 UTC  Minor    Rescue configuration is not sent
```

show system alarms (Fan Tray)

```
user@host> show system alarms
4 alarms currently active
Alarm time          Class    Description
2010-11-11 20:27:38 UTC  Major    Side Fan Tray 7 Failure
2010-11-11 20:27:13 UTC  Minor    Side Fan Tray 7 Overspeed
2010-11-11 20:27:13 UTC  Major    Side Fan Tray 5 Failure
2010-11-11 20:27:13 UTC  Major    Side Fan Tray 0 Failure
```

show system alarms (QFX Series)

```
user@switches> show system alarms
2 alarms currently active
Alarm time Class Description
2005-02-24 17:29:34 UTC Minor Rescue configuration is not sent
```

show system processes

Syntax	<pre>show system processes <brief detail extensive summary> <health (pid <i>process-identifier</i> process-name <i>process-name</i>)> <providers> <resource-limits (brief detail) <i>process-name</i>> <wide></pre>
Syntax (EX Series Switches)	<pre>show system processes <all-members> <brief detail extensive summary> <health (pid <i>process-identifier</i> process-name <i>process-name</i>)> <local> <member <i>member-id</i>> <providers> <resource-limits (brief detail) <i>process-name</i>> <wide></pre>
Syntax (MX Series Routers)	<pre>show system processes <all-members> <brief detail extensive summary> <health (pid <i>process-identifier</i> process-name <i>process-name</i>)> <local> <member <i>member-id</i>> <providers> <resource-limits (brief detail) <i>process-name</i>> <wide></pre>
Syntax (QFX Series)	<pre>show system processes <brief detail extensive summary > <health (pid <i>process-identifier</i> process-name <i>process-name</i>)> <interconnect-device <i>name</i>> <node-group <i>name</i>> <providers> <resource-limits> <wide></pre>
Syntax (TX Matrix Routers)	<pre>show system processes <brief detail extensive summary> <all-chassis all-lcc lcc <i>number</i> scc> <wide></pre>
Syntax (TX Matrix Plus Routers)	<pre>show system processes <brief detail extensive summary> <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i>> <wide></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option sfc introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>

Description	Display information about software processes that are running on the router or switch and that have controlling terminals.
Options	<p>none—Display standard information about system processes.</p> <p>brief detail extensive summary—(Optional) Display the specified level of detail.</p> <p>adaptive-services—(Optional) Display the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.</p> <p>alarm-control—(Optional) Display the process to configure the system alarm.</p> <p>all-chassis—(TX Matrix and TX Matrix Plus routers only) (Optional) Display standard system process information about all the T640 routers (in a routing matrix based on the TX Matrix router) or all the T1600 routers (in a routing matrix based on the TX Matrix Plus router) in the chassis.</p> <p>all-lcc—(TX Matrix and TX Matrix Plus routers only) (Optional) Display standard system process information for all T640 routers (or line-card chassis) connected to the TX Matrix router. Display standard system process information for all T1600 routers (or line-card chassis) connected to the TX Matrix Plus router.</p> <p>all-members—(EX4200 switches and MX Series routers only) (Optional) Display standard system process information for all members of the Virtual Chassis configuration.</p> <p>ancpd-service—Display the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.</p> <p>application-identification—Display the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.</p> <p>audit-process—(Optional) Display the RADIUS accounting process.</p> <p>auto-configuration—Display the Interface Auto-Configuration process.</p> <p>bootp—Display the process that enables a router, switch, or interface to act as a Dynamic Host Configuration Protocol (DHCP) or bootstrap protocol (BOOTP) relay agent. DHCP relaying is disabled.</p> <p>captive-portal-content-delivery—Display the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.</p> <p>ce-l2tp-service—(Optional) (M10, M10i, M7i, and MX Series routers only) Display the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.</p>

cfm—Display Ethernet Operations, Administration, and Maintenance (OAM) connectivity fault management (CFM) process, which can be used to monitor the physical link between two switches.

chassis-control—(Optional) Display the chassis management process.

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

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

craft-control—Display the process for the I/O of the craft interface.

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

datapath-trace-service—Display the packet path tracing process.

dhcp-service—(EX Series switches and MX Series routers only) (Optional) Display the Dynamic Host Configuration Protocol process, which enables a DHCP server to allocate network IP addresses and deliver configuration settings to client hosts without user intervention.

diameter-service—(Optional) Display the diameter process.

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

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

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

ethernet-connectivity-fault-management—Display the process that provides IEEE 802.1ag OAM connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.

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

event-processing—(Optional) Display the event process (eventd).

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

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

health (*pid* ***process-identifier*** | ***process-name*** ***process-name***)—(Optional) Display process health information, either by process id (PID) or by process name.

iccp-service—Display the Inter-Chassis Communication Protocol (ICCP) process.

idp-policy—Display the intrusion detection and prevention (IDP) protocol process.

ilmi—Display the Integrated Local Management Interface (ILMI) protocol process, which provides bidirectional exchange of management information between two ATM interfaces across a physical connection.

inet-process—Display the IP multicast family process.

init—Display the process that initializes the USB modem.

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

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

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

l2cpd-service—Display the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.

lACP—(Optional) Display the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG, and to enable the transmission and reception processes for the link to function in an orderly manner.

lcc number—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display standard system process information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display standard system process information for a specific T1600 router that is connected to the TX Matrix Plus router. Replace ***number*** with a value from 0 through 3.

local—(EX4200 switches and MX Series routers only) (Optional) Display standard system process information for the local Virtual Chassis member.

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

logical-system-mux—Display the logical router multiplexer process (lrmuxd), which manages the multiple instances of the routing protocols process (rpd) on a machine running logical routers.

mac-validation—Display the MAC validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

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

mib-process—(Optional) Display the MIB II process, which provides the router's MIB II agent.

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

moundd-service—(EX Series switches and MX Series routers only) (Optional) Display the service for NFS mounts requests.

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

mspd—(Optional) Display the Multiservice process.

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

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

neighbor-liveness—Display the process, which specifies the maximum length of time that the router waits for its neighbor to re-establish an LDP session.

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

ntp—Display the Network Time Protocol (NTP) process, which provides the mechanisms to synchronize time and coordinate time distribution in a large, diverse network.

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

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

periodic-packet-services—Display the Periodic packet management process, which is responsible for processing a variety of time-sensitive periodic tasks so that other processes can more optimally direct their resources.

pfe—Display the Packet Forwarding Engine management process.

pgcp-service—(Optional) Display the pgcpd service process running on the Routing Engine.

pgm—Display the Pragmatic General Multicast (PGM) protocol process, which enables a reliable transport layer for multicast applications.

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

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

ppp-service—Display the Universal edge PPP process, which is the encapsulation protocol process for transporting IP traffic across universal edge routers.

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

process-monitor—Display the process health monitor process (pmond).

providers—(Optional) Display provider processes.

redundancy-interface-process—(Optional) Display the ASP redundancy process.

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

resource-cleanup—Display the resource cleanup process.

resource-limits (brief | detail) *process-name*—(Optional) Display process resource limits.

routing—(Optional) Display the routing protocol process.

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

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

scc—(TX Matrix routers only) (Optional) Display standard system process information for the TX Matrix router (or switch-card chassis).

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

secure-neighbor-discovery—(EX Series switches and MX Series routers only) (Optional) Display the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.

send—(Optional) Display the Secure Neighbor Discovery Protocol (SEND) process, which provides support for protecting Neighbor Discovery Protocol (NDP) messages.

service-deployment—(Optional) Display the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.

sfc *number*—(TX Matrix Plus routers only) (Optional) Display system process information for the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

snmp—Display the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.

sonet-aps—Display the SONET Automatic Protection Switching (APS) process, which monitors any SONET interface that participates in APS.

static-subscribers—(Optional) Display the Static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.

tunnel-oamd—(Optional) Display the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.

vrrp—(EX Series switches and MX Series routers only) (Optional) Display the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

watchdog—Display the watchdog timer process, which enables the watchdog timer when Junos OS encounters a problem.

wide—(Optional) Display process information that might be wider than 80 columns.

Additional Information By default, when you issue the **show system processes** command on a TX Matrix or TX Matrix Plus master Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) master Routing Engines connected to it. Likewise, if you issue the same command on the TX Matrix or TX Matrix Plus backup Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) backup Routing Engines that are connected to it.

Required Privilege Level

view

Related Documentation

- *List of Junos OS Processes*

List of Sample Output

[show system processes on page 817](#)
[show system processes brief on page 817](#)
[show system processes detail on page 818](#)
[show system processes extensive on page 818](#)

[show system processes lcc wide \(TX Matrix Routing Matrix\) on page 819](#)
[show system processes summary on page 820](#)
[show system processes \(TX Matrix Plus Router\) on page 820](#)
[show system processes sfc \(TX Matrix Plus Router\) on page 827](#)
[show system processes lcc wide \(TX Matrix Plus Routing Matrix\) on page 830](#)
[show system processes \(QFX Series\) on page 831](#)

Output Fields [Table 65 on page 815](#) describes the output fields for the **show system processes** command. Output fields are listed in the approximate order in which they appear.

Table 65: show system processes Output Fields

Field Name	Field Description	Level of Output
last pid	Last process identifier assigned to the process.	brief extensive summary
load averages	Three load averages followed by the current time.	brief extensive summary
processes	Number of existing processes and the number of processes in each state (sleeping , running , starting , zombies , and stopped).	brief extensive summary
Mem	Information about physical and virtual memory allocation.	brief extensive summary
Swap	Information about physical and virtual memory allocation.	brief extensive summary
PID	Process identifier.	detail extensive summary
TT	Control terminal name.	none detail

Table 65: show system processes Output Fields (*continued*)

Field Name	Field Description	Level of Output
STAT	<p>Symbolic process state. The state is given by a sequence of letters. The first letter indicates the run state of the process:</p> <ul style="list-style-type: none"> • D—In disk or other short-term, uninterruptible wait • I—Idle (sleeping longer than about 20 seconds) • R—Runnable • S—Sleeping for less than 20 seconds • T—Stopped • Z—Dead (zombie) • + —The process is in the foreground process group of its control terminal. • <—The process has raised CPU scheduling priority. • >—The process has specified a soft limit on memory requirements and is currently exceeding that limit; such a process is not swapped. • A—The process requested random page replacement. • E—The process is trying to exit. • L—The process has pages locked in core. • N—The process has reduced CPU scheduling priority. • S—The process requested first-in, first-out (FIFO) page replacement. • s—The process is a session leader. • V—The process is temporarily suspended. • W—The process is swapped out. • X—The process is being traced or debugged. 	none detail
UID	User identifier.	detail
USERNAME	Process owner.	extensive summary
PPID	Parent process identifier.	detail
CPU	<p>(D)—Short-term CPU usage.</p> <p>(E and S)—Raw (unweighted) CPU usage. The value of this field is used to sort the processes in the output.</p>	detail extensive summary
RSS	Resident set size.	detail
WCHAN	Symbolic name of the wait channel.	detail
STARTED	Local time when the process started running.	detail
PRI	Current priority of the process. A lower number indicates a higher priority.	detail extensive summary
NI or NICE	UNIX "niceness" value. A lower number indicates a higher priority.	detail extensive summary
SIZE	Total size of the process (text, data, and stack), in kilobytes.	extensive summary

Table 65: show system processes Output Fields (*continued*)

Field Name	Field Description	Level of Output
RES	Current amount of resident memory, in kilobytes.	extensive summary
STATE	Current state of the process (for example, sleep , wait , run , idle , zombie , or stop).	extensive summary
TIME	(S)—Number of system and user CPU seconds that the process has used. (None, D, and E)—Total amount of time that the command has been running.	detail extensive summary
WCPU	Weighted CPU usage.	extensive summary
COMMAND	Command that is currently running.	detail extensive summary

Sample Output

show system processes

```

user@host> show system processes
PID  TT  STAT      TIME COMMAND
  0  ??  DLs      0:00.70 (swapper)
  1  ??  Is       0:00.35 /sbin/init --
  2  ??  DL       0:00.00 (pagedaemon)
  3  ??  DL       0:00.00 (vmdaemon)
  4  ??  DL       0:42.37 (update)
  5  ??  DL       0:00.00 (if_jnx)
 80  ??  Ss       0:14.66 syslogd -s
 96  ??  Is       0:00.01 portmap
128  ??  Is       0:02.70 cron
173  ??  Is       0:02.24 /usr/local/sbin/sshd (sshd1)
189  ??  S        0:03.80 /sbin/watchdog -t180
190  ??  I        0:00.03 /usr/sbin/tnetd -N
191  ??  S        2:24.76 /sbin/ifd -N
192  ??  S<       0:55.44 /usr/sbin/xntpd -N
195  ??  S        0:53.11 /usr/sbin/snmpd -N
196  ??  S        1:15.73 /usr/sbin/mib2d -N
198  ??  I        0:00.75 /usr/sbin/inetd -N
2677 ??  I        0:00.01 /usr/sbin/mgd -N
2712 ??  Ss       0:00.24 rlogind
2735 ??  R        0:00.00 /bin/ps -ax
1985 p0-  S        0:07.41 ./rpd -N
2713 p0  Is       0:00.24 -tcsh (tcsh)
2726 p0  S+       0:00.07 cli

```

show system processes brief

```

user@host> show system processes brief
last pid:  543;  load averages:  0.00,  0.00,  0.00    18:29:47
37 processes:  1 running, 36 sleeping

Mem: 25M Active, 3976K Inact, 19M Wired, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free

```

show system processes detail

```

user@host> show system processes detail

```

PID	UID	PPID	CPU	PRI	NI	RSS	WCHAN	STARTED	TT	STAT	TIME	COMMAND
3151	1049	3129	2	28	0	672	-	1:13PM	p0	R+	0:00.00	ps -ax -r
1	0	0	0	10	0	376	wait	1:51PM	??	Is	0:00.29	/sbin/ini
2	0	0	0	-18	0	12	psleep	1:51PM	??	DL	0:00.00	(pagedae
3	0	0	0	28	0	12	psleep	1:51PM	??	DL	0:00.00	(vmdaemon
4	0	0	0	28	0	12	update	1:51PM	??	DL	0:07.15	(update)
5	0	0	0	2	0	12	pfesel	1:51PM	??	IL	0:02.90	(if_pfe)
27	0	1	0	10	0	17936	mfsidl	1:51PM	??	Is	0:00.46	mfs /dev/
81	0	1	0	2	0	496	select	1:52PM	??	Ss	0:31.21	syslogd -
119	1	1	0	2	0	492	select	1:52PM	??	Is	0:00.00	portmap
134	0	1	0	2	0	580	select	1:52PM	??	S	0:02.95	amd -p -a
151	0	1	0	18	0	532	pause	1:52PM	??	Is	0:00.34	cron
183	0	1	0	2	0	420	select	1:52PM	??	Ss	0:00.07	/usr/loca
206	0	1	0	18	0	72	pause	1:52PM	??	S	0:00.51	/sbin/wat
207	0	1	0	2	0	520	select	1:52PM	??	I	0:00.16	/usr/sbin
208	0	1	0	2	0	536	select	1:52PM	??	S	0:08.21	/sbin/dcd
210	0	1	255	2	-12	740	select	1:52PM	??	S<	0:05.83	/usr/sbin
211	0	1	0	2	0	376	select	1:52PM	??	S	0:00.03	/usr/sbin
215	0	1	0	2	0	548	select	1:52PM	??	I	0:00.50	/usr/sbin
219	0	1	0	3	0	540	ttyin	1:52PM	v0	Is+	0:00.02	/usr/libe
220	0	1	0	3	0	540	ttyin	1:52PM	v1	Is+	0:00.01	/usr/libe
221	0	1	0	3	0	540	ttyin	1:52PM	v2	Is+	0:00.01	/usr/libe
222	0	1	0	3	0	540	ttyin	1:52PM	v3	Is+	0:00.01	/usr/libe
735	0	1	0	2	0	468	select	2:47PM	??	S	0:19.14	/usr/sbin
736	0	1	0	2	0	212	select	2:47PM	??	S	0:14.13	/usr/sbin
1380	0	1	0	3	0	888	ttyin	7:32PM	d0	Is+	0:00.46	bash
3019	0	207	0	2	0	636	select	10:49AM	??	Ss	0:02.93	tnp.chass
3122	0	1380	0	2	0	1764	select	12:33PM	d0	S	0:00.77	./rpd -N
3128	0	215	0	2	0	580	select	12:45PM	??	Ss	0:00.12	rlogind
3129	1049	3128	0	18	0	944	pause	12:45PM	p0	Ss	0:00.14	-tcsh (tc
0	0	0	0	-18	0	0	sched	1:51PM	??	DLs	0:00.10	(swapper

show system processes extensive

```

user@host> show system processes extensive

```

last pid: 544; load averages: 0.00, 0.00, 0.00 18:30:33

37 processes: 1 running, 36 sleeping

Mem: 25M Active, 3968K Inact, 19M Wired, 8346K Buf, 202M Free

Swap: 528M Total, 64K Used, 528M Free

PID	USERNAME	PRI	NICE	SIZE	RES	STATE	TIME	WCPU	CPU	COMMAND
544	root	30	0	604K	768K	RUN	0:00	0.00%	0.00%	top
3	root	28	0	0K	12K	psleep	0:00	0.00%	0.00%	vm Daemon
4	root	28	0	0K	12K	update	0:03	0.00%	0.00%	update
528	aviva	18	0	660K	948K	pause	0:00	0.00%	0.00%	tcsh
204	root	18	0	300K	544K	pause	0:00	0.00%	0.00%	csch
131	root	18	0	332K	532K	pause	0:00	0.00%	0.00%	cron
186	root	18	0	196K	68K	pause	0:00	0.00%	0.00%	watchdog
27	root	10	0	512M	16288K	mfsidl	0:00	0.00%	0.00%	mount_mfs
1	root	10	0	620K	344K	wait	0:00	0.00%	0.00%	init
304	root	3	0	884K	900K	ttyin	0:00	0.00%	0.00%	bash
200	root	3	0	180K	540K	ttyin	0:00	0.00%	0.00%	getty
203	root	3	0	180K	540K	ttyin	0:00	0.00%	0.00%	getty
202	root	3	0	180K	540K	ttyin	0:00	0.00%	0.00%	getty
201	root	3	0	180K	540K	ttyin	0:00	0.00%	0.00%	getty
194	root	2	0	2248K	1640K	select	0:11	0.00%	0.00%	rpd
205	root	2	0	964K	800K	select	0:12	0.00%	0.00%	tnp.chassisd
189	root	2	-12	352K	740K	select	0:03	0.00%	0.00%	xntpd


```

114 root      2   0   296K   612K select   0:00   0.00%   0.00% amd
188 root      2   0   780K   600K select   0:00   0.00%   0.00% dcd
527 root      2   0   176K   580K select   0:00   0.00%   0.00% rlogind
195 root      2   0   212K   552K select   0:00   0.00%   0.00% inetd
187 root      2   0   192K   532K select   0:00   0.00%   0.00% tnetd
 83 root      2   0   188K   520K select   0:00   0.00%   0.00% syslogd
538 root      2   0  1324K   516K select   0:00   0.00%   0.00% mgd
 99 daemon    2   0   176K   492K select   0:00   0.00%   0.00% portmap
163 root      2   0   572K   420K select   0:00   0.00%   0.00% nsrexecd
192 root      2   0   560K   400K select   0:10   0.00%   0.00% snmpd
191 root      2   0  1284K   376K select   0:00   0.00%   0.00% mgd
537 aviva     2   0   636K   364K select   0:00   0.00%   0.00% cli
193 root      2   0   312K   204K select   0:07   0.00%   0.00% mib2d
  5 root      2   0      0K    12K pfesel   0:00   0.00%   0.00% if_pfe
  2 root     -18   0      0K    12K psleep   0:00   0.00%   0.00% pagedaemon
  0 root     -18   0      0K     0K sched    0:00   0.00%   0.00% swapper

```

show system processes lcc wide (TX Matrix Routing Matrix)

```

user@host> show system processes lcc 2 wide
lcc2-re0:

```

```

-----
PID TT  STAT      TIME COMMAND
  0 ??  DLs      0:00.00 (swapper)
  1 ??  ILs      0:00.10 /sbin/preinit -- (init)
  2 ??  DL       0:00.00 (pagedaemon)
  3 ??  DL       0:00.00 (vmdaemon)
  4 ??  DL       0:00.00 (bufdaemon)
  5 ??  DL       0:00.04 (syncer)
  6 ??  DL       0:00.00 (netdaemon)
  7 ??  IL       0:00.00 (if_pic_listen)
  8 ??  IL       0:00.00 (scs_housekeeping)
  9 ??  IL       0:00.00 (if_pfe_listen)
 10 ??  DL       0:00.00 (vmuncachedaemon)
 11 ??  SL       0:00.02 (cb_poll)
 172 ??  ILs      0:00.21 mfs -o noauto /dev/ad1s1b /tmp (newfs)
2909 ??  Is       0:00.00 pccardd
2932 ??  Ss       0:00.07 syslogd -r -s
3039 ??  Is       0:00.00 cron
3217 ??  I        0:00.00 /sbin/watchdog -d
3218 ??  I        0:00.02 /usr/sbin/tnetd -N
3221 ??  S        0:00.11 /usr/sbin/alarmd -N
3222 ??  S        0:00.85 /usr/sbin/craftd -N
3223 ??  S        0:00.05 /usr/sbin/mgd -N
3224 ??  I        0:00.02 /usr/sbin/inetd -N
3225 ??  I        0:00.00 /usr/sbin/tnp.sntpd -N
3226 ??  I        0:00.01 /usr/sbin/tnp.sntpc -N
3228 ??  I        0:00.01 /usr/sbin/smartd -N
3231 ??  I        0:00.01 /usr/sbin/eccd -N
3425 ??  S        0:00.09 /usr/sbin/dfwd -N
3426 ??  S        0:00.19 /sbin/dcd -N
3427 ??  I        0:00.04 /usr/sbin/pfed -N
3430 ??  S        0:00.10 /usr/sbin/ksyncd -N
3482 ??  S        1:53.63 /usr/sbin/chassisd -N
4285 ??  SL       0:00.01 (peer proxy)
4286 ??  SL       0:00.00 (peer proxy)
4303 ??  Ss       0:00.00 mgd: (mgd) (root) (mgd)
4304 ??  R        0:00.00 /bin/ps -ax -ww
3270 d0  Is+      0:00.00 /usr/libexec/getty std.9600 ttyd0

```

show system processes summary

```
user@host> show system processes summary
last pid: 543; load averages: 0.00, 0.00, 0.00 18:29:47
37 processes: 1 running, 36 sleeping
```

```
Mem: 25M Active, 3976K Inact, 19M Wired, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free
```

PID	USERNAME	PRI	NICE	SIZE	RES	STATE	TIME	WCPU	CPU	COMMAND
527	root	2	0	176K	580K	select	0:00	0.04%	0.04%	rlogind
543	root	30	0	604K	768K	RUN	0:00	0.00%	0.00%	top

show system processes (TX Matrix Plus Router)

```
user@host> show system processes
sfc0-re0:
```

```
-----
PID  TT  STAT      TIME COMMAND
 0  ??  Wls      0:00.00 [swapper]
 1  ??  ILs      0:00.18 /packages/mnt/jbase/sbin/init --
 2  ??  DL       0:00.20 [g_event]
 3  ??  DL       0:00.39 [g_up]
 4  ??  DL       0:00.32 [g_down]
 5  ??  DL       0:00.00 [thread taskq]
 6  ??  DL       0:00.09 [kqueue taskq]
 7  ??  DL       0:00.01 [pagedaemon]
 8  ??  DL       0:00.00 [vmdaemon]
 9  ??  DL       0:06.63 [pagezero]
10  ??  DL       0:00.00 [ktrace]
11  ??  RL      310:52.98 [idle]
12  ??  WL       0:11.03 [swi2: net]
13  ??  WL       0:27.58 [swi7: clock sio]
14  ??  WL       0:00.00 [swi6: vm]
15  ??  DL       0:03.02 [yarrow]
16  ??  WL       0:00.00 [swi9: +]
17  ??  WL       0:00.00 [swi8: +]
18  ??  WL       0:00.00 [swi5: cambio]
19  ??  WL       0:00.00 [swi9: task queue]
20  ??  WL       0:11.41 [irq16: uhci0 uhci*]
21  ??  DL       0:00.00 [usb0]
22  ??  DL       0:00.00 [usbtask]
23  ??  WL       0:39.51 [irq17: uhci1 uhci*]
24  ??  DL       0:00.00 [usb1]
25  ??  WL       0:00.00 [irq18: uhci2 uhci*]
26  ??  DL       0:00.83 [usb2]
27  ??  DL       0:00.00 [usb3]
28  ??  DL       0:00.00 [usb4]
29  ??  DL       0:00.00 [usb5]
30  ??  DL       0:00.73 [usb6]
31  ??  DL       0:00.00 [usb7]
32  ??  WL       0:00.00 [irq14: ata0]
33  ??  WL       0:00.00 [irq15: ata1]
34  ??  WL       0:00.00 [irq1: atkbd0]
35  ??  WL       0:00.00 [swi0: sio]
36  ??  WL       0:00.00 [irq11: isab0]
37  ??  WL       0:00.00 [swi3: ip6opt ipopt]
38  ??  WL       0:00.00 [swi4: ip6mismatch+]
39  ??  WL       0:00.00 [swi1: ipfwd]
40  ??  DL       0:00.02 [bufdaemon]
41  ??  DL       0:00.02 [vn1ru]
```

```

42 ?? DL 0:00.39 [syncer]
43 ?? DL 0:00.05 [softdepflush]
44 ?? DL 0:00.00 [netdaemon]
45 ?? DL 0:00.02 [vmuncachedaemon]
46 ?? DL 0:00.00 [if_pic_listen]
47 ?? DL 0:00.35 [vmkmemdaemon]
48 ?? DL 0:00.00 [cb_poll]
49 ?? DL 0:00.06 [if_pfe_listen]
50 ?? DL 0:00.00 [scs_housekeeping]
51 ?? IL 0:00.00 [kern_dump_proc]
52 ?? IL 0:00.00 [nfsiod 0]
53 ?? IL 0:00.00 [nfsiod 1]
54 ?? IL 0:00.00 [nfsiod 2]
55 ?? IL 0:00.00 [nfsiod 3]
56 ?? DL 0:00.37 [schedcpu]
57 ?? DL 0:00.56 [md0]
79 ?? DL 0:02.58 [md1]
100 ?? DL 0:00.03 [md2]
118 ?? DL 0:00.01 [md3]
139 ?? DL 0:00.95 [md4]
160 ?? DL 0:00.12 [md5]
181 ?? DL 0:00.00 [md6]
217 ?? DL 0:00.02 [md7]
227 ?? DL 0:00.05 [md8]
1341 ?? SL 0:01.34 [bcmTX]
1342 ?? SL 0:01.68 [bcmXGS3AsyncTX]
1343 ?? SL 0:41.40 [bcmLINK.0]
1345 ?? SL 0:33.83 [bcmLINK.1]
1350 ?? Is 0:00.01 /usr/sbin/cron
1502 ?? S 0:00.01 /sbin/watchdog -t-1
1503 ?? S 0:00.86 /usr/libexec/bslockd -mp -N
1504 ?? S 0:00.01 /usr/sbin/tnetd -N
1507 ?? S 0:01.32 /usr/sbin/alarmd -N
1508 ?? S 0:14.54 /usr/sbin/craftd -N
1509 ?? S 0:01.19 /usr/sbin/mgd -N
1512 ?? I 0:00.05 /usr/sbin/inetd -N
1513 ?? S 0:00.10 /usr/sbin/tnp.snptd -N
1517 ?? S 0:00.11 /usr/sbin/smartd -N
1525 ?? S 0:01.10 /usr/sbin/idpd -N
1526 ?? S 0:01.43 /usr/sbin/license-check -U -M -p 10 -i 10
1527 ?? I 0:00.01 /usr/libexec/getty Pc ttyv0
1616 ?? DL 0:00.30 [peer proxy]
1617 ?? DL 0:00.32 [peer proxy]
1618 ?? DL 0:00.34 [peer proxy]
1619 ?? DL 0:00.30 [peer proxy]
2391 ?? Is 0:00.01 telnetd
7331 ?? Ss 0:00.03 telnetd
9538 ?? DL 0:01.16 [jsr_kkcm]
9613 ?? DL 0:00.18 [peer proxy]
23781 ?? Ss 0:00.01 telnetd
23926 ?? Ss 0:00.01 mgd: (mgd) (regress)/dev/tty2 (mgd)
36867 ?? S 0:03.14 /usr/sbin/rpd -N
36874 ?? S 0:00.08 /usr/sbin/lmpd
36876 ?? S 0:00.17 /usr/sbin/lacpd -N
36877 ?? S 0:00.15 /usr/sbin/bfdd -N
36878 ?? S 0:05.05 /usr/sbin/ppmd -N
36907 ?? S 0:25.07 /usr/sbin/chassisd -N
37775 ?? S 0:00.01 /usr/sbin/bdbrepd -N
45727 ?? S 0:00.02 /usr/sbin/xntpd -j -N -g (ntpd)
45729 ?? S 0:00.38 /usr/sbin/l2ald -N
45730 ?? S< 0:00.12 /usr/sbin/apsd -N

```

```

45731 ?? SN      0:00.10 /usr/sbin/sampled -N
45732 ?? S       0:00.03 /usr/sbin/ilmid -N
45733 ?? S       0:00.09 /usr/sbin/rmopd -N
45734 ?? S       0:00.30 /usr/sbin/cosd
45735 ?? I       0:00.00 /usr/sbin/rtspd -N
45736 ?? S       0:00.06 /usr/sbin/fsad -N
45737 ?? S       0:00.05 /usr/sbin/rdd -N
45738 ?? S       0:00.10 /usr/sbin/pppd -N
45739 ?? S       0:00.05 /usr/sbin/dfcd -N
45740 ?? S       0:00.07 /usr/sbin/lfmd -N
45741 ?? S       0:00.01 /usr/sbin/implsoamd -N
45742 ?? I       0:00.01 /usr/sbin/sendd -N
45743 ?? S       0:00.08 /usr/sbin/appidd -N
45744 ?? S       0:00.05 /usr/sbin/mspd -N
45745 ?? S       0:00.25 /usr/sbin/jdiameterd -N
45746 ?? S       0:00.10 /usr/sbin/pfed -N
45747 ?? S       0:00.19 /usr/sbin/lpdfd -N
45748 ?? S       0:00.63 /sbin/dcd -N
45750 ?? S       0:00.45 /usr/sbin/mib2d -N
45751 ?? S       0:00.15 /usr/sbin/dfwd -N
45752 ?? S       0:00.15 /usr/sbin/irsd -N
45764 ?? S       0:20.59 /usr/sbin/snmpd -N
56479 ?? Ss      0:00.00 mgd: (mgd) (root) (mgd)
56480 ?? R       0:00.00 /bin/ps -ax
1142 d0- I       0:00.01 /usr/sbin/usbd -N
1160 d0- S       0:29.17 /usr/sbin/eventd -N -r -s -A
6527 d0 Is+      0:00.00 /usr/libexec/getty std.9600 ttyd0
2392 p1 Is       0:00.00 login [pam] (login)
2393 p1 I        0:00.00 -csh (csh)
2394 p1 I        0:00.00 su -
2395 p1 I+       0:00.01 -su (csh)
23782 p2 Is       0:00.00 login [pam] (login)
23881 p2 I        0:00.00 -csh (csh)
23925 p2 S+      0:00.03 cli
7332 p3 Is       0:00.00 login [pam] (login)
7333 p3 I        0:00.00 -csh (csh)
23780 p3 S+      0:00.02 telnet aj

```

lcc0-re0:

```

-----
PID  TT  STAT      TIME COMMAND
  0  ??  WLS      0:00.00 [swapper]
  1  ??  ILs      0:00.16 /packages/mnt/jbase/sbin/init --
  2  ??  DL       0:00.01 [g_event]
  3  ??  DL       0:00.16 [g_up]
  4  ??  DL       0:00.11 [g_down]
  5  ??  DL       0:00.00 [thread taskq]
  6  ??  DL       0:00.00 [kqueue taskq]
  7  ??  DL       0:00.00 [pagedaemon]
  8  ??  DL       0:00.00 [vmdaemon]
  9  ??  DL       0:01.77 [pagezero]
 10  ??  DL       0:00.00 [ktrace]
 11  ??  RL      17:22.31 [idle]
 12  ??  WL       0:00.32 [swi2: net]
 13  ??  WL       0:01.21 [swi7: clock sio]
 14  ??  WL       0:00.00 [swi6: vm]
 15  ??  DL       0:00.10 [yarrow]
 16  ??  WL       0:00.00 [swi9: +]
 17  ??  WL       0:00.00 [swi8: +]
 18  ??  WL       0:00.00 [swi5: cambio]
 19  ??  WL       0:00.00 [swi9: task queue]

```

```

20 ?? WL 0:02.73 [irq10: bcm0 uhci1*]
21 ?? WL 0:00.02 [irq11: cb0 uhci0+*]
22 ?? DL 0:00.00 [usb0]
23 ?? DL 0:00.00 [usbtask]
24 ?? DL 0:00.00 [usb1]
25 ?? DL 0:00.05 [usb2]
26 ?? DL 0:00.00 [usb3]
27 ?? DL 0:00.00 [usb4]
28 ?? DL 0:00.00 [usb5]
29 ?? DL 0:00.04 [usb6]
30 ?? DL 0:00.00 [usb7]
31 ?? WL 0:00.00 [irq14: ata0]
32 ?? WL 0:00.00 [irq15: ata1]
33 ?? WL 0:00.00 [irq1: atkbd0]
34 ?? WL 0:00.00 [swi0: sio]
35 ?? WL 0:00.00 [swi3: ip6opt ipopt]
36 ?? WL 0:00.00 [swi4: ip6mismatch+]
37 ?? WL 0:00.00 [swi1: ipfwd]
38 ?? DL 0:00.00 [bufdaemon]
39 ?? DL 0:00.00 [vnlru]
40 ?? DL 0:00.01 [syncer]
41 ?? DL 0:00.00 [softdepflush]
42 ?? DL 0:00.00 [netdaemon]
43 ?? DL 0:00.00 [vmuncachedaemon]
44 ?? DL 0:00.00 [if_pic_listen]
45 ?? DL 0:00.02 [vmkmemdaemon]
46 ?? DL 0:00.01 [cb_poll]
47 ?? DL 0:00.00 [if_pfe_listen]
48 ?? DL 0:00.00 [scs_housekeeping]
49 ?? IL 0:00.00 [kern_dump_proc]
50 ?? IL 0:00.00 [nfsiod 0]
51 ?? IL 0:00.00 [nfsiod 1]
52 ?? IL 0:00.00 [nfsiod 2]
53 ?? IL 0:00.00 [nfsiod 3]
54 ?? DL 0:00.01 [schedcpu]
55 ?? DL 0:00.73 [md0]
77 ?? DL 0:03.54 [md1]
98 ?? DL 0:00.37 [md2]
116 ?? DL 0:00.02 [md3]
137 ?? DL 0:00.56 [md4]
158 ?? DL 0:00.15 [md5]
179 ?? DL 0:00.00 [md6]
215 ?? DL 0:00.03 [md7]
225 ?? DL 0:00.03 [md8]
1078 ?? DL 0:00.00 [jsr_kkcm]
1363 ?? SL 0:00.09 [bcmTX]
1364 ?? SL 0:00.10 [bcmXGS3AsyncTX]
1365 ?? SL 0:03.08 [bcmLINK.0]
1370 ?? Is 0:00.00 /usr/sbin/cron
1522 ?? S 0:00.00 /sbin/watchdog -t-1
1523 ?? S 0:00.05 /usr/libexec/bslockd -mp -N
1524 ?? I 0:00.01 /usr/sbin/tnetd -N
1526 ?? S 0:04.98 /usr/sbin/chassisd -N
1527 ?? S 0:00.04 /usr/sbin/alarmd -N
1528 ?? I 0:00.40 /usr/sbin/craftd -N
1529 ?? S 0:00.08 /usr/sbin/mgd -N
1532 ?? I 0:00.04 /usr/sbin/inetd -N
1533 ?? I 0:00.00 /usr/sbin/tnp.sntpd -N
1534 ?? I 0:00.00 /usr/sbin/tnp.sntpc -N
1536 ?? S 0:00.01 /usr/sbin/smartd -N
1540 ?? I 0:00.07 /usr/sbin/jcsd -N

```

```

1541 ?? S      0:00.11 /usr/sbin/idpd -N
1542 ?? I      0:00.00 /usr/libexec/getty Pc ttyv0
2089 ?? DL    0:00.01 [peer proxy]
2090 ?? DL    0:00.01 [peer proxy]
2091 ?? DL    0:00.01 [peer proxy]
2657 ?? S      0:00.02 /usr/sbin/dfwd -N
2658 ?? S      0:00.02 /sbin/dcd -N
2659 ?? S      0:00.05 /usr/sbin/snmpd -N
2660 ?? S      0:00.01 /usr/sbin/mib2d -N
2661 ?? S      0:00.01 /usr/sbin/pfed -N
2662 ?? S      0:00.01 /usr/sbin/irsd -N
2667 ?? S      0:00.13 /usr/sbin/ksyncd -N
2690 ?? Ss    0:00.00 mgd: (mgd) (root) (mgd)
2691 ?? R      0:00.00 /bin/ps -ax
1164 d0- S     0:00.00 /usr/sbin/usbd -N
1182 d0- S     0:00.34 /usr/sbin/eventd -N -r -s -A
1543 d0 Is+   0:00.00 /usr/libexec/getty std.9600 ttyd0

```

lcc1-re0:

```

-----
PID  TT  STAT      TIME COMMAND
  0  ??  Wls    0:00.00 [swapper]
  1  ??  ILs    0:00.17 /packages/mnt/jbase/sbin/init --
  2  ??  DL     0:00.01 [g_event]
  3  ??  DL     0:00.16 [g_up]
  4  ??  DL     0:00.11 [g_down]
  5  ??  DL     0:00.00 [thread taskq]
  6  ??  DL     0:00.00 [kqueue taskq]
  7  ??  DL     0:00.00 [pagedaemon]
  8  ??  DL     0:00.00 [vmdaemon]
  9  ??  DL     0:01.77 [pagezero]
 10  ??  DL     0:00.00 [ktrace]
 11  ??  RL    17:22.83 [idle]
 12  ??  WL     0:00.35 [swi2: net]
 13  ??  WL     0:01.20 [swi7: clock sio]
 14  ??  WL     0:00.00 [swi6: vm]
 15  ??  DL     0:00.10 [yarrow]
 16  ??  WL     0:00.00 [swi9: +]
 17  ??  WL     0:00.00 [swi8: +]
 18  ??  WL     0:00.00 [swi5: cambio]
 19  ??  WL     0:00.00 [swi9: task queue]
 20  ??  WL     0:02.87 [irq10: bcm0 uhci1*]
 21  ??  WL     0:00.02 [irq11: cb0 uhci0+*]
 22  ??  DL     0:00.00 [usb0]
 23  ??  DL     0:00.00 [usbtask]
 24  ??  DL     0:00.00 [usb1]
 25  ??  DL     0:00.05 [usb2]
 26  ??  DL     0:00.00 [usb3]
 27  ??  DL     0:00.00 [usb4]
 28  ??  DL     0:00.00 [usb5]
 29  ??  DL     0:00.04 [usb6]
 30  ??  DL     0:00.00 [usb7]
 31  ??  WL     0:00.00 [irq14: ata0]
 32  ??  WL     0:00.00 [irq15: ata1]
 33  ??  WL     0:00.00 [irq1: atkbd0]
 34  ??  WL     0:00.00 [swi0: sio]
 35  ??  WL     0:00.00 [swi3: ip6opt ipopt]
 36  ??  WL     0:00.00 [swi4: ip6mismatch+]
 37  ??  WL     0:00.00 [swi1: ipfwd]
 38  ??  DL     0:00.00 [bufdaemon]
 39  ??  DL     0:00.00 [vn1ru]

```

```

40 ?? DL 0:00.01 [syncer]
41 ?? DL 0:00.00 [softdepflush]
42 ?? DL 0:00.00 [netdaemon]
43 ?? DL 0:00.00 [vmuncachedaemon]
44 ?? DL 0:00.00 [if_pic_listen]
45 ?? DL 0:00.02 [vmkmemdaemon]
46 ?? DL 0:00.01 [cb_poll]
47 ?? DL 0:00.00 [if_pfe_listen]
48 ?? DL 0:00.00 [scs_housekeeping]
49 ?? IL 0:00.00 [kern_dump_proc]
50 ?? IL 0:00.00 [nfsiod 0]
51 ?? IL 0:00.00 [nfsiod 1]
52 ?? IL 0:00.00 [nfsiod 2]
53 ?? IL 0:00.00 [nfsiod 3]
54 ?? DL 0:00.02 [schedcpu]
55 ?? DL 0:00.75 [md0]
77 ?? DL 0:03.40 [md1]
98 ?? DL 0:00.37 [md2]
116 ?? DL 0:00.02 [md3]
137 ?? DL 0:00.56 [md4]
158 ?? DL 0:00.15 [md5]
179 ?? DL 0:00.00 [md6]
215 ?? DL 0:00.03 [md7]
225 ?? DL 0:00.03 [md8]
1052 ?? DL 0:00.00 [jsr_kkcm]
1337 ?? SL 0:00.09 [bcmTX]
1338 ?? SL 0:00.10 [bcmXGS3AsyncTX]
1339 ?? SL 0:03.10 [bcmLINK.0]
1344 ?? Is 0:00.00 /usr/sbin/cron
1496 ?? S 0:00.00 /sbin/watchdog -t-1
1497 ?? S 0:00.05 /usr/libexec/bslockd -mp -N
1498 ?? I 0:00.01 /usr/sbin/tnetd -N
1500 ?? S 0:04.97 /usr/sbin/chassisd -N
1501 ?? S 0:00.04 /usr/sbin/alarmd -N
1502 ?? I 0:00.40 /usr/sbin/craftd -N
1503 ?? S 0:00.08 /usr/sbin/mgd -N
1506 ?? I 0:00.04 /usr/sbin/inetd -N
1507 ?? I 0:00.00 /usr/sbin/tnp.snmpd -N
1508 ?? I 0:00.00 /usr/sbin/tnp.sntpc -N
1510 ?? S 0:00.01 /usr/sbin/smartd -N
1514 ?? I 0:00.07 /usr/sbin/jcsd -N
1515 ?? S 0:00.18 /usr/sbin/idpd -N
1516 ?? I 0:00.00 /usr/libexec/getty Pc ttyv0
2068 ?? DL 0:00.01 [peer proxy]
2069 ?? DL 0:00.01 [peer proxy]
2070 ?? DL 0:00.01 [peer proxy]
2666 ?? S 0:00.02 /sbin/dcd -N
2667 ?? S 0:00.01 /usr/sbin/irsd -N
2668 ?? S 0:00.01 /usr/sbin/pfed -N
2669 ?? S 0:00.05 /usr/sbin/snmpd -N
2670 ?? S 0:00.01 /usr/sbin/mib2d -N
2671 ?? S 0:00.02 /usr/sbin/dfwd -N
2675 ?? S 0:00.13 /usr/sbin/ksyncd -N
2699 ?? Ss 0:00.00 mgd: (mgd) (root) (mgd)
2700 ?? R 0:00.00 /bin/ps -ax
1138 d0- S 0:00.00 /usr/sbin/usbd -N
1156 d0- S 0:00.37 /usr/sbin/eventd -N -r -s -A
1517 d0 Is+ 0:00.00 /usr/libexec/getty std.9600 ttyd0

```

```
lcc2-re0:
```

PID	TT	STAT	TIME	COMMAND
0	??	Wls	0:00.00	[swapper]
1	??	ILs	0:00.18	/packages/mnt/jbase/sbin/init --
2	??	DL	0:00.01	[g_event]
3	??	DL	0:00.17	[g_up]
4	??	DL	0:00.12	[g_down]
5	??	DL	0:00.00	[thread taskq]
6	??	DL	0:00.00	[kqueue taskq]
7	??	DL	0:00.00	[pagedaemon]
8	??	DL	0:00.00	[vmdaemon]
9	??	DL	0:01.77	[pagezero]
10	??	DL	0:00.00	[ktrace]
11	??	RL	17:19.13	[idle]
12	??	WL	0:00.36	[swi2: net]
13	??	WL	0:01.20	[swi7: clock sio]
14	??	WL	0:00.00	[swi6: vm]
15	??	DL	0:00.13	[yarrow]
16	??	WL	0:00.00	[swi9: +]
17	??	WL	0:00.00	[swi8: +]
18	??	WL	0:00.00	[swi5: cambio]
19	??	WL	0:00.00	[swi9: task queue]
20	??	WL	0:03.03	[irq10: bcm0 uhci1*]
21	??	WL	0:00.02	[irq11: cb0 uhci0+*]
22	??	DL	0:00.00	[usb0]
23	??	DL	0:00.00	[usbtask]
24	??	DL	0:00.00	[usb1]
25	??	DL	0:00.05	[usb2]
26	??	DL	0:00.00	[usb3]
27	??	DL	0:00.00	[usb4]
28	??	DL	0:00.00	[usb5]
29	??	DL	0:00.04	[usb6]
30	??	DL	0:00.00	[usb7]
31	??	WL	0:00.00	[irq14: ata0]
32	??	WL	0:00.00	[irq15: ata1]
33	??	WL	0:00.00	[irq1: atkbd0]
34	??	WL	0:00.00	[swi0: sio]
35	??	WL	0:00.00	[swi3: ip6opt ipopt]
36	??	WL	0:00.00	[swi4: ip6mismatch+]
37	??	WL	0:00.00	[swi1: ipfwd]
38	??	DL	0:00.00	[bufdaemon]
39	??	DL	0:00.00	[vn1ru]
40	??	DL	0:00.01	[syncer]
41	??	DL	0:00.00	[softdepflush]
42	??	DL	0:00.00	[netdaemon]
43	??	DL	0:00.00	[vmuncachedaemon]
44	??	DL	0:00.00	[if_pic_listen]
45	??	DL	0:00.02	[vmkmemdaemon]
46	??	DL	0:00.01	[cb_poll]
47	??	DL	0:00.00	[if_pfe_listen]
48	??	DL	0:00.00	[scs_housekeeping]
49	??	IL	0:00.00	[kern_dump_proc]
50	??	IL	0:00.00	[nfsiod 0]
51	??	IL	0:00.00	[nfsiod 1]
52	??	IL	0:00.00	[nfsiod 2]
53	??	IL	0:00.00	[nfsiod 3]
54	??	DL	0:00.02	[schedcpu]
55	??	DL	0:00.75	[md0]
77	??	DL	0:03.48	[md1]
98	??	DL	0:00.59	[md2]
116	??	DL	0:00.02	[md3]
137	??	DL	0:00.56	[md4]


```

158 ?? DL 0:00.15 [md5]
179 ?? DL 0:00.00 [md6]
215 ?? DL 0:00.03 [md7]
225 ?? DL 0:00.03 [md8]
1052 ?? DL 0:00.00 [jsr_kkcm]
1337 ?? SL 0:00.09 [bcmTX]
1338 ?? SL 0:00.10 [bcmXGS3AsyncTX]
1339 ?? SL 0:03.22 [bcmLINK.0]
1344 ?? Is 0:00.00 /usr/sbin/cron
1496 ?? S 0:00.00 /sbin/watchdog -t-1
1497 ?? S 0:00.05 /usr/libexec/bslockd -mp -N
1498 ?? S 0:00.01 /usr/sbin/tnetd -N
1500 ?? R 0:05.17 /usr/sbin/chassisd -N
1501 ?? S 0:00.04 /usr/sbin/alarmd -N
1502 ?? I 0:00.39 /usr/sbin/craftd -N
1503 ?? S 0:00.08 /usr/sbin/mgd -N
1506 ?? I 0:00.05 /usr/sbin/inetd -N
1507 ?? I 0:00.00 /usr/sbin/tnp.sntpd -N
1508 ?? I 0:00.00 /usr/sbin/tnp.sntpc -N
1510 ?? S 0:00.01 /usr/sbin/smardd -N
1514 ?? I 0:00.07 /usr/sbin/jcsd -N
1515 ?? S 0:00.17 /usr/sbin/idpd -N
1516 ?? I 0:00.00 /usr/libexec/getty Pc ttyv0
2591 ?? DL 0:00.01 [peer proxy]
2592 ?? DL 0:00.01 [peer proxy]
2593 ?? DL 0:00.01 [peer proxy]
2597 ?? DL 0:00.00 [peer proxy]
3192 ?? S 0:00.01 /usr/sbin/irsd -N
3193 ?? S 0:00.05 /usr/sbin/snmpd -N
3194 ?? S 0:00.02 /sbin/dcd -N
3195 ?? S 0:00.01 /usr/sbin/pfed -N
3196 ?? S 0:00.01 /usr/sbin/mib2d -N
3197 ?? S 0:00.02 /usr/sbin/dfwd -N
3198 ?? S 0:00.13 /usr/sbin/ksyncd -N
3228 ?? Ss 0:00.00 mgd: (mgd) (root) (mgd)
3229 ?? R 0:00.00 /bin/ps -ax
1138 d0- S 0:00.00 /usr/sbin/usbd -N
1156 d0- S 0:00.42 /usr/sbin/eventd -N -r -s -A
1517 d0 Is+ 0:00.00 /usr/libexec/getty std.9600 ttyd0
...

```

show system processes sfc (TX Matrix Plus Router)

```

user@host> show system processes sfc 0
sfc0-re0:

```

```

-----
PID TT STAT TIME COMMAND
0 ?? Wls 0:00.00 [swapper]
1 ?? SLs 0:00.18 /packages/mnt/jbase/sbin/init --
2 ?? DL 0:00.20 [g_event]
3 ?? DL 0:00.39 [g_up]
4 ?? DL 0:00.32 [g_down]
5 ?? DL 0:00.00 [thread taskq]
6 ?? DL 0:00.09 [kqueue taskq]
7 ?? DL 0:00.01 [pagedaemon]
8 ?? DL 0:00.00 [vmdaemon]
9 ?? DL 0:06.63 [pagezero]
10 ?? DL 0:00.00 [ktrace]
11 ?? RL 312:09.00 [idle]
12 ?? WL 0:11.07 [swi2: net]
13 ?? WL 0:27.70 [swi7: clock sio]

```

```

14 ?? WL 0:00.00 [swi6: vm]
15 ?? DL 0:03.03 [yarrow]
16 ?? WL 0:00.00 [swi9: +]
17 ?? WL 0:00.00 [swi8: +]
18 ?? WL 0:00.00 [swi5: cambio]
19 ?? WL 0:00.00 [swi9: task queue]
20 ?? WL 0:11.46 [irq16: uhci0 uhci*]
21 ?? DL 0:00.00 [usb0]
22 ?? DL 0:00.00 [usbtask]
23 ?? WL 0:39.63 [irq17: uhci1 uhci*]
24 ?? DL 0:00.00 [usb1]
25 ?? WL 0:00.00 [irq18: uhci2 uhci*]
26 ?? DL 0:00.84 [usb2]
27 ?? DL 0:00.00 [usb3]
28 ?? DL 0:00.00 [usb4]
29 ?? DL 0:00.00 [usb5]
30 ?? DL 0:00.73 [usb6]
31 ?? DL 0:00.00 [usb7]
32 ?? WL 0:00.00 [irq14: ata0]
33 ?? WL 0:00.00 [irq15: ata1]
34 ?? WL 0:00.00 [irq1: atkbd0]
35 ?? WL 0:00.00 [swi0: sio]
36 ?? WL 0:00.00 [irq11: isab0]
37 ?? WL 0:00.00 [swi3: ip6opt ipopt]
38 ?? WL 0:00.00 [swi4: ip6mismatch+]
39 ?? WL 0:00.00 [swi1: ipfwd]
40 ?? DL 0:00.02 [bufdaemon]
41 ?? DL 0:00.02 [vn1ru]
42 ?? DL 0:00.39 [syncer]
43 ?? DL 0:00.05 [softdepflush]
44 ?? DL 0:00.00 [netdaemon]
45 ?? DL 0:00.02 [vmuncachedaemon]
46 ?? DL 0:00.00 [if_pic_listen]
47 ?? DL 0:00.35 [vmkmemdaemon]
48 ?? DL 0:00.00 [cb_poll]
49 ?? DL 0:00.06 [if_pfe_listen]
50 ?? DL 0:00.00 [scs_housekeeping]
51 ?? IL 0:00.00 [kern_dump_proc]
52 ?? IL 0:00.00 [nfsiod 0]
53 ?? IL 0:00.00 [nfsiod 1]
54 ?? IL 0:00.00 [nfsiod 2]
55 ?? IL 0:00.00 [nfsiod 3]
56 ?? DL 0:00.37 [schedcpu]
57 ?? DL 0:00.56 [md0]
79 ?? DL 0:02.58 [md1]
100 ?? DL 0:00.03 [md2]
118 ?? DL 0:00.01 [md3]
139 ?? DL 0:00.95 [md4]
160 ?? DL 0:00.12 [md5]
181 ?? DL 0:00.00 [md6]
217 ?? DL 0:00.02 [md7]
227 ?? DL 0:00.05 [md8]
1341 ?? SL 0:01.35 [bcmTX]
1342 ?? SL 0:01.69 [bcmXGS3AsyncTX]
1343 ?? SL 0:41.57 [bcmLINK.0]
1345 ?? SL 0:33.97 [bcmLINK.1]
1350 ?? Is 0:00.01 /usr/sbin/cron
1502 ?? S 0:00.01 /sbin/watchdog -t-1
1503 ?? S 0:00.86 /usr/libexec/bslockd -mp -N
1504 ?? I 0:00.01 /usr/sbin/tnetd -N
1507 ?? S 0:01.32 /usr/sbin/alarmd -N

```

```

1508 ?? S      0:14.54 /usr/sbin/craftd -N
1509 ?? S      0:01.20 /usr/sbin/mgd -N
1512 ?? S      0:00.05 /usr/sbin/inetd -N
1513 ?? S      0:00.10 /usr/sbin/tnp.sntpd -N
1517 ?? S      0:00.11 /usr/sbin/smartd -N
1525 ?? S      0:01.11 /usr/sbin/idpd -N
1526 ?? S      0:01.43 /usr/sbin/license-check -U -M -p 10 -i 10
1527 ?? I      0:00.01 /usr/libexec/getty Pc ttyv0
1616 ?? DL     0:00.30 [peer proxy]
1617 ?? DL     0:00.32 [peer proxy]
1618 ?? DL     0:00.34 [peer proxy]
1619 ?? DL     0:00.30 [peer proxy]
2391 ?? Is     0:00.01 telnetd
7331 ?? Ss     0:00.03 telnetd
9538 ?? DL     0:01.16 [jsr_kkcm]
9613 ?? DL     0:00.18 [peer proxy]
23781 ?? Ss     0:00.01 telnetd
23926 ?? Ss     0:00.03 mgd: (mgd) (regress)/dev/tty2 (mgd)
36867 ?? S      0:03.14 /usr/sbin/rpd -N
36874 ?? S      0:00.08 /usr/sbin/lmpd
36876 ?? S      0:00.17 /usr/sbin/lacpd -N
36877 ?? S      0:00.15 /usr/sbin/bfdd -N
36878 ?? S      0:05.05 /usr/sbin/ppmd -N
36907 ?? S      0:26.63 /usr/sbin/chassisd -N
37775 ?? S      0:00.01 /usr/sbin/bdbrepd -N
45727 ?? S      0:00.02 /usr/sbin/xntpd -j -N -g (ntpd)
45729 ?? S      0:00.40 /usr/sbin/l2ald -N
45730 ?? S<     0:00.13 /usr/sbin/apsd -N
45731 ?? SN     0:00.10 /usr/sbin/sampled -N
45732 ?? S      0:00.03 /usr/sbin/ilmid -N
45733 ?? S      0:00.09 /usr/sbin/rmopd -N
45734 ?? S      0:00.31 /usr/sbin/cosd
45735 ?? I      0:00.00 /usr/sbin/rtspd -N
45736 ?? S      0:00.06 /usr/sbin/fsad -N
45737 ?? S      0:00.05 /usr/sbin/rdd -N
45738 ?? S      0:00.10 /usr/sbin/pppd -N
45739 ?? S      0:00.05 /usr/sbin/dfcd -N
45740 ?? S      0:00.08 /usr/sbin/lfmd -N
45741 ?? S      0:00.01 /usr/sbin/mpiisoamd -N
45742 ?? I      0:00.01 /usr/sbin/sendd -N
45743 ?? S      0:00.08 /usr/sbin/appidd -N
45744 ?? S      0:00.05 /usr/sbin/mspd -N
45745 ?? S      0:00.27 /usr/sbin/jdiameterd -N
45746 ?? S      0:00.10 /usr/sbin/pfed -N
45747 ?? S      0:00.19 /usr/sbin/lpdfd -N
45748 ?? S      0:00.64 /sbin/dcd -N
45750 ?? S      0:00.46 /usr/sbin/mib2d -N
45751 ?? S      0:00.16 /usr/sbin/dfwd -N
45752 ?? S      0:00.15 /usr/sbin/irsd -N
45764 ?? S      0:20.60 /usr/sbin/snmpd -N
56481 ?? Ss     0:00.02 telnetd
56548 ?? Rs     0:00.19 mgd: (mgd) (regress)/dev/tty0 (mgd)
56577 ?? Ss     0:00.00 mgd: (mgd) (root) (mgd)
56578 ?? R      0:00.00 /bin/ps -ax
1142 d0- S      0:00.01 /usr/sbin/usbd -N
1160 d0- S      0:29.71 /usr/sbin/eventd -N -r -s -A
6527 d0 Is+    0:00.00 /usr/libexec/getty std.9600 ttyd0
56482 p0 Is     0:00.00 login [pam] (login)
56483 p0 S      0:00.01 -csh (csh)
56547 p0 S+     0:00.02 cli
2392 p1 Is     0:00.00 login [pam] (login)

```

```

2393 p1 I      0:00.00 -csh (csh)
2394 p1 I      0:00.00 su -
2395 p1 I+     0:00.01 -su (csh)
23782 p2 Is    0:00.00 login [pam] (login)
23881 p2 I      0:00.00 -csh (csh)
23925 p2 S+    0:00.03 cli
7332 p3 Is    0:00.00 login [pam] (login)
7333 p3 I      0:00.00 -csh (csh)
23780 p3 S+    0:00.02 telnet aj

```

show system processes lcc wide (TX Matrix Plus Routing Matrix)

```

user@host> show system processes lcc 2 wide
lcc2-re0:

```

```

-----
PID  TT  STAT    TIME PROVIDER  COMMAND
0   ??  WLS    0:00.00 (null)    [swapper]
1   ??  ILs    0:00.19    /packages/mnt/jbase/sbin/init --
2   ??  DL     0:00.02    [g_event]
3   ??  DL     0:00.19    [g_up]
4   ??  DL     0:00.13    [g_down]
5   ??  DL     0:00.00    [thread taskq]
6   ??  DL     0:00.00    [kqueue taskq]
7   ??  DL     0:00.00    [pagedaemon]
8   ??  DL     0:00.00    [vmdaemon]
9   ??  DL     0:01.77    [pagezero]
10  ??  DL     0:00.00    [ktrace]
11  ??  RL     20:33.81    [idle]
12  ??  WL     0:00.38    [swi2: net]
13  ??  WL     0:01.43    [swi7: clock sio]
14  ??  WL     0:00.00    [swi6: vm]
15  ??  DL     0:00.14    [yarrow]
16  ??  WL     0:00.00    [swi9: +]
17  ??  WL     0:00.00    [swi8: +]
18  ??  WL     0:00.00    [swi5: cambio]
19  ??  WL     0:00.00    [swi9: task queue]
20  ??  WL     0:03.18    [irq10: bcm0 uhci1*]
21  ??  WL     0:00.03    [irq11: cb0 uhci0+*]
22  ??  DL     0:00.00    [usb0]
23  ??  DL     0:00.00    [usbtask]
24  ??  DL     0:00.00    [usb1]
25  ??  DL     0:00.06    [usb2]
26  ??  DL     0:00.00    [usb3]
27  ??  DL     0:00.00    [usb4]
28  ??  DL     0:00.00    [usb5]
29  ??  DL     0:00.05    [usb6]
30  ??  DL     0:00.00    [usb7]
31  ??  WL     0:00.00    [irq14: ata0]
32  ??  WL     0:00.00    [irq15: ata1]
33  ??  WL     0:00.00    [irq1: atkbd0]
34  ??  WL     0:00.00    [swi0: sio]
35  ??  WL     0:00.00    [swi3: ip6opt ipopt]
36  ??  WL     0:00.00    [swi4: ip6mismatch+]
37  ??  WL     0:00.00    [swi1: ipfwd]
38  ??  DL     0:00.00    [bufdaemon]
39  ??  DL     0:00.00    [vn1ru]
40  ??  DL     0:00.02    [syncer]
41  ??  DL     0:00.01    [softdepflush]
42  ??  DL     0:00.00    [netdaemon]
43  ??  DL     0:00.00    [vmuncachedaemon]
44  ??  DL     0:00.00    [if_pic_listen]

```

```

45 ?? DL 0:00.03 [vmkmemdaemon]
46 ?? DL 0:00.01 [cb_poll]
47 ?? DL 0:00.00 [if_pfe_listen]
48 ?? DL 0:00.00 [scs_housekeeping]
49 ?? IL 0:00.00 [kern_dump_proc]
50 ?? IL 0:00.00 [nfsiod 0]
51 ?? IL 0:00.00 [nfsiod 1]
52 ?? IL 0:00.00 [nfsiod 2]
53 ?? IL 0:00.00 [nfsiod 3]
54 ?? DL 0:00.02 [schedcpu]
55 ?? DL 0:00.75 [md0]
77 ?? DL 0:03.84 [md1]
98 ?? DL 0:00.59 [md2]
116 ?? DL 0:00.02 [md3]
137 ?? DL 0:00.72 [md4]
158 ?? DL 0:00.15 [md5]
179 ?? DL 0:00.00 [md6]
215 ?? DL 0:00.03 [md7]
225 ?? DL 0:00.03 [md8]
1052 ?? DL 0:00.00 [jsr_kkcm]
1337 ?? SL 0:00.11 [bcmTX]
1338 ?? SL 0:00.12 [bcmXGS3AsyncTX]
1339 ?? SL 0:03.82 [bcmLINK.0]
1344 ?? Is 0:00.00 /usr/sbin/cron
1496 ?? I 0:00.00 /sbin/watchdog -t-1
1497 ?? S 0:00.06 /usr/libexec/bslockd -mp -N
1498 ?? I 0:00.01 /usr/sbin/tnetd -N
1500 ?? S 0:09.93 /usr/sbin/chassisd -N
1501 ?? S 0:00.05 /usr/sbin/alarmd -N
1502 ?? I 0:00.39 /usr/sbin/craftd -N
1503 ?? S 0:00.09 /usr/sbin/mgd -N
1506 ?? I 0:00.05 /usr/sbin/inetd -N
1507 ?? I 0:00.00 /usr/sbin/tnp.sntpd -N
1508 ?? I 0:00.00 /usr/sbin/tnp.sntpc -N
1510 ?? S 0:00.01 /usr/sbin/smartd -N
1514 ?? I 0:00.07 /usr/sbin/jcsd -N
1515 ?? S 0:00.17 /usr/sbin/idpd -N
1516 ?? I 0:00.00 /usr/libexec/getty Pc ttyv0
2591 ?? DL 0:00.01 [peer proxy]
2592 ?? DL 0:00.01 [peer proxy]
2593 ?? DL 0:00.01 [peer proxy]
2597 ?? DL 0:00.01 [peer proxy]
3192 ?? S 0:00.02 /usr/sbin/irsd -N
3193 ?? S 0:00.05 /usr/sbin/snmpd -N
3194 ?? S 0:00.04 /sbin/dcd -N
3195 ?? I 0:00.01 /usr/sbin/pfed -N
3196 ?? S 0:00.02 /usr/sbin/mib2d -N
3197 ?? I 0:00.03 /usr/sbin/dfwd -N
3198 ?? S 0:00.15 /usr/sbin/ksyncd -N
3559 ?? Ss 0:00.00 mgd: (mgd) (root) (mgd)
3560 ?? R 0:00.00 /bin/ps -ax -jpw
1138 d0- S 0:00.00 /usr/sbin/usbd -N
1156 d0- S 0:00.50 /usr/sbin/eventd -N -r -s -A
1517 d0 Is+ 0:00.00 /usr/libexec/getty std.9600 ttyd0

```

show system processes (QFX Series)

```

user@switch> show system processes
PID TT STAT TIME COMMAND
0 ?? Wls -2341043:-31.01 [swapper]
1 ?? SLs 0:01.34 /packages/mnt/jbase/sbin/init --

```

```

 2 ?? DL      2:48.31 [g_event]
 3 ?? DL      1:47.44 [g_up]
 4 ?? DL      1:37.82 [g_down]
 5 ?? DL      0:00.00 [kdm_tcp_poller]
 6 ?? DL      0:00.00 [thread taskq]
 7 ?? DL      0:04.86 [kqueue taskq]
 9 ?? DL      0:03.94 [pagedaemon]
10 ?? DL      0:00.00 [ktrace]
11 ?? RL      0:00.00 [idle: cpu31]
12 ?? RL      0:00.00 [idle: cpu30]
13 ?? RL      0:00.00 [idle: cpu29]
14 ?? RL      0:00.00 [idle: cpu28]
15 ?? RL      0:00.00 [idle: cpu27]
16 ?? RL      0:00.00 [idle: cpu26]
17 ?? RL      0:00.00 [idle: cpu25]
18 ?? RL      0:00.00 [idle: cpu24]
19 ?? RL      0:00.00 [idle: cpu23]
20 ?? RL      0:00.00 [idle: cpu22]
21 ?? RL      0:00.00 [idle: cpu21]
22 ?? RL      0:00.00 [idle: cpu20]
23 ?? RL      0:00.00 [idle: cpu19]
24 ?? RL      0:00.00 [idle: cpu18]
25 ?? RL      0:00.00 [idle: cpu17]
26 ?? RL      0:00.00 [idle: cpu16]
27 ?? RL      0:00.00 [idle: cpu15]
28 ?? RL      0:00.00 [idle: cpu14]
29 ?? RL      0:00.00 [idle: cpu13]
30 ?? RL      0:00.00 [idle: cpu12]
31 ?? RL      0:00.00 [idle: cpu11]
32 ?? RL      0:00.00 [idle: cpu10]
33 ?? RL      0:00.00 [idle: cpu9]
34 ?? RL      18184:07.25 [idle: cpu8]
35 ?? RL      0:00.00 [idle: cpu7]
36 ?? RL      17862:11.31 [idle: cpu6]
37 ?? RL      19343:45.16 [idle: cpu5]
38 ?? RL      5192:38.30 [idle: cpu4]
39 ?? RL      0:00.00 [idle: cpu3]
40 ?? RL      19278:02.24 [idle: cpu2]
41 ?? RL      19291:00.72 [idle: cpu1]
42 ?? RL      18910:31.21 [idle: cpu0]
43 ?? WL      19:03.74 [swi2: net]
44 ?? WL      261:43.82 [swi7: clock sio]
45 ?? WL      0:00.00 [swi6: vm]
46 ?? DL      2:18.57 [yarrow]
47 ?? WL      0:00.00 [swi9: +]
48 ?? WL      0:00.00 [swi8: +]
49 ?? WL      0:12.36 [swi5: cambio]
50 ?? WL      0:00.00 [swi9: task queue]
51 ?? WL      0:00.00 [swi0: sio]
52 ?? WL      0:32.40 [irq39: ehci0]
53 ?? DL      0:00.21 [usb0]
54 ?? DL      0:00.00 [usbtask]
55 ?? WL      0:00.00 [irq22: xlr_lbus0]
56 ?? WL      0:00.00 [irq38: xlr_lbus0]
57 ?? WL      0:00.00 [swi3: ip6opt ipopt]
58 ?? WL      0:00.00 [swi4: ip6mismatch+]
59 ?? WL      0:00.00 [swi1: ipfwd]
60 ?? DL      0:18.65 [pagezero]
61 ?? DL      0:18.59 [bufdaemon]
62 ?? DL      1:10.44 [vnlr_u_mem]
63 ?? DL      1:51.66 [syncer]

```

```

64 ?? DL      0:20.22 [vn1ru]
65 ?? DL      0:40.48 [softdepflush]
66 ?? DL      0:00.00 [netdaemon]
67 ?? DL      20:47.67 [vmkmemdaemon]
68 ?? DL      0:00.00 [if_pfe_listen]
69 ?? SL      0:02.80 [kdm_checkkcore]
70 ?? SL      0:03.34 [kdm_savekcore]
71 ?? SL      0:04.31 [kdm_livekcore]
72 ?? SL      0:06.14 [kdm_logger]
73 ?? SL      0:04.31 [kdm_kdb]
74 ?? SL      0:00.02 [devrt_kernel_thread]
75 ?? DL      0:21.54 [vmuncachedaemon]
76 ?? DL      0:00.00 [if_pic_listen0]
77 ?? SL      0:00.00 [nfsiod 0]
78 ?? SL      0:00.00 [nfsiod 1]
79 ?? SL      0:00.00 [nfsiod 2]
80 ?? SL      0:00.00 [nfsiod 3]
81 ?? WL      5:59.98 [irq13: +]
82 ?? RL      105:06.81 [pkt_sender: cpu0]
83 ?? DL      0:03.62 [md0]
95 ?? DL      0:37.04 [md1]
115 ?? DL     0:06.01 [md2]
135 ?? DL     0:00.75 [md3]
155 ?? DL     0:21.17 [md4]
175 ?? DL     0:01.90 [md5]
195 ?? DL     0:06.26 [md6]
231 ?? DL     0:00.01 [md7]
755 ?? Ss     0:04.17 /usr/sbin/cron
847 ?? S      0:00.10 /usr/sbin/tnetd -N
849 ?? S      0:06.82 /usr/sbin/mgd -N
850 ?? S      0:00.32 /usr/sbin/inetd -N
852 ?? S      1:05.34 /usr/sbin/dhcpd -N
853 ?? S      0:00.18 /usr/sbin/inetd -p /var/run/inetd_4.pid -N -JU __juni
855 ?? L      1181:02.21 /usr/sbin/dc-pfe -N (pafxpc)
857 ?? S      17:55.86 /usr/sbin/vccpd -N
896 ?? S      93:43.45 /usr/sbin/chassism -N
953 ?? S      0:02.89 /sbin/watchdog -t-1
954 ?? S      3:34.00 /sbin/dcd -N
955 ?? S      10:30.13 /usr/sbin/chassisd -N
956 ?? DL     0:00.21 [peer proxy]
957 ?? S      4:07.43 /usr/sbin/alarmd -N
958 ?? S      0:31.69 /usr/sbin/craftd -N
959 ?? S      0:55.16 /usr/sbin/mib2d -N
960 ?? S      3:40.64 /usr/sbin/rpd -N
961 ?? S      0:00.03 /usr/sbin/tnp.snmpd -N
962 ?? S      0:51.94 /usr/sbin/pfed -N
963 ?? S      0:47.31 /usr/sbin/rmopd -N
964 ?? S      0:33.65 /usr/sbin/cosd
965 ?? S      1:48.41 /usr/sbin/ppmd -N
966 ?? S      0:07.18 /usr/sbin/dfwd -N
967 ?? S      1:02.56 /usr/sbin/bfdd -N
968 ?? S      0:00.63 /usr/sbin/rdd -N
969 ?? S      0:40.61 /usr/sbin/dfcd -N
971 ?? S      0:07.81 /usr/sbin/bdbrepd -N
972 ?? S      0:00.28 /usr/sbin/sendd -N
973 ?? S      1:37.69 /usr/sbin/xntpd -j -N -g -JU __juniper_private4__ (nt
974 ?? S      5:56.28 /usr/sbin/snmpd -N -JU __juniper_private4__
975 ?? S      16:46.82 /usr/sbin/jdiameterd -N
976 ?? S      2:34.13 /usr/sbin/eswd -N
977 ?? S      1:03.05 /usr/sbin/sflowd -N
978 ?? S      0:22.30 /usr/sbin/fcd -N

```

```
979 ?? S      1:07.01 /usr/sbin/vccpdf -N
982 ?? S      0:25.25 /usr/sbin/mcsnoopd -N
983 ?? S      3:45.68 /usr/sbin/rpdf -N
1043 ?? S      0:37.87 /usr/sbin/lacpd -N
1048 ?? DL     0:01.29 [peer proxy]
1111 ?? WL     0:00.00 [swi2: FMNITHRD+]
1112 ?? DL     0:00.03 [peer proxy]
12816 ?? S     15:35.32 /usr/sbin/sfid -N
30893 ?? Ss    0:00.65 sshd: tlewis@tty0 (sshd)
30897 ?? Ss    0:00.15 mgd: (mgd) (tlewis)/dev/tty0 (mgd)
30905 ?? Ss    0:00.64 sshd: tlewis@tty1 (sshd)
30909 ?? Ss    0:00.15 mgd: (mgd) (tlewis)/dev/tty1 (mgd)
30910 ?? Ss    0:01.26 sshd: tcheng@tty2 (sshd)
30914 ?? Ss    0:00.80 mgd: (mgd) (tcheng)/dev/tty2 (mgd)
30937 ?? R     0:00.03 /bin/ps -ax
661 d0- S     0:21.24 /usr/sbin/eventd -N -r -s -A
860 d0 Ss+    0:00.07 /usr/libexec/getty std.9600 ttyd0
30896 p0 Ss+   0:00.55 -cli (cli)
30908 p1 Ss+   0:00.50 -cli (cli)
30913 p2 Ss+   0:00.85 -cli (cli)
```


show system statistics

Syntax	show system statistics
Syntax (EX Series Switches)	show system statistics <all-members> <local> <member <i>member-id</i> >
Syntax (TX Matrix Router)	show system statistics <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show system statistics <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Router)	show system statistics <all-members> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show system statistics
Release Information	Command introduced before JUNOS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in JUNOS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display system-wide protocol-related statistics.
Options	none —Display system statistics for all the following protocols: <ul style="list-style-type: none"> • arp—Address Resolution Protocol • bridge—IEEE 802.1 Bridging • clns—Connectionless Network Service • esis—End System-to-Intermediate System • ethoamcfm—Ethernet OAM protocol for connectivity fault management • ethoamlfm—Ethernet OAM protocol for link fault management • icmp—Internet Control Message Protocol • icmp6—Internet Control Message Protocol version 6 • igmp—Internet Group Management Protocol • ip—Internet Protocol version 4 • ip6—Internet Protocol version 6 • mpls—Multiprotocol Label Switching • rdp—Reliable Datagram Protocol

- **tcp**—Transmission Control Protocol
- **tnp**—Trivial Network Protocol
- **ttp**—TNP Tunneling Protocol
- **tudp**—Trivial User Datagram Protocol
- **udp**—User Datagram Protocol
- **vpls**—Virtual Private LAN Service

all-chassis—(TX Matrix and TX Matrix Plus routers only) (Optional) Display system statistics for a protocol for all the routers in the chassis.

all-lcc—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system statistics for a protocol for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for a protocol for all T1600 routers (or line-card chassis) connected to the TX Matrix Plus router

all-members—(EX4200 switches and MX Series routers only) (Optional) Display system statistics for a protocol for all members of the Virtual Chassis configuration.

lcc number—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system statistics for a protocol for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for a protocol for a specific T1600 router that is connected to the TX Matrix Plus router. Replace **number** with a value from 0 through 3.

local—(EX4200 switches and MX Series routers only) (Optional) Display system statistics for a protocol for the local Virtual Chassis member.

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

scc—(TX Matrix routers only) (Optional) Display system statistics for a protocol for the TX Matrix router (or switch-card chassis).

sfc number—(TX Matrix Plus routers only) (Optional) Display system statistics for a protocol for the TX Matrix Plus router (or switch-fabric chassis). Replace **number** with 0.

Additional Information By default, when you issue the **show system statistics** command on a TX Matrix or TX Matrix Plus master Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) master Routing Engines connected to it. Likewise, if you issue the same command on the TX Matrix or TX Matrix Plus backup Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) backup Routing Engines that are connected to it.

Required Privilege Level view

List of Sample Output [show system statistics on page 837](#)
[show system statistics \(EX Series Switches\) on page 844](#)
[show system statistics \(TX Matrix Router\) on page 853](#)
[show system statistics \(QFX Series\) on page 860](#)

Sample Output

show system statistics

```
user@host> show system statistics
ip:
    3682087 total packets received
    0 bad header checksums
    0 with size smaller than minimum
    0 with data size < data length
    0 with header length < data size
    0 with data length < header length
    0 with incorrect version number
    0 packets destined to dead next hop
    0 fragments received
    0 fragments dropped (dup or out of space)
    0 fragments dropped (queue overflow)
    0 fragments dropped after timeout
    0 fragments dropped due to over limit
    0 packets reassembled ok
    3664774 packets for this host
    17316 packets for unknown/unsupported protocol
    0 packets forwarded
    0 packets not forwardable
    0 redirects sent
    6528 packets sent from this host
    0 packets sent with fabricated ip header
    0 output packets dropped due to no bufs
    0 output packets discarded due to no route
    0 output datagrams fragmented
    0 fragments created
    0 datagrams that can't be fragmented
    0 packets with bad options
    1123 packets with options handled without error
    0 strict source and record route options
    0 loose source and record route options
    0 record route options
    0 timestamp options
    0 timestamp and address options
    0 timestamp and prespecified address options
    0 option packets dropped due to rate limit
    1123 router alert options
    0 multicast packets dropped (no iflist)
    0 packets dropped (src and int don't match)
icmp:
    0 drops due to rate limit
    0 calls to icmp_error
    0 errors not generated because old message was icmp
Output histogram:
    echo reply: 75
    0 messages with bad code fields
    0 messages less than the minimum length
    0 messages with bad checksum
```

```
0 messages with bad source address
0 messages with bad length
0 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
    echo: 75
    router advertisement: 130
75 message responses generated
tcp:
3844 packets sent
    3618 data packets (1055596 bytes)
    0 data packets (0 bytes) retransmitted
    0 resends initiated by MTU discovery
    205 ack-only packets (148 packets delayed)
    0 URG only packets
    0 window probe packets
    0 window update packets
    1079 control packets
5815 packets received
    3377 acks (for 1055657 bytes)
    24 duplicate acks
    0 acks for unsent data
    2655 packets (15004 bytes) received in-sequence
    1 completely duplicate packet (0 bytes)
    0 old duplicate packets
    0 packets with some dup. data (0 bytes duped)
    0 out-of-order packets (0 bytes)
    0 packets (0 bytes) of data after window
    0 window probes
    7 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
1 connection request
32 connection accepts
0 bad connection attempts
0 listen queue overflows
33 connections established (including accepts)
30 connections closed (including 0 drops)
    27 connections updated cached RTT on close
    27 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
0 embryonic connections dropped
3374 segments updated rtt (of 3220 attempts)
0 retransmit timeouts
    0 connections dropped by rexmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
344 keepalive timeouts
    0 keepalive probes sent
    0 connections dropped by keepalive
1096 correct ACK header predictions
1314 correct data packet header predictions
32 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    32 completed
    0 bucket overflow
    0 cache overflow
```

```

        0 reset
        0 stale
        0 aborted
        0 badack
        0 unreach
        0 zone failures
    0 cookies sent
    0 cookies received
    0 ACKs sent in response to in-window but not exact RSTs
    0 ACKs sent in response to in-window SYNs on established connections
    0 rcv packets dropped by TCP due to bad address
    0 out-of-sequence segment drops due to insufficient memory
    1058 RST packets
    0 ICMP packets ignored by TCP
    0 send packets dropped by TCP due to auth errors
    0 rcv packets dropped by TCP due to auth errors
udp:
    3658884 datagrams received
    0 with incomplete header
    0 with bad data length field
    0 with bad checksum
    3657342 dropped due to no socket
    3657342 broadcast/multicast datagrams dropped due to no socket
    0 dropped due to full socket buffers
    0 not for hashed pcb
    4291311496 delivered
    1551 datagrams output
ipsec:
    0 inbound packets processed successfully
    0 inbound packets violated process security policy
    0 inbound packets with no SA available
    0 invalid inbound packets
    0 inbound packets failed due to insufficient memory
    0 inbound packets failed getting SPI
    0 inbound packets failed on AH replay check
    0 inbound packets failed on ESP replay check
    0 inbound AH packets considered authentic
    0 inbound AH packets failed on authentication
    0 inbound ESP packets considered authentic
    0 inbound ESP packets failed on authentication
    0 outbound packets processed successfully
    0 outbound packets violated process security policy
    0 outbound packets with no SA available
    0 invalid outbound packets
    0 outbound packets failed due to insufficient memory
    0 outbound packets with no route
igmp:
    17186 messages received
    0 messages received with too few bytes
    0 messages received with bad checksum
    0 membership queries received
    0 membership queries received with invalid field(s)
    0 membership reports received
    0 membership reports received with invalid field(s)
    0 membership reports received for groups to which we belong
    0 membership reports sent
arp:
    44181302 datagrams received
    2 ARP requests received
    2028 ARP replies received
    3156 resolution requests received

```

```
0 unrestricted proxy requests
0 received proxy requests
0 proxy requests not proxied
0 with bogus interface
787 with incorrect length
712 for non-IP protocol
0 with unsupported op code
0 with bad protocol address length
0 with bad hardware address length
0 with multicast source address
7611 with multicast target address
0 with my own hardware address
14241699 for an address not on the interface
0 with a broadcast source address
0 with source address duplicate to mine
29929250 which were not for me
0 packets discarded waiting for resolution
6 packets sent after waiting for resolution
17812 ARP requests sent
2 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry

ip6:
0 total packets received
0 with size smaller than minimum
0 with data size < data length
0 with bad options
0 with incorrect version number
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped after timeout
0 fragments that exceeded limit
0 packets reassembled ok
0 packets for this host
0 packets forwarded
0 packets not forwardable
0 redirects sent
0 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs, etc.
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets that violated scope rules
0 multicast packets which we don't join
Mbuf statistics:
0 packets whose headers are not continuous
0 tunneling packets that can't find gif
0 packets discarded due to too many headers
0 failures of source address selection
0 forward cache hit
0 forward cache miss
0 packets destined to dead next hop
0 option packets dropped due to rate limit
0 packets dropped (src and int don't match)
0 packets dropped due to bad protocol

icmp6:
0 calls to icmp_error
0 errors not generated because old message was icmp error or so
```

```

0 errors not generated because rate limitation
0 messages with bad code fields
0 messages < minimum length
0 bad checksums
0 messages with bad length
Histogram of error messages to be generated:
    0 no route
    0 administratively prohibited
    0 beyond scope
    0 address unreachable
    0 port unreachable
    0 packet too big
    0 time exceed transit
    0 time exceed reassembly
    0 erroneous header field
    0 unrecognized next header
    0 unrecognized option
    0 redirect
    0 unknown
0 message responses generated
0 messages with too many ND options
ipsec6:
0 inbound packets processed successfully
0 inbound packets violated process security policy
0 inbound packets with no SA available
0 invalid inbound packets
0 inbound packets failed due to insufficient memory
0 inbound packets failed getting SPI
0 inbound packets failed on AH replay check
0 inbound packets failed on ESP replay check
0 inbound AH packets considered authentic
0 inbound AH packets failed on authentication
0 inbound ESP packets considered authentic
0 inbound ESP packets failed on authentication
0 outbound packets processed successfully
0 outbound packets violated process security policy
0 outbound packets with no SA available
0 invalid outbound packets
0 outbound packets failed due to insufficient memory
0 outbound packets with no route
c1n1:
0 total packets received
0 packets delivered
0 too small
0 bad header length
0 bad checksum
0 bad version
0 unknown or unsupported protocol
0 bogus sdl size
0 no free memory in socket buffer
0 send packets discarded
0 sbappend failure
0 mcopy failure
0 address fields were not reasonable
0 segment information forgotten
0 forwarded packets
0 total packets sent
0 output packets discarded
0 non-forwarded packets
0 packets fragmented
0 fragments sent

```

```
0 fragments discarded
0 fragments timed out
0 fragmentation prohibited
0 packets reconstructed
0 packets destined to dead nexthop
0 packets discarded due to no route
0 Error pdu rate drops
0 ER pdu generation failure
esis:
0 total pkts received
0 total packets consumed by protocol
0 pdus received with bad checksum
0 pdus received with bad version number
0 pdus received with bad type field
0 short pdus received
0 bogus sdl size
0 bad header length
0 unknown or unsupported protocol
0 no free memory in socket buffer
0 send packets discarded
0 sbappend failure
0 mcopy failure
0 ISO family not configured
tnp:
146776365 unicast packets received
0 broadcast packets received
0 fragmented packets received
0 hello packets dropped
0 fragments dropped
0 fragment reassembly queue flushes
0 hello packets received
0 control packets received
49681642 rdp packets received
337175 udp packets received
96757548 tunnel packets received
0 input packets discarded with no protocol
98397591 unicast packets sent
0 broadcast packets sent
0 fragmented packets sent
0 hello packets dropped
0 fragments dropped
0 hello packets sent
0 control packets sent
49681642 rdp packets sent
337175 udp packets sent
48378774 tunnel packets sent
0 packets sent with unknown protocol
rdp:
49681642 input packets
0 discards for bad checksum
0 discards bad sequence number
0 refused connections
2031964 acks received
0 dropped due to full socket buffers
49692 retransmits
49681642 output packets
24815968 acks sent
28 connects
0 closes
22783990 keepalives received
22783990 keepalives sent
```



```

tudp:
    337175 datagrams received
    0 with incomplete header
    0 with bad data length field
    0 with bad checksum
    0 dropped due to no socket
    0 broadcast/multicast datagrams dropped due to no socket
    0 dropped due to full socket buffers
    337175 delivered
    337175 datagrams output

ttp:
    398749 packets sent
    0 packets sent while unconnected
    0 packets sent while interface down
    0 packets sent couldn't get buffer
    0 packets sent couldn't find neighbor
    44696687 L2 packets received
    0 unknown L3 packets received
    3682087 IPv4 L3 packets received
    0 MPLS L3 packets received
    0 MPLS->IPv4 L3 packets received
    0 IPv4->MPLS L3 packets received
    0 IPv6 L3 packets received
    0 ARP L3 packets received
    0 CLNP L3 packets received
    0 TNP L3 packets received
    0 NULL L3 packets received
    0 cyclotron cycle L3 packets received
    0 cyclotron send L3 packets received
    0 packets received while unconnected
    0 packets received from unknown ifl
    0 input packets couldn't get buffer
    0 input packets with bad type
    0 input packets with discard type
    0 Input packets with too many tlvs
    0 Input packets with bad tlv header
    70633 Input packets with bad tlv type
    68877 Input packets dropped based on tlv result
    0 input packets for which rt lookup is bypassed

mpls:
    0 total mpls packets received
    0 packets forwarded
    0 packets dropped
    0 with header too small
    0 after tagging, can't fit link MTU
    0 with IPv4 explicit NULL tag
    0 with IPv4 explicit NULL cksum errors
    0 with router alert tag
    0 lsp ping packets (ttl-expired/router alert)
    0 with ttl expired
    0 with tag encoding error
    0 packets discarded, no route

vpls:
    0 total packets received
    0 with size smaller than minimum
    0 with incorrect version number
    0 packets for this host
    0 packets with no logical interface
    0 packets with no family
    0 packets with no route table
    0 packets with no auxiliary table

```

```
0 packets with no corefacing entry
0 packets with no CE-facing entry
0 mac route learning requests
0 mac routes learnt
0 requests to learn an existing route
0 learning requests while learning disabled on interface
0 learning requests over capacity
0 mac routes moved
0 requests to move static route
0 mac route aging requests
0 mac routes aged
0 bogus address in aging requests
0 requests to age static route
0 requests to re-ageout aged route
0 requests involving multiple peer FEs
0 aging acks from PFE
0 aging non-acks from PFE
0 aging requests timed out waiting on FEs
0 aging requests over max-rate
0 errors finding peer FEs
```

show system statistics (EX Series Switches)

```
user@host> show system statistics
```

```
Tcp:
```

```
571779 packets sent
    21517 data packets (1797102 bytes)
    2 data packets retransmitted (20 bytes)
    0 resends initiated by MTU discovery
    3708 ack only packets (531 packets delayed)
    0 URG only packets
    1 window probe packets
    1 window update packets
    1093063 control packets
1132541 packets received
    20961 acks(for 1796102 bytes)
    5861 duplicate acks
    0 acks for unsent data
    19556 packets received in-sequence(232079 bytes)
    3018 completely duplicate packets(0 bytes)
    0 old duplicate packets
    4 packets with some duplicate data(4 bytes duped)
    2 out-of-order packets(2 bytes)
    0 packets of data after window(0 bytes)
    0 window probes
    39 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
546519 connection requests
78 connection accepts
0 bad connection attempts
0 listen queue overflows
100 connections established (including accepts)
546596 connections closed (including 6 drops)
    47 connections updated cached RTT on close
    47 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
546497 embryonic connections dropped
20453 segments updated rtt(of 566914 attempts)
```

```

2 retransmit timeouts
    0 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
3028 keepalive timeouts
    3027 keepalive probes sent
    1 connections dropped by keepalive
7515 correct ACK header predictions
12258 correct data packet header predictions
78 syncache entries added
    0 retransmitted
    0 dupsyn
    4 dropped
    78 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
1 SACK recovery episodes
1 segment retransmits in SACK recovery episodes
1 byte retransmits in SACK recovery episodes
71 SACK options (SACK blocks) received
1 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
546544 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing

udp:
147 datagrams received
0 with incomplete header
0 with bad data length field
0 with bad checksum
9 dropped due to no socket
0 broadcast/multicast datagrams dropped due to no socket
0 dropped due to full socket buffers
0 not for hashed pcb
138 delivered
0 datagrams output

ip:
73704 total packets received
0 bad header checksums
0 with size smaller than minimum
0 with data size < data length
0 with header length < data size
0 with data length < header length
0 with incorrect version number
0 packets destined to dead next hop
0 fragments received
0 fragments dropped (dup or out of space)

```

```

0 fragments dropped (queue overflow)
0 fragments dropped after timeout
0 fragments dropped due to over limit
0 packets reassembled ok
1133057 packets for this host
0 packets for unknown/unsupported protocol
40146 packets forwarded
0 packets not forwardable
40146 redirects sent
1121700 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets with bad options
0 packets with options handled without error
0 strict source and record route options
0 loose source and record route options
0 record route options
0 timestamp options
0 timestamp and address options
0 timestamp and prespecified address options
0 option packets dropped due to rate limit
0 router alert options
0 multicast packets dropped (no iflist)
0 packets dropped (src and int don't match)
0 transit re packets dropped on mgmt i/f
0 packets used first nexthop in ecmp unilist
0 incoming ttpoip packets received
0 incoming ttpoip packets dropped
0 outgoing TTPoIP packets sent
0 outgoing TTPoIP packets dropped

icmp:
0 drops due to rate limit
9 calls to icmp_error
0 errors not generated because old message was icmp
Output histogram:
    295 echo reply
    9 destination unreachable
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
0 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
    295 echo
295 message responses generated

igmp:
0 messages received
0 messages received with too few bytes
0 messages received with bad checksum
0 membership queries received
0 membership queries received with invalid fields
0 membership reports received
0 membership reports received with invalid fields
0 membership reports received for groups to which we belong
0 Membership reports sent

```

```

raw_if:
    0 RAW packets transmitted
    0 PPPOE packets transmitted
    0 ISDN packets transmitted
    0 DIALER packets transmitted
    0 PPP packets transmitted to pppd
    0 PPP packets transmitted to jppd
    0 IGMPv2 packets transmitted
    13 output drops due to tx error
    0 MPU packets transmitted
    0 PPPOE packets received
    0 ISDN packets received
    0 DIALER packets received
    0 PPP packets received from pppd
    0 MPU packets received
    0 PPP packets received from jppd
    0 IGMPv2 packets received
    0 Input drops due to bogus protocol
    0 input drops due to no mbufs available
    0 input drops due to no space in socket
    0 input drops due to no socket

arp:
    186413 datagrams received
    88 ARP requests received
    88 ARP replies received
    0 resolution request received
    0 unrestricted proxy requests
    0 restricted proxy requests
    0 received proxy requests
    0 proxy requests not proxied
    0 restricted proxy requests not proxied
    0 datagrams with bogus interface
    0 datagrams with incorrect length
    0 datagrams for non-IP protocol
    0 datagrams with unsupported op code
    0 datagrams with bad protocol address length
    0 datagrams with bad hardware address length
    0 datagrams with multicast source address
    0 datagrams with multicast source address
    0 datagrams with my own hardware address
    164 datagrams for an address not on the interface
    0 datagrams with a broadcast source address
    0 datagrams with source address duplicate to mine
    186065 datagrams which were not for me
    0 packets discarded waiting for resolution
    0 packets sent after waiting for resolution
    50 ARP requests sent
    88 ARP replies sent
    0 requests for memory denied
    0 requests dropped on entry
    0 requests dropped during retry
    0 requests dropped due to interface deletion
    0 requests on unnumbered interfaces
    0 new requests on unnumbered interfaces
    0 replies for from unnumbered interfaces
    0 requests on unnumbered interface with non-subnetted donor
    0 replies from unnumbered interface with non-subnetted donor

ip6:
    0 total packets received
    0 packets with size smaller than minimum
    0 packets with data size < data length

```

```
0 packets with bad options
0 packets with incorrect version number
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped after timeout
0 fragments that exceeded limit
0 packets reassembled ok
0 packets for this host
0 packets forwarded
0 packets not forwardable
0 redirects sent
0 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs, etc.
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets that violated scope rules
0 multicast packets which we don't join
0 packets whose headers are not continuous
0 tunneling packets that can't find gif
0 packets discarded due to too many headers
0 failures of source address selection
0 forward cache hit
0 forward cache miss
0 Packets destined to dead next hop
0 option packets dropped due to rate limit
0 Packets dropped (src and int don't match)
0 packets dropped due to bad protocol
0 transit re packet(null) dropped on mgmt i/f

icmp6:
0 Calls to icmp_error
0 Errors not generated because old message was icmp error
0 Errors not generated because rate limitation
0 Messages with bad code fields
0 Messages < minimum length
0 Bad checksums
0 Messages with bad length
    0 No route
    0 Administratively prohibited
    0 Beyond scope
    0 Address unreachable
    0 Port unreachable
    0 packet too big
    0 Time exceed transit
    0 Time exceed reassembly
    0 Erroneous header field
    0 Unrecognized next header
    0 Unrecognized option
    0 redirect
    0 Unknown
0 Message responses generated
0 Messages with too many ND options

pfkey:
0 Requests sent from userland
0 Bytes sent from userland
histogram by message type:
    0 reserved
    0 dump
0 Messages with invalid length field
0 Messages with invalid version field
```

```

0 Messages with invalid message type field
0 Messages too short
0 Messages with memory allocation failure
0 Messages with duplicate extension
0 Messages with invalid extension type
0 Messages with invalid sa type
0 Messages with invalid address extension
0 Requests sent to userland
0 Bytes sent to userland
histogram by message type:
    0 reserved
    0 dump
0 Messages toward single socket
0 Messages toward all sockets
0 Messages toward registered sockets
0 Messages with memory allocation failure
cInl:
0 Total packets received
0 Packets delivered
0 Too small packets
0 Packets with bad header length
0 Packets with bad checksum
0 Bad version packets
0 Unknown or unsupported protocol packets
0 Packets with bogus sdl size
0 No free memory in socket buffer
0 Send packets discarded
0 Sbappend failure
0 Mcopy failure
0 Address fields were not reasonable
0 Segment information forgotten
0 Forwarded packets
0 Total packets sent
0 Output packets discarded
0 Non-forwarded packets
0 Packets fragmented
0 Fragments sent
0 Fragments discarded
0 Fragments timed out
0 Fragmentation prohibited
0 Packets reconstructed
0 Packets destined to dead nexthop
0 Packets discarded due to no route
0 Error pdu rate drops
    0 ER pdu generation failure
esis:
0 Total pkts received
0 Total packets consumed by protocol
0 Pdus received with bad checksum
0 Pdus received with bad version number
0 Pdus received with bad type field
0 Short pdus received
0 Pdus with bogus sdl size
0 Pdus with bad header length
0 Pdu with unknown or unsupported protocol
0 No free memory in socket buffer
0 Send packets discarded
0 Sbappend failure
0 Mcopy failure
0 ISO family not configured
tnp:

```

```
0 Unicast packets received
0 Broadcast packets received
0 Fragmented packets received
0 Hello packets dropped
0 Fragments dropped
0 Fragment reassembly queue flushes
0 Packets with tnp src address collision received
0 Hello packets received
0 Control packets received
0 Rdp packets received
0 Udp packets received
0 Tunnel packets received
0 Input packets discarded with no protocol
0 Packets of version unspecified received
0 Packets of version 1 received
0 Packets of version 2 received
0 Packets of version 3 received
0 Unicast packets sent
0 Broadcast packets sent
0 Fragmented packets sent
0 Hello packets dropped
0 Fragments dropped
0 Hello packets sent
0 Control packets sent
0 Rdp packets sent
0 Udp packets sent
0 Tunnel packets sent
0 Packets sent with unknown protocol
0 Packets of version unspecified sent
0 Packets of version 1 sent
0 Packets of version 2 sent
0 Packets of version 3 sent
rdp:
0 Input packets
0 Packets discarded for bad checksum
0 Packets discarded due to bad sequence number
0 Refused connections
0 Acks received
0 Packets dropped due to full socket buffers
0 Retransmits
0 Output packets
0 Acks sent
0 Connects
0 Closes
0 Keepalives received
0 Keepalives sent
tudp:
67 Datagrams received
0 Datagrams with incomplete header
0 Datagrams with bad data length field
0 Datagrams with bad checksum
0 Datagrams dropped due to no socket
0 Broadcast/multicast datagrams dropped due to no socket
0 Datagrams dropped due to full socket buffers
67 Delivered
68 Datagrams output
ttp:
0 Packets sent
0 Packets sent while unconnected
0 Packets sent while interface down
0 Packets sent couldn't get buffer
```



```

0 Packets sent couldn't find neighbor
0 L2 packets received
0 Unknown L3 packets received
0 IPv4 L3 packets received
0 MPLS L3 packets received
0 MPLS->IPv4 L3 packets received
0 IPv4->MPLS L3 packets received
0 IPv6 L3 packets received
0 ARP L3 packets received
0 CLNP L3 packets received
0 TNP L3 packets received
0 NULL L3 packets received
0 Cyclotron cycle L3 packets received
0 Cyclotron send L3 packets received
0 Packets received while unconnected
0 Packets received from unknown ifl
0 Input packets couldn't get buffer
0 Input packets with bad type
0 Input packets with discard type
0 Input packets with too many tlvs
0 Input packets with bad tlv header
70633 Input packets with bad tlv type
68877 Input packets dropped based on tlv result
0 Input packets for which rt lookup is bypassed

mpls:
0 Total MPLS packets received
0 Packets forwarded
0 Packets dropped
0 Packets with header too small
0 After tagging, packets can't fit link MTU
0 Packets with IPv4 explicit NULL tag
0 Packets with IPv4 explicit NULL cksum errors
0 Packets with router alert tag
0 LSP ping packets (ttl-expired/router alert)
0 Packets with ttl expired
0 Packets with tag encoding error
0 Packets discarded due to no route
0 Packets used first nexthop in ecmp unilist

vpls:
0 Total packets received
0 Packets with size smaller than minimum
0 Packets with incorrect version number
0 Packets for this host
0 Packets with no logical interface
0 Packets with no family
0 Packets with no route table
0 Packets with no auxiliary table
0 Packets with no corefacing entry
0 packets with no CE-facing entry
0 MAC route learning requests
0 MAC routes learnt
0 Requests to learn an existing route
0 Learning requests while learning disabled on interface
0 Learning requests over capacity
0 MAC routes moved
0 Requests to move static route
0 MAC route aging requests
0 MAC routes aged
0 Bogus address in aging requests
0 Requests to age static route
0 Requests to re-ageout aged route

```

```
0 Requests involving multiple peer FEs
0 Aging acks from PFE
0 Aging non-acks from PFE
0 Aging requests timed out waiting on FEs
0 Aging requests over max-rate
0 Errors finding peer FEs
0 Unsupported platform
0 Packets dropped due to no l3 route table
0 Packets dropped due to no local ifl
0 Packets punted
0 Packets dropped due to no socket

bridge:
Input:
0 packets received
0 packets forwarded
0 packets failed to forward
0 packets dropped
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with stp state lookup failures
0 packets dropped due to stp blocked/listening
0 packets dropped due to stp learning
0 packets with src MAC learning failures
0 packets with input control processing failures
Forward:
0 packets sent successfully
0 packets with send failures
0 packets forwarded to l3 interface
0 packets with l3 send failures
0 packets discarded
0 packets with l2ifl store failures
0 packets with ifl mismatch failures
0 packets with packet duplication failures
0 packets with tag lookup failures
0 packets with no route for DMAC
0 packets with no route table
0 packets with no nexthop
0 packets with dead nexthop
0 packets with eof reached error
Learning:
0 MACs learned
0 packets sent to l3 interface
0 packets with l3 send failures
0 packets hit holdq while learning
0 MAC moves
0 packets discarded
0 packets with no route for SMAC
0 packets with no nexthop
0 packets with dead nexthop
0 packets dropped due to no resolve route
0 packets with l3 ifd lookup failures
0 packets with l3 ifl lookup failures
0 packets with l3 invalid rnh
0 packets with no route for SMAC in clone learning
0 packets with no nexthop in clone learning
0 packets with dead nexthop in clone learning
0 packets dropped due to no resolve nh in clone learning
Output:
0 packets forwarded
0 packets failed to forward
0 packets with vmember lookup failures
```

```

    0 packets with vlan lookup failures
0 packets with input control processing failures
Send:
0 packets sent successfully
0 packets with send failures
0 packets dropped due to interface down
0 packets with dev output failures
0 blocked ifl discards
0 packets with tag lookup failures
0 packets with stp state lookup failures
0 packets with tag insertion failures
0 packets with tag removal failures
Flood:
0 packets flooded
0 flood failures
IGMP:
0 packets sent successfully
0 packets with send failures
0 packets forwarded
0 packets failed to forward
0 packets with mpull failures
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with ifl lookup failures
0 packets with tag lookup failures
Misc:
0 packets with size smaller than minimum
0 packets with double tags
0 packets with no ifl
0 packets with no family
0 packets with no route table

```

show system statistics (TX Matrix Router)

```

user@host> show system statistics
sfc0-re0:

```

```

-----
Tcp:
    361694 packets sent
        326507 data packets (103237236 bytes)
        2343 data packets retransmitted (2673324 bytes)
        0 resends initiated by MTU discovery
        33857 ack only packets (31613 packets delayed)
        0 URG only packets
        14 window probe packets
        387 window update packets
        1108 control packets
    345879 packets received
        298207 acks(for 103141728 bytes)
        438 duplicate acks
        0 acks for unsent data
        204578 packets received in-sequence(13820995 bytes)
        6 completely duplicate packets(18 bytes)
        0 old duplicate packets
        0 packets with some duplicate data(0 bytes duped)
        0 out-of-order packets(0 bytes)
        0 packets of data after window(0 bytes)
        0 window probes
        899 window update packets
        166 packets received after close
        0 discarded for bad checksums

```

```
        0 discarded for bad header offset fields
        0 discarded because packet too short
406 connection requests
233 connection accepts
0 bad connection attempts
0 listen queue overflows
616 connections established (including accepts)
911 connections closed (including 41 drops)
    346 connections updated cached RTT on close
    346 connections updated cached RTT variance on close
    200 connections updated cached ssthresh on close
23 embryonic connections dropped
298155 segments updated rtt(of 287216 attempts)
1163 retransmit timeouts
    27 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
5 keepalive timeouts
    5 keepalive probes sent
    0 connections dropped by keepalive
69922 correct ACK header predictions
34993 correct data packet header predictions
233 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    233 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
23 SACK recovery episodes
68 segment retransmits in SACK recovery episodes
71542 byte retransmits in SACK recovery episodes
158 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
259 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing
```

lcc0-re0:

Tcp:

```
346 packets sent
    222 data packets (22894 bytes)
    0 data packets retransmitted (0 bytes)
    0 resends initiated by MTU discovery
    80 ack only packets (12 packets delayed)
```

```

    0 URG only packets
    0 window probe packets
    5 window update packets
    42 control packets
358 packets received
    268 acks(for 22939 bytes)
    9 duplicate acks
    0 acks for unsent data
    203 packets received in-sequence(33820 bytes)
    0 completely duplicate packets(0 bytes)
    0 old duplicate packets
    0 packets with some duplicate data(0 bytes duped)
    0 out-of-order packets(0 bytes)
    0 packets of data after window(0 bytes)
    0 window probes
    6 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
13 connection requests
18 connection accepts
0 bad connection attempts
0 listen queue overflows
31 connections established (including accepts)
35 connections closed (including 2 drops)
    3 connections updated cached RTT on close
    3 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
0 embryonic connections dropped
268 segments updated rtt(of 247 attempts)
0 retransmit timeouts
    0 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
0 keepalive timeouts
    0 keepalive probes sent
    0 connections dropped by keepalive
0 correct ACK header predictions
42 correct data packet header predictions
18 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    18 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 SACK recovery episodes
0 segment retransmits in SACK recovery episodes
0 byte retransmits in SACK recovery episodes
0 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow

```

- 0 ACKs sent in response to in-window but not exact RSTs
- 0 ACKs sent in response to in-window SYNs on established connections
- 0 rcv packets dropped by TCP due to bad address
- 0 out-of-sequence segment drops due to insufficient memory
- 5 RST packets
- 0 ICMP packets ignored by TCP
- 0 send packets dropped by TCP due to auth errors
- 0 rcv packets dropped by TCP due to auth errors
- 0 outgoing segments dropped due to policing

lcc1-re0:

Tcp:

- 348 packets sent
 - 223 data packets (22895 bytes)
 - 0 data packets retransmitted (0 bytes)
 - 0 resends initiated by MTU discovery
 - 81 ack only packets (13 packets delayed)
 - 0 URG only packets
 - 0 window probe packets
 - 5 window update packets
 - 42 control packets
- 360 packets received
 - 269 acks(for 22940 bytes)
 - 9 duplicate acks
 - 0 acks for unsent data
 - 203 packets received in-sequence(33820 bytes)
 - 0 completely duplicate packets(0 bytes)
 - 0 old duplicate packets
 - 0 packets with some duplicate data(0 bytes duped)
 - 0 out-of-order packets(0 bytes)
 - 0 packets of data after window(0 bytes)
 - 0 window probes
 - 6 window update packets
 - 0 packets received after close
 - 0 discarded for bad checksums
 - 0 discarded for bad header offset fields
 - 0 discarded because packet too short
- 13 connection requests
- 18 connection accepts
- 0 bad connection attempts
- 0 listen queue overflows
- 31 connections established (including accepts)
- 36 connections closed (including 2 drops)
 - 3 connections updated cached RTT on close
 - 3 connections updated cached RTT variance on close
 - 0 connections updated cached ssthresh on close
- 0 embryonic connections dropped
- 269 segments updated rtt(of 248 attempts)
- 0 retransmit timeouts
 - 0 connections dropped by retransmit timeout
- 0 persist timeouts
 - 0 connections dropped by persist timeout
- 0 keepalive timeouts
 - 0 keepalive probes sent
 - 0 connections dropped by keepalive
- 0 correct ACK header predictions
- 43 correct data packet header predictions
- 18 syncache entries added
 - 0 retransmitted
 - 0 dupsyn

```

    0 dropped
    18 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 SACK recovery episodes
0 segment retransmits in SACK recovery episodes
0 byte retransmits in SACK recovery episodes
0 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
5 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing

```

1cc2-re0:

 Tcp:

```

405 packets sent
    271 data packets (23926 bytes)
    0 data packets retransmitted (0 bytes)
    0 resends initiated by MTU discovery
    86 ack only packets (13 packets delayed)
    0 URG only packets
    0 window probe packets
    5 window update packets
    46 control packets
418 packets received
    321 acks(for 23975 bytes)
    9 duplicate acks
    0 acks for unsent data
    234 packets received in-sequence(34403 bytes)
    0 completely duplicate packets(0 bytes)
    0 old duplicate packets
    0 packets with some duplicate data(0 bytes duped)
    0 out-of-order packets(0 bytes)
    0 packets of data after window(0 bytes)
    0 window probes
    7 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
15 connection requests
19 connection accepts
0 bad connection attempts
0 listen queue overflows
34 connections established (including accepts)

```

```
39 connections closed (including 2 drops)
    4 connections updated cached RTT on close
    4 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
0 embryonic connections dropped
321 segments updated rtt(of 299 attempts)
0 retransmit timeouts
    0 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
0 keepalive timeouts
    0 keepalive probes sent
    0 connections dropped by keepalive
0 correct ACK header predictions
48 correct data packet header predictions
19 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    19 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 SACK recovery episodes
0 segment retransmits in SACK recovery episodes
0 byte retransmits in SACK recovery episodes
0 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
5 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing
```

lcc3-re0:

Tcp:

```
346 packets sent
    221 data packets (22895 bytes)
    0 data packets retransmitted (0 bytes)
    0 resends initiated by MTU discovery
    81 ack only packets (13 packets delayed)
    0 URG only packets
    0 window probe packets
    5 window update packets
    42 control packets
360 packets received
    267 acks(for 22940 bytes)
    9 duplicate acks
```



```

    0 acks for unsent data
    203 packets received in-sequence(33820 bytes)
    0 completely duplicate packets(0 bytes)
    0 old duplicate packets
    0 packets with some duplicate data(0 bytes duped)
    0 out-of-order packets(0 bytes)
    0 packets of data after window(0 bytes)
    0 window probes
    6 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
13 connection requests
18 connection accepts
0 bad connection attempts
0 listen queue overflows
31 connections established (including accepts)
35 connections closed (including 2 drops)
    3 connections updated cached RTT on close
    3 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
0 embryonic connections dropped
267 segments updated rtt(of 246 attempts)
0 retransmit timeouts
    0 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
0 keepalive timeouts
    0 keepalive probes sent
    0 connections dropped by keepalive
0 correct ACK header predictions
43 correct data packet header predictions
18 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    18 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 SACK recovery episodes
0 segment retransmits in SACK recovery episodes
0 byte retransmits in SACK recovery episodes
0 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
5 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors

```

0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing

show system statistics (QFX Series)

```
user@switch> show system statistics
Tcp:
571779 packets sent
21517 data packets (1797102 bytes)
2 data packets retransmitted (20 bytes)
0 resends initiated by MTU discovery
3708 ack only packets (531 packets delayed)
0 URG only packets
1 window probe packets
1 window update packets
1093063 control packets
1132541 packets received
20961 acks(for 1796102 bytes)
5861 duplicate acks
0 acks for unsent data
19556 packets received in-sequence(232079 bytes)
3018 completely duplicate packets(0 bytes)
0 old duplicate packets
4 packets with some duplicate data(4 bytes duped)
2 out-of-order packets(2 bytes)
0 packets of data after window(0 bytes)
0 window probes
39 window update packets
0 packets received after close
0 discarded for bad checksums
0 discarded for bad header offset fields
0 discarded because packet too short
546519 connection requests
78 connection accepts
0 bad connection attempts
0 listen queue overflows
100 connections established (including accepts)
546596 connections closed (including 6 drops)
47 connections updated cached RTT on close
47 connections updated cached RTT variance on close
0 connections updated cached ssthresh on close
546497 embryonic connections dropped
20453 segments updated rtt(of 566914 attempts)
2 retransmit timeouts
0 connections dropped by retransmit timeout
0 persist timeouts
0 connections dropped by persist timeout
3028 keepalive timeouts
3027 keepalive probes sent
1 connections dropped by keepalive
7515 correct ACK header predictions
12258 correct data packet header predictions
78 syncache entries added
0 retransmitted
0 dupsyn
4 dropped
78 completed
0 bucket overflow
0 cache overflow
0 reset
0 stale
```

```
0 aborted
0 badack
0 unreach
0 zone failures
0 cookies sent
0 cookies received
1 SACK recovery episodes
1 segment retransmits in SACK recovery episodes
1 byte retransmits in SACK recovery episodes
71 SACK options (SACK blocks) received
1 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
546544 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing
udp:
147 datagrams received
0 with incomplete header
0 with bad data length field
0 with bad checksum
9 dropped due to no socket
0 broadcast/multicast datagrams dropped due to no socket
0 dropped due to full socket buffers
0 not for hashed pcb
138 delivered
0 datagrams output
ip:
73704 total packets received
0 bad header checksums
0 with size smaller than minimum
0 with data size < data length
0 with header length < data size
0 with data length < header length
0 with incorrect version number
0 packets destined to dead next hop
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped (queue overflow)
0 fragments dropped after timeout
0 fragments dropped due to over limit
0 packets reassembled ok
1133057 packets for this host
0 packets for unknown/unsupported protocol
40146 packets forwarded
0 packets not forwardable
40146 redirects sent
1121700 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets with bad options
0 packets with options handled without error
```

```
0 strict source and record route options
0 loose source and record route options
0 record route options
0 timestamp options
0 timestamp and address options
0 timestamp and prespecified address options
0 option packets dropped due to rate limit
0 router alert options
0 multicast packets dropped (no iflist)
0 packets dropped (src and int don't match)
0 transit re packets dropped on mgmt i/f
0 packets used first nexthop in ecmp unilist
0 incoming ttpoip packets received
0 incoming ttpoip packets dropped
0 outgoing TTPoIP packets sent
0 outgoing TTPoIP packets dropped
icmp:
0 drops due to rate limit
9 calls to icmp_error
0 errors not generated because old message was icmp
Output histogram:
295 echo reply
9 destination unreachable
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
0 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
295 echo
295 message responses generated
igmp:
0 messages received
0 messages received with too few bytes
0 messages received with bad checksum
0 membership queries received
0 membership queries received with invalid fields
0 membership reports received
0 membership reports received with invalid fields
0 membership reports received for groups to which we belong
0 Membership reports sent
raw_if:
0 RAW packets transmitted
0 PPPOE packets transmitted
0 ISDN packets transmitted
0 DIALER packets transmitted
0 PPP packets transmitted to pppd
0 PPP packets transmitted to jppd
0 IGMPv2 packets transmitted
13 output drops due to tx error
0 MPU packets transmitted
0 PPPOE packets received
0 ISDN packets received
0 DIALER packets received
0 PPP packets received from pppd
0 MPU packets received
0 PPP packets received from jppd
0 IGMPv2 packets received
0 Input drops due to bogus protocol
```

```
0 input drops due to no mbufs available
0 input drops due to no space in socket
0 input drops due to no socket
arp:
186413 datagrams received
88 ARP requests received
88 ARP replies received
0 resolution request received
0 unrestricted proxy requests
0 restricted proxy requests
0 received proxy requests
0 proxy requests not proxied
0 restricted proxy requests not proxied
0 datagrams with bogus interface
0 datagrams with incorrect length
0 datagrams for non-IP protocol
0 datagrams with unsupported op code
0 datagrams with bad protocol address length
0 datagrams with bad hardware address length
0 datagrams with multicast source address
0 datagrams with multicast source address
0 datagrams with my own hardware address
164 datagrams for an address not on the interface
0 datagrams with a broadcast source address
0 datagrams with source address duplicate to mine
186065 datagrams which were not for me
0 packets discarded waiting for resolution
0 packets sent after waiting for resolution
50 ARP requests sent
88 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry
0 requests dropped due to interface deletion
0 requests on unnumbered interfaces
0 new requests on unnumbered interfaces
0 replies for from unnumbered interfaces
0 requests on unnumbered interface with non-subnetted donor
0 replies from unnumbered interface with non-subnetted donor
ip6:
0 total packets received
0 packets with size smaller than minimum
0 packets with data size < data length
0 packets with bad options
0 packets with incorrect version number
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped after timeout
0 fragments that exceeded limit
0 packets reassembled ok
0 packets for this host
0 packets forwarded
0 packets not forwardable
0 redirects sent
0 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs, etc.
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets that violated scope rules
```

```
0 multicast packets which we don't join
0 packets whose headers are not continuous
0 tunneling packets that can't find gif
0 packets discarded due to too many headers
0 failures of source address selection
0 forward cache hit
0 forward cache miss
0 Packets destined to dead next hop
0 option packets dropped due to rate limit
0 Packets dropped (src and int don't match)
0 packets dropped due to bad protocol
0 transit re packet(null) dropped on mgmt i/f
icmp6:
0 Calls to icmp_error
0 Errors not generated because old message was icmp error
0 Errors not generated because rate limitation
0 Messages with bad code fields
0 Messages < minimum length
0 Bad checksums
0 Messages with bad length
0 No route
0 Administratively prohibited
0 Beyond scope
0 Address unreachable
0 Port unreachable
0 packet too big
0 Time exceed transit
0 Time exceed reassembly
0 Erroneous header field
0 Unrecognized next header
0 Unrecognized option
0 redirect
0 Unknown
0 Message responses generated
0 Messages with too many ND options
pfkey:
0 Requests sent from userland
0 Bytes sent from userland
histogram by message type:
0 reserved
0 dump
0 Messages with invalid length field
0 Messages with invalid version field
0 Messages with invalid message type field
0 Messages too short
0 Messages with memory allocation failure
0 Messages with duplicate extension
0 Messages with invalid extension type
0 Messages with invalid sa type
0 Messages with invalid address extension
0 Requests sent to userland
0 Bytes sent to userland
histogram by message type:
0 reserved
0 dump
0 Messages toward single socket
0 Messages toward all sockets
0 Messages toward registered sockets
0 Messages with memory allocation failure
c1n1:
0 Total packets received
```

```
0 Packets delivered
0 Too small packets
0 Packets with bad header length
0 Packets with bad checksum
0 Bad version packets
0 Unknown or unsupported protocol packets
0 Packets with bogus sdl size
0 No free memory in socket buffer
0 Send packets discarded
0 Sbappend failure
0 Mcopy failure
0 Address fields were not reasonable
0 Segment information forgotten
0 Forwarded packets
0 Total packets sent
0 Output packets discarded
0 Non-forwarded packets
0 Packets fragmented
0 Fragments sent
0 Fragments discarded
0 Fragments timed out
0 Fragmentation prohibited
0 Packets reconstructed
0 Packets destined to dead nexthop
0 Packets discarded due to no route
0 Error pdu rate drops
0 ER pdu generation failure
esis:
0 Total pkts received
0 Total packets consumed by protocol
0 Pdus received with bad checksum
0 Pdus received with bad version number
0 Pdus received with bad type field
0 Short pdus received
0 Pdus with bogus sdl size
0 Pdus with bad header length
0 Pdus with unknown or unsupported protocol
0 No free memory in socket buffer
0 Send packets discarded
0 Sbappend failure
0 Mcopy failure
0 ISO family not configured
tnp:
0 Unicast packets received
0 Broadcast packets received
0 Fragmented packets received
0 Hello packets dropped
0 Fragments dropped
0 Fragment reassembly queue flushes
0 Packets with tnp src address collision received
0 Hello packets received
0 Control packets received
0 Rdp packets received
0 Udp packets received
0 Tunnel packets received
0 Input packets discarded with no protocol
0 Packets of version unspecified received
0 Packets of version 1 received
0 Packets of version 2 received
0 Packets of version 3 received
0 Unicast packets sent
```

```
0 Broadcast packets sent
0 Fragmented packets sent
0 Hello packets dropped
0 Fragments dropped
0 Hello packets sent
0 Control packets sent
0 Rdp packets sent
0 Udp packets sent
0 Tunnel packets sent
0 Packets sent with unknown protocol
0 Packets of version unspecified sent
0 Packets of version 1 sent
0 Packets of version 2 sent
0 Packets of version 3 sent
rdp:
0 Input packets
0 Packets discarded for bad checksum
0 Packets discarded due to bad sequence number
0 Refused connections
0 Acks received
0 Packets dropped due to full socket buffers
0 Retransmits
0 Output packets
0 Acks sent
0 Connects
0 Closes
0 Keepalives received
0 Keepalives sent
tudp:
67 Datagrams received
0 Datagrams with incomplete header
0 Datagrams with bad data length field
0 Datagrams with bad checksum
0 Datagrams dropped due to no socket
0 Broadcast/multicast datagrams dropped due to no socket
0 Datagrams dropped due to full socket buffers
67 Delivered
68 Datagrams output
ttp:
0 Packets sent
0 Packets sent while unconnected
0 Packets sent while interface down
0 Packets sent couldn't get buffer
0 Packets sent couldn't find neighbor
0 L2 packets received
0 Unknown L3 packets received
0 IPv4 L3 packets received
0 MPLS L3 packets received
0 MPLS->IPv4 L3 packets received
0 IPv4->MPLS L3 packets received
0 IPv6 L3 packets received
0 ARP L3 packets received
0 CLNP L3 packets received
0 TNP L3 packets received
0 NULL L3 packets received
0 Cyclotron cycle L3 packets received
0 Cyclotron send L3 packets received
0 Packets received while unconnected
0 Packets received from unknown ifl
0 Input packets couldn't get buffer
0 Input packets with bad type
```



```

0 Input packets with discard type
0 Input packets with too many tlvs
0 Input packets with bad tlv header
70633 Input packets with bad tlv type
68877 Input packets dropped based on tlv result0 Input packets for which rt lookup
  is bypassed
mpls:
0 Total MPLS packets received
0 Packets forwarded
0 Packets dropped
0 Packets with header too small
0 After tagging, packets can't fit link MTU
0 Packets with IPv4 explicit NULL tag
0 Packets with IPv4 explicit NULL cksum errors
0 Packets with router alert tag
0 LSP ping packets (ttl-expired/router alert)
0 Packets with ttl expired
0 Packets with tag encoding error
0 Packets discarded due to no route
0 Packets used first nexthop in ecmp unilist
vpls:
0 Total packets received
0 Packets with size smaller than minimum
0 Packets with incorrect version number
0 Packets for this host
0 Packets with no logical interface
0 Packets with no family
0 Packets with no route table
582 Copyright © 2010, Juniper Networks, Inc.
0 Packets with no auxiliary table
0 Packets with no corefacing entry
0 packets with no CE-facing entry
0 MAC route learning requests
0 MAC routes learnt
0 Requests to learn an existing route
0 Learning requests while learning disabled on interface
0 Learning requests over capacity
0 MAC routes moved
0 Requests to move static route
0 MAC route aging requests
0 MAC routes aged
0 Bogus address in aging requests
0 Requests to age static route
0 Requests to re-ageout aged route
0 Requests involving multiple peer FEs
0 Aging acks from PFE
0 Aging non-acks from PFE
0 Aging requests timed out waiting on FEs
0 Aging requests over max-rate
0 Errors finding peer FEs
0 Unsupported platform
0 Packets dropped due to no l3 route table
0 Packets dropped due to no local ifl
0 Packets punted
0 Packets dropped due to no socket
bridge:
Input:
0 packets received
0 packets forwarded
0 packets failed to forward
0 packets dropped

```

```
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with stp state lookup failures
0 packets dropped due to stp blocked/listening
0 packets dropped due to stp learning
0 packets with src MAC learning failures
0 packets with input control processing failures
Forward:
0 packets sent successfully
0 packets with send failures
0 packets forwarded to l3 interface
0 packets with l3 send failures
0 packets discarded
0 packets with l2ifl store failures
0 packets with ifl mismatch failures
0 packets with packet duplication failures
0 packets with tag lookup failures
0 packets with no route for DMAC
0 packets with no route table
0 packets with no nexthop
0 packets with dead nexthop
0 packets with eof reached error
Learning:
0 MACs learned
0 packets sent to l3 interface
0 packets with l3 send failures
0 packets hit holdq while learning
0 MAC moves
0 packets discarded
0 packets with no route for SMAC
0 packets with no nexthop
0 packets with dead nexthop
0 packets dropped due to no resolve route
0 packets with l3 ifd lookup failures
0 packets with l3 ifl lookup failures
0 packets with l3 invalid rnh
0 packets with no route for SMAC in clone learning
0 packets with no nexthop in clone learning
0 packets with dead nexthop in clone learning
0 packets dropped due to no resolve nh in clone learning
Output:
0 packets forwarded
0 packets failed to forward
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with input control processing failures
Send:
0 packets sent successfully
0 packets with send failures
0 packets dropped due to interface down
0 packets with dev output failures
0 blocked ifl discards
0 packets with tag lookup failures
0 packets with stp state lookup failures
0 packets with tag insertion failures
0 packets with tag removal failures
Flood:
0 packets flooded
0 flood failures
IGMP:
0 packets sent successfully
```

```
0 packets with send failures
0 packets forwarded
0 packets failed to forward
0 packets with mpull failures
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with ifl lookup failures
0 packets with tag lookup failures
Misc:
0 packets with size smaller than minimum
0 packets with double tags
0 packets with no ifl
0 packets with no family
0 packets with no route table
```

show system uptime

Syntax	show system uptime
Syntax (EX Series Switches)	show system uptime <all-members> <local> <member <i>member-id</i> >
Syntax (QFX Series)	show system uptime <director-group <i>name</i> > <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-group <i>name</i> >
Syntax (TX Matrix Router)	show system uptime <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show system uptime <detail> <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Router)	show system uptime <all-members> <invoke-on> <local> <member <i>member-id</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in JUNOS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the current time and information about how long the router or switch, router or switch software, and routing protocols have been running.
Options	none —Show time since the system rebooted and processes started. all-chassis —(TX Matrix and TX Matrix Plus routers only) (Optional) Show time since the system rebooted and processes started on all the routers in the chassis. all-lcc —(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show time since the system rebooted and processes started for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, show time since the system rebooted and processes started for all T1600 routers (or line-card chassis) connected to the TX Matrix Plus router. all-members —(EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on all members of the Virtual Chassis configuration.

director-group *name*—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Director group.

infrastructure *name*—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the fabric control Routing Engine and fabric manager Routing Engine.

interconnect-device *name*—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Interconnect device.

invoke-on—(MX Series routers only) (Optional) Display the time since the system rebooted and processes started on the master Routing Engine, backup Routing Engine, or both, on a router with two Routing Engines.

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show time since the system rebooted and processes started for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, show time since the system rebooted and processes started for a specific T1600 router that is connected to the TX Matrix Plus router. Replace ***number*** with a value from 0 through 3.

local—(EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on the local Virtual Chassis member.

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

node-group *name*—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Node group.

scc—(TX Matrix routers only) (Optional) Show time since the system rebooted and processes started for the TX Matrix router (or switch-card chassis).

sfc *number*—(TX Matrix Plus routers only) (Optional) Show time since the system rebooted and processes started for the TX Matrix Plus router (or switch-fabric chassis). Replace ***number*** with 0.

Additional Information By default, when you issue the **show system uptime** command on a TX Matrix or TX Matrix Plus master Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on a TX Matrix router) or T1600 (in a routing matrix based on a TX Matrix Plus router) master Routing Engines connected to it. Likewise, if you issue the same command on the TX Matrix or TX Matrix Plus backup Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on a TX Matrix router) or T1600 (in a routing matrix based on a TX Matrix Plus router) backup Routing Engines that are connected to it.

Required Privilege Level view

- Related Documentation**
- *Monitoring System Process Information*
 - *Monitoring System Properties*
 - *10-Gigabit Ethernet LAN/WAN PIC with XFP (T640 Router)*

List of Sample Output

[show system uptime on page 872](#)
[show system uptime all-lcc \(TX Matrix Router\) on page 872](#)
[show system uptime all-lcc \(TX Matrix Plus Router\) on page 873](#)
[show system uptime \(QFX Series\) on page 873](#)

Output Fields [Table 66 on page 872](#) describes the output fields for the **show system uptime** command. Output fields are listed in the approximate order in which they appear.

Table 66: show system uptime Output Fields

Field Name	Field Description
Current time	Current system time in UTC.
System booted	Date and time when the Routing Engine on the router or switch was last booted and how long it has been running.
Protocols started	Date and time when the routing protocols were last started and how long they have been running.
Last configured	Date and time when a configuration was last committed. Also shows the name of the user who issued the last commit command.
time and up	Current time, in the local time zone, and how long the router or switch has been operational.
users	Number of users logged in to the router or switch.
load averages	Load averages for the last 1 minute, 5 minutes, and 15 minutes.

Sample Output

show system uptime

```
user@host> show system uptime
Current time:      1998-10-13 19:45:47 UTC
System booted:    1998-10-12 20:51:41 UTC (22:54:06 ago)
Protocols started: 1998-10-13 19:33:45 UTC (00:12:02 ago)
Last configured:  1998-10-13 19:33:45 UTC (00:12:02 ago) by abc
12:45PM up 22:54, 2 users, load averages: 0.07, 0.02, 0.01
```

show system uptime all-lcc (TX Matrix Router)

```
user@host> show system uptime all-lcc
lcc0-re0:
-----
Current time: 2004-09-13 09:55:35 PDT
System booted: 2004-09-13 03:13:55 PDT (06:41:40 ago)
Last configured: 2004-09-13 03:17:48 PDT (06:37:47 ago) by root
9:55AM PDT up 6:42, 1 user, load averages: 0.02, 0.03, 0.00
lcc2-re0:
```

```

-----
Current time: 2004-09-13 09:55:35 PDT
System booted: 2004-09-12 03:23:43 PDT (1d 06:31 ago)
Last configured: 2004-09-13 03:05:36 PDT (06:49:59 ago) by root
9:55AM PDT up 1 day, 6:32, 1 user, load averages: 0.02, 0.01, 0.00

```

show system uptime all-lcc (TX Matrix Plus Router)

```

user@host> show system uptime all-lcc
sfc0-re0:

```

```

-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:33 PDT (17:44:57 ago)
Protocols started: 2009-05-24 06:40:30 PDT (17:44:00 ago)
Last configured: 2009-05-24 06:33:27 PDT (17:51:03 ago) by gregdo
12:24AM up 17:45, 2 users, load averages: 0.07, 0.05, 0.01

```

```
lcc0-re0:
```

```

-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:46 PDT (17:44:44 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:47 PDT (17:43:43 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

```

```
lcc1-re0:
```

```

-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:38 PDT (17:44:52 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:18 PDT (17:44:12 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

```

```
lcc2-re0:
```

```

-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:48 PDT (17:44:42 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:44 PDT (17:43:46 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

```

```
lcc3-re0:
```

```

-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:44 PDT (17:44:46 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:08 PDT (17:44:22 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

```

show system uptime (QFX Series)

```

user@switch> show system uptime
Current time: 2010-08-27 03:12:30 PDT
System booted: 2010-08-13 17:11:54 PDT (1w6d 10:00 ago)
Protocols started: 2010-08-13 17:13:56 PDT (1w6d 09:58 ago)
Last configured: 2010-08-26 05:54:00 PDT (21:18:30 ago) by regress
3:12AM up 13 days, 10:01, 3 users, load averages: 0.00, 0.00, 0.00

```


CHAPTER 5

System Services

- [Overview on page 875](#)
- [Configuration on page 911](#)
- [Administration on page 1157](#)
- [Troubleshooting on page 1245](#)

Overview

- [DHCP Local Server on page 875](#)
- [DHCP Relay Agent on page 899](#)

DHCP Local Server

- [Extended DHCP Local Server Overview on page 876](#)
- [DHCPv6 Local Server Overview on page 880](#)
- [DHCP Local Server Handling of Client Information Request Messages on page 881](#)
- [DHCP Duplicate Client Differentiation Using Client Subinterface Overview on page 882](#)
- [Group-Specific DHCP Local Server Options on page 883](#)
- [Understanding Dynamic Reconfiguration of Extended DHCP Local Server Clients on page 883](#)
- [DHCP Snooping Support on page 887](#)
- [DHCP Auto Logout Overview on page 888](#)
- [Address-Assignment Pools Overview on page 890](#)
- [Use of DHCP Option 50 and DHCPv6 IA_NA Option to Request a Specific IP Address on page 891](#)
- [Multiple Address Assignment for DHCPv6 Clients on page 891](#)
- [Centrally Configured Opaque DHCP Options on page 893](#)
- [Graceful Routing Engine Switchover on page 897](#)
- [Port Number Requirements for DHCP Firewall Filters on page 898](#)

Extended DHCP Local Server Overview

You can enable the router or switch to function as an extended DHCP local server and configure the extended DHCP local server options on the router (or switch). The extended DHCP local server provides an IP address and other configuration information in response to a client request. The DHCP local server supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication or DHCP client authentication. You can configure dynamic profile and authentication support on a global basis or for a specific group of interfaces.

The extended DHCP local server enhances traditional DHCP server operation by utilizing centralized address-assignment pools. The address-assignment pools are managed independently of the DHCP local server and can be shared by different client applications.

You can also configure the extended DHCP local server to support IPv6 clients. Both DHCP local server and DHCPv6 local server support the specific address request feature, which enables you to assign a particular address to a client. See [“DHCPv6 Local Server Overview” on page 880](#) for information about the DHCPv6 local server feature.



NOTE: You cannot configure the extended DHCP local server and extended DHCP relay on the same interface.

To configure the extended DHCP local server on the router (or switch), you include the **dhcp-local-server** statement at the **[edit system services]** hierarchy level. See the *[edit system services dhcp-local-server] Hierarchy Level* for the complete DHCP local server syntax.

This overview covers:

- [Interaction Among the DHCP Client, Extended DHCP Local Server, and Address-Assignment Pools on page 876](#)
- [Providing DHCP Client Configuration Information on page 877](#)
- [Minimal Configuration for Clients on page 878](#)
- [DHCP Local Server and Address-Assignment Pools on page 878](#)
- [DHCP Liveness Detection on page 879](#)

Interaction Among the DHCP Client, Extended DHCP Local Server, and Address-Assignment Pools

The pattern of interaction between the DHCP local server, the DHCP client, and address-assignment pools is the same regardless of whether the software installation is on a router or a switch. Technically, the codes operates in the same manner, regardless of the hardware platform. However, there are some difference in the details of usage.

- On routers—In a typical carrier edge network configuration, the DHCP client is on the subscriber's computer, and the DHCP local server is configured on the router.

- On switches—In a typical network configuration, the DHCP client is on an access device, such as a personal computer, and the DHCP local server is configured on the switch.

The following steps provide a high-level description of the interaction among the DHCP local server, DHCP client, and address-assignment pools:

1. The DHCP client sends a discover packet to one or more DHCP local servers in the network to obtain configuration parameters and an IP address for the subscriber (or DHCP client).
2. Each DHCP local server that receives the discover packet then searches its address-assignment pool for the client address and configuration options. Each local server creates an entry in its internal client table to keep track of the client state, then sends a DHCP offer packet to the client.
3. On receipt of the offer packet, the DHCP client selects the DHCP local server from which to obtain configuration information and sends a request packet indicating the DHCP local server selected to grant the address and configuration information.
4. The selected DHCP local server sends an acknowledgement packet to the client that contains the client address lease and configuration parameters. The server also installs the host route and ARP entry, and then monitors the lease state.

Providing DHCP Client Configuration Information

When the extended DHCP application receives a response from an external authentication server, the response might include information in addition to the IP address and subnet mask. The extended DHCP application uses the information from the authentication grant for the response the DHCP application sends to the DHCP client. The DHCP application can either send the information in its original form or the application might merge the information with local configuration specifications. For example, if the authentication grant includes an address pool name and a local configuration specifies DHCP attributes for that pool, the extended DHCP application merges the authentication results and the attributes in the reply that the server sends to the client.

A local configuration is optional — a client can be fully configured by the external authentication service. However, if the external authentication service does not provide client configuration, you must configure the local address-assignment pool to provide the configuration for the client. When a local configuration specifies options, the extended DHCP application adds the local configuration options to the offer PDU the server sends to the client. If the two sets of options overlap, the options in the authentication response from the external service take precedence.

When you use RADIUS to provide the authentication, the additional information might be in the form of RADIUS attributes and Juniper Networks VSAs. [Table 67 on page 878](#) lists the information that RADIUS might include in the authentication grant. See *RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework* for a complete list of RADIUS attributes and Juniper Networks VSAs that the extended DHCP applications supports for subscriber access management or DHCP management.

Table 67: Information in Authentication Grant

Attribute Number	Attribute Name	Description
RADIUS attribute 8	Framed-IP-Address	Client IP address
RADIUS attribute 9	Framed-IP-Netmask	Subnet mask for client IP address (DHCP option 1)
Juniper Networks VSA 26-4	Primary-DNS	Primary domain server (DHCP option 6)
Juniper Networks VSA 26-5	Secondary-DNS	Secondary domain server (DHCP option 6)
Juniper Networks VSA 26-6	Primary-WINS	Primary WINS server (DHCP option 44)
Juniper Networks VSA 26-7	Secondary-WINS	Secondary WINS server (DHCP option 44)
RADIUS attribute 27	Session-Timeout	Lease time
RADIUS attribute 88	Framed-Pool	Address assignment pool name
Juniper Networks VSA 26-109	DHCP-Guided-Relay-Server	DHCP relay server

Minimal Configuration for Clients

The extended DHCP local server provides a minimal configuration to the DHCP client if the client does not have DHCP option 55 configured. The server provides the subnet mask of the address-assignment pool that is selected for the client. In addition to the subnet mask, the server provides the following values to the client if the information is configured in the selected address-assignment pool:

- **router**—A router (or switch) located on the client's subnet. This statement is the equivalent of DHCP option 3.
- **domain name**—The name of the domain in which the client searches for a DHCP server host. This is the default domain name that is appended to hostnames that are not fully qualified. This is equivalent to DHCP option 15.
- **domain name server**—A Domain Name System (DNS) name server that is available to the client to resolve hostname-to-client mappings. This is equivalent to DHCP option 6.

DHCP Local Server and Address-Assignment Pools

In the traditional DHCP server operation, the client address pool and client configuration information reside on the DHCP server. With the extended DHCP local server, the client address and configuration information reside in centralized address-assignment pools, which are managed independently of the DHCP local server and which can be shared by different client applications.

The extended DHCP local server also supports advanced pool matching and the use of named address ranges. You can also configure the local server to use DHCP option 82 information in the client PDU to determine which named address range to use for a particular client. The client configuration information, which is configured in the address-assignment pool, includes user-defined options, such as boot server, grace period, and lease time.

Configuring the DHCP environment that includes the extended DHCP local server requires two independent configuration operations, which you can complete in any order. In one operation, you configure the extended DHCP local server on the router and specify how the DHCP local server determines which address-assignment pool to use. In the other operation, you configure the address-assignment pools used by the DHCP local server. The address-assignment pools contain the IP addresses, named address ranges, and configuration information for DHCP clients. See *Configuring Address-Assignment Pools* for details about creating and using address-assignment pools.



NOTE: The extended DHCP local server and the address-assignment pools used by the server must be configured in the same logical system and routing instance.

DHCP Liveness Detection

Liveness detection for DHCP subscriber IP (or DHCP client IP) sessions utilizes an active liveness detection protocol to institute liveness detection checks for relevant clients. Clients are expected to respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, then the liveness detection check fails and a failure action is implemented. You can configure



NOTE: DHCP liveness detection either globally or per DHCP group.

Related Documentation

- [Configuring Address-Assignment Pools](#)
- [Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 949](#)
- [Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview](#)
- [Using External AAA Authentication Services with DHCP on page 923](#)
- [Use of DHCP Option 50 and DHCPv6 IA_NA Option to Request a Specific IP Address on page 891](#)
- [Graceful Routing Engine Switchover on page 897](#)
- [Subscriber Management Unified ISSU Support](#)
- [Tracing Extended DHCP Operations on page 1245](#)
- [Verifying and Managing DHCP Local Server Configuration on page 1158](#)

- [Example: Minimum Extended DHCP Local Server Configuration on page 912](#)
- [Example: Extended DHCP Local Server Configuration with Optional Pool Matching on page 912](#)
- [Example: Configuring a DHCP Firewall Filter to Protect the Routing Engine](#)

DHCPv6 Local Server Overview

The DHCPv6 local server enhances the extended DHCP local server by providing support for IPv6. When a DHCPv6 client logs in, the DHCPv6 local server uses the AAA service framework to interact with the RADIUS server. The RADIUS server, which is configured independently of DHCP, authenticates the client and supplies the IPv6 prefix and client configuration parameters.

You can configure DHCPv6 local server to communicate the following attributes to the AAA service framework and RADIUS at login time:

- Client username
- Client password



NOTE: The client username, which uniquely identifies a subscriber or a DHCP client, must be present in the configuration in order for DHCPv6 local server to use RADIUS authentication.

Based on the attributes that the DHCPv6 local server provides, RADIUS returns the information listed in [Table 68 on page 880](#) to configure the client:

Table 68: RADIUS Attributes and VSAs for DHCPv6 Local Server

Attribute Number	Attribute Name	Description
27	Session-Timeout	Lease time, in seconds. If not supplied, the lease does not expire
123	Delegated-IPv6-Prefix	Prefix that is delegated to the client
26-143	Max-Clients-Per-Interface	Maximum number of clients allowed per interface

The DHCPv6 local server is compatible with the extended DHCP local server and the extended DHCP relay agent, and can be enabled on the same interface as either the extended DHCP local server or DHCP relay agent.

The DHCPv6 local server provides many of the same features as the extended DHCP local server, including:

- Configuration for a specific interface or for a group of interfaces
- Site-specific usernames and passwords
- Numbered Ethernet interfaces
- Statically configured CoS and filters
- AAA directed login
- Use of the IA_NA option to assign a specific address to a client

To configure the extended DHCPv6 local server on the router (or switch), you include the **dhcpv6** statement at the **[edit system services dhcp-local-server]** hierarchy level. See the *[edit system services dhcp-local-server] Hierarchy Level* for the complete DHCP local server syntax, including the DHCPv6 syntax.

You can also include the **dhcpv6** statement at the following hierarchy levels:

- **[edit logical-systems *logical-system-name* system services dhcp-local-server]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server]**
- **[edit routing-instances *routing-instance-name* system services dhcp-local-server]**

Related Documentation

- [Extended DHCP Local Server Overview on page 876](#)
- [Using External AAA Authentication Services with DHCP on page 923](#)
- [Grouping Interfaces with Common DHCP Configurations on page 926](#)
- [Group-Specific DHCP Local Server Options on page 883](#)
- [Overriding Default DHCP Local Server Configuration Settings on page 928](#)
- [Configuring Passwords for Usernames on page 945](#)
- [Creating Unique Usernames for DHCP Clients on page 946](#)
- [Use of DHCP Option 50 and DHCPv6 IA_NA Option to Request a Specific IP Address on page 891](#)
- [Verifying and Managing DHCPv6 Local Server Configuration on page 1158](#)
- [Example: Extended DHCPv6 Local Server Configuration](#)

DHCP Local Server Handling of Client Information Request Messages

DHCP clients that already have externally provided addresses may solicit further configuration information from a DHCP server by sending a DHCP information request that indicates what information is desired. By default, DHCP local server and DHCPv6 local server ignore any DHCP information requests that they receive. You can override this default behavior to enable processing of these messages. Include the **process-inform**

statement at the **[edit system services dhcp-local-server overrides]** or **[edit system services dhcp-local-server dhcpv6 overrides]** hierarchy level.

By default, DHCP relay and DHCP relay proxy automatically forward DHCP information request messages without modification if the messages are received on an interface configured for a DHCP server group. DHCP relay and relay proxy drop information request messages received on any other interfaces. You cannot disable this default DHCP relay and relay proxy behavior.

The information requested by these clients has typically been configured with the **dhcp-attributes** statement for an address pool defined by the **address-assignment pool pool-name** statement at the **[edit access]** hierarchy level.

When you enable processing of DHCP information requests, you can optionally specify the name of the pool from which the local server retrieves the requested configuration information for the client. If you do not specify a local pool, then the local server requests that AAA selects and returns only the name of the relevant pool.

DHCP local server responds to the client with a DHCP acknowledgment message that includes the requested information—if it is available. DHCPv6 local server responds in the same manner but uses a DHCP reply message. No subscriber management or DHCP-management is applied as a result of the DHCP information request message.



NOTE: PPP interfaces are not supported on EX Series switches.

When DHCPv6 is configured over PPP interfaces, the PPP RADIUS authentication data can be used to select the pool from which the response information is taken. Additionally other RADIUS attributes can also be inserted into the DHCPv6 reply message. If an overlap exists between RADIUS attributes and local pool attributes, the RADIUS values are used instead of the local configuration data. If no RADIUS information is received from the underlying PPP interface, then the behavior is the same as described previously for non-PPP interfaces.

**Related
Documentation**

- [Overriding Default DHCP Local Server Configuration Settings on page 928](#)
- [Enabling Processing of Client Information Requests on page 933](#)

DHCP Duplicate Client Differentiation Using Client Subinterface Overview

In some network environments, client IDs and MAC addresses might not be unique, resulting in duplicate clients. For example, two network adapters might be manufactured with the same hardware address, resulting in a duplicate MAC address among the DHCP clients attached to the router (or switch). A duplicate DHCP client occurs when a client attempts to get a lease, and that client has the same client ID or the same MAC address as an existing DHCP client.

When DHCP server receives a request from a new client that has a duplicate ID or MAC address, DHCP server terminates the address lease for the existing client and returns the address to its original address pool. DHCP server then assigns a new address and lease to the new client.

By default, both DHCP local server and DHCP relay use the subnet information to differentiate between duplicate clients. However, in some cases, this level of differentiation is not adequate. For example, when multiple subinterfaces share the same underlying loopback interface with the same preferred source address, the interfaces appear to be on the same subnet. In this situation, the default configuration prevents duplicate clients.

You can provide greater differentiation between duplicate clients by configuring DHCP to consider the client subinterface when duplicate clients occur. In this optional configuration, DHCP uniquely identifies:

- The subnet on which the client resides
- The subinterface on which the client resides
- The client within the subnet

**Related
Documentation**

- [Configuring DHCP Duplicate Client Support on page 925](#)
- [Guidelines for Configuring Support for DHCP Duplicate Clients on page 925](#)

Group-Specific DHCP Local Server Options

You can include the following statements at the **[edit system services dhcp-local-server group group-name]** hierarchy level to set group-specific DHCP local server configuration options, and at the **[edit system services dhcp-local-server]** hierarchy level to set global DHCP local server configuration options. Statements configured at the **[edit system services dhcp-local-server group group-name]** hierarchy level apply only to the named group of interfaces, and override any global DHCP local server settings configured with the same statements at the **[edit system services dhcp-local-server]** hierarchy level.

DHCPv6 local server supports the same set of statements with the exception of the **dynamic-profile** statement.

- **authentication**—Configure the parameters the router sends to the external AAA server.
- **dynamic-profile**—Specify the dynamic profile that is attached to a group of interfaces.
- **interface**—Specify one or more interfaces, or a range of interfaces, that are within the specified group.
- **overrides**—Override the default configuration settings for the extended DHCP local server. For information, see “[Overriding Default DHCP Local Server Configuration Settings](#)” on page 928.

**Related
Documentation**

- [Grouping Interfaces with Common DHCP Configurations on page 926](#)

Understanding Dynamic Reconfiguration of Extended DHCP Local Server Clients

Dynamic reconfiguration of clients enables the extended DHCP local server to initiate a client update without waiting for the client to initiate a request.

Default Client/Server Interaction

Typically the DHCP client initiates all of the basic DHCP client/server interactions. The DHCP server sends information to a client only in response to a request from that client. This behavior does not enable a client to be quickly updated with its network address and configuration in the event of server changes:



NOTE: Technically, the DHCP client/server interactions are the same on routers and switches. However, the primary usage of this technology on the routers is for subscriber management. The switches are not used for subscriber management. Therefore, this topic provides two sample scenarios. The actions are the same, but the implementation details are different.

- On routers—Suppose a service provider restructures its addressing scheme or changes the server IP addresses that it provided to clients. Without dynamic reconfiguration, the service provider typically clears the DHCP server binding table, but cannot inform the DHCP clients that their bindings have been cleared. Consequently, the DHCP client operates as though its IP address is still valid, but it is now unable to communicate over the access network, resulting in an outage. The DHCP local server needs to wait for the client to send a message to renew its lease or rebind to the server. In response, the server sends a NAK message to the client to force it to begin the DHCP connection process again. Alternatively, the provider can wait for customers to make a service call about the network failures and then instruct them to power cycle their customer premises equipment to reinitiate the connection. Neither of these actions is timely or convenient for customers.
- On switches—Suppose you restructure the addressing scheme or change the server IP addresses that the DHCP server provides to clients. Without dynamic reconfiguration, the network typically clears the DHCP server binding table, but cannot inform the DHCP clients that their bindings have been cleared. Consequently, the DHCP client operates as though its IP address is still valid, but it is now unable to communicate over the access network, resulting in an outage. The DHCP local server needs to wait for the client to send a message to renew its lease or rebind to the server. In response, the server sends a NAK message to the client to force it to begin the DHCP connection process again. Alternatively, you can wait for users to notify you of the network failures and then instruct them to power cycle their equipment to reinitiate the connection. Neither of these actions is timely or convenient for users.

Dynamic Client/Server Interaction for DHCPv4

Dynamic reconfiguration for DHCPv4 is available through a partial implementation of RFC 3203, *DHCP Reconfigure Extension* for DHCPv4. It enables the DHCPv4 local server to send a message to the client to force reconfiguration.

The server sends a `forcerenew` message to a DHCPv4 client, initiating a message exchange. In response, DHCPv4 clients that support the `forcerenew` message then send a lease renewal message to the server. The server rejects the lease renewal request and sends a NAK to the client, causing the client to reinitiate the DHCP connection. A successful reconnection results in the reconfiguration of the DHCP client. Only the

exchange of `forcerenew`, `renew`, and `NAK` messages is supported from RFC 3202. DHCP relay and DHCP relay proxy do not participate in the client reconfiguration or react to `forcerenew` messages other than to forward them to the client.

When the local server state machine starts the reconfiguration process on a bound client, the client transitions to the reconfiguring state and the local server sends a `forcerenew` message to the client. Because the client was in the bound state before entering the reconfiguring state, all subscriber services or DHCP-managed services, such as forwarding and statistics, continue to work. Client statistics are not maintained in the interval between a successful reconfiguration and the subsequent client binding. When the server responds to the client renewal request with a `NAK`, the client entry is removed from the binding table and final statistics are reported. New statistics are collected when the client sends a `discover` message to establish a new session.

Dynamic Client/Server Interaction for DHCPv6

Dynamic reconfiguration for DHCPv6 is available through a partial implementation of RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*. It enables the DHCPv6 local server to send a message to the client to force reconfiguration.

DHCPv6 servers send `reconfigure` messages to DHCPv6 clients, initiating a message exchange. In response, DHCPv6 clients that support the `reconfigure` message transition to the renewing state and send a `renew` message to the server. The server returns a reply message with a lifetime of zero (0). The client transitions to the init state and sends a `solicit` message. The server sends an `advertise` message to indicate that it is available for service. The client sends a request for configuration parameters, which the server then includes in its reply. DHCP relay and DHCP relay proxy do not participate in the client reconfiguration or react to `reconfigure` messages other than to forward them to the client.

When a DHCPv6 server is triggered to initiate reconfiguration on a bound DHCPv6 client, the client transitions to the `reconfigure` state. All subscriber services, such as forwarding and statistics, continue to work. The server then sends the `reconfigure` message to the client. If the DHCPv6 client is already in the `reconfigure` state, the DHCPv6 server ignores the reconfiguration trigger. For clients in any state other than bound or `reconfigure`, the server clears the binding state of the client, as if the `clear dhcpv6 server binding` command had been issued.

Dynamic Configuration Options

You can enable dynamic reconfiguration for all DHCP clients or only the DHCP clients serviced by a specified group of interfaces, and you can modify the behavior accordingly.

- To enable dynamic reconfiguration with default reconfiguration values for all DHCP clients, include the `reconfigure` statement at the `[edit system services dhcp-local-server]` hierarchy level for DHCPv4 clients, and at the `[edit system services dhcp-local-server dhcpv6]` hierarchy level for DHCPv6 clients.
- Alternatively, to enable dynamic reconfiguration for only the DHCP clients serviced by a specified group of interfaces, include the `reconfigure` statement at the `[edit system services dhcp-local-server group group-name]` hierarchy level for DHCPv4 clients, and at the `[edit system services dhcp-local-server dhcpv6 group group-name]` hierarchy level for DHCPv6 clients.

You can optionally modify the behavior of the reconfiguration process by including the appropriate statements at the **[edit system services dhcp-local-server reconfigure]** hierarchy level for all DHCPv4 clients, and at the **[edit system services dhcp-local-server dhcpv6 reconfigure]** hierarchy level for all DHCPv6 clients. To override this global configuration for only the DHCP clients serviced by a specified group of interfaces, you can include the statements with different values at the **[edit system services dhcp-local-server group group-name reconfigure]** hierarchy level for DHCPv4 clients, and at the **[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]** hierarchy level for DHCPv6 clients.

Include the **attempts** statement to specify how many times the local server sends the **forcerenew** or **reconfigure** message to initiate client reconfiguration. Include the **timeout** statement to set the interval between the first and second attempts. The interval between each subsequent attempt doubles the previous value. For example, if the first value is 2, the first retry is attempted 2 seconds after the first attempt fails. The second retry is attempted 4 seconds after the first retry fails. The third retry is attempted 8 seconds after the second retry fails, and so on.

By default, the DHCP client's original configuration is restored if all of the reconfiguration attempts fail. Include the **clear-on-abort** statement to delete the client instead.

You can configure an authentication token by including the **token** statement. The DHCP local server then includes this token inside the authentication option when it sends **forcerenew** or **reconfigure** messages. If the service provider has previously configured the DHCP client with this token, then the client can compare that token against the newly received token, and reject the message if the tokens do not match. This functionality corresponds to RFC 3118, *Authentication for DHCP Messages*, section 4.

In the event of a RADIUS-initiated disconnect (RID), the client is deleted by default. You can configure the client to be reconfigured instead of deleted by including the **radius-disconnect** statement. The client is deleted if all attempts to reconfigure the client fail.

For the DHCPv6 server only, you can include the **strict** statement. By default, the server accepts solicit messages from clients that do not support server-initiated reconfiguration. Including this statement causes the server to discard solicit messages from nonsupporting clients; consequently the server does not bind these clients.

You can force the local server to initiate the reconfiguration process for clients by issuing the **request dhcp server reconfigure** command for DHCPv4 clients, and the **request dhcpv6 server reconfigure** command for DHCPv6 clients. Command options determine whether reconfiguration is then attempted for all clients or specified clients.

Events that take place while a reconfiguration is in process take precedence over the reconfiguration. [Table 69 on page 887](#) lists the actions taken in response to several different events.

Table 69: Action Taken for Events That Occur During a Reconfiguration

Event	Action
Server receives a discover (DHCPv4) or solicit (DHCPv6) message from the client.	Server drops packet and deletes client.
Server receives a request, renew, rebind, or init-reboot message from the client.	DHCPv4—Server sends NAK message and deletes client. DHCPv6—Server drops packet and deletes client. Server replies to renew message with lease time of zero (0).
Server receives a release or decline message from the client.	Server deletes client.
The client lease times out.	Server deletes client.
The clear dhcp server binding command is issued.	Server deletes client.
The request dhcp server reconfigure (DHCPv4) or request dhcpv6 server reconfigure (DHCPv6) command is issued.	Command is ignored.
GRES or DHCP restart occurs.	Reconfiguration process is halted.

Related Documentation

- [Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 936](#)

DHCP Snooping Support

DHCP snooping provides DHCP security on the router by filtering incoming messages. When DHCP snooping is enabled, the router differentiates between trusted and untrusted interfaces, and forwards messages from trusted sources while rejecting the untrusted messages.

In Junos OS, DHCP snooping is enabled in a routing instance when you configure either the **dhcp-relay** statement at the **[edit forwarding-options]** hierarchy level, or the **dhcp-local-server** statement at the **[edit system services]** hierarchy level in that routing instance. However, depending on the Junos OS release, the router processes the snooped packets differently, as described in the following list:

- In Junos OS Release 10.0 and earlier, the router processes snooped packets normally.
- In Junos OS Release 10.1 and later, the router discards snooped packets by default. To enable normal processing of snooped packets in Junos OS Release 10.1 and later, you must explicitly configure the **allow-snooped-clients** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level.

You can configure DHCP snooping support for the following:

- DHCPv4 relay agent—Override the router's (or switch's) default snooping configuration and specify that DHCP snooping is enabled or disabled globally, for a named group of interfaces, or for a specific interface within a named group.

In a separate procedure, you can set a global configuration to specify whether the DHCPv4 relay agent forwards or drops snooped packets for all interfaces, only configured interfaces, or only nonconfigured interfaces. The router also uses the global DHCP relay agent snooping configuration to determine whether to forward or drop snooped BOOTREPLY packets.

- DHCPv6 relay agent—As you can with snooping support for the DHCPv4 relay agent, you can override the default DHCPv6 relay agent snooping configuration on the router to explicitly enable or disable snooping support globally, for a named group of interfaces, or for a specific interface with a named group of interfaces.

In multi-relay topologies where more than one DHCPv6 relay agent is between the DHCPv6 client and the DHCPv6 server, snooping enables intervening DHCPv6 relay agents between the client and the server to correctly receive and process the unicast traffic from the client and forward it to the server. The DHCPv6 relay agent snoops incoming unicast DHCPv6 packets by setting up a filter with UDP port 547 (the DHCPv6 UDP server port) on a per-forwarding table basis. The DHCPv6 relay agent then processes the packets intercepted by the filter and forwards the packets to the DHCPv6 server.

Unlike the DHCPv4 relay agent, the DHCPv6 relay agent does not support global configuration of forwarding support for DHCPv6 snooped packets.

- DHCP local server—Configure whether DHCP local server forwards or drops snooped packets for all interfaces, only configured interfaces, or only nonconfigured interfaces.

**Related
Documentation**

- [Configuring DHCP Snooping for DHCP Relay Agent on page 963](#)
- [Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server on page 944](#)
- [Example: Configuring DHCP Snooping Support for DHCP Relay Agent on page 921](#)

DHCP Auto Logout Overview

This topic provides an introduction to the optional DHCP auto logout feature and includes the following sections:

- [Auto Logout Overview on page 888](#)
- [How DHCP Identifies and Releases Clients on page 889](#)
- [Option 60 and Option 82 Requirements on page 890](#)

Auto Logout Overview

Auto logout is an optional configuration for DHCP local server and DHCP relay agent that improves the efficiency of DHCP IP address assignment. Auto logout enables IP addresses to be immediately released and returned to the address pool when the addresses are no

longer used by DHCP clients. DHCP can then assign the addresses to other clients. Without auto logout, an IP address is blocked for the entire lease period, and DHCP must wait until the address lease time expires before reusing the address.

Auto logout is particularly useful when DHCP uses long lease times for IP address assignments and to help avoid allocating duplicate IP addresses for a single client.

For example (on the routers), you might have an environment that includes set-top boxes (STB) that are often upgraded or replaced. Each time a STB is changed, the new STB repeats the DHCP discover process to obtain client configuration information and an IP address. DHCP views the new STB as a completely new client and assigns a new IP address—the previous IP address assigned to the client (the old STB) remains blocked and unavailable until the lease expires. If auto logout is configured in this situation, DHCP recognizes that the new STB is actually the same client and then immediately releases the original IP address. DHCP relay agent acts as a proxy client for auto logout and sends a DHCP release message to the DHCP server.

How DHCP Identifies and Releases Clients

The auto logout feature requires that DHCP explicitly identify clients. By default, DHCP local server and DHCP relay agent identify clients based on MAC address or Client Identifier. However, in some cases this type of identification might not be sufficient. For example, in the previous STB example, each STB has a different MAC address, so DHCP incorrectly assumes that an upgraded or replacement STB is a new client.

In order to explicitly identify clients, auto logout uses a secondary identification method when the primary identification method is unsuccessful—the primary method is considered unsuccessful if the MAC address or Client Identifier does not match that of an existing client. The secondary identification method is based on the DHCP option 60 and option 82 information in DHCP discover messages.

Both the primary and secondary identification methods use subnet information to differentiate between clients. The primary identification method differentiates between two clients with the same MAC address (or same Client Identifier) if the clients are on different subnets. Similarly, the secondary identification method considers two clients as different if they have the same option 60 and option 82 information, but different subnets.

DHCP local server and DHCP relay agent perform the following operations when auto logout is enabled and the secondary identification method identifies a duplicate client (that is, the discover packet is from an existing client).

- DHCP local server immediately releases the existing address.
- DHCP relay agent immediately releases the existing client and then sends a DHCP release packet to the DHCP server. Sending the release packet ensures that DHCP relay and the DHCP server are synchronized.

If the DHCP relay receives a DISCOVER message from an existing client, the DHCP relay forwards the DISCOVER message to the DHCP server. The DHCP relay preserves the binding if the client's existing IP address is returned by the DHCP server. This behavior is not applicable if the proxy-mode override or client-discover-match functionality are enabled.



NOTE: If the DHCP relay agent is in snoop mode, DHCP relay releases the client but does not send a release packet to the DHCP server if the discover packet is for a passive client (a client added as a result of snooped packets) or if the discover packet is a snooped packet.

Option 60 and Option 82 Requirements

DHCP local server requires that the received discover packet include both DHCP option 60 and option 82. If either option is missing, DHCP local server cannot perform the secondary identification method and auto logout is not used.

DHCP relay agent requires that the received discover packet contain DHCP option 60. DHCP relay determines the option 82 value based on the guidelines provided in [“DHCP Relay Agent Option 82 Value for Auto Logout” on page 962](#).

Related Documentation

- [Automatically Logging Out DHCP Clients on page 932](#)
- [DHCP Relay Agent Option 82 Value for Auto Logout on page 962](#)
- [Clearing DHCP Bindings for Subscriber Access](#)

Address-Assignment Pools Overview

The address-assignment pool feature supports subscriber management and DHCP management functionality by enabling you to create IPv4 and IPv6 address pools that different client applications can share. For example, multiple client applications, such as DHCP, can use an address-assignment pool to provide addresses for their particular clients. Client applications can acquire addresses for either authenticated or unauthenticated clients.

Address-assignment pools support both dynamic and static address assignment. In dynamic address assignment, a client is automatically assigned an address from the address-assignment pool. In static address assignment, which is supported for IPv4 pools only, you reserve an address that is then always used by a particular client. Addresses that are reserved for static assignment are removed from the dynamic address pool and cannot be assigned to other clients.

You can configure named address ranges within an address-assignment pool. A named range is a subset of the overall address range. A client application can use named ranges to manage address assignment based on client-specific criteria. For example, for IPv4 address-assignment pools, you might create a named range that is based on a specific DHCP option 82 value. Then, when a DHCP client request matches the specified option 82 value, an address from the specified range is assigned to the client.

You can link address-assignment pools together to provide backup pools for address assignment. When the primary pool is fully allocated, the router or switch automatically switches to the linked, or secondary, pool and begins allocating addresses from that pool.

You can also explicitly identify that an address-assignment pool is used for ND/RA.

Related Documentation

- [Configuring Address-Assignment Pools](#)
- [DNS Address Assignment Precedence](#)
- [Address-Assignment Pools Licensing Requirements](#)
- [Example: Configuring an Address-Assignment Pool](#)
- [Configuring an Extended DHCP Server with DHCPv6 on EX Series Switches \(CLI Procedure\)](#)

Use of DHCP Option 50 and DHCPv6 IA_NA Option to Request a Specific IP Address

Subscriber management or DHCP management enables you to specify that DHCP local server assign a particular address to a client. For example, if a client is disconnected, you might use this capability to assign the same address that the client was using prior to being disconnected. If the requested address is available, DHCP assigns it to the client. If the address is unavailable, the DHCP local server offers another address, based on the address allocation process.

Both DHCP local server and DHCPv6 local server support the specific address request feature. DHCP local server uses DHCP option 50 in DHCP DISCOVER messages to request a particular address, while DHCPv6 local server uses the IA_NA option (Identity Association for Non-Temporary Addresses) in DHCPv6 SOLICIT messages.



NOTE: Subscriber management (DHCP management) supports only one address for each of the DHCPv6 IA_NA or IA_PD address types. If the DHCPv6 client requests more than one address for a given type, the DHCPv6 local server uses only the first address and ignores the other addresses.

Multiple Address Assignment for DHCPv6 Clients

Subscriber management (on the routers) or DHCP management (on the switches) enables you to assign multiple addresses to a single DHCPv6 client. Multiple address support is enabled by default, and is activated when the DHCPv6 local server receives a DHCPv6 Solicit message from a subscriber (or DHCP client) that contains multiple addresses.

For example, if you are implementing this feature on the routers, you might use the multiple address assignment feature in a networking environment in which a customer premises equipment (CPE) device requires a host address and a delegated prefix. In such an environment, you can configure subscriber management to assign both a DHCPv6 IA_NA (Identity Association for Non-Temporary Addresses) and an IA_PD (Identity Association for Prefix Delegation) address to the client (the CPE device).

- [Multiple Address Assignment Using Local Address Pools or RADIUS on page 892](#)
- [Junos OS Predefined Variable for Multiple DHCPv6 Address Assignment on page 892](#)

Multiple Address Assignment Using Local Address Pools or RADIUS

You can use either local address pools or RADIUS when assigning multiple addresses to a DHCP client. When at least one address is successfully allocated, the router or switch creates a subscriber (or DHCP client) entry and binds the entry to the assigned address. If both addresses are successfully allocated, the router (or switch) creates a single subscriber (or DHCP client) entry and binds both addresses to that entry.

You can also configure a delegated address pool, which explicitly specifies the address pool that subscriber management (or DHCP management) uses to assign IPv6 prefixes for subscribers (or DHCP clients).

Junos OS Predefined Variable for Multiple DHCPv6 Address Assignment



NOTE: EX Series switches do not support demux.

(On the routers only) Subscriber management provides a predefined variable that you can use to dynamically configure DHCPv6 multiple address assignment. You apply the Junos OS predefined variable, **\$junos-subscriber-ipv6-multi-address**, as a demux source address in a dynamic profile. When the dynamic profile is attached to a subscriber, the variable is expanded to include both the host and prefix addresses. You use this variable instead of the **\$junos-subscriber-ipv6-address** variable, which supports a single IPv6 address.

You include the **\$junos-subscriber-ipv6-multi-address** variable at the **[edit dynamic-profile profile-name interfaces interface-name unit logical-unit-number family inet6 demux-source]** hierarchy level.

Related Documentation

- [Specifying the Delegated Address Pool for IPv6 Prefix Assignment on page 934](#)
- [Junos OS Predefined Variables](#)

Centrally Configured Opaque DHCP Options

Subscriber management (on the routers) or DHCP management (on the switches) enables you to centrally configure DHCP options on a RADIUS server and then distribute the options on a per-subscriber or per DHCP-client basis. This method results in RADIUS-sourced DHCP options—the DHCP options originate at the RADIUS server and are sent to the subscriber (or DHCP client). This differs from the traditional client-sourced method (also called DHCP-sourced) of configuring DHCP options, in which the options originate at the client and are sent to the RADIUS server. The subscriber management (DHCP management) RADIUS-sourced DHCP options are also considered to be *opaque*, because DHCP local server performs minimal processing and error checking for the DHCP options string before passing the options to the subscriber (DHCP client).

Subscriber management (or DHCP management) uses Juniper Networks VSA 26-55 (DHCP-Options) to distribute the RADIUS-sourced DHCP options. The RADIUS server includes VSA 26-55 in the Access-Accept message that the server returns during subscriber authentication or DHCP client authentication. The RADIUS server sends the Access-Accept message to the RADIUS client, and then on to DHCP local server for return to the DHCP subscriber. The RADIUS server can include multiple instances of VSA 26-55 in a single Access-Accept message. The RADIUS client concatenates the multiple instances and uses the result as a single instance.

There is no CLI configuration required to enable subscriber management (DHCP management) to use the centrally configured DHCP options—the procedure is triggered by the presence of VSA 26-55 in the RADIUS Access-Accept message.

When building the offer packet for the DHCP client, DHCP local server uses the following sequence:

1. Processes any RADIUS-configured parameters that are passed as separate RADIUS attributes; for example, RADIUS attribute 27 (Session Timeout).
2. Processes any client-sourced parameters; for example, RADIUS attributes 53 (DHCP Message Type) and 54 (Server Identifier).
3. Appends (without performing any processing) the opaque DHCP options string contained in the VSA 26-55 received from the RADIUS server.

In addition to supporting central configuration of DHCP options directly on the RADIUS server (RADIUS-sourced options), subscriber management (DHCP management) also supports the traditional client-sourced options configuration, in which the router's (switch's) DHCP component sends the options to the RADIUS server. The client-sourced DHCP options method is supported for both DHCP local server and DHCP relay agent; however, the RADIUS-sourced central configuration method is supported on DHCP local server only. Both the RADIUS-sourced and client-sourced methods support DHCPv4 and DHCPv6 subscribers (clients).




NOTE: You can use the RADIUS-sourced and client-sourced methods simultaneously on DHCP local server. However, you must ensure that the central configuration method does not include options that override client-sourced DHCP options, because this can create unpredictable results.

- [Data Flow for RADIUS-Sourced DHCP Options on page 895](#)
- [Multiple VSA 26-55 Instances Configuration on page 896](#)
- [DHCP Options That Cannot Be Centrally Configured on page 896](#)

Data Flow for RADIUS-Sourced DHCP Options

Figure 4 on page 895 shows the procedure subscriber management (DHCP management) uses when configuring DHCP options for subscribers (DHCP clients).

Figure 4: DHCP Options Data Flow



The following general sequence describes the data flow when subscriber management (DHCP management) uses RADIUS-sourced DHCP options and VSA 26-55 to configure a DHCP subscriber (client):

1. The subscriber (DHCP client) sends a DHCP discover message (or DHCPv6 solicit message) to the DHCP local server. The message includes client-sourced DHCP options.
2. The DHCP local server initiates authentication with the Junos OS RADIUS client.
3. The RADIUS client sends an Access-Request message on behalf of the subscriber (DHCP client) to the external RADIUS server. The message includes the subscriber's (DHCP client's) client-sourced DHCP options.
4. The external RADIUS server responds by sending an Access-Accept message to the RADIUS client. The Access-Accept message includes the RADIUS-sourced opaque DHCP options in VSA 26-55.
5. The RADIUS client sends the DHCP options string to DHCP local server. If there are multiple VSA 26-55 instances, the RADIUS client first assembles them into a single options string.
6. DHCP local server processes all options into the DHCP offer (or DHCPv6 reply) message, except for the RADIUS-sourced VSA 26-55 DHCP options. After processing all other options, DHCP local server then appends the unmodified VSA 26-55 DHCP options to the message and sends the message to the subscriber (DHCP client).

7. The subscriber (DHCP client) is configured with the DHCP options.
8. The following operations occur after the subscriber (DHCP client) receives the DHCP options:
 - Accounting—The RADIUS client sends Acct-Start and Interim-Accounting requests to the RADIUS server, including the RADIUS-sourced DHCP options in VSA 26-55. By default, the DHCP options are included in accounting requests.
 - Renewal—When the subscriber (DHCP client) renews, the cached DHCP options value is returned in the DHCP renew (or DHCPv6 ACK) message. The originally assigned DHCP options cannot be modified during a renew cycle.
 - Logout—When the subscriber (DHCP client) logs out, the RADIUS client sends an Acct-Stop message to the RADIUS server, including the RADIUS-sourced VSA 26-55.

Multiple VSA 26-55 Instances Configuration

VSA 26-55 supports a maximum size of 247 bytes. If your RADIUS-sourced DHCP options field is greater than 247 bytes, you must break the field up and manually configure multiple instances of VSA 26-55 for the RADIUS server to return. When using multiple instances for an options field, you must place the instances in the packet in the order in which the fragments are to be reassembled by the RADIUS client. The fragments can be of any size of 247 bytes or less.



BEST PRACTICE: For ease of configuration and management of your DHCP options, you might want to have one DHCP option per VSA 26-55 instance, regardless of the size of the option field.

When the RADIUS client returns a reassembled opaque options field in an accounting request to the RADIUS server, the client uses 247-byte fragments. If you had originally created instances of fewer than 247 bytes, the returned fragments might not be the same as you originally configured on the RADIUS server.



NOTE: If you are configuring Steel-Belted Radius (SBR) to support multiple VSA 26-55 instances, ensure that you specify VSA 26-55 with the RO flags in the Subscriber Management RADIUS dictionary file. The R value indicates a multivalued reply attribute and the O value indicates an ordered attribute.

DHCP Options That Cannot Be Centrally Configured

Table 70 on page 896 shows the DHCP options that you must not centrally configure on the RADIUS server.

Table 70: Unsupported Opaque DHCP Options

DHCP Option	Option Name	Comments
Option 0	Pad Option	Not supported.

Table 70: Unsupported Opaque DHCP Options (*continued*)

DHCP Option	Option Name	Comments
Option 51	IP Address Lease Time	Value is provided by RADIUS attribute 27 (Session-Timeout).
Option 52	Option Overload	Not supported.
Option 53	DHCP Message Type	Value is provided by DHCP local server.
Option 54	Server Identifier	Value is provided by DHCP local server.
Option 55	Parameter Request List	Value is provided by DHCP local server.
Option 255	End	Value is provided by DHCP local server.
–	DHCP magic cookie	Not supported.

**Related
Documentation**

- *Monitoring DHCP Options Configured on RADIUS Servers*

Graceful Routing Engine Switchover

For EX Series switches, only extended DHCP local server maintains the state of active DHCP client leases. The DHCP local server supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication. You can configure dynamic profile and authentication support on a global basis or for a specific group of interfaces. The extended DHCP local server also supports the use of Junos address-assignment pools or external authorities, such as RADIUS, to provide the client address and configuration information.

For MX Series routers, the extended DHCP local server and the DHCP relay agent applications both maintain the state of active DHCP client leases in the session database. The extended DHCP application can recover this state if the DHCP process fails or is manually restarted, thus preventing the loss of active DHCP clients in either of these circumstances. However, the state of active DHCP client leases is lost if a power failure occurs or if the kernel stops operating (for example, when the router is reloaded) on a single Routing Engine.

You can enable graceful switchover support on both EX Series switches and MX Series routers. To enable graceful switchover support for the extended DHCP local server or extended DHCP relay agent on a switch, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level. To enable graceful Routing Engine switchover support on MX Series routers, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level. You cannot disable graceful Routing Engine

switchover support for the extended DHCP application when the router is configured to support graceful Routing Engine switchover.

For more information about using graceful Routing Engine switchover, see the *Junos OS High Availability Configuration Guide*.

**Related
Documentation**

- [Extended DHCP Local Server Overview on page 876](#)
- [Extended DHCP Relay Agent Overview on page 900](#)
- *Subscriber Management Unified ISSU Support*

Port Number Requirements for DHCP Firewall Filters

When you configure a firewall filter to perform some action on DHCP packets at the Routing Engine, such as protecting the Routing Engine by allowing only proper DHCP packets, you must specify both port 67 (bootps) and port 68 (bootpc) for both the source and destination. The firewall filter acts at both the line cards and the Routing Engine.

This requirement applies to both DHCP local server and DHCP relay, but it applies only when DHCP is provided by the `jdhcpd` process. MX Series routers, M120 routers, and M320 routers use `jdhcpd`. For DHCP relay, that means the configuration is required only at the `[edit forwarding-options dhcp-relay]` hierarchy level and not at the `[edit forwarding-options helpers bootp]` hierarchy level.

DHCP packets received on the line cards are encapsulated by `jdhcpd` with a new UDP header where their source and destination addresses are set to port 68 before being forwarded to the Routing Engine.

For DHCP relay and DHCP proxy, packets sent to the DHCP server from the router have both the source and destination UDP ports set to 67. The DHCP server responds using the same ports. However, when the line card receives these DHCP response packets, it changes both port numbers from 67 to 68 before passing the packets to the Routing Engine. Consequently the filter needs to accept port 67 for packets relayed from the client to the server, and port 68 for packets relayed from the server to the client.

Failure to include both port 67 and port 68 as described here results in most DHCP packets not being accepted.

For complete information about configuring firewall filters in general, see *Junos OS Firewall Filters and Traffic Policers Configuration Guide*.

**Related
Documentation**

- *Example: Configuring a DHCP Firewall Filter to Protect the Routing Engine*
- [Extended DHCP Local Server Overview on page 876](#)
- [Extended DHCP Relay Agent Overview on page 900](#)
- *Dynamic Firewall Filters Overview*

DHCP Relay Agent

- [Extended DHCP Relay Agent Overview on page 900](#)
- [DHCP Relay Proxy Overview on page 903](#)
- [DHCPv6 Relay Agent Overview on page 905](#)
- [DHCP Duplicate Client Differentiation Using Client Subinterface Overview on page 906](#)
- [Group-Specific DHCP Relay Options on page 906](#)
- [DHCP Snooping Support on page 907](#)
- [DHCP Auto Logout Overview on page 908](#)
- [Graceful Routing Engine Switchover on page 910](#)
- [Port Number Requirements for DHCP Firewall Filters on page 911](#)

Extended DHCP Relay Agent Overview

You can configure extended DHCP relay options on the router or on the switch and enable the router (or switch) to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.

DHCP relay supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication or DHCP client authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.



NOTE: The PTX Series Packet Transport Switches do not support authentication for DHCP relay agents.

On the routers, you can use DHCP relay in carrier edge applications such as video/IPTV to obtain configuration parameters, including an IP address, for your subscribers. For more information about how to use the DHCP relay agent in a video/IPTV application, see the *Junos OS Feature Guides*.

On the switches, you can use DHCP relay to obtain configuration parameters including an IP address for DHCP clients.



NOTE: The extended DHCP relay agent options configured with the `dhcp-relay` statement are incompatible with the DHCP/BOOTP relay agent options configured with the `bootp` statement. As a result, you cannot enable both the extended DHCP relay agent and the DHCP/BOOTP relay agent on the router at the same time.

For information about the DHCP/BOOTP relay agent, see the *Routing Policy Configuration Guide*.

You can also configure the extended DHCP relay agent to support IPv6 clients. See “[DHCPv6 Relay Agent Overview](#)” on [page 905](#) for information about the DHCPv6 relay agent feature.

To configure the extended DHCP relay agent on the router (or switch), include the `dhcp-relay` statement at the `[edit forwarding-options]` hierarchy level. See the *[edit forwarding-options dhcp-relay] Hierarchy Level* for the complete DHCP relay agent syntax.

You can also include the `dhcp-relay` statement at the following hierarchy levels:

- `[edit logical-systems logical-system-name forwarding-options]`
- `[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options]`
- `[edit routing-instances routing-instance-name forwarding-options]`

This overview covers:

- [Interaction Among the DHCP Relay Agent, DHCP Client, and DHCP Servers on page 901](#)
- [DHCP Liveness Detection on page 902](#)

Interaction Among the DHCP Relay Agent, DHCP Client, and DHCP Servers

The pattern of interaction among the DHCP Relay agent, DHCP client, and DHCP servers is the same regardless of whether the software installation is on a router or a switch. However, there are some difference in the details of usage.

On routers—In a typical carrier edge network configuration, the DHCP client is on the subscriber's computer, and the DHCP relay agent is configured on the router between the DHCP client and one or more DHCP servers.

On switches—In a typical network configuration, the DHCP client is on an access device such as a personal computer and the DHCP relay agent is configured on the switch between the DHCP client and one or more DHCP servers.

The following steps describe, at a high level, how the DHCP client, DHCP relay agent, and DHCP server interact in a configuration that includes two DHCP servers.

1. The DHCP client sends a discover packet to find a DHCP server in the network from which to obtain configuration parameters for the subscriber (or DHCP client), including an IP address.
2. The DHCP relay agent receives the discover packet and forwards copies to each of the two DHCP servers. The DHCP relay agent then creates an entry in its internal client table to keep track of the client's state.
3. In response to receiving the discover packet, each DHCP server sends an offer packet to the client. The DHCP relay agent receives the offer packets and forwards them to the DHCP client.
4. On receipt of the offer packets, the DHCP client selects the DHCP server from which to obtain configuration information. Typically, the client selects the server that offers the longest lease time on the IP address.
5. The DHCP client sends a request packet that specifies the DHCP server from which to obtain configuration information.
6. The DHCP relay agent receives the request packet and forwards copies to each of the two DHCP servers.
7. The DHCP server requested by the client sends an acknowledgement (ACK) packet that contains the client's configuration parameters.
8. The DHCP relay agent receives the ACK packet and forwards it to the client.
9. The DHCP client receives the ACK packet and stores the configuration information.
10. If configured to do so, the DHCP relay agent installs a host route and Address Resolution Protocol (ARP) entry for this client.
11. After establishing the initial lease on the IP address, the DHCP client and the DHCP server use unicast transmission to negotiate lease renewal or release. The DHCP relay

agent “snoops” on all of the packets unicast between the client and the server that pass through the router (or switch) to determine when the lease for this client has expired or been released. This process is referred to as *lease shadowing* or *passive snooping*.

DHCP Liveness Detection

Liveness detection for DHCP subscriber or DHCP client IP sessions utilizes an active liveness detection protocol to institute liveness detection checks for relevant clients. Clients are expected to respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, then the liveness detection check fails and a failure action is implemented.



NOTE: DHCP liveness detection either globally or per DHCP group.

Related Documentation

- [DHCPv6 Relay Agent Overview on page 905](#)
- [Access and Access-Internal Routes for Subscriber Management](#)
- [Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview](#)
- [Using External AAA Authentication Services with DHCP on page 923](#)
- [DHCP Relay Proxy Overview on page 903](#)
- [Graceful Routing Engine Switchover on page 897](#)
- [Subscriber Management Unified ISSU Support](#)
- [Verifying and Managing DHCP Relay Configuration on page 1159](#)
- [Tracing Extended DHCP Operations on page 1245](#)
- [Example: Minimum DHCP Relay Agent Configuration on page 916](#)
- [Example: DHCP Relay Agent Configuration with Multiple Clients and Servers](#)
- [Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 917](#)
- [Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing](#)
- [Example: Configuring a DHCP Firewall Filter to Protect the Routing Engine](#)

DHCP Relay Proxy Overview

DHCP relay proxy mode is an enhancement to extended DHCP relay. DHCP relay proxy supports all DHCP relay features while providing additional features and benefits.

Normally, extended DHCP relay operates as a helper application for DHCP operations. Except for the ability to add DHCP relay agent options and the gateway address (giaddr) to DHCP packets, DHCP relay is transparent to DHCP clients and DHCP servers, and simply forwards messages between DHCP clients and servers.

When you configure DHCP relay to operate in proxy mode, the relay is no longer transparent. In proxy mode, DHCP relay conceals DHCP server details from DHCP clients, which interact with a DHCP relay in proxy mode as though it is the DHCP server. For DHCP servers there is no change, because proxy mode has no effect on how the DHCP server interacts with the DHCP relay.

DHCP relay proxy provides the following benefits:

- DHCP server isolation and DoS protection—DHCP clients are unable to detect the DHCP servers, learn DHCP server addresses, or determine the number of servers that are providing DHCP support. Server isolation also provides denial-of-service (DoS) protection for the DHCP servers.
- Multiple lease offer selection—DHCP relay proxy receives lease offers from multiple DHCP servers and selects a single offer to send to the DHCP client, thereby reducing traffic in the network. Currently, the DHCP relay proxy selects the first offer received.
- Support for both numbered and unnumbered Ethernet interfaces—For DHCP clients connected through Ethernet interfaces, when the DHCP client obtains an address, the DHCP relay proxy adds an access internal host route specifying that interface as the outbound interface. The route is automatically removed when the lease time expires or when the client releases the address.
- Logical system support—DHCP relay proxy can be configured in a logical system, whereas a non-proxy mode DHCP relay cannot.



NOTE: Extended DHCP relay proxy is not supported for the J Series Services Routers DHCP server. Also, you cannot configure both DHCP relay proxy and extended DHCP local server on the same interface.

Interaction Among DHCP Relay Proxy, DHCP Client, and DHCP Servers

The DHCP relay agent is configured on the router (or switch), which operates between the DHCP client and one or more DHCP servers.

The following steps provide a high-level description of how DHCP relay proxy interacts with DHCP clients and DHCP servers.

1. The DHCP client sends a discover packet to locate a DHCP server in the network from which to obtain configuration parameters for the subscriber.
2. The DHCP relay proxy receives the discover packet from the DHCP client and forwards copies of the packet to each supporting DHCP server. The DHCP relay proxy then creates a client table entry to keep track of the client state.
3. In response to the discover packet, each DHCP server sends an offer packet to the client, which the DHCP relay proxy receives. The DHCP relay proxy does the following:
 - a. Selects the first offer received as the offer to sent to the client
 - b. Replaces the DHCP server address with the address of the DHCP relay proxy
 - c. Forwards the offer to the DHCP client.
4. The DHCP client receives the offer from the DHCP relay proxy.
5. The DHCP client sends a request packet that indicates the DHCP server from which to obtain configuration information—the request packet specifies the address of the DHCP relay proxy.
6. The DHCP relay proxy receives the request packet and forwards copies, which include the address of selected server, to all supporting DHCP servers.
7. The DHCP server requested by the client sends an acknowledgement (ACK) packet that contains the client configuration parameters.
8. The DHCP relay proxy receives the ACK packet, replaces the DHCP server address with its own address, and forwards the packet to the client.
9. The DHCP client receives the ACK packet and stores the configuration information.
10. If configured to do so, the DHCP relay proxy installs a host route and Address Resolution Protocol (ARP) entry for the DHCP client.
11. After the initial DHCP lease is established, the DHCP relay proxy receives all lease renewals and lease releases from the DHCP client and forwards them to the DHCP server.

**Related
Documentation**

- [Extended DHCP Relay Agent Overview on page 900](#)
- [Enabling DHCP Relay Proxy Mode on page 977](#)
- [Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 981](#)

DHCPv6 Relay Agent Overview

The DHCPv6 relay agent enhances the extended DHCP relay agent by providing support in an IPv6 network. The DHCPv6 relay agent passes messages between the DHCPv6 client and the DHCPv6 server, similar to the way DHCP relay agent supports an IPv4 network.

When a DHCPv6 client logs in, the DHCPv6 relay agent uses the AAA service framework to interact with the RADIUS server to provide authentication and accounting. The RADIUS server, which is configured independently of DHCP, authenticates the client and supplies the IPv6 prefix and client configuration parameters, such as session timeout and the maximum number of clients allowed per interface.



NOTE: The PTX Series Packet Transport Switches do not support authentication for DHCPv6 relay agents.

The DHCPv6 relay agent is compatible with the extended DHCP local server and the extended DHCP relay agent, and can be enabled on the same interface as either the extended DHCP local server or DHCP relay agent.

To configure the DHCPv6 relay agent on the router (or switch), you include the **dhcpv6** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level.

You can also include the **dhcpv6** statement at the following hierarchy levels:

- **[edit logical-systems *logical-system-name* forwarding-options dhcp-relay]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay]**
- **[edit routing-instances *routing-instance-name* forwarding-options dhcp-relay]**

Related Documentation

- [Using External AAA Authentication Services with DHCP on page 923](#)
- [Grouping Interfaces with Common DHCP Configurations on page 926](#)
- [Group-Specific DHCP Relay Options on page 906](#)
- [Overriding the Default DHCP Relay Configuration Settings on page 955](#)
- [Configuring Passwords for Usernames on page 945](#)
- [Creating Unique Usernames for DHCP Clients on page 946](#)
- [Verifying and Managing DHCPv6 Local Server Configuration on page 1158](#)
- [Example: Extended DHCPv6 Local Server Configuration](#)

DHCP Duplicate Client Differentiation Using Client Subinterface Overview

In some network environments, client IDs and MAC addresses might not be unique, resulting in duplicate clients. For example, two network adapters might be manufactured with the same hardware address, resulting in a duplicate MAC address among the DHCP clients attached to the router (or switch). A duplicate DHCP client occurs when a client attempts to get a lease, and that client has the same client ID or the same MAC address as an existing DHCP client.

When DHCP server receives a request from a new client that has a duplicate ID or MAC address, DHCP server terminates the address lease for the existing client and returns the address to its original address pool. DHCP server then assigns a new address and lease to the new client.

By default, both DHCP local server and DHCP relay use the subnet information to differentiate between duplicate clients. However, in some cases, this level of differentiation is not adequate. For example, when multiple subinterfaces share the same underlying loopback interface with the same preferred source address, the interfaces appear to be on the same subnet. In this situation, the default configuration prevents duplicate clients.

You can provide greater differentiation between duplicate clients by configuring DHCP to consider the client subinterface when duplicate clients occur. In this optional configuration, DHCP uniquely identifies:

- The subnet on which the client resides
- The subinterface on which the client resides
- The client within the subnet

Related Documentation

- [Configuring DHCP Duplicate Client Support on page 925](#)
- [Guidelines for Configuring Support for DHCP Duplicate Clients on page 925](#)

Group-Specific DHCP Relay Options

You can include the following statements at the **[edit forwarding-options dhcp-relay group group-name]** hierarchy level to set group-specific DHCP relay agent configuration options. Group-specific statements apply only to the named group of interfaces, and override any global DHCP relay agent settings for the same statement.

Include the statements at the **[edit forwarding-options dhcp-relay dhcpv6 group group-name]** hierarchy level to configure group-specific options for DHCPv6 relay agent.

- **active-server-group**—Configure an active server group to apply a common DHCP relay agent configuration to a named group of DHCP server addresses. For information, see [“Configuring Active Server Groups” on page 977](#).
- **authentication**—Configure the parameters the router (or switch) sends to the external AAA server.
- **dynamic-profile**—Specify the dynamic profile that is attached to a group of interfaces.

- **interface**—Specify one or more interfaces, or a range of interfaces, that are within the specified group.
- **liveness-detection**—Configure bidirectional failure detection timers and authentication criteria for static routes. For more information, see [“DHCP Liveness Detection Overview” on page 980](#).
- **overrides**—Override the default configuration settings for the extended DHCP relay agent. For information, see [“Overriding the Default DHCP Relay Configuration Settings” on page 955](#).
- **relay-agent-interface-id**—(DHCPv6 only) Insert the DHCPv6 Relay Agent Interface-ID option (option 18) in DHCPv6 packets destined for the DHCPv6 server.
- **relay-option**—Configure selective processing, which uses DHCP options in client packets to identify and filter client traffic, and to specify the action DHCP relay agent takes with the traffic. For more information, see *Using DHCP Option Information to Selectively Process DHCP Client Traffic*.
- **relay-option-82**—(DHCPv4 only) Enable or disable the insertion of option 82 information in packets destined for a DHCP server. For information, see [“Enabling and Disabling Insertion of Option 82 Information” on page 973](#).
- **service-profile**—Specify the default subscriber service, (or default profile) which is activated when the subscriber (or DHCP client) logs in and no other service is activated by a RADIUS server or a provisioning server. For more information, see *Default Subscriber Service Overview*.

Related Documentation

- [Grouping Interfaces with Common DHCP Configurations on page 926](#)

DHCP Snooping Support

DHCP snooping provides DHCP security on the router by filtering incoming messages. When DHCP snooping is enabled, the router differentiates between trusted and untrusted interfaces, and forwards messages from trusted sources while rejecting the untrusted messages.

In Junos OS, DHCP snooping is enabled in a routing instance when you configure either the **dhcp-relay** statement at the **[edit forwarding-options]** hierarchy level, or the **dhcp-local-server** statement at the **[edit system services]** hierarchy level in that routing instance. However, depending on the Junos OS release, the router processes the snooped packets differently, as described in the following list:

- In Junos OS Release 10.0 and earlier, the router processes snooped packets normally.
- In Junos OS Release 10.1 and later, the router discards snooped packets by default. To enable normal processing of snooped packets in Junos OS Release 10.1 and later, you must explicitly configure the **allow-snooped-clients** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level.

You can configure DHCP snooping support for the following:

- DHCPv4 relay agent—Override the router's (or switch's) default snooping configuration and specify that DHCP snooping is enabled or disabled globally, for a named group of interfaces, or for a specific interface within a named group.

In a separate procedure, you can set a global configuration to specify whether the DHCPv4 relay agent forwards or drops snooped packets for all interfaces, only configured interfaces, or only nonconfigured interfaces. The router also uses the global DHCP relay agent snooping configuration to determine whether to forward or drop snooped BOOTREPLY packets.

- DHCPv6 relay agent—As you can with snooping support for the DHCPv4 relay agent, you can override the default DHCPv6 relay agent snooping configuration on the router to explicitly enable or disable snooping support globally, for a named group of interfaces, or for a specific interface with a named group of interfaces.

In multi-relay topologies where more than one DHCPv6 relay agent is between the DHCPv6 client and the DHCPv6 server, snooping enables intervening DHCPv6 relay agents between the client and the server to correctly receive and process the unicast traffic from the client and forward it to the server. The DHCPv6 relay agent snoops incoming unicast DHCPv6 packets by setting up a filter with UDP port 547 (the DHCPv6 UDP server port) on a per-forwarding table basis. The DHCPv6 relay agent then processes the packets intercepted by the filter and forwards the packets to the DHCPv6 server.

Unlike the DHCPv4 relay agent, the DHCPv6 relay agent does not support global configuration of forwarding support for DHCPv6 snooped packets.

- DHCP local server—Configure whether DHCP local server forwards or drops snooped packets for all interfaces, only configured interfaces, or only nonconfigured interfaces.

**Related
Documentation**

- [Configuring DHCP Snooping for DHCP Relay Agent on page 963](#)
- [Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server on page 944](#)
- [Example: Configuring DHCP Snooping Support for DHCP Relay Agent on page 921](#)

DHCP Auto Logout Overview

This topic provides an introduction to the optional DHCP auto logout feature and includes the following sections:

- [Auto Logout Overview on page 908](#)
- [How DHCP Identifies and Releases Clients on page 909](#)
- [Option 60 and Option 82 Requirements on page 910](#)

Auto Logout Overview

Auto logout is an optional configuration for DHCP local server and DHCP relay agent that improves the efficiency of DHCP IP address assignment. Auto logout enables IP addresses to be immediately released and returned to the address pool when the addresses are no

longer used by DHCP clients. DHCP can then assign the addresses to other clients. Without auto logout, an IP address is blocked for the entire lease period, and DHCP must wait until the address lease time expires before reusing the address.

Auto logout is particularly useful when DHCP uses long lease times for IP address assignments and to help avoid allocating duplicate IP addresses for a single client.

For example (on the routers), you might have an environment that includes set-top boxes (STB) that are often upgraded or replaced. Each time a STB is changed, the new STB repeats the DHCP discover process to obtain client configuration information and an IP address. DHCP views the new STB as a completely new client and assigns a new IP address—the previous IP address assigned to the client (the old STB) remains blocked and unavailable until the lease expires. If auto logout is configured in this situation, DHCP recognizes that the new STB is actually the same client and then immediately releases the original IP address. DHCP relay agent acts as a proxy client for auto logout and sends a DHCP release message to the DHCP server.

How DHCP Identifies and Releases Clients

The auto logout feature requires that DHCP explicitly identify clients. By default, DHCP local server and DHCP relay agent identify clients based on MAC address or Client Identifier. However, in some cases this type of identification might not be sufficient. For example, in the previous STB example, each STB has a different MAC address, so DHCP incorrectly assumes that an upgraded or replacement STB is a new client.

In order to explicitly identify clients, auto logout uses a secondary identification method when the primary identification method is unsuccessful—the primary method is considered unsuccessful if the MAC address or Client Identifier does not match that of an existing client. The secondary identification method is based on the DHCP option 60 and option 82 information in DHCP discover messages.

Both the primary and secondary identification methods use subnet information to differentiate between clients. The primary identification method differentiates between two clients with the same MAC address (or same Client Identifier) if the clients are on different subnets. Similarly, the secondary identification method considers two clients as different if they have the same option 60 and option 82 information, but different subnets.

DHCP local server and DHCP relay agent perform the following operations when auto logout is enabled and the secondary identification method identifies a duplicate client (that is, the discover packet is from an existing client).

- DHCP local server immediately releases the existing address.
- DHCP relay agent immediately releases the existing client and then sends a DHCP release packet to the DHCP server. Sending the release packet ensures that DHCP relay and the DHCP server are synchronized.

If the DHCP relay receives a DISCOVER message from an existing client, the DHCP relay forwards the DISCOVER message to the DHCP server. The DHCP relay preserves the binding if the client's existing IP address is returned by the DHCP server. This behavior is not applicable if the proxy-mode override or client-discover-match functionality are enabled.



NOTE: If the DHCP relay agent is in snoop mode, DHCP relay releases the client but does not send a release packet to the DHCP server if the discover packet is for a passive client (a client added as a result of snooped packets) or if the discover packet is a snooped packet.

Option 60 and Option 82 Requirements

DHCP local server requires that the received discover packet include both DHCP option 60 and option 82. If either option is missing, DHCP local server cannot perform the secondary identification method and auto logout is not used.

DHCP relay agent requires that the received discover packet contain DHCP option 60. DHCP relay determines the option 82 value based on the guidelines provided in [“DHCP Relay Agent Option 82 Value for Auto Logout” on page 962](#).

Related Documentation

- [Automatically Logging Out DHCP Clients on page 932](#)
- [DHCP Relay Agent Option 82 Value for Auto Logout on page 962](#)
- [Clearing DHCP Bindings for Subscriber Access](#)

Graceful Routing Engine Switchover

For EX Series switches, only extended DHCP local server maintains the state of active DHCP client leases. The DHCP local server supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication. You can configure dynamic profile and authentication support on a global basis or for a specific group of interfaces. The extended DHCP local server also supports the use of Junos address-assignment pools or external authorities, such as RADIUS, to provide the client address and configuration information.

For MX Series routers, the extended DHCP local server and the DHCP relay agent applications both maintain the state of active DHCP client leases in the session database. The extended DHCP application can recover this state if the DHCP process fails or is manually restarted, thus preventing the loss of active DHCP clients in either of these circumstances. However, the state of active DHCP client leases is lost if a power failure occurs or if the kernel stops operating (for example, when the router is reloaded) on a single Routing Engine.

You can enable graceful switchover support on both EX Series switches and MX Series routers. To enable graceful switchover support for the extended DHCP local server or extended DHCP relay agent on a switch, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level. To enable graceful Routing Engine switchover support on MX Series routers, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level. You cannot disable graceful Routing Engine switchover support for the extended DHCP application when the router is configured to support graceful Routing Engine switchover.

For more information about using graceful Routing Engine switchover, see the *Junos OS High Availability Configuration Guide*.

**Related
Documentation**

- [Extended DHCP Local Server Overview on page 876](#)
- [Extended DHCP Relay Agent Overview on page 900](#)
- *Subscriber Management Unified ISSU Support*

Port Number Requirements for DHCP Firewall Filters

When you configure a firewall filter to perform some action on DHCP packets at the Routing Engine, such as protecting the Routing Engine by allowing only proper DHCP packets, you must specify both port 67 (bootps) and port 68 (bootpc) for both the source and destination. The firewall filter acts at both the line cards and the Routing Engine.

This requirement applies to both DHCP local server and DHCP relay, but it applies only when DHCP is provided by the `jdhcpd` process. MX Series routers, M120 routers, and M320 routers use `jdhcpd`. For DHCP relay, that means the configuration is required only at the `[edit forwarding-options dhcp-relay]` hierarchy level and not at the `[edit forwarding-options helpers bootp]` hierarchy level.

DHCP packets received on the line cards are encapsulated by `jdhcpd` with a new UDP header where their source and destination addresses are set to port 68 before being forwarded to the Routing Engine.

For DHCP relay and DHCP proxy, packets sent to the DHCP server from the router have both the source and destination UDP ports set to 67. The DHCP server responds using the same ports. However, when the line card receives these DHCP response packets, it changes both port numbers from 67 to 68 before passing the packets to the Routing Engine. Consequently the filter needs to accept port 67 for packets relayed from the client to the server, and port 68 for packets relayed from the server to the client.

Failure to include both port 67 and port 68 as described here results in most DHCP packets not being accepted.

For complete information about configuring firewall filters in general, see *Junos OS Firewall Filters and Traffic Policers Configuration Guide*.

**Related
Documentation**

- *Example: Configuring a DHCP Firewall Filter to Protect the Routing Engine*
- [Extended DHCP Local Server Overview on page 876](#)
- [Extended DHCP Relay Agent Overview on page 900](#)
- *Dynamic Firewall Filters Overview*

Configuration

- [DHCP Local Server Examples on page 912](#)
- [DHCP Relay Agent Examples on page 916](#)
- [Configuration Tasks for DHCP Local Server on page 923](#)

- [Configuration Tasks for DHCP Relay Agent on page 950](#)
- [DHCP Local Server Configuration Statements on page 983](#)
- [DHCP Relay Agent Configuration Statements on page 1069](#)

DHCP Local Server Examples

- [Example: Minimum Extended DHCP Local Server Configuration on page 912](#)
- [Example: Extended DHCP Local Server Configuration with Optional Pool Matching on page 912](#)
- [Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913](#)

Example: Minimum Extended DHCP Local Server Configuration

This example shows the minimum configuration you need to use for the extended DHCP local server on the router or switch:

```
[edit system services]
dhcp-local-server {
  group group_one {
    interface fe-0/0/2.0;
  }
}
```



NOTE: The interface type in this topic is just an example. The **fe-** interface type is not supported by EX Series switches.

This example creates the server group named **group_one**, and specifies that the DHCP local server is enabled on interface **fe-0/0/2.0** within the group. The DHCP local server uses the default pool match configuration of **ip-address-first**.

Related Documentation

- [Extended DHCP Local Server Overview on page 876](#)

Example: Extended DHCP Local Server Configuration with Optional Pool Matching

This example shows an extended DHCP local server configuration that includes optional IPv4 address-assignment pool matching and interface groups. For pool matching, this configuration specifies that the DHCP local server first check the response from an external authentication authority (for example, RADIUS) and use the Framed-IPv6-Pool attribute to determine the address-assignment pool to use for the client address. If no external authority match is found, the DHCP local server then uses **ip-address-first** matching together with the option 82 information to match the named address range for client IPv4 address assignment. The option 82 matching must also be included in the address-assignment pool configuration.

```
[edit system services]
dhcp-local-server {
  group group_one {
    interface fe-0/0/2.0;
    interface fe-0/0/2.1;
  }
}
```

```

}
group group_two {
  interface fe-0/0/3.0;
  interface fe-0/0/3.1;
}
pool-match-order {
  external-authority
  ip-address-first;
  option-82;
}
}

```



NOTE: The interface type in this topic is just an example. The fe- interface type is not supported by EX Series switches.

Related Documentation

- [Extended DHCP Local Server Overview on page 876](#)
- [Address-Assignment Pools Overview on page 890](#)

Example: Configuring Group Liveness Detection for DHCP Local Server Clients

This example shows how to configure group liveness detection for DHCP local server subscribers or DHCP clients using Bidirectional Forwarding Detection (BFD) as the liveness detection method.

- [Requirements on page 913](#)
- [Overview on page 913](#)
- [Configuration on page 914](#)

Requirements

- Juniper Networks MX Series routers
- Juniper Networks EX Series switches
- Junos OS Release 12.1R1 or later
- Junos OS Release 12.3R2 or later for EX Series switches
- Configure DHCP local server. See [“Extended DHCP Local Server Overview” on page 876](#).

Overview

In this example, you configure group liveness detection for DHCP local server subscribers (clients) by completing the following operations:

1. Enable liveness detection for DHCP local server subscriber (or DHCP client) groups.
2. Specify BFD as the liveness detection method for all dynamically created DHCP local server subscribers (clients).
3. Configure BFD-specific statements to define how the protocol behaves.
4. Configure the action the router (switch) takes when a liveness detection failure occurs.



NOTE: This example explains how to configure liveness detection for a DHCPv4 network. Liveness detection is also supported for DHCPv6 configurations. To configure DHCPv6 liveness detection, include the `liveness-detection` statement, and any subsequent configuration statements, at the `[edit system services dhcp-local-server dhcpv6]` or `[edit system services dhcp-local-server dhcpv6 group group-name]` hierarchy level.

Configuration

Step-by-Step Procedure

To configure group liveness detection for DHCP local server:

1. Specify that you want to configure liveness detection.

```
[edit system services dhcp-local-server ]
user@host# edit liveness-detection
```
2. Specify that you want to configure liveness detection for a specific DHCP local server group.

```
[edit system services dhcp-local-server liveness-detection]
user@host# edit group local_group_1
```
3. Specify that you want to configure the liveness detection method.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection]
user@host# edit method
```
4. Specify BFD as the liveness detection method that you want DHCP to use.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
method]
user@host# edit bfd
```
5. Configure the detection time threshold (in milliseconds) at which a trap is produced.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
method bfd]
user@host# set detection-time threshold 30000
```
6. Configure the time (in milliseconds) for which BFD holds a session up notification.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
method bfd]
user@host# set holddown-interval 50
```
7. Configure the BFD minimum transmit and receive interval (in milliseconds).



NOTE: You do not need to configure the BFD minimum transmit and receive interval if you configure the `minimum-interval` for the BFD `transmit-interval` statement and the `minimum-receive-interval`.

```
[edit system services dhcp-local-servergroup local_group_1 liveness-detection method
bfd]
user@host# set minimum-interval 45000
```


8. Configure the minimum receive interval (in milliseconds).



NOTE: You do not need to configure the BFD minimum receive interval if you configure the BFD minimum transmit and receive interval.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
method bfd]
user@host# set minimum-receive-interval 60000
```

9. Configure a multiplier value for the detection time.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
method bfd]
user@host# set multiplier 100
```

10. Disable the ability for BFD interval timers to change or adapt to network situations.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
method bfd]
user@host# set no-adaptation
```

11. Configure the BFD session mode.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
method bfd]
user@host# set session-mode automatic
```

12. Configure the threshold and minimum interval for the BFD transmit interval.



NOTE: You do not need to configure the transmit interval values if you have already configured the minimum transmit and receive interval for BFD.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
method bfd]
user@host# set transmit-interval threshold 60000 minimum-interval 45000
```

13. Configure the BFD protocol version you want to detect.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection
method bfd]
user@host# set version automatic
```

14. Configure the action the router (switch) takes when a liveness detection failure occurs. In this example, the failure action is to clear the client session only when a liveness detection failure occurs and the local interface is detected as being up.

```
[edit system services dhcp-local-server group local_group_1 liveness-detection]
user@host# edit failure-action action
```

Results From configuration mode, confirm your configuration by entering the **show system** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it.

```
[edit]
regress@montag# show system
services {
  dhcp-local-server {
    group local_group_1 {
      liveness-detection {
        failure-action clear-binding-if-interface-up;
        method {
          bfd {
            version automatic;
            minimum-interval 45000;
            minimum-receive-interval 60000;
            multiplier 100;
            no-adaptation;
            transmit-interval {
              minimum-interval 45000;
              threshold 60000;
            }
            detection-time {
              threshold 30000;
            }
            session-mode automatic;
            holddown-interval 50;
          }
        }
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Related Documentation

- [Extended DHCP Local Server Overview on page 876](#)
- [DHCP Liveness Detection Overview on page 980](#)
- [Configuring Detection of DHCP Local Server Client Connectivity on page 941](#)

DHCP Relay Agent Examples

- [Example: Minimum DHCP Relay Agent Configuration on page 916](#)
- [Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings on page 917](#)
- [Example: Configuring DHCP Snooping Support for DHCP Relay Agent on page 921](#)

Example: Minimum DHCP Relay Agent Configuration

This example shows the minimum configuration you need to use the extended DHCP relay agent on the router or switch:

```
[edit forwarding-options]
dhcp-relay {
  server-group {
    test 10.0.2.1;
  }
}
```

```

active-server-group test;
group all {
    interface fe-0/0/2.0;
}
}

```



NOTE: The interface type in this topic is just an example. The `fe-` interface type is not supported by EX Series switches.

This example creates a server group and an active server group named **test** with IP address 10.0.2.1. The DHCP relay agent configuration is applied to a group named **all**. Within this group, the DHCP relay agent is enabled on interface `fe-0/0/2.0`.

Related Documentation

- [Extended DHCP Relay Agent Overview on page 900](#)

Example: Configuring DHCP Relay Agent Selective Traffic Processing Based on DHCP Option Strings

This example shows how to configure DHCP relay agent to use DHCP option strings to selectively identify, filter, and process client traffic.

- [Requirements on page 917](#)
- [Overview on page 917](#)
- [Configuration on page 918](#)
- [Verification on page 920](#)

Requirements

This example uses the following hardware and software components:

- MX Series 3D Universal Edge Routers or EX Series Switches
- Junos OS Release 12.3 or later or Junos OS Release 12.3R2 for EX Series switches

Before you configure DHCP relay agent selective processing support, be sure you:

- Configure DHCP relay agent.

See [“Extended DHCP Relay Agent Overview” on page 900](#).

- (Optional) Configure a named DHCP local server group if you want to forward client traffic to a server group.

See [“Grouping Interfaces with Common DHCP Configurations” on page 926](#).

Overview

In this example, you configure DHCP relay agent to use DHCP option strings in client packets to selectively identify, filter, and process client traffic. To configure selective processing, you perform the following procedures:

1. Identify the client traffic—Specify the DHCP option that DHCP relay agent uses to identify the client traffic you want to process. The option you specify matches the option in the client traffic.
2. Configure a default action—Specify the default processing action, which DHCP relay uses for identified client traffic that does not satisfy any configured match criteria.
3. Create match filters and associate an action with each filter—Specify match criteria that filter the client traffic. The criteria can be an exact match or a partial match with the option string in the client traffic. Associate a processing action with each match criterion.

Configuration

To configure DHCP relay agent selective processing based on DHCP option information, perform these tasks:

- [Configuring DHCP Relay Agent To Selectively Process Client Traffic Based on DHCP Option Strings on page 918](#)
- [Results on page 919](#)

CLI Quick Configuration

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

```
set forwarding-options dhcp-relay relay-option option-number 60
set forwarding-options dhcp-relay relay-option equals ascii video-gold forward-only
set forwarding-options dhcp-relay relay-option equals ascii video-bronze local-server-group
  servergroup-15
set forwarding-options dhcp-relay relay-option starts-with hexadecimal ffff
  local-server-group servergroup-east
set forwarding-options dhcp-relay relay-option default-action drop
```

Configuring DHCP Relay Agent To Selectively Process Client Traffic Based on DHCP Option Strings

Step-by-Step Procedure

To configure DHCP relay selective processing:

1. Specify that you want to configure DHCP relay agent support.

```
[edit forwarding-options]
user@host# edit dhcp-relay
```
2. Specify the DHCP option that DHCP relay agent uses to identify incoming client traffic.

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option option-number 60
```
3. Configure a default action, which DHCP relay agent uses when the incoming client traffic does not satisfy any configured match criteria.

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option default-action drop
```

4. Configure an exact match condition and associated action that DHCP relay uses to process the identified client traffic.

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option equals ascii video-gold forward-only
```

5. Configure a second exact match condition and associated action that DHCP relay uses to process client traffic.

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option equals ascii video-bronze local-server-group
servergroup-15
```

6. Configure a partial match criteria and associated action that DHCP relay uses to process client traffic.

```
[edit forwarding-options dhcp-relay]
user@host# set relay-option starts-with hexadecimal ffff local-server-group
servergroup-east
```

Results

From configuration mode, confirm the results of your configuration by issuing the **show** statement at the **[edit forwarding-options]** hierarchy level. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit forwarding-options]
user@host# show
dhcp-relay {
  relay-option {
    option-number 60;
    equals {
      ascii video-gold {
        forward-only;
      }
    }
    equals {
      ascii video-bronze {
        local-server-group servergroup-15;
      }
    }
    default-action {
      drop;
    }
    starts-with {
      hexadecimal ffff {
        local-server-group servergroup-east;
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To verify the status of DHCP relay agent selective traffic processing, perform this task:

- [Verifying the Status of DHCP Relay Agent Selective Traffic Processing on page 920](#)

Verifying the Status of DHCP Relay Agent Selective Traffic Processing

Purpose Verify the DHCP relay agent selective traffic processing status.

Action Display statistics for DHCP relay agent.

```
user@host> show dhcp relay statistics
Packets dropped:
  Total                30
  Bad hardware address  1
  Bad opcode            1
  Bad options           3
  Invalid server address 5
  No available addresses 1
  No interface match    2
  No routing instance match 9
  No valid local address 4
  Packet too short      2
  Read error            1
  Send error            1
  Option 60             1
  Option 82             2

Messages received:
  BOOTREQUEST          116
  DHCPDECLINE           0
  DHCPDISCOVER          11
  DHCPINFORM            0
  DHCPRELEASE           0
  DHCPREQUEST          105

Messages sent:
  BOOTREPLY             0
  DHCPOFFER             2
  DHCPACK                1
  DHCPNAK                0
  DHCPFORCERENEW         0

Packets forwarded:
  Total                4
  BOOTREQUEST           2
  BOOTREPLY             2
```

Meaning The **Packets forwarded** field in the **show dhcp relay statistics** command output displays the number of client packets that have been forwarded as a result of the selective traffic processing configuration. In this example, the output indicates the total number of packets that DHCP relay agent has forwarded, as well as a breakdown for the number of **BOOTREQUEST** and **BOOTREPLY** packets forwarded.

- Related Documentation**
- [Extended DHCP Relay Agent Overview on page 900](#)
 - [DHCP Options and Selective Traffic Processing Overview](#)
 - [Using DHCP Option Information to Selectively Process DHCP Client Traffic](#)
 - [Displaying a Count of DHCP Packets That Are Dropped or Forwarded During Selective Processing That Is Based on DHCP Option Strings](#)
 - [Example: Configuring DHCP and DHCPv6 Relay Agent Group-Level Selective Traffic Processing](#)

Example: Configuring DHCP Snooping Support for DHCP Relay Agent

This example shows how to configure DHCP snooping support for DHCP relay agent.

- [Requirements on page 921](#)
- [Overview on page 921](#)
- [Configuration on page 921](#)

Requirements

- Configure DHCP relay agent. See “[Extended DHCP Relay Agent Overview](#)” on page 900.

Overview

In this example, you configure DHCP snooping support for DHCP relay agent by completing the following operations:

- Override the default DHCP snooping configuration and enable DHCP snooping support for the interfaces in group **frankfurt**.
- Configure DHCP relay agent to forward snooped packets to only configured interfaces.



NOTE: By default, DHCP snooping is enabled globally in Junos OS Release 10.0 and earlier and disabled globally in Junos OS Release 10.1 and later.

Configuration

Step-by-Step Procedure

To configure DHCP relay support for DHCP snooping:

1. Specify that you want to configure DHCP relay agent.

```
[edit]
user@host# edit forwarding-options dhcp-relay
```
2. Specify the named group of interfaces on which DHCP snooping is supported.

```
[edit forwarding-options dhcp-relay]
user@host# edit group frankfurt
```
3. Specify the interfaces that you want to include in the group. DHCP relay agent considers these as the configured interfaces when determining whether to forward or drop traffic.

```
[edit forwarding-options dhcp-relay group frankfurt]
user@host# set interface fe-1/0/1.3 upto fe-1/0/1.9
```

4. Specify that you want to override the default configuration for the group.

```
[edit forwarding-options dhcp-relay group frankfurt]
user@host# edit overrides
```

5. Enable DHCP snooping support for the group.

```
[edit forwarding-options dhcp-relay group frankfurt overrides]
user@host# set allow-snooped-clients
```

6. Return to the `[edit forwarding-options dhcp-relay]` hierarchy level to configure the forwarding action and specify that DHCP relay agent forward snooped packets on only configured interfaces:

```
[edit forwarding-options dhcp-relay group frankfurt overrides]
user@host# up 2
```

7. Enable DHCP snooped packet forwarding for DHCP relay agent.

```
[edit forwarding-options dhcp-relay]
user@host# edit forward-snooped-clients
```

8. Specify that snooped packets are forwarded on only configured interfaces (the interfaces in group `frankfurt`).

```
[edit forwarding-options dhcp-relay forward-snooped-clients]
user@host# set configured-interfaces
```

Results From configuration mode, confirm your configuration by entering the **show forwarding-options** command. If the output does not display the intended configuration, repeat the instructions in this example to correct it. The following output also shows a range of configured interfaces in group `frankfurt`.

```
[edit]
regress@montag# show forwarding-options
dhcp-relay {
  forward-snooped-clients configured-interfaces;
  group frankfurt {
    overrides {
      allow-snooped-clients;
    }
    interface fe-1/0/1.3 {
      upto fe-1/0/1.9;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

**Related
Documentation**

- [DHCP Snooping Support on page 887](#)
- [Configuring DHCP Snooping for DHCP Relay Agent on page 963](#)

Configuration Tasks for DHCP Local Server

- [Using External AAA Authentication Services with DHCP on page 923](#)
- [Guidelines for Configuring Support for DHCP Duplicate Clients on page 925](#)
- [Configuring DHCP Duplicate Client Support on page 925](#)
- [Grouping Interfaces with Common DHCP Configurations on page 926](#)
- [Guidelines for Configuring Interface Ranges on page 927](#)
- [Overriding Default DHCP Local Server Configuration Settings on page 928](#)
- [Specifying the Maximum Number of DHCP Clients Per Interface on page 929](#)
- [Disabling ARP Table Population on page 930](#)
- [Automatically Logging Out DHCP Clients on page 932](#)
- [Enabling Processing of Client Information Requests on page 933](#)
- [Specifying the Delegated Address Pool for IPv6 Prefix Assignment on page 934](#)
- [Enabling DHCPv6 Rapid Commit Support on page 935](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 935](#)
- [Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 936](#)
- [Configuring Dynamic Reconfiguration Attempts for DHCP Clients on page 937](#)
- [Configuring Deletion of the Client When Dynamic Reconfiguration Fails on page 938](#)
- [Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect on page 938](#)
- [Configuring a Token for DHCP Local Server Authentication on page 939](#)
- [Preventing Binding of Clients That Do Not Support Reconfigure Messages on page 939](#)
- [Requesting DHCP Local Server to Initiate Reconfiguration of Client Bindings on page 940](#)
- [Configuring Detection of DHCP Local Server Client Connectivity on page 941](#)
- [Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 942](#)
- [Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server on page 944](#)
- [Configuring Passwords for Usernames on page 945](#)
- [Creating Unique Usernames for DHCP Clients on page 946](#)
- [Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 949](#)

Using External AAA Authentication Services with DHCP

The extended DHCP local server, including DHCPv6 local server, and the extended DHCP relay agent, including DHCPv6 relay agent, support the use of external AAA authentication services, such as RADIUS, to authenticate DHCP clients. When the extended DHCP local server or relay agent receives a discover PDU from a client, the extended DHCP application contacts the AAA server to authenticate the DHCP client. The extended DHCP application

can obtain client addresses and DHCP configuration options from the external AAA authentication server.



NOTE: This section uses the term *extended DHCP application* to refer to both the extended DHCP local server and the extended DHCP relay agent.

The external authentication feature also supports AAA directed logout. If the external AAA service supports a user logout directive, the extended DHCP application honors the logout and responds as though it were requested by a CLI management command. All of the client state information and allocated resources are deleted at logout. The extended DHCP application supports directed logout using the list of configured authentication servers you specify with the **authentication-server** statement at the **[edit access profile profile-name]** hierarchy level.

You can configure either global authentication support or group-specific support.

You must configure the **username-include** statement to enable the use of authentication. The **password** statement is not required and does not cause DHCP to use authentication if the **username-include** statement is not included.

To configure DHCP local server and DHCP relay agent authentication support:

1. Specify that you want to configure authentication options.

- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# edit authentication
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit authentication
```

- For DHCPv6 local server:

```
[edit system services dhcp-local-server dhcpv6]
user@host# edit authentication
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit authentication
```

2. (Optional) Configure a password that authenticates the username to the external authentication service.

See [“Configuring Passwords for Usernames” on page 945](#).

3. (Optional) Configure optional features to create a unique username.

See [“Creating Unique Usernames for DHCP Clients” on page 946](#).

Related Documentation

- [Extended DHCP Local Server Overview on page 876](#)
- [Extended DHCP Relay Agent Overview on page 900](#)

- [DHCPv6 Local Server Overview on page 880](#)
- [DHCPv6 Relay Agent Overview on page 905](#)

Guidelines for Configuring Support for DHCP Duplicate Clients

This topic describes the guidelines for configuring DHCP to include the client subinterface in order to distinguish between duplicate clients (clients with the same MAC address or client ID) in a subscriber access or in a DHCP-managed environment.

When configuring DHCP duplicate client support, consider the following guidelines:

- The optional DHCP duplicate client support feature is used for DHCPv4 clients. For DHCPv6, client identification is independent of MAC address.
- For DHCP relay agent configuration:
 - DHCP relay must be configured to insert option 82, regardless of whether or not the incoming packet has option 82.
 - Option 82 must include the Agent Circuit ID suboption (suboption 1).
 - Option 82 must be the interface name, not the interface description.
 - DHCP server must echo option 82 in the server's reply. This is required because of the following:
 - The giaddr inserted by DHCP relay is the same for duplicate clients on different subinterfaces. The DHCP local server uses option 82 when allocating the IP address.
 - DHCP relay uses the echoed option 82 to learn the client subinterface and to construct the client key.
- For the Layer 3 wholesale model:
 - The wholesaler and retailer logical system/routing instances must have the same **duplicate-clients-on-interface** statement configuration.
 - For DHCP relay, the wholesaler and the retailer routing contexts must both be configured with the Agent Circuit ID suboption (suboption 1) in option 82.

Related Documentation

- [DHCP Duplicate Client Differentiation Using Client Subinterface Overview on page 882](#)
- [Configuring DHCP Duplicate Client Support on page 925](#)

Configuring DHCP Duplicate Client Support

You can optionally configure DHCP local server and DHCP relay to include a client subinterface when distinguishing between two clients that have the same MAC address or client ID. The configuration is a global setting for each logical system/routing instance.

To configure DHCP local server to include the client subinterface:

1. Specify that you want to configure DHCP local server.

```
[edit system services]
user@host# edit dhcp-local-server
```

2. Configure the optional duplicate client support.

```
[edit system services dhcp-local-server]
user@host# set duplicate-clients-on-interface
```

To configure DHCP relay agent to include the client subinterface:

1. Specify that you want to configure DHCP relay agent.

```
[edit forwarding-options]
user@host# edit dhcp-relay
```

2. Configure the optional duplicate client support.

```
[edit system services dhcp-relay]
user@host# set duplicate-clients-on-interface
```

**Related
Documentation**

- [DHCP Duplicate Client Differentiation Using Client Subinterface Overview on page 882](#)
- [Guidelines for Configuring Support for DHCP Duplicate Clients on page 925](#)

Grouping Interfaces with Common DHCP Configurations

You use the group feature to group together a set of interfaces and then apply a common DHCP configuration to the named interface group. The extended DHCP local server, DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent all support interface groups.

The following steps create a DHCP local server group; the steps are similar for the DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent.

To configure a DHCP local server interface group:

1. Specify that you want to configure DHCP local server.

```
[edit system services]
user@host# edit dhcp-local-server
```

2. Create the group and assign a name.

```
[edit system services dhcp-local-server]
user@host# edit group boston
```

3. Specify the names of one or more interfaces on which the extended DHCP application is enabled. You can repeat the `interface interface-name` statement to specify multiple interfaces within the group, but you cannot use the same interface in more than one group.

```
[edit system services dhcp-local-server group boston]
user@host# set interface fe-1/0/1.1
user@host# set interface fe-1/0/1.2
```

4. (Optional) You can use the `upto` option to specify a range of interfaces for a group.

```
[edit system services dhcp-local-server group boston]
user@host# set interface fe-1/0/1.3 upto fe-1/0/1.9
```

5. (Optional) You can use the **exclude** option to exclude a specific interface or a specified range of interfaces from the group. For example:

```
[edit system services dhcp-local-server group boston]
user@host# set interface fe-1/0/1.1 upto fe-1/0/1.102
user@host# set interface fe-1/0/1.6 exclude
user@host# set interface fe-1/0/1.70 upto fe-1/0/1.80 exclude
```

Related Documentation

- [Extended DHCP Local Server Overview on page 876](#)
- [Extended DHCP Relay Agent Overview on page 900](#)
- [DHCPv6 Local Server Overview on page 880](#)
- [DHCPv6 Relay Agent Overview on page 905](#)
- [Group-Specific DHCP Local Server Options on page 883](#)
- [Group-Specific DHCP Relay Options on page 906](#)
- [Guidelines for Configuring Interface Ranges on page 927](#)

Guidelines for Configuring Interface Ranges

This topic describes guidelines to consider when configuring interface ranges for named interface groups for DHCP local server and DHCP relay. The guidelines refer to the following configuration statement:

```
user@host# set interface interface-name upto upto-interface-name
```

- The start subunit, **interface *interface-name***, serves as the key for the stanza. The remaining configuration settings are considered attributes.
- If the subunit is not included, an implicit **.0** subunit is enforced. The implicit subunit is applied to all interfaces when autoconfiguration is enabled. For example, **interface ge-2/2/2** is treated as **interface ge-2/2/2.0**.
- Ranged entries contain the **upto** option, and the configuration applies to all interfaces within the specified range. The start of a ranged entry must be less than the end of the range. Discrete entries apply to a single interface, except in the case of autoconfiguration, in which a **0** (zero) subunit acts as a wildcard.
- Interface stanzas defined within the same router or switch context are dependent and can constrain each other—both DHCP local server and DHCP relay are considered. Interface stanzas defined across different router (switch) contexts are independent and do not constrain one another.
- Each interface stanza, whether discrete or ranged, has a unique start subunit across a given router context. For example, the following configuration is not allowed within the same group because **ge-1/0/0.10** is the start subunit for both.

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.10
```

- Two groups cannot share interface space. For example, the following configuration is not allowed because the three stanzas share the same space and interfere with one another—interface **ge-1/0/0.26** is common to all three.

```
dhcp-relay group diamond interface ge-1/0/0.10 upto ge-1/0/0.30
dhcp-local-server group ruby interface ge-1/0/0.25
dhcp-relay group sapphire interface ge-1/0/0.25 upto ge-1/0/0.35
```

- Two ranges cannot overlap, either within a group or across groups. Overlapping occurs when two interface ranges share common subunit space but neither range is a proper subset of the other. The following ranges overlap:

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.20 upto ge-1/0/0.40
```

- A range can contain multiple nested ranges. A nested range is a proper subset of another range. When ranges are nested, the smallest matching range applies.

In the following example, the three ranges nest properly:

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.12 upto ge-1/0/0.15 exclude
interface ge-1/0/0.25 upto ge-1/0/0.29 exclude
```

- Discrete interfaces take precedence over ranges. In the following example, interface **ge-1/0/0.20** takes precedence and enforces an interface client limit of 5.

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.15 upto ge-1/0/0.25 exclude
interface ge-1/0/0.20 overrides interface-client-limit 5
```

Related Documentation

- [Grouping Interfaces with Common DHCP Configurations on page 926](#)

Overriding Default DHCP Local Server Configuration Settings

Subscriber management enables you to override certain default DHCP and DHCPv6 local server configuration settings. You can override settings at the global level, for a named group of interfaces, or for a specific interface within a named group.

- To override global default DHCP local server configuration options, include the **overrides** statement and its subordinate statements at the **[edit system services dhcp-local-server]** or **[edit system services dhcp-local-server dhcpv6]** hierarchy level.
- To override DHCP local server configuration options for a named group of interfaces, include the statements at the **[edit system services dhcp-local-server group group-name]** or **[edit system services dhcp-local-server dhcpv6 group]** hierarchy level.
- To override DHCP local server configuration options for a specific interface within a named group of interfaces, include the statements at the **[edit system services dhcp-local-server group group-name interface]** or **[edit system services dhcp-local-server dhcpv6 group group-name interface]** hierarchy level.

To override default DHCP local server configuration settings:

1. Specify that you want to configure override options.

Global override:

```
[edit system services dhcp-local-server]
user@host# edit overrides
```

Group level override:

```
[edit system services dhcp-local-server]
user@host# edit group boston overrides
```

Per-interface override:

```
[edit system services dhcp-local-server]
user@host# edit group boston overrides interface fe-1/0/1.1
```

2. (Optional) Override the maximum number of DHCP clients allowed per interface.
See [“Specifying the Maximum Number of DHCP Clients Per Interface” on page 929](#).
3. (Optional) Override ARP table population in distrusted environments.
See [“Disabling ARP Table Population” on page 930](#).
4. (Optional) Configure DHCP client auto logout.
See [“Automatically Logging Out DHCP Clients” on page 932](#).
5. (Optional) Enable processing of information requests from clients.
See [“Enabling Processing of Client Information Requests” on page 933](#).
6. (Optional, DHCPv6 only) Specify a delegated pool name to use for DHCPv6 multiple address assignment.
See [“Specifying the Delegated Address Pool for IPv6 Prefix Assignment” on page 934](#).
7. (Optional, DHCPv6 only) Enable DHCPv6 rapid commit support.
See [“Enabling DHCPv6 Rapid Commit Support” on page 935](#).
8. (Optional) Delete DHCP override settings.
See [“Deleting DHCP Local Server and DHCP Relay Override Settings” on page 935](#).

**Related
Documentation**

- [Group-Specific DHCP Local Server Options on page 883](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 935](#)

Specifying the Maximum Number of DHCP Clients Per Interface

By default, there is no limit to the number of DHCP local server or DHCP relay clients allowed on an interface. However, you can override the default setting and specify the maximum number of clients allowed per interface, in the range 1 through 500,000. When the number of clients on the interface reaches the specified limit, no additional DHCP Discover PDUs or DHCPv6 Solicit PDUs are accepted. When the number of clients subsequently drops below the limit, new clients are again accepted.



NOTE: The maximum number of DHCP (and DHCPv6) local server clients or DHCP (and DHCPv6) relay clients can also be specified by Juniper Networks VSA 26-143 during client login. The VSA-specified value always takes precedence if the `interface-client-limit` statement specifies a different number.

If the VSA-specified value differs with each client login, DHCP uses the largest limit set by the VSA until there are no clients on the interface.

To configure the maximum number of DHCP clients allowed per interface:

1. Specify that you want to configure override options.

- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# edit overrides
```

- For DHCPv6 local server:

```
[edit system services dhcp-local-server dhcpv6]
user@host# edit overrides
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit overrides
```

2. Configure the maximum number of clients allowed per interface. (DHCP local server, DHCPv6 local server, DHCP relay agent and DHCPv6 relay agent all support the `interface-client-limit` statement.)

```
[edit system services dhcp-local-server overrides]
user@host# set interface-client-limit number
```

Related Documentation

- [Overriding Default DHCP Local Server Configuration Settings on page 928](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 935](#)
- [Extended DHCP Local Server Overview on page 876](#)
- [Extended DHCP Relay Agent Overview on page 900](#)

Disabling ARP Table Population

By default, DHCP populates the ARP table with the MAC address of a client when the client binding is established. However, you may choose to use the DHCP `no-arp` statement to hide the subscriber MAC address or the DHCP client MAC address, as it appears in ARP table entries.

When running in a trusted environment (that is, when not using the **no-arp** statement), DHCP populates the ARP table with unique MAC addresses contained within the DHCP PDU for each DHCP client:

Table 71: ARP Table in Trusted Environment

IP Address	MAC Address
Client 1 IP Address	MAC A
Client 2 IP Address	MAC B
Client 3 IP Address	MAC C

In distrusted environments, you can specify the **no-arp** statement to hide the MAC addresses of clients. When you specify the **no-arp** statement, DHCP does not automatically populate the ARP table with MAC address information from the DHCP PDU for each client. Instead, the system performs an ARP to obtain the MAC address of each client and obtains the MAC address of the immediately attached device (for example, a DSLAM). DHCP populates the ARP table with the same interface MAC address (for example, MAC X from a DSLAM interface) for each client:

Table 72: ARP Table in Distrusted Environment

IP Address	MAC Address
Client 1 IP Address	MAC X
Client 2 IP Address	MAC X
Client 3 IP Address	MAC X

To disable ARP table population:

- Specify that you want to configure override options.
 - For DHCP local server:


```
[edit system services dhcp-local-server]
user@host# edit overrides
```
 - For DHCP relay:


```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```
- Disable ARP table population with client-specific information. (DHCP local server and DHCP relay agent both support the **no-arp** statement.)
 - For DHCP local server:


```
[edit system services dhcp-local-server overrides]
user@host# set no-arp
```
 - For DHCP relay:

```
[edit forwarding-options dhcp-relay overrides]
user@host# set no-arp
```

**Related
Documentation**

- [Overriding Default DHCP Local Server Configuration Settings on page 928](#)
- [Extended DHCP Local Server Overview on page 876](#)
- [DHCPv6 Local Server Overview on page 880](#)
- [Extended DHCP Relay Agent Overview on page 900](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 935](#)

Automatically Logging Out DHCP Clients

You can configure the extended DHCP local server and extended DHCP relay to automatically log out DHCP clients. Auto logout immediately releases an existing client when DHCP receives a discover packet that has the same DHCP option 60 and DHCP option 82 information as the existing client. DHCP then releases the existing client IP address without waiting for the normal lease expiration.



NOTE: When the existing client is released, the new client undergoes the normal authentication process. The new client might not receive the same IP address as the original client.

To configure DHCP client auto logout:

1. Specify that you want to configure override options.
 - For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# edit overrides
```
 - For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```
2. Enable auto logout. (DHCP local server and DHCP relay agent both support the **client-discover-match** statement.)
 - For DHCP local server:

```
[edit system services dhcp-local-server overrides]
user@host# set client-discover-match
```
 - For DHCP relay:

```
[edit forwarding-options dhcp-relay overrides]
user@host# set client-discover-match
```



NOTE: If you change the auto logout configuration, existing clients continue to use the auto logout setting that was configured when they logged in. New clients use the new setting.

Related Documentation

- [DHCP Auto Logout Overview on page 888](#)
- [DHCP Relay Agent Option 82 Value for Auto Logout on page 962](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 935](#)
- [Extended DHCP Local Server Overview on page 876](#)
- [Extended DHCP Relay Agent Overview on page 900](#)

Enabling Processing of Client Information Requests

By default, DHCP local server and DHCPv6 local server do not respond to information request messages from the client. You can enable DHCP local server and DHCPv6 local server to process these messages and respond to them with an acknowledgment (ack or reply message, respectively) and the requested information.

DHCP relay agent automatically forwards the information request messages without modification to the configured server group by means of the interfaces configured for the respective server group. The messages are dropped if they are received on an unconfigured interface. DHCP relay proxy also supports forwarding these messages. You cannot disable forwarding of the information request messages.

Configure one or more local address pools if you want to use a local pool rather than one provided by AAA. See *Configuring an Address-Assignment Pool Name and Addresses*. For processing information request messages, the address configuration is not necessary. For DHCP local server, you must specify the IPv4 family; for DHCPv6 local server, you must specify the IPv6 family.

See *Configuring DHCP Client-Specific Attributes* for details about how to configure the information sought by clients that send information request messages.

To enable processing of DHCP client information request messages:

1. Specify that you want to configure override options.
 - For DHCP local server:


```
[edit system services dhcp-local-server overrides]
user@host# set process-inform
```
 - For DHCPv6 local server:


```
[edit system services dhcp-local-server dhcpv6 overrides]
user@host# set process-inform
```
2. (Optional) Specify a pool name from which DHCP information is returned to the client.
 - For DHCP local server:


```
[edit system services dhcp-local-server overrides process-inform]
```

```
user@host# set pool pool-name
```

- For DHCPv6 local server:

```
[edit system services dhcp-local-server dhcpv6 overrides process-inform]
```

```
user@host# set pool pool-name
```

Related Documentation

- [Overriding Default DHCP Local Server Configuration Settings on page 928](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 935](#)
- [Extended DHCP Local Server Overview on page 876](#)
- [Extended DHCP Relay Agent Overview on page 900](#)

Specifying the Delegated Address Pool for IPv6 Prefix Assignment

You can explicitly specify a delegated address pool:

- On routers—Subscriber management uses the pool to assign IPv6 prefixes for subscribers. You can specify the delegated address pool globally, for a specific group of interfaces, or for a particular interface.
- On switches—DHCP management uses the pool to assign IPv6 prefixes for DHCP clients. You can specify the delegated address pool globally, for a specific group of interfaces, or for a particular interface.



NOTE: You can also use by Juniper Networks VSA 26-161 to specify the delegated address pool. The VSA-specified value always takes precedence over the delegated-address statement.

To configure the delegated address pool for DHCPv6 local server:

1. Specify that you want to configure override options.

```
[edit system services dhcp-local-server dhcpv6]  
user@host# edit overrides
```

2. Configure the delegated address pool.

```
[edit system services dhcp-local-server dhcpv6 overrides]  
user@host# set delegated-pool paris-cable-12
```

Related Documentation

- [Overriding Default DHCP Local Server Configuration Settings on page 928](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 935](#)
- [Extended DHCP Local Server Overview on page 876](#)
- [Extended DHCP Relay Agent Overview on page 900](#)

Enabling DHCPv6 Rapid Commit Support

You can configure the extended DHCPv6 local server to support the DHCPv6 Rapid Commit option (DHCPv6 option 14). When rapid commit is enabled on the extended DHCPv6 local server, the server recognizes the Rapid Commit option in Solicit messages sent from the DHCPv6 client. (DHCPv6 clients are configured separately to include the DHCPv6 Rapid Commit option in the Solicit messages.) The server and client then use a two-message exchange (Solicit and Reply) to configure clients, rather than the default four-method exchange (Solicit, Advertise, Request, and Reply). The two-message exchange provides faster client configuration, and is beneficial in environments in which networks are under a heavy load.

You can configure the DHCPv6 local server to support the Rapid Commit option globally, for a specific group, or for a specific interface. By default, Rapid Commit support is disabled on the DHCPv6 local server.

To configure the DHCPv6 local server to support the DHCPv6 Rapid Commit option:

1. Specify that you want to configure override options.

```
[edit system services dhcp-local-server dhcpv6]
user@host# edit overrides
```

2. Enable rapid commit support.

```
[edit system services dhcp-local-server dhcpv6 overrides]
user@host# set rapid-commit
```

Related Documentation

- [Overriding Default DHCP Local Server Configuration Settings on page 928](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 935](#)
- [Extended DHCP Local Server Overview on page 876](#)

Deleting DHCP Local Server and DHCP Relay Override Settings

You can delete override settings for DHCP local server and DHCP relay globally, for a named group, or for a specific interface within a named group. You can delete a specific override setting or all overrides.

- To delete a specific DHCP override setting at a particular hierarchy level, include the **overrides** statement with the appropriate subordinate statements. For example, to delete the DHCP local server override **no-arp** setting for a group named **marin20**:

```
[edit system services dhcp-local-server]
user@host# delete group marin20 overrides no-arp
```

- To delete all DHCP override settings at a hierarchy level, include the **overrides** statement without any subordinate statements. For example, to delete all DHCP relay overrides for interface **fxp0.0**, which is in group **marin20**:

```
[edit forwarding-options dhcp-relay]
user@host# delete group marin20 interface fxp0.0 overrides
```

Related Documentation

- [Overriding Default DHCP Local Server Configuration Settings on page 928](#)
- [Extended DHCP Local Server Overview on page 876](#)
- [Extended DHCP Relay Agent Overview on page 900](#)

Configuring Extended DHCP Local Server Dynamic Client Reconfiguration

The DHCP local server can initiate reconfiguration of its clients to avoid extended outages because of server configuration changes. In addition to requesting that the DHCP local server initiate reconfiguration, you can specify the reconfiguration behavior.

To configure dynamic reconfiguration of DHCP clients:

1. Enable dynamic reconfiguration with default values for all clients.

For DHCPv4:

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

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6]
user@host# set reconfigure
```

2. (Optional) Override the global configuration for a particular group of clients.

For DHCPv4:

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

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 group group-name]
user@host# set reconfigure
```

3. (Optional) Configure how the server attempts reconfiguration.

See [“Configuring Dynamic Reconfiguration Attempts for DHCP Clients” on page 937](#).

4. (Optional) Configure the response to a failed reconfiguration.

See [“Configuring Deletion of the Client When Dynamic Reconfiguration Fails” on page 938](#).

5. (Optional) Configure the behavior in response to a RADIUS-initiated disconnect.

See [“Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect” on page 938](#).

6. (Optional) Configure a token for rudimentary server authentication.

See [“Configuring a Token for DHCP Local Server Authentication” on page 939](#).

7. (Optional) Initiate reconfiguration of some or all client bindings.

See [“Requesting DHCP Local Server to Initiate Reconfiguration of Client Bindings”](#) on page 940.

8. (Optional) Prevent DHCPv6 clients from binding if they do not support reconfigure messages.

See [“Preventing Binding of Clients That Do Not Support Reconfigure Messages”](#) on page 939.

Configuring Dynamic Reconfiguration Attempts for DHCP Clients

You can configure how many attempts the local server makes to initiate reconfiguration of the DHCP client by sending `forcerenew` messages. You can also specify how long the server waits between attempts. By default, eight attempts are made and the initial interval is two seconds.

Each successive attempt doubles the interval between attempts. For example, if the first value is 2, the first retry is attempted 2 seconds after the first attempt fails. The second retry is attempted 4 seconds after the first retry fails. The third retry is attempted 8 seconds after the second retry fails, and so on. A group configuration takes precedence over a DHCP local server configuration.

(Optional) To configure DHCP local server reconfiguration behavior for all DHCP clients:

1. Specify the number of reconfiguration attempts.

For DHCPv4:

```
[edit system services dhcp-local-server reconfigure]
user@host# set attempts 5
```

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 reconfigure]
user@host# set attempts 5
```

2. Specify the interval between reconfiguration attempts.

For DHCPv4:

```
[edit system services dhcp-local-server reconfigure]
user@host# set timeout 8
```

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 reconfigure]
user@host# set timeout 8
```

To override the global configuration for a particular group of clients, include the statements at the `[edit system services dhcp-local-server group group-name reconfigure]` hierarchy level or the `[edit system services dhcpv6 dhcp-local-server group group-name reconfigure]` hierarchy level.

Related Documentation

- [Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 936](#)
- [attempts on page 999](#)
- [timeout on page 1062](#)

Configuring Deletion of the Client When Dynamic Reconfiguration Fails

You can configure the local server to delete the client when the maximum number of reconfiguration attempts has been made without success. By default, the client's original configuration is restored.

(Optional) To configure the DHCP local server to delete the client when reconfiguration is not successful, for all clients:

- Specify the client deletion.

For DHCPv4:

```
[edit system services dhcp-local-server reconfigure]
user@host# set clear-on-abort
```

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 reconfigure]
user@host# set clear-on-abort
```

To override the global configuration for a particular group of clients, include the statement at the `[edit system services dhcp-local-server group group-name reconfigure]` hierarchy level or the `[edit system services dhcpv6 dhcp-local-server group group-name reconfigure]` hierarchy level.

Related Documentation

- [Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 936](#)
- [clear-on-abort on page 1003](#)

Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect

You can configure the local server to reconfigure the client when the client receives a RADIUS-initiated disconnect. By default, the client is deleted when a RADIUS-initiated disconnect is received.

(Optional) To configure the DHCP local server to reconfigure the client instead of deleting the client when a RADIUS-initiated disconnect is received, for all clients:

- Specify the RADIUS-initiated disconnect trigger.

For DHCPv4:

```
[edit system services dhcp-local-server reconfigure trigger]
user@host# set radius-disconnect
```

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 reconfigure trigger]
user@host# set radius-disconnect
```

To override the global configuration for a particular group of clients, include the statement at the `[edit system services dhcp-local-server group group-name reconfigure trigger]`

hierarchy level or the `[edit system services dhcpv6 dhcp-local-server group group-name reconfigure trigger]` hierarchy level.

- Related Documentation**
- [Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 936](#)
 - [radius-disconnect on page 1051](#)
 - [trigger on page 1065](#)

Configuring a Token for DHCP Local Server Authentication

You can configure the local server to include a constant, unencoded token in the DHCP forcerenew message as part of the authentication option it sends to clients. The client compares the received token with a token already configured on the client. If the tokens do not match, the DHCP client discards the forcerenew message. Use of the token provides rudimentary protection against inadvertently instantiated DHCP servers.

(Optional) To configure the DHCP local server to include a token in the forcerenew message sent to the client, for all clients:

- Specify the token.

For DHCPv4:

```
[edit system services dhcp-local-server reconfigure]
user@host# set token 8ysIU9E32k8r
```

For DHCPv6:

```
[edit system services dhcp-local-server dhcpv6 reconfigure]
user@host# set token 8ysIU9E32k8r
```

To override the global configuration for a particular group of clients, include the statement at the `[edit system services dhcp-local-server group group-name reconfigure]` hierarchy level or the `[edit system services dhcpv6 dhcp-local-server group group-name reconfigure]` hierarchy level.

- Related Documentation**
- [Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 936](#)
 - [token on page 1063](#)

Preventing Binding of Clients That Do Not Support Reconfigure Messages

The DHCPv6 client and server negotiate the use of reconfigure messages. When the client can accept reconfigure messages from the server, then the client includes the Reconfigure Accept option in both solicit and request messages sent to the server.

By default, the DHCPv6 server accepts solicit messages from clients regardless of whether they support reconfiguration. You can specify that the server require clients to accept reconfigure messages. In this case, the DHCPv6 server includes the Reconfigure Accept option in both advertise and reply messages when reconfiguration is configured for the client interface. Solicit messages from nonsupporting clients are discarded and the clients are not allowed to bind.

(Optional) To configure the DHCPv6 local server to require that all clients accept reconfiguration:

- Specify strict reconfiguration.

```
[edit system services dhcp-local-server dhcpv6 reconfigure]  
user@host# set strict
```

To override the global configuration for a particular group of clients, include the statement at the `[edit system services dhcp-local-server dhcpv6 group group-name reconfigure]` hierarchy level.

The `show dhcpv6 server statistics` command displays a count of solicit messages that the server has discarded.

**Related
Documentation**

- [Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 936](#)
- [strict on page 1059](#)

Requesting DHCP Local Server to Initiate Reconfiguration of Client Bindings

You can request that the DHCP local server initiate reconfiguration of all of clients or only specified clients.

To request reconfiguration of all clients:

- Specify the `all` option.

For DHCPv4:

```
user@host> request dhcp server reconfigure all
```

For DHCPv6:

```
user@host> request dhcpv6 server reconfigure all
```

You can use any of the following methods to request reconfiguration of specific clients:

- Specify the IP address of the DHCP client.

For DHCPv4:

```
user@host> request dhcp server reconfigure 192.168.27.3
```

For DHCPv6:

```
user@host> request dhcpv6 server reconfigure 2001:bd8:1111:2222::
```

- Specify the client ID of a DHCPv6 client.

```
user@host> request dhcpv6 server reconfigure  
LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
```

- Specify the session ID of a DHCPv6 client.

```
user@host> request dhcpv6 server reconfigure 5
```

- Specify the MAC address of a DHCPv4 client.

```
user@host> request dhcp server reconfigure 12:23:34:45:56:67
```

- Specify an interface; reconfiguration is attempted for all clients on this interface.

```
user@host> request dhcp server reconfigure interface fe-0/0/0.100
```

- Specify a logical system; reconfiguration is attempted for all clients or the specified clients in this logical system.

```
user@host> request dhcp server reconfigure all logical-system ls-bldg5
```

- Specify a routing instance; reconfiguration is attempted for all clients or the specified clients in this routing instance.

```
user@host> request dhcp server reconfigure all routing-instance ri-boston
```

Related Documentation

- [Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 936](#)
- [request dhcp server reconfigure on page 1168](#)

Configuring Detection of DHCP Local Server Client Connectivity

Liveness detection for DHCP subscriber IP sessions or DHCP client IP sessions utilizes an active liveness detection protocol to institute liveness detection checks for relevant clients. Clients must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, then the liveness detection check fails and a failure action is implemented.



NOTE: You can also configure DHCP liveness detection for DHCP relay.

To configure liveness detection for DHCP local server:

1. Specify that you want to configure liveness detection.

- For DHCP global configuration:

```
[edit system services dhcp-local-server]
user@host# edit liveness-detection
```

- For DHCP group configuration:

```
[edit system services dhcp-local-server group group-name]
user@host# edit liveness-detection
```



NOTE: Liveness detection is also supported for DHCPv6 configurations. To configure DHCPv6 liveness detection, include the [liveness-detection](#) statement, and any subsequent configuration statements, at the `[edit system services dhcp-local-server dhcpv6]` or `[edit system services dhcp-local-server dhcpv6 group group-name]` hierarchy level.

2. Specify that you want to configure the liveness detection method.

- For DHCP global configuration:

```
[edit system services dhcp-local-server liveness-detection]
user@host# edit method
```

- For DHCP group configuration:

```
[edit system services dhcp-local-server group group-name liveness-detection]
user@host# edit method
```

3. Specify the liveness detection method that you want DHCP to use.



NOTE: In this release, the only method supported for liveness detection is Bidirectional Forwarding Detection (BFD).

- For DHCP global configuration:

```
[edit system services dhcp-local-server liveness-detection method]
user@host# edit bfd
```

- For DHCP group configuration:

```
[edit system services dhcp-local-server group group-name liveness-detection method]
user@host# edit bfd
```

4. Configure the liveness detection method as desired.

See [“Example: Configuring Group Liveness Detection for DHCP Local Server Clients” on page 913](#) for an example of how to configure DHCPv4 groups for DHCP local server liveness detection.

5. Configure the action the router takes when a liveness detection failure occurs.

- For DHCP global configuration:

```
[edit system services dhcp-local-server liveness-detection]
user@host# edit failure-action action
```

- For DHCP group configuration:

```
[edit system services dhcp-local-server group group-name liveness-detection]
user@host# edit failure-action action
```

Related Documentation

- [DHCP Liveness Detection Overview on page 980](#)
- [Extended DHCP Local Server Overview on page 876](#)
- [Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 981](#)
- [Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913](#)
- [Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients](#)

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

This topic describes how to attach a dynamic profile to a DHCP subscriber interface or a DHCP client interface. When a DHCP subscriber or DHCP client logs in, the specified

dynamic profile is instantiated and the services defined in the profile are applied to the interface.

This topic contains the following sections:

- [Attaching a Dynamic Profile to All DHCP Subscriber or All DHCP Client Interfaces on page 943](#)
- [Attaching a Dynamic Profile to a Group of DHCP Subscriber Interfaces or a Group of DHCP Client Interfaces on page 943](#)

Attaching a Dynamic Profile to All DHCP Subscriber or All DHCP Client Interfaces

To attach a dynamic profile to all DHCP subscriber or all DHCP client interfaces:

1. At the DHCP configuration hierarchy, use the **dynamic-profile** statement to specify the name of the dynamic profile to attach to all interfaces.
 - For DHCP local server:


```
[edit system services dhcp-local-server]
user@host# set dynamic-profile vod-profile-22
```
 - For DHCP relay agent:


```
[edit forwarding-options dhcp-relay]
user@host# set dynamic-profile vod-profile-west
```
2. (Routers only) Optionally, you can configure the attribute to use when attaching the specified profile.

You can include either the **aggregate-clients** option to enable multiple DHCP subscribers to share the same VLAN logical interface, or the **use-primary** option to specify that the primary dynamic profile is used. The **aggregate-clients** option does not apply to demux subscriber interfaces. The two options are mutually exclusive.

- To enable multiple subscribers to share the same VLAN logical interface:

```
[edit system services dhcp-local-server dynamic-profile]
user@host# set aggregate-clients merge
```

- To use the primary dynamic profile:

```
[edit forwarding-options dhcp-relay dynamic-profile]
user@host# set use-primary subscriber_profile
```

Attaching a Dynamic Profile to a Group of DHCP Subscriber Interfaces or a Group of DHCP Client Interfaces

Before you begin:

- Configure the interface group.

See [“Grouping Interfaces with Common DHCP Configurations” on page 926](#).

To attach a dynamic profile to a group of interfaces:

1. At the DHCP configuration hierarchy, specify the name of the interface group and the dynamic profile to attach to the group.

- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# set group boston dynamic-profile vod-profile-42
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# set group quebec dynamic-profile vod-profile-east
```

2. (Routers only) Optionally, you can configure the attribute to use when attaching the specified profile.

You can include either the **aggregate-clients** option to enable multiple DHCP subscribers to share the same VLAN logical interface, or the **use-primary** option to specify that the primary dynamic profile is used. The **aggregate-clients** option does not apply to demux subscriber interfaces. The two options are mutually exclusive.

- To enable multiple subscribers to share the same VLAN logical interface:

```
[edit system services dhcp-local-server dynamic-profile]
user@host# set aggregate-clients merge
```

- To use the primary dynamic profile:

```
[edit forwarding-options dhcp-relay dynamic-profile]
user@host# set use-primary subscriber_profile
```

Related Documentation

- *Dynamic Profiles Overview*
- *Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview*
- *Example: Configuring Dynamic Subscriber Interfaces on IP Demux Interfaces*

Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server

You can configure how DHCP local server handles DHCP snooped packets. Depending on the configuration, DHCP local server either forwards or drops the snooped packets it receives.

Table 73 on page 945 indicates the action the router takes for DHCP local server snooped packets.



NOTE: Configured interfaces are those interfaces that have been configured with the **group** statement in the [edit system services dhcp-local-server] hierarchy. Non-configured interfaces are those that are in the logical system/routing instance but have not been configured by the **group** statement.

Table 73: Actions for DHCP Local Server Snooped Packets

forward-snooped-clients Configuration	Action on Configured Interfaces	Action on Non-Configured Interfaces
forward-snooped-clients not configured	dropped	dropped
all-interfaces	forwarded	forwarded
configured-interfaces	forwarded	dropped
non-configured-interfaces	dropped	forwarded

To configure DHCP snooped packet forwarding for DHCP local server:

1. Specify that you want to configure DHCP local server.

```
[edit]
user@host# edit system services dhcp-local-server
```

2. Enable DHCP snooped packet forwarding for DHCP local server.

```
[edit system services dhcp-local-server]
user@host# edit forward-snooped-clients
```

3. Specify the interfaces that are supported for snooped packet forwarding.

```
[edit system services dhcp-local-server forward-snooped-clients]
user@host# set (all-interfaces | configured-interfaces | non-configured-interfaces)
```

For example, to configure DHCP local server to forward DHCP snooped packets on only configured interfaces:

```
[edit]
system {
  services {
    dhcp-local-server {
      forward-snooped-clients configured-interfaces;
    }
  }
}
```

Related Documentation

- [DHCP Snooping Support on page 887](#)

Configuring Passwords for Usernames

You can configure an optional password that the extended DHCP application presents to the external AAA authentication service to authenticate the specified username.

To configure a password that authenticates the username:

1. Specify that you want to configure authentication options.
 - For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# edit authentication
```

- For DHCPv6 local server:

```
[edit system services dhcp-local-server dhcpv6]
user@host# edit authentication
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit authentication
```

2. Configure the password. (DHCP local server, DHCPv6 local server, and DHCP relay agent all support the **password** statement.)

```
[edit system services dhcp-local-server authentication]
user@host# set password myPassword1234
```

Related Documentation

- [Extended DHCP Local Server Overview on page 876](#)
- [DHCPv6 Local Server Overview on page 880](#)
- [Extended DHCP Relay Agent Overview on page 900](#)
- [Using External AAA Authentication Services with DHCP on page 923](#)
- For information about supported characters in passwords, see “Configuring Special Requirements for Plain-Text Passwords” in the *Junos OS System Basics Configuration Guide*

Creating Unique Usernames for DHCP Clients

You can configure the extended DHCP application to include additional information in the username that is passed to the external AAA authentication service when the DHCP client logs in. This additional information enables you to construct usernames that uniquely identify subscribers (DHCP clients).



NOTE: If you do not include a username in the authentication configuration, the router (or switch) does not perform authentication; however, the IP address is provided by the local pool if it is configured.

When you use the DHCPv6 local server, you must configure authentication and the client username; otherwise client login fails.

The following list describes the optional information that you can include as part of the username:

- **circuit-type**—The circuit type used by the DHCP client, for example **enet**.
- **client-id**—The client identifier option (option 1). (DHCPv6 local server DHCPv6 relay agent only)

- **delimiter**—The delimiter character that separates components that make up the concatenated username. The default delimiter is a period (.). The semicolon (;) is not supported as a delimiter character.
- **domain-name**—The client domain name as a string. The router adds the @ delimiter to the username.
- **interface-name**—The interface name, including the interface device and associated VLAN IDs.
- **logical-system-name**—The name of the logical system, if the receiving interface is in a logical system.
- **mac-address**—The client MAC address, in a string of the format *xxxx.xxxx.xxxx*. (Not supported for DHCPv6 local server)
- **option-60**—The portion of the option 60 payload that follows the length field. (Not supported for DHCPv6 local server)
- **option-82 <circuit-id> <remote-id>**—The specified contents of the option 82 payload. (Not supported for DHCPv6 local server)
 - **circuit-id**—The payload of the Agent Circuit ID suboption.
 - **remote-id**—The payload of the Agent Remote ID suboption.
 - Both **circuit-id** and **remote-id**—The payloads of both suboptions, in the format: **circuit-id[delimiter]remote-id**.
 - Neither **circuit-id** or **remote-id**—The raw payload of the option 82 from the PDU is concatenated to the username.



NOTE: For DHCP relay agent, the option 82 value used in creating the username is based on the option 82 value that is encoded in the outgoing (relayed) PDU.

- **relay-agent-interface-id**—The Interface-ID option (option 18). (DHCPv6 local server only)
- **relay-agent-remote-id**—The DHCPv6 Relay Agent Remote-ID option (option 37). (DHCPv6 local server only)
- **relay-agent-subscriber-id**—(On routers only) The DHCPv6 Relay Agent Subscriber-ID option (option 38). (DHCPv6 local server only)
- **routing-instance-name**—The name of the routing instance, if the receiving interface is in a routing instance.
- **user-prefix**—A string indicating the user prefix.

The router (switch) creates the unique username by including the specified additional information in the following order, with the fields separated by a delimiter.

For DHCP local server and DHCP relay agent:

```
user-prefix[delimiter]mac-address[delimiter]logical-system-name[delimiter]  
routing-instance-name[delimiter]circuit-type[delimiter]interface-name[delimiter]option-82[delimiter]  
option-60@domain-name
```

For DHCPv6 local server:

```
user-prefix[delimiter]logical-system-name[delimiter]routing-instance-name[delimiter]  
circuit-type[delimiter]interface-name[delimiter]relay-agent-remote-id[delimiter]  
relay-agent-subscriber-id[delimiter]relay-agent-interface-id[delimiter]client-id@domain-name
```

To configure a unique username:

1. Specify that you want to configure authentication.

- For DHCP local server:

```
[edit system services dhcp-local-server]  
user@host# edit authentication
```

- For DHCPv6 local server:

```
[edit system services dhcp-local-server dhcpv6]  
user@host# edit authentication
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]  
user@host# edit authentication
```

2. Specify that you want to include optional information in the username. (DHCP local server, DHCPv6 local server, and DHCP relay agent all support the **username-include** statement.)

```
[edit system services dhcp-local-server authentication]  
user@host# set username-include
```

3. (Optional) Specify the optional information you want to include in the username.

```
[edit system services dhcp-local-server authentication username-include]  
user@host# set username-include circuit-type  
user@host# set username-include domain-name isp55.com  
user@host# set username-include mac-address  
user@host# set username-include user-prefix wallybrown
```

The previous **username-include** configuration produces this unique username:

wallybrown.0090.1a01.1234.enet@isp55.com

Related Documentation

- [Extended DHCP Local Server Overview on page 876](#)
- [DHCPv6 Local Server Overview on page 880](#)
- [Extended DHCP Relay Agent Overview on page 900](#)
- [Using External AAA Authentication Services with DHCP on page 923](#)

Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use

You can specify the match order in which the extended DHCP local server uses the client data to determine the address-assignment pool that provides the IP address and configuration for a DHCP client. You use the **pool-match-order** statement to specify the match order. If you do not specify the **pool-match-order**, the router (or switch) uses the default **ip-address-first** matching to select the address pool. After DHCP local server determines the address assignment pool to use, the server performs the matching based on the criteria you specified in the pool configuration.

In the default **ip-address-first** matching, the server selects the address-assignment pool to use by matching the IP address in the client DHCP request with the network address of the address-assignment pool. If the client request contains the gateway IP address (giaddr), the local server matches the giaddr to the address-assignment pool's address. If there is no giaddr in the request, then the DHCP local server matches the IP address of the receiving interface to the address of the address-assignment pool.

In **external-authority** matching, the DHCP local server receives the address assignment from an external authority, such as RADIUS or Diameter. If RADIUS is the external authority, the DHCP local server uses the Framed-IPv6-Pool attribute (RADIUS attribute 100) to select the pool. If Diameter is the external authority, the server uses the Diameter counterpart of the Framed-IPv6-Pool attribute to determine the pool.

For IPv4 address-assignment pools, you can optionally configure the extended DHCP local server to match the DHCP relay agent information option (option 82) in the client DHCP packets to a named range in the address-assignment pool used for the client. Named ranges are subsets within the overall address-assignment pool address range, which you can configure when you create the address-assignment pool.



NOTE: To use the DHCP local server option 82 matching feature with an IPv4 address-assignment pool, you must ensure that the **option-82** statement is included in the **dhcp-attributes** statement for the address-assignment pool.

To configure the matching order the extended DHCP local server uses to determine the address-assignment pool used for a client:

1. Access the **pool-match-order** configuration.

```
[edit system services dhcp-local-server]
user@host# edit pool-match-order
```

2. Specify the pool matching methods in the order in which the router (switch) performs the methods. You can specify the methods in any order. All methods are optional—the router (switch) uses the **ip-address-first** method by default.

- Configure the router (switch) to use an external addressing authority.

```
[edit system services dhcp-local-server pool-match-order]
user@host# set external-authority
```

- Configure the router (switch) to use the ip-address-first method.

```
[edit system services dhcp-local-server pool-match-order]  
user@host# set ip-address-first
```
- (IPv4 address-assignment pools only) Specify the option 82 matching method.

```
[edit system services dhcp-local-server pool-match-order]  
user@host# set option-82
```

**Related
Documentation**

- [Address-Assignment Pools Overview on page 890](#)
- [Configuring Address-Assignment Pools](#)
- [Extended DHCP Local Server Overview on page 876](#)
- [Example: Extended DHCP Local Server Configuration with Optional Pool Matching on page 912](#)

Configuration Tasks for DHCP Relay Agent

- [Using External AAA Authentication Services with DHCP on page 951](#)
- [Configuring DHCP Duplicate Client Support on page 952](#)
- [Grouping Interfaces with Common DHCP Configurations on page 953](#)
- [Guidelines for Configuring Interface Ranges on page 954](#)
- [Overriding the Default DHCP Relay Configuration Settings on page 955](#)
- [Overwriting giaddr Information on page 957](#)
- [Replacing the DHCP Relay Request and Release Packet Source Address on page 957](#)
- [Overriding Option 82 Information on page 957](#)
- [Using Layer 2 Unicast Transmission for DHCP Packets on page 958](#)
- [Trusting Option 82 Information on page 958](#)
- [Disabling ARP Table Population on page 959](#)
- [Specifying the Maximum Number of DHCP Clients Per Interface on page 960](#)
- [Automatically Logging Out DHCP Clients on page 961](#)
- [DHCP Relay Agent Option 82 Value for Auto Logout on page 962](#)
- [Configuring DHCP Snooping for DHCP Relay Agent on page 963](#)
- [Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 964](#)
- [Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 969](#)
- [Sending Release Messages When Clients Are Deleted on page 971](#)
- [Disabling Automatic Binding of Stray DHCP Requests on page 972](#)
- [Enabling and Disabling Insertion of Option 82 Information on page 973](#)
- [Configuring Server Groups on page 976](#)
- [Configuring Active Server Groups on page 977](#)

- [Enabling DHCP Relay Proxy Mode on page 977](#)
- [Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 978](#)
- [Inserting DHCPv6 Interface-ID Option \(Option 18\) In DHCPv6 Packets on page 979](#)
- [DHCP Liveness Detection Overview on page 980](#)
- [Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 981](#)
- [Disabling DHCP Relay on page 983](#)

Using External AAA Authentication Services with DHCP

The extended DHCP local server, including DHCPv6 local server, and the extended DHCP relay agent, including DHCPv6 relay agent, support the use of external AAA authentication services, such as RADIUS, to authenticate DHCP clients. When the extended DHCP local server or relay agent receives a discover PDU from a client, the extended DHCP application contacts the AAA server to authenticate the DHCP client. The extended DHCP application can obtain client addresses and DHCP configuration options from the external AAA authentication server.



NOTE: This section uses the term *extended DHCP application* to refer to both the extended DHCP local server and the extended DHCP relay agent.

The external authentication feature also supports AAA directed logout. If the external AAA service supports a user logout directive, the extended DHCP application honors the logout and responds as though it were requested by a CLI management command. All of the client state information and allocated resources are deleted at logout. The extended DHCP application supports directed logout using the list of configured authentication servers you specify with the **authentication-server** statement at the **[edit access profile profile-name]** hierarchy level.

You can configure either global authentication support or group-specific support.

You must configure the **username-include** statement to enable the use of authentication. The **password** statement is not required and does not cause DHCP to use authentication if the **username-include** statement is not included.

To configure DHCP local server and DHCP relay agent authentication support:

1. Specify that you want to configure authentication options.

- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# edit authentication
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit authentication
```

- For DHCPv6 local server:

```
[edit system services dhcp-local-server dhcpv6]
user@host# edit authentication
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit authentication
```

2. (Optional) Configure a password that authenticates the username to the external authentication service.

See [“Configuring Passwords for Usernames” on page 945](#).

3. (Optional) Configure optional features to create a unique username.

See [“Creating Unique Usernames for DHCP Clients” on page 946](#).

Related Documentation

- [Extended DHCP Local Server Overview on page 876](#)
- [Extended DHCP Relay Agent Overview on page 900](#)
- [DHCPv6 Local Server Overview on page 880](#)
- [DHCPv6 Relay Agent Overview on page 905](#)

Configuring DHCP Duplicate Client Support

You can optionally configure DHCP local server and DHCP relay to include a client subinterface when distinguishing between two clients that have the same MAC address or client ID. The configuration is a global setting for each logical system/routing instance.

To configure DHCP local server to include the client subinterface:

1. Specify that you want to configure DHCP local server.

```
[edit system services]
user@host# edit dhcp-local-server
```

2. Configure the optional duplicate client support.

```
[edit system services dhcp-local-server]
user@host# set duplicate-clients-on-interface
```

To configure DHCP relay agent to include the client subinterface:

1. Specify that you want to configure DHCP relay agent.

```
[edit forwarding-options]
user@host# edit dhcp-relay
```

2. Configure the optional duplicate client support.

```
[edit system services dhcp-relay]
user@host# set duplicate-clients-on-interface
```

Related Documentation

- [DHCP Duplicate Client Differentiation Using Client Subinterface Overview on page 882](#)
- [Guidelines for Configuring Support for DHCP Duplicate Clients on page 925](#)

Grouping Interfaces with Common DHCP Configurations

You use the group feature to group together a set of interfaces and then apply a common DHCP configuration to the named interface group. The extended DHCP local server, DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent all support interface groups.

The following steps create a DHCP local server group; the steps are similar for the DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent.

To configure a DHCP local server interface group:

1. Specify that you want to configure DHCP local server.

```
[edit system services]
user@host# edit dhcp-local-server
```

2. Create the group and assign a name.

```
[edit system services dhcp-local-server]
user@host# edit group boston
```

3. Specify the names of one or more interfaces on which the extended DHCP application is enabled. You can repeat the `interface interface-name` statement to specify multiple interfaces within the group, but you cannot use the same interface in more than one group.

```
[edit system services dhcp-local-server group boston]
user@host# set interface fe-1/0/1.1
user@host# set interface fe-1/0/1.2
```

4. (Optional) You can use the `upto` option to specify a range of interfaces for a group.

```
[edit system services dhcp-local-server group boston]
user@host# set interface fe-1/0/1.3 upto fe-1/0/1.9
```

5. (Optional) You can use the `exclude` option to exclude a specific interface or a specified range of interfaces from the group. For example:

```
[edit system services dhcp-local-server group boston]
user@host# set interface fe-1/0/1.1 upto fe-1/0/1.102
user@host# set interface fe-1/0/1.6 exclude
user@host# set interface fe-1/0/1.70 upto fe-1/0/1.80 exclude
```

Related Documentation

- [Extended DHCP Local Server Overview on page 876](#)
- [Extended DHCP Relay Agent Overview on page 900](#)
- [DHCPv6 Local Server Overview on page 880](#)
- [DHCPv6 Relay Agent Overview on page 905](#)
- [Group-Specific DHCP Local Server Options on page 883](#)
- [Group-Specific DHCP Relay Options on page 906](#)
- [Guidelines for Configuring Interface Ranges on page 927](#)

Guidelines for Configuring Interface Ranges

This topic describes guidelines to consider when configuring interface ranges for named interface groups for DHCP local server and DHCP relay. The guidelines refer to the following configuration statement:

```
user@host# set interface interface-name upto upto-interface-name
```

- The start subunit, **interface *interface-name***, serves as the key for the stanza. The remaining configuration settings are considered attributes.
- If the subunit is not included, an implicit **.0** subunit is enforced. The implicit subunit is applied to all interfaces when autoconfiguration is enabled. For example, **interface ge-2/2/2** is treated as **interface ge-2/2/2.0**.
- Ranged entries contain the **upto** option, and the configuration applies to all interfaces within the specified range. The start of a ranged entry must be less than the end of the range. Discrete entries apply to a single interface, except in the case of autoconfiguration, in which a **0** (zero) subunit acts as a wildcard.
- Interface stanzas defined within the same router or switch context are dependent and can constrain each other—both DHCP local server and DHCP relay are considered. Interface stanzas defined across different router (switch) contexts are independent and do not constrain one another.
- Each interface stanza, whether discrete or ranged, has a unique start subunit across a given router context. For example, the following configuration is not allowed within the same group because **ge-1/0/0.10** is the start subunit for both.

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.10
```

- Two groups cannot share interface space. For example, the following configuration is not allowed because the three stanzas share the same space and interfere with one another—interface **ge-1/0/0.26** is common to all three.

```
dhcp-relay group diamond interface ge-1/0/0.10 upto ge-1/0/0.30
dhcp-local-server group ruby interface ge-1/0/0.26
dhcp-relay group sapphire interface ge-1/0/0.25 upto ge-1/0/0.35
```

- Two ranges cannot overlap, either within a group or across groups. Overlapping occurs when two interface ranges share common subunit space but neither range is a proper subset of the other. The following ranges overlap:

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.20 upto ge-1/0/0.40
```

- A range can contain multiple nested ranges. A nested range is a proper subset of another range. When ranges are nested, the smallest matching range applies.

In the following example, the three ranges nest properly:

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.12 upto ge-1/0/0.15 exclude
interface ge-1/0/0.25 upto ge-1/0/0.29 exclude
```


- Discrete interfaces take precedence over ranges. In the following example, interface **ge-1/0/0.20** takes precedence and enforces an interface client limit of 5.

```
interface ge-1/0/0.10 upto ge-1/0/0.30
interface ge-1/0/0.15 upto ge-1/0/0.25 exclude
interface ge-1/0/0.20 overrides interface-client-limit 5
```

Related Documentation

- [Grouping Interfaces with Common DHCP Configurations on page 926](#)

Overriding the Default DHCP Relay Configuration Settings

You can override the default DHCP and DHCPv6 relay agent configuration settings at the global level, for a named group of interfaces, or for a specific interface within a named group.

- To override global default DHCP relay agent configuration options, include the **overrides** statement and its subordinate statements at the **[edit forwarding-options dhcp-relay]** hierarchy level.
- To override DHCP relay configuration options for a named group of interfaces, include the statements at the **[edit forwarding-options dhcp-relay group group-name]** hierarchy level.
- To override DHCP relay configuration options for a specific interface within a named group of interfaces, include the statements at the **[edit forwarding-options dhcp-relay group group-name interface]** hierarchy level.
- To configure overrides for DHCPv6 relay, use the supported statements at the **[edit forwarding-options dhcp-relay dhcpv6]** hierarchy level.

To override default DHCP relay agent configuration settings:

- Specify that you want to configure override options.

Global override:

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

Group-level override:

```
[edit forwarding-options dhcp-relay]
user@host# edit group boston overrides
```

Per-interface override:

```
[edit forwarding-options dhcp-relay]
user@host# edit group boston interface fe-1/0/1.2 overrides
```

- (DHCPv4 only) Enable DHCP relay proxy mode.
See [“Enabling DHCP Relay Proxy Mode” on page 977](#).
- (DHCPv4 only) Overwrite the giaddr in DHCP packets that the DHCP relay agent forwards.
See [“Overwriting giaddr Information” on page 957](#).

4. (DHCPv4 only) Replace the IP source address in DHCP relay request and release packets with the gateway IP address (giaddr).
See [“Replacing the DHCP Relay Request and Release Packet Source Address” on page 957.](#)
5. (DHCPv4 only) Override the DHCP relay agent information option (option 82) in DHCP packets.
See [“Overriding Option 82 Information” on page 957.](#)
6. (DHCPv4 only) Override the setting of the broadcast bit in DHCP request packets and use the Layer 2 unicast transmission method.
See [“Using Layer 2 Unicast Transmission for DHCP Packets” on page 958.](#)
7. (DHCPv4 only) Trust DHCP client packets that have a giaddr of 0 and that contain option 82 information.
See [“Trusting Option 82 Information” on page 958.](#)
8. (DHCPv4 only) Override the ARP table population in distrusted environments.
See [“Disabling ARP Table Population” on page 930.](#)
9. (DHCPv4 and DHCPv6) Override the maximum number of DHCP clients allowed per interface.
See [“Specifying the Maximum Number of DHCP Clients Per Interface” on page 929.](#)
10. (DHCPv4 only) Configure client auto logout.
See [“DHCP Auto Logout Overview” on page 888.](#)
11. (DHCPv4 and DHCPv6) Enable or disable support for DHCP snooped clients on interfaces.
See [“Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent” on page 964.](#)
12. (DHCPv4 and DHCPv6) Send release messages to the DHCP server when clients are deleted.
See [“Sending Release Messages When Clients Are Deleted” on page 971.](#)
13. (DHCPv4 only) Disable the DHCP relay agent on specific interfaces.
See [“Disabling DHCP Relay” on page 983.](#)
14. (DHCPv4 and DHCPv6) Disable automatic binding of stray DHCP requests.
See [“Disabling Automatic Binding of Stray DHCP Requests” on page 972.](#)

**Related
Documentation**

- [Group-Specific DHCP Relay Options on page 906](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 935](#)

Overwriting giaddr Information

You can configure the DHCP relay agent to change the gateway IP address (giaddr) field in packets that it forwards between a DHCP client and a DHCP server.

To overwrite the giaddr of every DHCP packet with the giaddr of the DHCP relay agent before forwarding the packet to the DHCP server:

1. Specify that you want to configure override options.

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

2. Specify that the giaddr of DHCP packets is overwritten.

```
[edit forwarding-options dhcp-relay overrides]
user@host# set always-write-giaddr
```

Related Documentation

- [Extended DHCP Relay Agent Overview on page 900](#)
- [Overriding the Default DHCP Relay Configuration Settings on page 955](#)

Replacing the DHCP Relay Request and Release Packet Source Address

You can configure the DHCP relay agent to replace request and release packets with the gateway IP address (giaddr) before forwarding the packet to the DHCP server.

To replace the source address with giaddr:

1. Specify that you want to configure override options.

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

2. Specify that you want to replace the IP source address in DHCP relay request and release packets with the gateway IP address (giaddr).

```
[edit forwarding-options dhcp-relay overrides]
user@host# set replace-ip-source-with giaddr
```

Related Documentation

- [Extended DHCP Relay Agent Overview on page 900](#)
- [Overriding the Default DHCP Relay Configuration Settings on page 955](#)

Overriding Option 82 Information

You can configure the DHCP relay agent to add or remove the DHCP relay agent information option (option 82) in DHCP packets.

This feature causes the DHCP relay agent to perform one of the following actions, depending on the configuration:

- If the DHCP relay agent is configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the DHCP packets and inserts the new values before forwarding the packets to the DHCP server.

- If the DHCP relay agent is not configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the packets, but does not add any new values before forwarding the packets to the DHCP server.

To override the default option 82 information in DHCP packets destined for a DHCP server:

1. Specify that you want to configure override options.

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

2. Specify that the option 82 information in DHCP packets is overwritten.

```
[edit forwarding-options dhcp-relay overrides]
user@host# set always-write-option-82
```

**Related
Documentation**

- [Extended DHCP Relay Agent Overview on page 900](#)
- [Overriding the Default DHCP Relay Configuration Settings on page 955](#)

Using Layer 2 Unicast Transmission for DHCP Packets

You can configure the DHCP relay agent to override the setting of the broadcast bit in DHCP request packets. DHCP relay agent then instead uses the Layer 2 unicast transmission method to send DHCP Offer reply packets and DHCP ACK reply packets from the DHCP server to DHCP clients during the discovery process.

To override the default setting of the broadcast bit in DHCP request packets:

1. Specify that you want to configure override options.

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

2. Specify that the DHCP relay agent uses the Layer 2 unicast transmission method.

```
[edit forwarding-options dhcp-relay overrides]
user@host# set layer2-unicast-replies
```

**Related
Documentation**

- [Extended DHCP Relay Agent Overview on page 900](#)
- [Overriding the Default DHCP Relay Configuration Settings on page 955](#)

Trusting Option 82 Information

By default, the DHCP relay agent treats client packets with a giaddr of 0 (zero) and option 82 information as if the packets originated at an untrusted source, and drops them without further processing. You can override this behavior and specify that the DHCP relay agent process DHCP client packets that have a giaddr of 0 (zero) and contain option 82 information.

To configure DHCP relay agent to trust option 82 information:

1. Specify that you want to configure override options.

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

- Specify that the DHCP relay agent process DHCP client packets with a giaddr of 0 and that contain option 82 information.

```
[edit forwarding-options dhcp-relay overrides]
user@host# set trust-option-82
```

Related Documentation

- [Extended DHCP Relay Agent Overview on page 900](#)
- [Overriding the Default DHCP Relay Configuration Settings on page 955](#)

Disabling ARP Table Population

By default, DHCP populates the ARP table with the MAC address of a client when the client binding is established. However, you may choose to use the DHCP **no-arp** statement to hide the subscriber MAC address or the DHCP client MAC address, as it appears in ARP table entries.

When running in a trusted environment (that is, when not using the **no-arp** statement), DHCP populates the ARP table with unique MAC addresses contained within the DHCP PDU for each DHCP client:

Table 74: ARP Table in Trusted Environment

IP Address	MAC Address
Client 1 IP Address	MAC A
Client 2 IP Address	MAC B
Client 3 IP Address	MAC C

In distrusted environments, you can specify the **no-arp** statement to hide the MAC addresses of clients. When you specify the **no-arp** statement, DHCP does not automatically populate the ARP table with MAC address information from the DHCP PDU for each client. Instead, the system performs an ARP to obtain the MAC address of each client and obtains the MAC address of the immediately attached device (for example, a DSLAM). DHCP populates the ARP table with the same interface MAC address (for example, MAC X from a DSLAM interface) for each client:

Table 75: ARP Table in Distrusted Environment

IP Address	MAC Address
Client 1 IP Address	MAC X
Client 2 IP Address	MAC X
Client 3 IP Address	MAC X

To disable ARP table population:

1. Specify that you want to configure override options.

- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# edit overrides
```

- For DHCP relay:

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

2. Disable ARP table population with client-specific information. (DHCP local server and DHCP relay agent both support the **no-arp** statement.)

- For DHCP local server:

```
[edit system services dhcp-local-server overrides]
user@host# set no-arp
```

- For DHCP relay:

```
[edit forwarding-options dhcp-relay overrides]
user@host# set no-arp
```

Related Documentation

- [Overriding Default DHCP Local Server Configuration Settings on page 928](#)
- [Extended DHCP Local Server Overview on page 876](#)
- [DHCPv6 Local Server Overview on page 880](#)
- [Extended DHCP Relay Agent Overview on page 900](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 935](#)

Specifying the Maximum Number of DHCP Clients Per Interface

By default, there is no limit to the number of DHCP local server or DHCP relay clients allowed on an interface. However, you can override the default setting and specify the maximum number of clients allowed per interface, in the range 1 through 500,000. When the number of clients on the interface reaches the specified limit, no additional DHCP Discover PDUs or DHCPv6 Solicit PDUs are accepted. When the number of clients subsequently drops below the limit, new clients are again accepted.



NOTE: The maximum number of DHCP (and DHCPv6) local server clients or DHCP (and DHCPv6) relay clients can also be specified by Juniper Networks VSA 26-143 during client login. The VSA-specified value always takes precedence if the `interface-client-limit` statement specifies a different number.

If the VSA-specified value differs with each client login, DHCP uses the largest limit set by the VSA until there are no clients on the interface.

To configure the maximum number of DHCP clients allowed per interface:

1. Specify that you want to configure override options.

- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# edit overrides
```

- For DHCPv6 local server:

```
[edit system services dhcp-local-server dhcpv6]
user@host# edit overrides
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit overrides
```

2. Configure the maximum number of clients allowed per interface. (DHCP local server, DHCPv6 local server, DHCP relay agent and DHCPv6 relay agent all support the **interface-client-limit** statement.)

```
[edit system services dhcp-local-server overrides]
user@host# set interface-client-limit number
```

Related Documentation

- [Overriding Default DHCP Local Server Configuration Settings on page 928](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 935](#)
- [Extended DHCP Local Server Overview on page 876](#)
- [Extended DHCP Relay Agent Overview on page 900](#)

Automatically Logging Out DHCP Clients

You can configure the extended DHCP local server and extended DHCP relay to automatically log out DHCP clients. Auto logout immediately releases an existing client when DHCP receives a discover packet that has the same DHCP option 60 and DHCP option 82 information as the existing client. DHCP then releases the existing client IP address without waiting for the normal lease expiration.



NOTE: When the existing client is released, the new client undergoes the normal authentication process. The new client might not receive the same IP address as the original client.

To configure DHCP client auto logout:

1. Specify that you want to configure override options.

- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# edit overrides
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

2. Enable auto logout. (DHCP local server and DHCP relay agent both support the **client-discover-match** statement.)

- For DHCP local server:

```
[edit system services dhcp-local-server overrides]
user@host# set client-discover-match
```

- For DHCP relay:

```
[edit forwarding-options dhcp-relay overrides]
user@host# set client-discover-match
```



NOTE: If you change the auto logout configuration, existing clients continue to use the auto logout setting that was configured when they logged in. New clients use the new setting.

Related Documentation

- [DHCP Auto Logout Overview on page 888](#)
- [DHCP Relay Agent Option 82 Value for Auto Logout on page 962](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 935](#)
- [Extended DHCP Local Server Overview on page 876](#)
- [Extended DHCP Relay Agent Overview on page 900](#)

DHCP Relay Agent Option 82 Value for Auto Logout

Table 76 on page 962 indicates how the DHCP relay agent determines the option 82 value used for the client auto logout feature. Depending on the configuration settings, DHCP relay agent takes the action indicated in the right column.

Table 76: DHCP Relay Agent Option 82 Value for Auto Logout

DHCP Relay Agent Configuration Settings				giaddr in non-snooped packet	Action Taken
DHCP Relay Configured with Option 82	Discover Packet Contains Option 82	Override "trust-option-82"	Override "always-write-option-82"		
No	No	—	—	—	No secondary search performed
No	Yes	Yes	—	—	Use option 82 from packet

Table 76: DHCP Relay Agent Option 82 Value for Auto Logout (*continued*)

DHCP Relay Agent Configuration Settings				giaddr in non-snooped packet	Action Taken
DHCP Relay Configured with Option 82	Discover Packet Contains Option 82	Override "trust-option-82"	Override "always-write-option-82"		
No	Yes	No	–	Zero	Drop packet
No	Yes	No	–	Non-zero	Use option 82 from packet
Yes	No	–	–	–	Use configured option 82
Yes	Yes	No	–	Zero	Drop packet
Yes	Yes	No	No	Non-zero	Use option 82 from packet
Yes	Yes	No	Yes	Non-zero	Overwrite the configured option 82
Yes	Yes	Yes	No	–	Use option 82 from packet
Yes	Yes	Yes	Yes	–	Overwrite the configured option 82

Related Documentation

- [DHCP Auto Logout Overview on page 888](#)
- [Automatically Logging Out DHCP Clients on page 932](#)

Configuring DHCP Snooping for DHCP Relay Agent

DHCP relay agent uses a two-part configuration to determine how to handle DHCP snooped packets. First, you enable or disable snooping support for DHCP relay agent and, optionally, override the default snooping configuration. Then you configure the forwarding action for snooped clients, which specifies whether DHCP relay agent forwards or drops snooped traffic.

To configure DHCP snooping for DHCP relay agent:

1. (DHCPv4 and DHCPv6) Enable or disable DHCP snooping. You can configure DHCP snooping globally, for a named group of interfaces, or for a specific interface.

See “Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent” on page 964.

2. (DHCPv4 only) Configure snooped packets forwarding support.

See [“Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent”](#) on page 969.

**Related
Documentation**

- [DHCP Snooping Support on page 887](#)
- [Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 964](#)
- [Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 969](#)
- [Example: Configuring DHCP Snooping Support for DHCP Relay Agent on page 921](#)

Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent

DHCP relay agent uses a two-part configuration to determine how to handle DHCP snooped packets. This topic describes the first procedure, in which you configure DHCP relay to either enable or disable support for snooped packets.

The second procedure, which applies only to DHCPv4 relay agent, is described in [“Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent”](#) on page 969, and configures the **forward-snooped-clients** statement, which determines whether the snooped packets are forwarded or dropped, depending on the type of interface.



NOTE: The router has a default global setting that specifies whether DHCP snooping support is enabled or disabled for DHCP relay. In Junos OS Release 10.0 and earlier, DHCP snooping is enabled by default. In Junos OS Release 10.1 and later, DHCP snooping is disabled by default.

You can override the default global DHCP snooping configuration and explicitly enable or disable DHCP snooping support. You can configure the explicit snooping support globally, for a group of interfaces, or for a specific interface in a group.

- To enable DHCP relay or DHCPv6 relay snooping support, include the **allow-snooped-clients** option in the **overrides** statement.
- To disable DHCP relay or DHCPv6 relay snooping support, include the **no-allow-snooped-clients** option in the **overrides** statement.

To enable or disable DHCP snooping support globally:

1. Specify that you want to configure DHCP relay agent.
 - For DHCP relay agent:

```
[edit]  
user@host# edit forwarding-options dhcp-relay
```
 - For DHCPv6 relay agent:

```
[edit]
user@host# edit forwarding-options dhcp-relay dhcpv6
```

2. Specify that you want to override the default configuration.

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit overrides
```

3. Enable or disable DHCP snooping support.

- To enable DHCP snooping:

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay overrides]
user@host# set allow-snooped-clients
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 overrides]
user@host# set allow-snooped-clients
```

- To disable DHCP snooping:

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay overrides]
user@host# set no-allow-snooped-clients
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 overrides]
user@host# set no-allow-snooped-clients
```

For example, to enable global DHCP snooping support :

```
forwarding-options {
  dhcp-relay {
    overrides {
      allow-snooped-clients;
    }
  }
}
```

To enable or disable DHCP snooping support for a group of interfaces:

1. Specify that you want to configure DHCP relay agent.

- For DHCP relay agent:

```
[edit]
user@host# edit forwarding-options dhcp-relay
```

- For DHCPv6 relay agent:

```
[edit]
user@host# edit forwarding-options dhcp-relay dhcpv6
```

2. Specify the named group.

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit group group-name
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit group group-name
```

3. Specify that you want to override the default configuration.

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay group group-name]
user@host# edit overrides
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name]
user@host# edit overrides
```

4. Enable or disable DHCP snooping support.

- To enable DHCP snooping:

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay group group-name overrides]
user@host# set allow-snooped-clients
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides]
user@host# set allow-snooped-clients
```

- To disable DHCP snooping:

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay group group-name overrides]
user@host# set no-allow-snooped-clients
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name overrides]
user@host# set no-allow-snooped-clients
```

For example, to enable DHCP snooping support on all interfaces in group **boston**:

```
forwarding-options {
  dhcp-relay {
    group boston {
      overrides {
        allow-snooped-clients;
      }
    }
  }
}
```

To enable or disable DHCP snooping support on a specific interface:

1. Specify that you want to configure DHCP relay agent.

- For DHCP relay agent:

```
[edit]
user@host# edit forwarding-options dhcp-relay
```

- For DHCPv6 relay agent:

```
[edit]
user@host# edit forwarding-options dhcp-relay dhcpv6
```

2. Specify the named group containing the interface.

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit group group-name
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit group group-name
```

3. Specify the interface for which you want to configure DHCP snooping.

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay group group-name]
user@host# edit interface interface-name
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name]
user@host# edit interface interface-name
```

4. Specify that you want to override the default configuration on the interface.

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay group group-name interface interface-name]
user@host# edit overrides
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface
interface-name]
user@host# edit overrides
```

5. Enable or disable DHCP snooping support.

- To enable DHCP snooping:
 - For DHCP relay agent:

```
[edit forwarding-options dhcp-relay group group-name interface interface-name overrides]
user@host# set allow-snooped-clients
```
 - For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name overrides]
user@host# set allow-snooped-clients
```
- To disable DHCP snooping:
 - For DHCP relay agent:

```
[edit forwarding-options dhcp-relay group group-name interface interface-name overrides]
user@host# set no-allow-snooped-clients
```
 - For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name interface interface-name overrides]
user@host# set no-allow-snooped-clients
```

For example, to disable DHCP snooping support on interface **ge-2/1/8.0** in group **boston**:

```
forwarding-options {
  dhcp-relay {
    group boston {
      interface ge-2/1/8.0 {
        overrides {
          no-allow-snooped-clients;
        }
      }
    }
  }
}
```

To enable DHCPv6 snooping support on interface **ge-3/2/1.1** in group **sunnyvale**:

```
forwarding-options {
  dhcp-relay {
    dhcpv6 {
      group sunnyvale {
        interface ge-3/2/1.1 {
          overrides {
            allow-snooped-clients;
          }
        }
      }
    }
  }
}
```

Related Documentation

- [Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent on page 969](#)
- [DHCP Snooping Support on page 887](#)
- [Overriding the Default DHCP Relay Configuration Settings on page 955](#)

Configuring DHCP Snooped Packets Forwarding Support for DHCP Relay Agent

You can configure how DHCP relay agent handles DHCP snooped packets. Depending on the configuration, DHCP relay agent either forwards or drops the snooped packets it receives.

DHCP relay uses a two-part configuration to determine how to handle DHCP snooped packets. This topic describes how you use the [forward-snooped-clients](#) statement to manage whether DHCP relay agent forwards or drops snooped packets, depending on the type of interface on which the packets are snooped. In the other part of the DHCP relay agent snooping configuration, which is described in “[Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent](#)” on page 964, you enable or disable the DHCP relay snooping feature.

[Table 77 on page 969](#) shows the action the router or switch takes on snooped packets when DHCP snooping is enabled by the [allow-snooped-clients](#) statement.

[Table 78 on page 970](#) shows the action the router (or switch) takes on snooped packets when DHCP snooping is disabled by the [no-allow-snooped-clients](#) statement.

The router or switch also uses the configuration of the DHCP relay agent forwarding support to determine how to handle snooped BOOTREPLY packets. [Table 79 on page 970](#) shows the action the router (or switch) takes for the snooped BOOTREPLY packets.



NOTE: Configured interfaces have been configured with the `group` statement in the `[edit forwarding-options dhcp-relay]` hierarchy. Non-configured interfaces are in the logical system/routing instance but have not been configured by the `group` statement.

Table 77: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Enabled

forward-snooped-clients Configuration	Action on Configured Interfaces	Action on Non-Configured Interfaces
<code>forward-snooped-clients</code> not configured	snooped packets result in subscriber (DHCP client) creation	dropped
<code>all-interfaces</code>	forwarded	forwarded
<code>configured-interfaces</code>	forwarded	dropped

Table 77: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Enabled (*continued*)

forward-snooped-clients Configuration	Action on Configured Interfaces	Action on Non-Configured Interfaces
non-configured-interfaces	snooped packets result in subscriber (DHCP client) creation	forwarded

Table 78: Actions for DHCP Relay Agent Snooped Packets When DHCP Snooping Is Disabled

forward-snooped-clients Configuration	Action on Configured Interfaces	Action on Non-Configured Interfaces
forward-snooped-clients not configured	dropped	dropped
all-interfaces	dropped	forwarded
configured-interfaces	dropped	dropped
non-configured-interfaces	dropped	forwarded

Table 79: Actions for Snooped BOOTREPLY Packets

forward-snooped-clients Configuration	Action
forward-snooped-clients not configured	snooped BOOTREPLY packets dropped if client is not found
forward-snooped-clients all configurations	snooped BOOTREPLY packets forwarded if client is not found

To configure DHCP snooped packet forwarding and BOOTREPLY snooped packet forwarding for DHCP relay agent:

- Specify that you want to configure DHCP relay agent.

```
[edit]
user@host# edit forwarding-options dhcp-relay
```
- Enable DHCP snooped packet forwarding.

```
[edit forwarding-options dhcp-relay]
user@host# edit forward-snooped-clients
```
- Specify the interfaces that are supported for snooped packet forwarding.

```
[edit forwarding-options dhcp-relay forward-snooped-clients]
user@host# set (all-interfaces | configured-interfaces | non-configured-interfaces)
```


For example, to configure DHCP relay agent to forward DHCP snooped packets on only configured interfaces:

```
[edit]
forwarding-options {
  dhcp-relay {
    forward-snooped-clients configured-interfaces;
  }
}
```

**Related
Documentation**

- [DHCP Snooping Support on page 887](#)
- [Enabling and Disabling DHCP Snooped Packets Support for DHCP Relay Agent on page 964](#)

[Sending Release Messages When Clients Are Deleted](#)

By default, when DHCP relay and relay proxy delete a client, they do not send a release message to the DHCP server. You can override the default behavior and configure DHCP relay and relay proxy to send a release message whenever they delete a client. The release message sent by DHCP relay and relay proxy includes option 82 information.



NOTE: In Junos OS Release 10.1 and earlier, DHCP relay sends a release message to the DHCP server when the `client-discover-match` statement is included as a DHCP relay override. In Junos OS Release 10.2 and later, you must include the `send-release-on-delete` statement to configure DHCP relay and relay proxy to send the release message when the `client-discover-match` statement is included.

You can use the `[edit forwarding-options dhcp-relay dhcpv6]` hierarchy level to override the default behavior for DHCPv6 relay agent.

To send a release message:

1. Specify that you want to configure override options.

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

- For DHCPv6 relay agent:

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit overrides
```

2. Specify that you want DHCP relay and relay proxy (or DHCPv6 relay agent) to send a release message when clients are deleted.

```
[edit forwarding-options dhcp-relay overrides]
user@host# set send-release-on-delete
```

- Related Documentation**
- [Extended DHCP Relay Agent Overview on page 900](#)
 - [Overriding the Default DHCP Relay Configuration Settings on page 955](#)

Disabling Automatic Binding of Stray DHCP Requests

DHCP requests that are received but have no entry in the database are known as stray requests. By default, DHCP relay, DHCP relay proxy, and DHCPv6 relay agent attempt to bind the requesting client by creating a database entry and forwarding the request to the DHCP server. If the server responds with an ACK, the client is bound and the ACK is forwarded to the client. If the server responds with a NAK, the database entry is deleted and the NAK is forwarded to the client. This behavior occurs regardless of whether authentication is configured.

You can override the default configuration at the global level, for a named group of interfaces, or for a specific interface within a named group. Overriding the default causes DHCP relay, DHCP relay proxy, and DHCPv6 relay agent to drop all stray requests instead of attempting to bind the clients.



NOTE: In Junos OS Release 10.4 and later, automatic binding of stray requests is enabled by default.

In Junos OS Release 10.3 and earlier releases, automatic binding of stray requests is disabled by default. In those releases, DHCP relay drops stray requests and forwards a NAK to the client when authentication is configured. Otherwise, DHCP relay attempts to bind the requesting client. In those releases, DHCP relay proxy always drops stray requests and forwards a NAK to the client, regardless of the authentication configuration.

- To disable automatic binding behavior, include the **no-bind-on-request** statement when you configure DHCP overrides at the global, group, or interface level.

```
[edit forwarding-options dhcp-relay overrides]
user@host# set no-bind-on-request
```

- To override the default behavior for DHCPv6 relay agent, configure the override at the **[edit forwarding-options dhcp-relay dhcpv6]** hierarchy level.

```
[edit forwarding-options dhcp-relay dhcpv6 overrides]
user@host# set no-bind-on-request
```

The following two examples show a configuration that disables automatic binding of stray requests for a group of interfaces and a configuration that disables automatic binding on a specific interface.

To disable automatic binding of stray requests on a group of interfaces:

1. Specify the named group.

```
[edit forwarding-options dhcp-relay]
user@host# edit group boston
```

2. Specify that you want to configure overrides.

```
[edit forwarding-options dhcp-relay group boston]
user@host# edit overrides
```

3. Disable automatic binding for the group.

```
[edit forwarding-options dhcp-relay group boston overrides]
user@host# set no-bind-on-request
```

To disable automatic binding of stray requests on a specific interface:

1. Specify the named group of which the interface is a member.

```
[edit forwarding-options dhcp-relay]
user@host# edit group boston
```

2. Specify the interface on which you want to disable automatic binding.

```
[edit forwarding-options dhcp-relay group boston]
user@host# edit interface fe-1/0/1.2
```

3. Specify that you want to configure overrides.

```
[edit forwarding-options dhcp-relay group boston interface fe-1/0/1.2]
user@host# edit overrides
```

4. Disable automatic binding on the interface.

```
[edit forwarding-options dhcp-relay group boston interface fe-1/0/1.2 overrides]
user@host# set no-bind-on-request
```

Related Documentation

- [Extended DHCP Relay Agent Overview on page 900](#)
- [Overriding the Default DHCP Relay Configuration Settings on page 955](#)

Enabling and Disabling Insertion of Option 82 Information

You can enable or disable support for the DHCP relay agent information option (option 82) in packets destined for a DHCP server. You can configure option 82 support globally or for a named group of interfaces.

To restore the default behavior (option 82 information is not inserted into DHCP packets), you use the **delete relay-option-82** statement.

To configure support for the DHCP relay agent information option 82, you use the **relay-option-82** statement.

The following sections describe the option 82 operations you can configure:

- [Configuring Agent Circuit ID Information on page 974](#)
- [Configuring an Option 82 Prefix on page 974](#)
- [Using a Textual Description in Option 82 on page 976](#)

Configuring Agent Circuit ID Information

You use the **relay-option-82** statement to enable insertion of option 82 information in DHCP packets. You must also specify at least the **circuit-id** statement to include the Agent Circuit ID suboption (suboption 1) of the DHCP relay agent information option.

If you specify the **circuit-id** statement, the format of the Agent Circuit ID information for Fast Ethernet (**fe**) or Gigabit Ethernet (**ge**) interfaces is one of the following, depending on your network configuration:

- For Fast Ethernet or Gigabit Ethernet interfaces that do not use virtual local area networks (VLANs) or stacked VLANs (S-VLANs):

`(fe | ge)-fpc/pic/port`

- For Fast Ethernet or Gigabit Ethernet interfaces that use VLANs:

`(fe | ge)-fpc/pic/port:vlan-id`

- For Fast Ethernet or Gigabit Ethernet interfaces that use S-VLANs:

`(fe | ge)-fpc/pic/port:svlan-id-vlan-id`

To enable insertion of option 82 information:

1. Specify that you want to configure option 82 support.

```
[edit forwarding-options dhcp-relay]
user@host# edit relay-option-82
```

2. Specify insertion of the Agent Circuit ID suboption.

```
[edit forwarding-options dhcp-relay relay-option-82]
user@host# set circuit-id
```

Configuring an Option 82 Prefix

You can include an optional prefix to the base option 82 information in DHCP packets destined for a DHCP server.

The prefix is separated from the option 82 Agent Circuit ID information by a colon (:), and can include any combination of the **host-name**, **logical-system-name**, and **routing-instance-name** options. The DHCP relay agent obtains the values for the **host-name**, **logical-system-name**, and **routing-instance-name** as follows:

- If you include the **host-name** option, the DHCP relay agent uses the hostname of the router configured with the **host-name** statement at the **[edit system]** hierarchy level.
- If you include the **logical-system-name** option, the DHCP relay agent uses the logical system name configured with the **logical-system** statement at the **[edit logical-system]** hierarchy level.
- If you include the **routing-instance-name** option, the DHCP relay agent uses the routing instance name configured with the **routing-instance** statement at the **[edit routing-instances]** hierarchy level or at the **[edit logical-system logical-system-name routing-instances]** hierarchy level.

If you include the hostname and either or both of the logical system name and the routing instance name in the prefix, the hostname is followed by a forward slash (/). If you include both the logical system name and the routing instance name in the prefix, these values are separated by a semicolon (;).

The following examples show several possible formats for the Agent Circuit ID information when you specify the **prefix** statement for Fast Ethernet (**fe**) or Gigabit Ethernet (**ge**) interfaces with S-VLANs.

- If you include only the hostname in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

hostname:(fe | ge)-fpc/pic/port:svlan-id-vlan-id

- If you include only the logical system name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

logical-system-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id

- If you include only the routing instance name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

routing-instance-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id

- If you include both the hostname and the logical system name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

host-name/logical-system-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id

- If you include both the logical system name and the routing instance name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

logical-system-name;routing-instance-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id

- If you include the hostname, logical system name, and routing instance name in the prefix for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs:

host-name/logical-system-name;routing-instance-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id

For Fast Ethernet or Gigabit Ethernet interfaces that use VLANs but not S-VLANs, only the **vlan-id** value appears in the Agent Circuit ID format. For Fast Ethernet or Gigabit Ethernet interfaces that do not use VLANs or S-VLANs, neither the **vlan-id** value nor the **svlan-id** value appears.

To configure an optional prefix with the option 82 information:

1. Specify that you want to configure option 82 support.

```
[edit forwarding-options dhcp-relay]
user@host# edit relay-option-82
```

2. Specify insertion of the Agent Circuit ID information.

```
[edit forwarding-options dhcp-relay relay-option-82]
user@host# edit circuit-id
```

3. Specify that the prefix is included in the option 82 information. In this example, the prefix includes the hostname and logical system name

```
[edit forwarding-options dhcp-relay relay-option-82 circuit-id]
user@host# set prefix host-name logical-system-name
```

Using a Textual Description in Option 82

By default, when DHCP option 82 is inserted into client packets, the Agent Circuit ID suboption includes the interface identifier. You can optionally configure that the Agent Circuit ID suboption include the textual description that is configured for the interface instead of the interface identifier. You can use the textual description for either the logical interface or the device interface.

You can include the textual interface description in the Agent Circuit ID suboption for static interfaces. The textual description is configured using the **description** statement at the **[edit interfaces interface-name]** hierarchy level. If you specify that the textual description is used and no description is configured for the interface, DHCP relay defaults to using the interface identifier.

To configure the DHCP relay option 82 suboption to include the textual interface description:

1. Specify that you want to configure option 82 support.

```
[edit forwarding-options dhcp-relay]
user@host# edit relay-option-82
```

2. Specify insertion of the Agent Circuit ID information.

```
[edit forwarding-options dhcp-relay relay-option-82]
user@host# edit circuit-id
```

3. Specify that the textual description is included in the option 82 information. In this example, the option 82 information includes the description used for the device interface.

```
[edit forwarding-options dhcp-relay relay-option-82 circuit-id]
user@host# set use-interface-description device
```

Configuring Server Groups

You can configure a named group of DHCP servers for use by the extended DHCP relay agent on the router or switch.

You specify the name of the DHCP server group and the IP addresses of one or more DHCP servers that belong to this group. You can configure a maximum of five IP addresses per named server group.

To configure a named server group:

1. Specify the name of the server group.

```
[edit forwarding-options dhcp-relay]
user@host# set server-group myServerGroup
```

2. Add the IP addresses of the DHCP servers belonging to the group.

```
[edit forwarding-options dhcp-relay server-group myServerGroup]
user@host# set 192.168.100.50
```

```
user@host# set 192.168.100.75
```

Related Documentation

- [Extended DHCP Relay Agent Overview on page 900](#)

Configuring Active Server Groups

You can configure an active server group. Using an active server group enables you to apply a common DHCP relay agent configuration to a named group of DHCP server addresses.

Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.

To configure an active server group:

- Specify the name of the active server group.

```
[edit forwarding-options dhcp-relay]
user@host# set active-server-group myServerGroup
```

To create an active server group as a global DHCP relay agent configuration option, include the **active-server-group** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level. To have the group apply only to a named group of interfaces, include the **active-server-group** statement at the **[edit forwarding-options dhcp-relay group group-name]** hierarchy level.

Including the **active-server-group** statement at the **[edit forwarding-options dhcp-relay group group-name]** hierarchy level (as a group-specific option) overrides the effect of including the **active-server-group** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level as a global option.

Related Documentation

- [Extended DHCP Relay Agent Overview on page 900](#)
- [Grouping Interfaces with Common DHCP Configurations on page 926](#)

Enabling DHCP Relay Proxy Mode

You can enable DHCP relay proxy mode on all interfaces or a group of interfaces.

To enable DHCP relay proxy mode:

1. Specify that you want to configure override options.

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

2. Enable DHCP relay proxy mode.

```
[edit forwarding-options dhcp-relay overrides]
user@host# set proxy-mode
```

Related Documentation

- [DHCP Relay Proxy Overview on page 903](#)
- [Overriding the Default DHCP Relay Configuration Settings on page 955](#)

Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces

This topic describes how to attach a dynamic profile to a DHCP subscriber interface or a DHCP client interface. When a DHCP subscriber or DHCP client logs in, the specified dynamic profile is instantiated and the services defined in the profile are applied to the interface.

This topic contains the following sections:

- [Attaching a Dynamic Profile to All DHCP Subscriber or All DHCP Client Interfaces on page 978](#)
- [Attaching a Dynamic Profile to a Group of DHCP Subscriber Interfaces or a Group of DHCP Client Interfaces on page 979](#)

Attaching a Dynamic Profile to All DHCP Subscriber or All DHCP Client Interfaces

To attach a dynamic profile to all DHCP subscriber or all DHCP client interfaces:

1. At the DHCP configuration hierarchy, use the **dynamic-profile** statement to specify the name of the dynamic profile to attach to all interfaces.
 - For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# set dynamic-profile vod-profile-22
```
 - For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# set dynamic-profile vod-profile-west
```
2. (Routers only) Optionally, you can configure the attribute to use when attaching the specified profile.

You can include either the **aggregate-clients** option to enable multiple DHCP subscribers to share the same VLAN logical interface, or the **use-primary** option to specify that the primary dynamic profile is used. The **aggregate-clients** option does not apply to demux subscriber interfaces. The two options are mutually exclusive.

- To enable multiple subscribers to share the same VLAN logical interface:

```
[edit system services dhcp-local-server dynamic-profile]
user@host# set aggregate-clients merge
```

- To use the primary dynamic profile:

```
[edit forwarding-options dhcp-relay dynamic-profile]
user@host# set use-primary subscriber_profile
```


Attaching a Dynamic Profile to a Group of DHCP Subscriber Interfaces or a Group of DHCP Client Interfaces

Before you begin:

- Configure the interface group.

See “[Grouping Interfaces with Common DHCP Configurations](#)” on page 926.

To attach a dynamic profile to a group of interfaces:

1. At the DHCP configuration hierarchy, specify the name of the interface group and the dynamic profile to attach to the group.

- For DHCP local server:

```
[edit system services dhcp-local-server]
user@host# set group boston dynamic-profile vod-profile-42
```

- For DHCP relay agent:

```
[edit forwarding-options dhcp-relay]
user@host# set group quebec dynamic-profile vod-profile-east
```

2. (Routers only) Optionally, you can configure the attribute to use when attaching the specified profile.

You can include either the **aggregate-clients** option to enable multiple DHCP subscribers to share the same VLAN logical interface, or the **use-primary** option to specify that the primary dynamic profile is used. The **aggregate-clients** option does not apply to demux subscriber interfaces. The two options are mutually exclusive.

- To enable multiple subscribers to share the same VLAN logical interface:

```
[edit system services dhcp-local-server dynamic-profile]
user@host# set aggregate-clients merge
```

- To use the primary dynamic profile:

```
[edit forwarding-options dhcp-relay dynamic-profile]
user@host# set use-primary subscriber_profile
```

Related Documentation

- *Dynamic Profiles Overview*
- *Dynamic Profile Attachment to DHCP Subscriber Interfaces Overview*
- *Example: Configuring Dynamic Subscriber Interfaces on IP Demux Interfaces*

Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets

You can configure DHCPv6 relay agent to insert the DHCPv6 Interface-ID (Option 18) in the packets that the relay sends to a DHCPv6 server. You can optionally include a prefix, which can include any combination of hostname, logical system name, and routing instance name. You can also specify that the packets include the textual interface description instead of the interface identifier.



NOTE: If you configure the optional Steps 2 or 3, and the specified information does not exist (for example, there is no interface description), DHCPv6 relay ignores the optional configuration and inserts the interface identifier in the packets.

To insert the DHCPv6 Interface-ID option (Option 18) in DHCPv6 packets :

1. Configure the DHCPv6 relay to include Option 18.

```
[edit forwarding-options dhcp-relay dhcpv6]
user@host# edit relay-agent-interface-id
```

2. (Optional) Specify the prefix to include in Option 18.

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id]
user@host# set prefix prefix
```

3. (Optional) Specify that the Option 18 include the textual description of the interface. You can specify either the **logical** interface description or the **device** interface description.

```
[edit forwarding-options dhcp-relay dhcpv6 relay-agent-interface-id]
user@host# set use-interface-description (logical | device)
```

Related Documentation

- [Extended DHCP Relay Agent Overview on page 900](#)

DHCP Liveness Detection Overview

Unlike PPP, DHCP does not define a native keepalive mechanism as part of either the DHCPv4 or DHCPv6 protocols. Without a keepalive mechanism, DHCP local server, DHCP relay, or DHCP relay proxy is unable to quickly detect if it has lost connectivity with a subscriber or a DHCP client; and it must rely on standard DHCP subscriber session or DHCP client session termination messages.

DHCP clients often do not send DHCP release messages prior to exiting the network. The discovery of their absence is dependent on existing DHCP lease time and release request mechanisms. These mechanisms are often considered insufficient when serving as session health checks for clients in a DHCP subscriber access or a DHCP-managed network. Because DHCP lease times are typically too long to provide an adequate response time for a session health failure, and configuring short DHCP lease times can pose an undue burden on control plane processing, implementing a DHCP liveness detection mechanism enables better monitoring of bound DHCP clients. When configured with a liveness detection protocol, if a given subscriber (or client) fails to respond to a configured number of consecutive liveness detection requests, the subscriber (or client) binding is deleted and its resources released.

DHCP liveness detection for DHCP subscriber IP or DHCP client IP sessions utilizes an active liveness detection protocol to institute liveness detection checks for relevant clients. Clients must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of

consecutive attempts, then the liveness detection check fails and a failure action is implemented.

Using DHCP liveness detection, IP sessions are acted upon as soon as liveness detection checks fail. This faster response time serves to:

- Provide more accurate time-based accounting of subscriber (or DHCP client) sessions.
- Better preserve router (switch) resources.
- Help to reduce the window of vulnerability to some security attacks.

Examples of liveness detection protocols include Bidirectional Forwarding Detection (BFD) for both DHCPv4 and DHCPv6 subscribers, IPv4 Address Resolution Protocol (ARP) for DHCPv4 subscribers, and IPv6 Neighbor Unreachability Detection for DHCPv6 subscribers.



NOTE: This release supports only BFD for DHCPv4 and DHCPv6 liveness detection.

When configuring BFD liveness detection, keep the following in mind:

- You can configure DHCPv4 and DHCPv6 liveness detection either globally or per DHCPv4 or DHCPv6 group.
- DHCPv4 or DHCPv6 subscriber access clients that do not support BFD are not affected by the liveness detection configuration. These clients can continue to access the network (once validated) even if BFD liveness detection is enabled on the router (or switch).
- When configured, DHCPv4 or DHCPv6 initiates liveness detection checks for relevant clients (that is, clients that support BFD) when those clients enter a bound state.
- After protocol-specific messages are initiated for a BFD client, they are periodically sent to the subscriber (or client) IP address of the client and responses to those liveness detection requests are expected within a configured amount of time.
- If liveness detection responses are not received from clients that support BFD within the configured amount of time for a configured number of consecutive attempts, the liveness detection check is deemed to have failed and a configured failure action is implemented.

Related Documentation

- [Configuring Detection of DHCP Local Server Client Connectivity on page 941](#)
- [Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 981](#)

Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity

Liveness detection for DHCP subscriber IP or DHCP client IP sessions utilizes an active liveness detection protocol to institute liveness detection checks for relevant clients. Clients must respond to liveness detection requests within a specified amount of time.

If the responses are not received within that time for a given number of consecutive attempts, then the liveness detection check fails and a failure action is implemented.



NOTE: You can also configure DHCP liveness detection for DHCP local server.

To configure liveness detection for DHCP relay:

1. Specify that you want to configure liveness detection.

- For DHCP global configuration:

```
[edit forwarding-options dhcp-relay]
user@host# edit liveness-detection
```

- For DHCP group configuration:

```
[edit forwarding-options dhcp-relay group group-name]
user@host# edit liveness-detection
```



NOTE: Liveness detection is also supported for DHCPv6 configurations. To configure DHCPv6 liveness detection, include the **liveness-detection** statement, and any subsequent configuration statements, at the [edit forwarding-options dhcp-relay dhcpv6] or [edit forwarding-options dhcp-relay dhcpv6 group *group-name*] hierarchy level.

2. (Optional) Specify that you want to use DHCP relay proxy mode.

```
[edit forwarding-options dhcp-relay group group-name]
user@host# set overrides proxy-mode
```

3. Specify that you want to configure the liveness detection method.

- For DHCP global configuration:

```
[edit forwarding-options dhcp-relay liveness-detection]
user@host# edit method
```

- For DHCP group configuration:

```
[edit forwarding-options dhcp-relay group group-name liveness-detection]
user@host# edit method
```

4. Specify the liveness detection method that you want DHCP to use.



NOTE: In this release, the only method supported for liveness detection is Bidirectional Forwarding Detection (BFD).

- For DHCP global configuration:

```
[edit forwarding-options dhcp-relay liveness-detection method]
user@host# edit bfd
```

- For DHCP group configuration:

```
[edit forwarding-options dhcp-relay group group-name liveness-detection method]
user@host# edit bfd
```

5. Configure the liveness detection method as desired.

See *Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients* for an example of how to globally configure DHCP relay liveness detection.

6. Configure the action the router takes when a liveness detection failure occurs.

- For DHCP global configuration:

```
[edit forwarding-options dhcp-relay liveness-detection]
user@host# edit failure-action action
```

- For DHCP group configuration:

```
[edit forwarding-options dhcp-relay group group-name liveness-detection]
user@host# edit failure-action action
```

Related Documentation

- [Extended DHCP Relay Agent Overview on page 900](#)
- [DHCP Liveness Detection Overview on page 980](#)
- [Configuring Detection of DHCP Local Server Client Connectivity on page 941](#)
- [Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913](#)
- [Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients](#)

Disabling DHCP Relay

You can disable DHCP relay on all interfaces or a group of interfaces.

To disable DHCP relay agent:

1. Specify that you want to configure override options.

```
[edit forwarding-options dhcp-relay]
user@host# edit overrides
```

2. Disable the DHCP relay agent.

```
[edit forwarding-options dhcp-relay overrides]
user@host# set disable-relay
```

Related Documentation

- [Extended DHCP Relay Agent Overview on page 900](#)
- [Deleting DHCP Local Server and DHCP Relay Override Settings on page 935](#)

DHCP Local Server Configuration Statements

- [\[edit system\] Hierarchy Level on page 983](#)

[edit system] Hierarchy Level

```
system {
  accounting {
```

```
destination {
  radius {
    server {
      server-address {
        accounting-port port-number;
        max-outstanding-requests
        port port-number;
        retry number;
        secret password;
        source-address address;
        timeout seconds;
      }
    }
  }
  tacplus {
    server {
      server-address {
        port port-number;
        secret password;
        single-connection;
        source-address address;
        timeout seconds;
      }
    }
  }
}
events [ change-log interactive-commands login ];
}
allow-6pe-traceroute;
allow-v4mapped-packets;
archival {
  configuration {
    archive-sites {
      ftp://<username>:<password>@<host>:<port>/<url-path>;
      scp://<username>:<password>@<host>:<port>/<url-path>;
    }
    transfer-interval interval;
    transfer-on-commit;
  }
}
arp {
  aging-timer minutes;
  gratuitous-arp-delay;
  gratuitous-arp-on-ifup;
  interfaces {
    logical-interface-name {
      aging-timer minutes;
    }
  }
  passive-learning;
  purging;
}
authentication-order [ authentication-methods ];
auto-configuration {
  traceoptions {
```

```

    file <filename> <files number> <match regular-expression> <size size>
      <world-readable | no-world-readable>;
    flag <all | auth | configuration | ;interfaces | io | rtsock | ui>
    level level;
    no-remote-trace;
  }
}
backup-router address <destination [ destination-addresses ]>;
commit {
  fast-synchronize;
  synchronize;
  server {
    commit-interval number;
    days-to-keep-error-logs number;
    maximum-aggregate-pool number;
    maximum-entries number;
    traceoptions {
      file <filename> <files number> <match regular-expression> <size size>
        <world-readable | no-world-readable>;
      flag <all | auth | configuration | ;interfaces | io | rtsock | ui>
      level level;
      no-remote-trace;
    }
  }
}
(compress-configuration-files | no-compress-configuration-files);
ddos-protection {
  global {
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    flow-detection;
    flow-report-rate;
    violation-report-rate;
  }
  protocols protocol-group (aggregate | packet-type) {
    bandwidth packets-per-second;
    burst size;
    disable-fpc;
    disable-logging;
    disable-routing-engine;
    fpc {
      bandwidth-scale percentage;
      burst-scale percentage;
      disable-fpc;
    }
    priority level;
    recover-time seconds;
    flow-detection {
      flow-detect-time detect-period;
      no-flow-logging;
      timeout-active-flows enable-period;
      flow-level-bandwidth;
      flow-level-control (all | keep-all | police);
      flow-detection-mode (always-on | automatic | disabled);
      physical-interface;
    }
  }
}

```

```
        flow-recover-time recover-period;  
        flow-timeout-time timeout-period;  
        subscriber;  
    }  
}  
traceoptions{  
    file filename <files number> <match regular-expression> <size maximum-file-size>  
        <world-readable | no-world-readable>;  
    flag flag;  
    level (all | error | info | notice | verbose | warning);  
    no-remote-trace;  
}  
}  
default-address-selection;  
diag-port-authentication (encrypted-password "password" | plain-text-password);  
dynamic-profile-options {  
    versioning;  
}  
domain-name domain-name;  
domain-search [ domain-list ];  
donot-disable-ip6op-ondad;  
extensions {  
    providers {  
        provider-id {  
            license-type license deployment-scope [ deployments ];  
        }  
    }  
}  
resource-limits {  
    package package-name {  
        resources {  
            cpu {  
                priority number;  
                time seconds;  
            }  
            file {  
                core-size bytes;  
                open number;  
                size bytes;  
            }  
            memory {  
                data-size bytes;  
                locked-in bytes;  
                resident-set-size bytes;  
                socket-buffers bytes;  
                stack-size bytes;  
            }  
        }  
    }  
}  
process process-ui-name {  
    resources {  
        cpu {  
            priority number;  
            time seconds;  
        }  
        file {  
            core-size bytes;  
        }  
    }  
}
```



```

        open number;
        size bytes;
    }
    memory {
        data-size bytes;
        locked-in bytes;
        resident-set-size bytes;
        socket-buffers bytes;
        stack-size bytes;
    }
}
}
}
fips {
    level level;
}
host-name hostname;
inet6-backup-router ipv6-address <destination address>;
internet-options {
    (gre-path-mtu-discovery | no-gre-path-mtu-discovery);
    icmpv4-rate-limit bucket-size number packet-rate rate;
    icmpv6-rate-limit bucket-size number packet-rate rate;
    (ipip-path-mtu-discovery | no-ipip-path-mtu-discovery);
    (ipv6-path-mtu-discovery | noipv6-path-mtu-discovery);
    ipv6-path-mtu-discovery-timeout;
    no-tcp-rfc1323-paws;
    no-tcp-rfc1323;
    (path-mtu-discovery | no-path-mtu-discovery);
    source-port upper-limit port-number;
    (source-quench | no-source-quench);
    tcp-drop-synfin-set;
}
kernel-replication;
license {
    autoupdate {
        url URL;
        password password;
    }
    renew before-expiration number;
    interval number
    traceoptions {
        file <filename> <files number> <size maximum-file-size> <world-readable |
            no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
location {
    altitude feet;
    building name;
    country-code code;
    floor number;
    hcoord horizontal-coordinate;
    lata service-area;
    latitude degrees;

```

```

longitude degrees;
npa-nxx number;
postal-code postal-code;
rack number;
vcoord vertical-coordinate;
}
login {
  announcement "text";
  class class-name {
    access-end "hh<:mm:<ss>>";
    access-start "hh<:mm:<ss>>";
    allow-commands "regular-expression";
    ( allow-configuration | allow-configuration-regexps ) "regular expression 1" "regular
      expression 2";
    allowed-days [ sunday monday tuesday wednesday thursday friday saturday ];
    configuration-breadcrumbs;
    deny-commands "regular-expression";
    ( deny-configuration | deny-configuration-regexps ) "regular expression 1" "regular
      expression 2";
    idle-timeout minutes;
    logical-system logical-system-name;
    login-alarms;
    login-script filename;
    login-tip;
    permissions [ permissions ];
    security-role [ security-role ];
  }
  deny-sources ( address address | apply-groups | apply-groups-except );
  message "text";
  password {
    change-type ( character-sets | set-transitions );
    format ( des | md5 | sha1 );
    maximum-length length;
    minimum-changes number;
    minimum-length length;
    minimum-lower-cases number;
    minimum-numeric number;
    minimum-punctuations number;
    minimum-upper-cases number;
  }
  retry-options {
    backoff-factor number;
    backoff-threshold number;
    maximum-time number;
    minimum-time number;
    tries-before-disconnect number;
  }
  user username {
    authentication {
      ( encrypted-password "password" | plain-text-password );
      load-key-file filename;
      ssh-dsa "public-key" <from hostname>;
      ssh-ecdsa "public-key" <from hostname>;
      ssh-rsa "public-key" <from hostname>;
    }
    class class-name;
  }
}

```

```

        full-name "complete-name";
        uid uid-value;
    }
}
max-configurations-on-flash number;
mirror-flash-on-disk;
name-server {
    address;
}
nd-maxmcast-solicit
nd-retransmit-timer
no-multicast-echo;
no-neighbor-learn;;
no-ping-record-route;
no-ping-time-stamp;
no-redirects;
no-redirects-ipv6;
ntp {
    authentication-key key-number type md5 value password;
    boot-server address;
    broadcast <address> <key key-number> <ttl value> <version value>;
    broadcast-client;
    multicast-client <address>;
    peer address <key key-number> <prefer> <version value>;
    server address <key key-number> <prefer> <version value>;
    source-address source-address;
    trusted-key [ key-numbers ];
}
pic-console-authentication {
    (encrypted-password "encrypted-password" | plain-text-password);
}
ports {
    auxiliary {
        disable;
        insecure;
        type (ansi | small-xterm | vt100 | xterm);
        port-type (mini-usb | rj45) ;
    }
}
console {
    disable;
    insecure;
    log-out-on-disconnect;
    type (ansi | small-xterm | vt100 | xterm);
}
}
processes {
    process-name (enable | disable) failover (alternate-media | other-routing-engine);
    command path;
    timeout seconds;
}
proxy {
    password password;
    port port-number;
    server (hostname | ip-address);
    username username;
}

```

```
}
radius-options {
  attributes {
    nas-ip-address address;
  }
  password-protocol mschap-v2;
}
radius-server {
  server-address {
    accounting-port port-number;
    max-outstanding-requests number;
    port port-number;
    retry number;
    secret password;
    source-address source-address;
    timeout seconds;
  }
}
root-authentication {
  (encrypted-password "password" | plain-text-password);
  load-key-file filename;
  ssh-dsa "public-key" <from hostname>;
  ssh-ecdsa "public-key" <from hostname>;
  ssh-rsa "public-key" <from hostname>;
}
(saved-core-context | no-saved-core-context);
saved-core-files number;
scripts {
  load-scripts-from-flash;
  commit {
    allow-transients;
    direct-access;
    file filename.xml {
      checksum (md5 | sha-256 | sha1) hash;
      optional;
      refresh;
      refresh-from url;
      source url;
    }
  }
  max-datasize
  refresh;
  refresh-from url;
  traceoptions {
    file <filename> <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
op {
  file filename.xml {
    arguments {
      argument-name {
        description descriptive-text;
      }
    }
  }
}
```

```

checksum (md5 | sha-256 | sha1) hash;
command filename-alias;
description descriptive-text;
refresh;
refresh-from url;
source url;
}
max-datasize
no-allow-url
refresh;
refresh-from url;
traceoptions {
    file <filename> <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag;
    no-remote-trace;
}
}
static-host-mapping {
    hostname {
        alias [ aliases ];
        inet [ addresses ];
        inet6 [ addresses ];
        sysid system-identifier;
    }
}
syslog {
    allow-duplicates;
    archive <binary-data | no-binary-data> <files number> <size size> <world-readable |
        no-world-readable>;
    console {
        any | authorization | change-log | conflict-log | daemon | dfc | external | firewall | ftp
            | interactive-commands | kernel | ntp | pfe | security | user) (alert | any | critical |
            emergency | error | info | none | notice | warning);
    }
    file filename {
        facility severity;
        allow-duplicates;
        any (alert | any | critical | emergency | error | info | none | notice | warning);
        archive <archive-sites {ftp-url <password password>}> <files number> <size size>
            <start-time "YYYY-MM-DD.hh:mm"> <transfer-interval minutes> <world-readable |
            no-world-readable>;
        authorization (alert | any | critical | emergency | error | info | none | notice | warning);
        change-log (alert | any | critical | emergency | error | info | none | notice | warning);
        conflict-log (alert | any | critical | emergency | error | info | none | notice | warning);
        daemon (alert | any | critical | emergency | error | info | none | notice | warning);
        dfc (alert | any | critical | emergency | error | info | none | notice | warning);
        explicit-priority;
        external (alert | any | critical | emergency | error | info | none | notice | warning);
        firewall (alert | any | critical | emergency | error | info | none | notice | warning);
        ftp (alert | any | critical | emergency | error | info | none | notice | warning);
        interactive-commands (alert | any | critical | emergency | error | info | none | notice
            | warning);
        kernel (alert | any | critical | emergency | error | info | none | notice | warning);
        match "regular-expression";
    }
}

```

```

ntp (alert | any | critical | emergency | error | info | none | notice | warning);
pfe (alert | any | critical | emergency | error | info | none | notice | warning);
security (alert | any | critical | emergency | error | info | none | notice | warning);
structured-data {
  brief
}
host (hostname | other-routing-engine | scc-master) {
  facility severity;
  authorization (alert | any | critical | emergency | error | info | none | notice | warning);
  change-log (alert | any | critical | emergency | error | info | none | notice | warning);
  conflict-log (alert | any | critical | emergency | error | info | none | notice | warning);
  daemon (alert | any | critical | emergency | error | info | none | notice | warning);
  dfc (alert | any | critical | emergency | error | info | none | notice | warning);
  explicit-priority;
  external (alert | any | critical | emergency | error | info | none | notice | warning);
  facility-override facility;
  firewall (alert | any | critical | emergency | error | info | none | notice | warning);
  ftp (alert | any | critical | emergency | error | info | none | notice | warning);
  interactive-commands (alert | any | critical | emergency | error | info | none | notice
    | warning);
  kernel (alert | any | critical | emergency | error | info | none | notice | warning);
  log-prefix string;
  match "regular-expression";
  ntp (alert | any | critical | emergency | error | info | none | notice | warning);
  pfe (alert | any | critical | emergency | error | info | none | notice | warning);
  security (alert | any | critical | emergency | error | info | none | notice | warning);
  source-address source-address;
  structured-data {
    brief
  }
  user (username | *) {
  }
  log-rotate-frequency minutes;
  server;
  source-address address;
  time-format (year | millisecond | year millisecond);
  user (username | *) {
    facility severity;
    match "regular-expression";
  }
}
tacplus-options {
  (exclude-cmd-attribute | no-cmd-attribute-value);
  service-name service-name;
}
tacplus-server {
  server-address {
    port port-number;
    secret password;
    single-connection;
    source-address source-address;
    timeout seconds;
  }
}
time-zone (GMT | GMT+hour-offset | GMT-hour-offset | zone-name);
tracing destination-override syslog host address;
use-imported-time-zones;

```

```

}
}
system {
  services {
    database-replication {
      traceoptions {
        file <filename> <files number> <match regular-expression>
          <size maximum-file-size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
      }
    }
    dhcp-local-server {
      authentication {
        password password;
        username-include {
          circuit-type;
          delimiter delimiter-character;
          domain-name domain-name;
          logical-system-name;
          mac-address;
          option-60;
          option-82 <circuit-id> <remote-id>;
          routing-instance-name;
          user-prefix user-prefix;
        }
      }
      duplicate-clients-on-interface;
      dynamic-profile (profile-name | junos-default-profile) <aggregate-clients <merge |
        replace> | use-primary primary-profile-name>;
      forward-snooped-clients (all-interfaces | configured-interfaces |
        non-configured-interfaces);
      group group-name {
        dynamic-profile (profile-name | junos-default-profile) <aggregate-clients <merge |
          replace> | use-primary primary-profile-name>;
        interface interface-name {
          exclude;
          overrides {
            ...same statements as at the [edit system services dhcp-local-server overrides]
              hierarchy level ...
          }
          trace;
          upto upto-interface-name;
        }
      }
      overrides {
        client-discover-match <option60-and-option82>;
        interface-client-limit number;
        no-arp;
        process-inform {
          pool pool-name;
        }
      }
      pool-match-order {
        external-authority;
        ip-address-first;
      }
    }
  }
}

```

```

    option-82;
}
reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {
        radius-disconnect;
    }
}
traceoptions {
    file <filename> <files number> <match regular-expression>
        <size maximum-file-size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
}
dhcpv4-profiles profile-name {
    bind-interface interface-name;
    dead-server-retry-interval interval-in-seconds;
    dead-server-successive-retry-attempt number-of-attempts;
    dhcp-server-selection-algorithm (highest-priority-server | round-robin);
    lease-time time-in-seconds;
    pool-name pool-name;
    retransmission-attempt number-of-attempts;
    retransmission-interval interval-in-seconds;
    servers ip-address {
        priority value;
    }
}
}
dhcpv6-profiles profile-name {
    bind-interface interface-name;
    lease-time time-in-seconds;
    pool-name pool-name;
    retransmission-attempt number-of-attempts;
    retransmission-interval interval-in-seconds;
}
traceoptions {
    file <filename> <files number> <match regular-expression>
        <size maximum-file-size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
}
}
finger {
    connection-limit limit;
    rate-limit limit;
}
}
flow-tap-dtcp {
    ssh {
        connection-limit limit;
        rate-limit limit;
    }
}
}
}

```



```

ftp {
    connection-limit limit;
    rate-limit limit;
}
local-policy-decision-function {
    statistics {
        aacl-statistics-profile profile-name {
            aacl-fields {
                address;
                all-fields;
                application;
                application-group;
                input-bytes;
                input-interface;
                input-packets;
                ipv6-address
                ipv6-prefix-length
                mask;
                output-bytes;
                output-packets;
                subscriber-name;
                timestamp;
                vrf-name;
            }
            file filename;
            record-type (delta | interim);
        }
        file filename {
            archive-sites {
                url;
            }
            files number;
            size bytes;
            transfer-interval minutes;
        }
        record-type (data | interim);
    }
    traceoptions {
        file <filename> <files number> <match regular-expression>
            <size maximum-file-size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
netconf {
    ssh {
        connection-limit limit;
        port port;
        rate-limit limit;
    }
    traceoptions {
        file <filename> <files number> <match regular-expression> <size size>
            <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
        on-demand;
    }
}

```

```
    }
  }
  outbound-ssh {
    client client-id {
      address {
        port port-number;
        retry number;
        timeout seconds;
      }
      device-id device-id;
      keep-alive {
        retry number;
        timeout seconds;
      }
      reconnect-strategy (in-order | sticky);
      secret secret;
      services netconf;
    }
    traceoptions {
      file <filename> <files number> <match regular-expression>
        <size maximum-file-size> <world-readable | no-world-readable>;
      flag flag;
      no-remote-trace;
    }
  }
  resource-monitor {
    resource-category jtree {
      resource-type free-dwords {
        low-watermark number;
        high-watermark number;
      }
      resource-type free-pages {
        low-watermark number;
        high-watermark number;
      }
    }
  }
  no-throttle;
  no-logging;
  high-threshold number;
  traceoptions {
    file filename <files number> <match regular-expression> <size maximum-file-size>
      <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
service-deployment {
  local-certificate certificate-name;
  servers {
    server-address {
      port port-number;
      security-options {
        (ssl3 | tls);
      }
      user username;
    }
  }
}
```

```

    }
    source-address source-address;
    traceoptions {
        flag flag;
    }
}
ssh {
    ciphers [ cipher-1 cipher-2 cipher-3 ...]
    client-alive-count-max seconds;
    client-alive-interval seconds;
    connection-limit limit;
    hostkey-algorithm limit;
    key-exchange limit;
    macs limit;
    max-sessions-per-connection number;
    no-tcp-forwarding;
    protocol-version [v1 v2];
    rate-limit limit;
    root-login (allow | deny | deny-password);
}
subscriber-management {
    enforce-strict-scale-limit-license;
    gres-route-flush-delay;
    maintain-subscriber {
        interface-delete;
    }
}
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size>
    <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
}
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size>
    <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
}
telnet {
    connection-limit limit;
    rate-limit limit;
}
tftp-server {
    connection-limit limit;
    rate-limit limit;
}
xnm-clear-text {
    connection-limit limit;
    rate-limit limit;
}
xnm-ssl {
    connection-limit limit;
    local-certificate certificate-name;
    rate-limit limit;
}
}

```

}

Related Documentation • *Notational Conventions Used in Junos OS Configuration Hierarchies*

attempts (DHCP Local Server)

Syntax	<code>attempts attempt-count;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure],</p> <p>[edit system services dhcp-local-server reconfigure],</p> <p>[edit system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.0.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit ... dhcpv6 ...] hierarchy levels introduced in Junos OS Release 10.4.</p>
Description	Configure how many attempts are made to reconfigure all DHCP clients or only the DHCP clients serviced by the specified group of interfaces before reconfiguration is considered to have failed. A group configuration takes precedence over a DHCP local server configuration.
Options	<p><i>attempt-count</i>—Maximum number of attempts.</p> <p>Range: 1 through 10</p> <p>Default: 8</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 936 • Configuring Dynamic Reconfiguration Attempts for DHCP Clients on page 937

authentication (DHCP Local Server)

Syntax	<pre> authentication { password <i>password-string</i>; username-include { circuit-type; client-id; delimiter <i>delimiter-character</i>; domain-name <i>domain-name-string</i>; interface-name; logical-system-name; mac-address; option-60; option-82 <circuit-id> <remote-id>; relay-agent-interface-id; relay-agent-remote-id; relay-agent-subscriber-id; routing-instance-name; user-prefix <i>user-prefix-string</i>; } }</pre>
Hierarchy Level	<pre> [edit system services dhcp-local-server], [edit system services dhcp-local-server dhcpv6], [edit system services dhcp-local-server dhcpv6 group group-name], [edit system services dhcp-local-server group group-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server ...], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...]</pre>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure the parameters the router sends to the external AAA server. A group configuration takes precedence over a global DHCP relay or DHCP local server configuration.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Using External AAA Authentication Services with DHCP on page 923

bfd

Syntax	<pre> bfd { version (0 1 automatic); minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; multiplier <i>number</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } detection-time { threshold <i>milliseconds</i>; } session-mode (automatic multihop singlehop); holddown-interval <i>milliseconds</i>; } </pre>
Hierarchy Level	<p>[edit system services dhcp-local-server liveness-detection method], [edit system services dhcp-local-server dhcpv6 liveness-detection method], [edit forwarding-options dhcp-relay liveness-detection method], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method]</p>
Release Information	<p>Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure Bidirectional Forwarding Detection (BFD) as the liveness detection method.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913 • Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

circuit-type (DHCP Local Server)

Syntax	circuit-type;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit system services dhcp-local-server authentication username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit system services dhcp-local-server group group-name authentication username-include]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Specify that the circuit type is concatenated with the username during the subscriber authentication or client authentication process.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Using External AAA Authentication Services with DHCP on page 923

clear-on-abort (DHCP Local Server)

Syntax	clear-on-abort;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure],</p> <p>[edit system services dhcp-local-server reconfigure],</p> <p>[edit system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.0.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit ... dhcpv6 ...] hierarchy levels introduced in Junos OS Release 10.4.</p>
Description	Delete all DHCP clients or only the DHCP clients serviced by the specified group of interfaces when reconfiguration fails; that is, when the maximum number of retry attempts have been made without success. A group configuration takes precedence over a DHCP local server configuration.
Default	Restores the original client configuration when reconfiguration fails.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 936 • Configuring Deletion of the Client When Dynamic Reconfiguration Fails on page 938

client-discover-match (DHCP Local Server)

Syntax	client-discover-match <option60-and-option82>;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> overrides],</p> <p>[edit system services dhcp-local-server overrides],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> overrides],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i> overrides]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure DHCP local server to use option 60 and option 82 information to uniquely identify DHCP subscribers or clients when primary subscriber or client identification fails. The statement always uses the option60-and-option82 option. Specifying the option has no effect.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Extended DHCP Local Server Overview on page 876• Overriding Default DHCP Local Server Configuration Settings on page 928

client-id (DHCP Local Server)

Syntax	client-id;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication username-include]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Specify that the DHCPv6 Client-ID option (option 1) in the client PDU name is concatenated with the username during the subscriber authentication or client authentication process.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Creating Unique Usernames for DHCP Clients on page 946

delegated-pool (DHCP Local Server)

Syntax	<code>delegated-pool <i>pool-name</i>;</code>
Hierarchy Level	<code>[edit system services dhcp-local-server dhcpv6 overrides],</code> <code>[edit system services dhcp-local-server dhcpv6 group group-name overrides],</code> <code>[edit system services dhcp-local-server dhcpv6 group group-name interface <i>interface-name</i></code> <code>overrides],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system</code> <code>services dhcp-local-server dhcpv6 ...],</code> <code>[edit logical-systems <i>logical-system-name</i> system services system services dhcp-local-server</code> <code>dhcpv6 ...],</code> <code>[edit routing-instances <i>routing-instance-name</i> system services system services</code> <code>dhcp-local-server dhcpv6 ...]</code>
Release Information	Statement introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify the address pool that assigns the IA_PD address. A pool specified by RADIUS VSA 26-161 takes precedence over the pool specified by this delegated-pool statement.
Options	<i>pool-name</i> —Name of the address-assignment pool.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Specifying the Delegated Address Pool for IPv6 Prefix Assignment on page 934• Overriding Default DHCP Local Server Configuration Settings on page 928

delimiter (DHCP Local Server)

Syntax	<code>delimiter <i>delimiter-character</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit system services dhcp-local-server authentication username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit system services dhcp-local-server group group-name authentication username-include]</p>
Release Information	Statement introduced in Junos OS Release 9.1.
Description	<p>Specify the character used as the delimiter between the concatenated components of the username.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Options	<i>delimiter-character</i> —Character that separates components that make up the concatenated username. You cannot use the semicolon (;) as a delimiter.

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

Related Documentation

- [Using External AAA Authentication Services with DHCP on page 923](#)

detection-time

Syntax `detection-time {
 threshold milliseconds;
}`

Hierarchy Level [edit system services dhcp-local-server liveness-detection method [bfd](#)],
[edit system services dhcp-local-server dhcpv6 liveness-detection method [bfd](#)],
[edit forwarding-options dhcp-relay liveness-detection method [bfd](#)], [edit forwarding-options
 dhcp-relay dhcpv6 liveness-detection method [bfd](#)],
[edit system services dhcp-local-server group *group-name* liveness-detection method [bfd](#)],
[edit system services dhcp-local-server dhcpv6 group *group-name* liveness-detection method
 [bfd](#)],
[edit forwarding-options dhcp-relay group *group-name* liveness-detection method [bfd](#)],
[edit forwarding-options dhcp-relay dhcpv6 group *group-name* liveness-detection method
 [bfd](#)]

Release Information Statement introduced in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Enable failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the one configured.

The remaining statement is explained separately.

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

Related Documentation

- [Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913](#)
- [Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients](#)

dhcp-local-server

```
Syntax  dhcp-local-server {
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                logical-system-name;
                mac-address;
                option-60;
                option-82 <circuit-id> <remote-id>;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        dhcpv6 {
            authentication {
                ...
            }
            group group-name {
                authentication {
                    ...
                }
                interface interface-name {
                    exclude;
                    liveness-detection {
                        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                        method {
                            bfd {
                                version (0 | 1 | automatic);
                                minimum-interval milliseconds;
                                minimum-receive-interval milliseconds;
                                multiplier number;
                                no-adaptation;
                                transmit-interval {
                                    minimum-interval milliseconds;
                                    threshold milliseconds;
                                }
                                detection-time {
                                    threshold milliseconds;
                                }
                            }
                            session-mode (automatic | multihop | singlehop);
                            holddown-interval milliseconds;
                        }
                    }
                }
            }
            overrides {
                interface-client-limit number;
                process-inform {
                    pool pool-name;
                }
            }
        }
    }
```

```
        rapid-commit;
    }
    service-profile dynamic-profile-name;
    trace;
    upto upto-interface-name;
}
overrides {
    delegated-pool;
    interface-client-limit number;
    process-inform {
        pool pool-name;
    }
    rapid-commit;
}
service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
overrides {
    delegated-pool;
    interface-client-limit number;
    process-inform {
        pool pool-name;
    }
    rapid-commit;
}
reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {
        radius-disconnect;
    }
}
service-profile dynamic-profile-name;
```



```

}
duplicate-clients-on-interface;
dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
  primary-profile-name>;
forward-snooped-clients (all-interfaces | configured-interfaces |
  non-configured-interfaces);
group group-name {
  authentication {
    ...
  }
}
dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
  primary-profile-name>;
interface interface-name {
  exclude;
  liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
      bfd {
        version (0 | 1 | automatic);
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
        detection-time {
          threshold milliseconds;
        }
        session-mode (automatic | multihop | singlehop);
        holddown-interval milliseconds;
      }
    }
  }
}
overrides {
  client-discover-match <option60-and-option82>;
  interface-client-limit number;
  no-arp;
  process-inform {
    pool pool-name;
  }
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
overrides {
  client-discover-match <option60-and-option82>;
  interface-client-limit number;
  no-arp;
  process-inform {
    pool pool-name;
  }
}
}
requested-ip-network-match subnet-mask


```

```

    service-profile dynamic-profile-name;
  }
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
    }
    session-mode (automatic | multihop | singlehop);
    holddown-interval milliseconds;
  }
}
}
overrides {
  client-discover-match <option60-and-option82>;
  interface-client-limit number;
  no-arp;
  process-inform {
    pool pool-name;
  }
}
pool-match-order {
  external-authority;
  ip-address-first;
  option-82;
}
reconfigure {
  attempts attempt-count;
  clear-on-abort;
  strict;
  timeout timeout-value;
  token token-value;
  trigger {
    radius-disconnect;
  }
}
}
requested-ip-network-match subnet-mask;
service-profile dynamic-profile-name;
}

```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services],
 [edit logical-systems *logical-system-name* system services],
 [edit routing-instances *routing-instance-name* system services],
 [edit system services]

Release Information	<p>Statement introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>requested-ip-network-match option introduced in Junos OS Release 12.2 (MX Series routers only).</p>
Description	<p>Configure Dynamic Host Configuration Protocol (DHCP) local server options on the router or switch and enable the router or switch to function as an extended DHCP local server. The DHCP local server receives DHCP request and reply packets from DHCP clients and then responds with an IP address and other optional configuration information to the client.</p> <p>The extended DHCP local server is incompatible with the DHCP server on J Series routers and so is not supported on J Series routers. Also, the DHCP local server and the DHCP/BOOTP relay server, which are configured under the [edit forwarding-options helpers] hierarchy level, cannot both be enabled on the router or switch at the same time. The extended DHCP local server is fully compatible with the extended DHCP relay feature.</p> <p>The dhcpv6 stanza configures the router or switch to support Dynamic Host Configuration Protocol for IPv6 (DHCPv6). The DHCPv6 local server is fully compatible with the extended DHCP local server and the extended DHCP relay feature.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 20px;">  <p>NOTE: When you configure the dhcp-local-server statement at the routing instance hierarchy level, you must use a routing instance type of virtual-router.</p> </div> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Extended DHCP Local Server Overview on page 876 • DHCPv6 Local Server Overview on page 880 • Configuring a DHCP Server on Switches (CLI Procedure)

dhcpv6 (DHCP Local Server)

```
Syntax  dhcpv6 {
    authentication {
        password password-string;
        username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name-string;
            logical-system-name;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix-string;
        }
    }
    group group-name {
        authentication {
            ...
        }
        interface interface-name {
            exclude;
            liveness-detection {
                failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                method {
                    bfd {
                        version (0 | 1 | automatic);
                        minimum-interval milliseconds;
                        minimum-receive-interval milliseconds;
                        multiplier number;
                        no-adaptation;
                        transmit-interval {
                            minimum-interval milliseconds;
                            threshold milliseconds;
                        }
                        detection-time {
                            threshold milliseconds;
                        }
                    }
                    session-mode (automatic | multihop | singlehop);
                    holddown-interval milliseconds;
                }
            }
        }
        overrides {
            interface-client-limit number;
            process-inform {
                pool pool-name;
            }
            rapid-commit;
        }
        service-profile dynamic-profile-name;
        trace;
        upto upto-interface-name;
    }
}
```

```

}
overrides {
  delegated-pool;
  interface-client-limit number;
  process-inform {
    pool pool-name;
  }
  rapid-commit;
}
service-profile dynamic-profile-name;
}
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
      session-mode (automatic | multihop | singlehop);
      holddown-interval milliseconds;
    }
  }
}
overrides {
  delegated-pool;
  interface-client-limit number;
  process-inform {
    pool pool-name;
  }
  rapid-commit;
  reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {
      radius-disconnect;
    }
  }
}
reconfigure {
  attempts attempt-count;
  clear-on-abort;
  strict;
  timeout timeout-value;
  token token-value;

```

```
trigger {  
    radius-disconnect;  
}  
}  
requested-ip-network-match subnet-mask;  
service-profile dynamic-profile-name;  
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services *dhcp-local-server*],
[edit logical-systems *logical-system-name* system services *dhcp-local-server*],
[edit routing-instances *routing-instance-name* system services *dhcp-local-server*],
[edit system services *dhcp-local-server*]

Release Information Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.3 for EX Series switches.

Description Configure DHCPv6 local server options on the router or switch and enable the router or switch to function as a server for the DHCP protocol for IP version 6 (IPv6). The DHCPv6 local server sends and receives packets using the IPv6 protocol and informs IPv6 of the routing requirements of router clients. The local server works together with the AAA service framework to control subscriber access (or DHCP client access) and accounting.

The DHCPv6 local server is fully compatible with the extended DHCP local server and DHCP relay agent.

The remaining statements are explained separately.

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

Related Documentation • [DHCPv6 Local Server Overview on page 880](#)

domain-name (DHCP Local Server)

Syntax	<code>domain-name <i>domain-name-string</i>;</code>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server authentication username-include], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 authentication username-include], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group group-name authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include], [edit system services dhcp], [edit system services dhcp-local-server authentication username-include], [edit system services dhcp-local-server dhcpv6 authentication username-include], [edit system services dhcp-local-server dhcpv6 group group-name authentication username-include], [edit system services dhcp-local-server group group-name authentication username-include] </pre>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Specify the domain name that is concatenated with the username during the subscriber authentication or DHCP client authentication process.
Options	<i>domain-name-string</i> —Domain name formatted string.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Using External AAA Authentication Services with DHCP on page 923](#)

duplicate-clients-on-interface (DHCP Local Server)

Syntax	duplicate-clients-on-interface;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server], [edit system services dhcp-local-server]
Release Information	Statement introduced in Junos OS Release 10.2. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure DHCP local server to include the client subinterface when distinguishing between duplicate DHCP clients (clients with the same MAC address or client ID) in the same subnet. By default, DHCP distinguishes clients by subnet. This feature is supported on DHCPv4 only.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	• Configuring DHCP Duplicate Client Support on page 925

dynamic-profile (DHCP Local Server)

Syntax	<pre>dynamic-profile <i>profile-name</i> { aggregate-clients (merge replace); use-primary <i>primary-profile-name</i>; }</pre>
Hierarchy Level	<pre>[edit system services dhcp-local-server], [edit system services dhcp-local-server dhcpv6], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i>], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> interface <i>interface-name</i>], [edit system services dhcp-local-server group <i>group-name</i>], [edit system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...]</pre>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Options aggregate-clients and use-primary introduced in Junos OS Release 9.3.</p> <p>Support at the [edit ... interface] hierarchy levels introduced in Junos OS Release 11.2.</p>
Description	Specify the dynamic profile that is attached to all interfaces, a named group of interfaces, or a specific interface.
Options	<p><i>profile-name</i>—Name of the dynamic profile.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 942 • Configuring a Default Subscriber Service

external-authority

Syntax	external-authority;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server pool-match-order], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server pool-match-order], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server pool-match-order], [edit system services dhcp-local-server pool-match-order]
Release Information	Statement introduced in Junos OS Release 10.0. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>Specify that an external authority (for example, RADIUS or Diameter) provides the address assignment.</p> <p>When RADIUS is the external authority, the router uses the Framed-IPv6-Pool attribute (RADIUS attribute 100) to select the pool. When Diameter is the external authority, the router uses the Diameter counterpart of RADIUS Framed-IPv6-Pool attribute.</p>
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 949• Extended DHCP Local Server Overview on page 876• Address-Assignment Pools Overview on page 890

failure-action

Syntax	failure-action (clear-binding clear-binding-if-interface-up log-only);
Hierarchy Level	[edit system services dhcp-local-server liveness-detection], [edit system services dhcp-local-server dhcpv6 liveness-detection], [edit forwarding-options dhcp-relay liveness-detection], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the action the router (or switch) takes when a liveness detection failure occurs.
Options	<p>clear-binding—The client session is cleared when a liveness detection failure occurs.</p> <p>clear-binding-if-interface-up—The client session is cleared only when a liveness detection failure occurs and the local interface is detected as being up.</p> <p>log-only—A message is logged to indicate the event; no action is taken and DHCP is left to manage the failure.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • DHCP Liveness Detection Overview on page 980 • Configuring Detection of DHCP Local Server Client Connectivity on page 941 • Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 981 • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913 • Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

forward-snooped-clients (DHCP Local Server)

Syntax	forward-snooped-clients (all-interfaces configured-interfaces non-configured-interfaces);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server], [edit system services dhcp-local-server]
Release Information	Statement introduced in Junos OS Release 10.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure how the DHCP local server handles DHCP snooped packets on specific interfaces.
Options	all-interfaces —Perform the action on all interfaces. configured-interfaces —Perform the action only on configured interfaces. non-configured-interfaces —Perform the action only on nonconfigured interfaces.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• DHCP Snooping Support on page 887• Configuring DHCP Snooped Packets Forwarding Support for DHCP Local Server on page 944

group (DHCP Local Server)

```
Syntax  group group-name {
        authentication {
            password password-string;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name-string;
                logical-system-name;
                mac-address;
                option-60;
                option-82 <circuit-id> <remote-id>;
                relay-agent-interface-id
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
            primary-profile-name>;
        interface interface-name {
            exclude;
            overrides {
                client-discover-match <option60-and-option82>;
                interface-client-limit number;
                no-arp;
                process-inform {
                    pool pool-name;
                }
                rapid-commit;
            }
            service-profile dynamic-profile-name;
            trace;
            upto upto-interface-name;
        }
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                }
                detection-time {
                    threshold milliseconds;
                }
            }
        }
    }
```

```
        session-mode(automatic | multihop | singlehop);
        holddown-interval milliseconds;
    }
}
overrides {
    client-discover-match <option60-and-option82>;
    delegated-pool;
    interface-client-limit number;
    no-arp;
    process-inform {
        pool pool-name;
    }
    rapid-commit;
}
reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {
        radius-disconnect;
    }
}
service-profile dynamic-profile-name;
}
```

Hierarchy Level	[edit system services dhcp-local-server], [edit system services dhcp-local-server dhcpv6], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server ...], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...]
Release Information	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 12.1 for EX Series switches.
Description	Configure a group of interfaces that have a common configuration, such as authentication parameters. A group must contain at least one interface.
Options	group-name —Name of the group. The remaining statements are explained separately.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.

- Related Documentation**
- [Extended DHCP Local Server Overview on page 876](#)
 - [Grouping Interfaces with Common DHCP Configurations on page 926](#)
 - [Using External AAA Authentication Services with DHCP on page 923](#)
 - [Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 942](#)
 - [Configuring a DHCP Server on Switches \(CLI Procedure\)](#)

holddown-interval

Syntax	<code>holddown-interval <i>milliseconds</i>;</code>
Hierarchy Level	[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the time (in milliseconds) for which Bidirectional Forwarding Detection (BFD) holds a session up notification.
Options	<i>milliseconds</i> —Interval specifying how long a BFD session must remain up before a state change notification is sent. Range: 0 through 255,000 Default: 0
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913 • Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

interface (DHCP Local Server)

Syntax `interface interface-name {
 exclude;
 overrides {
 client-discover-match <option60-and-option82>;
 interface-client-limit number;
 no-arp;
 rapid-commit;
 }
 service-profile dynamic-profile-name;
 trace;
 upto upto-interface-name;
 }`

Hierarchy Level [edit system services dhcp-local-server *group group-name*],
 [edit system services dhcp-local-server *dhcpv6 group group-name*]
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system
 services *dhcp-local-server ...*],
 [edit logical-systems *logical-system-name* system services *dhcp-local-server ...*],
 [edit routing-instances *routing-instance-name* system services *dhcp-local-server ...*]

Release Information Statement introduced in Junos OS Release 9.0.
 Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
 Options **upto** and **exclude** introduced in Junos OS Release 9.1.

Description Specify one or more interfaces, or a range of interfaces, that are within a specified group on which the DHCP local server is enabled. You can repeat the **interface *interface-name*** statement to specify multiple interfaces within a group, but you cannot specify the same interface in more than one group. Also, you cannot use an interface that is being used by the DHCP relay agent.



NOTE: DHCP values are supported in Integrated Routing and Bridging (IRB) configurations. When you configure an IRB interface in a network that is using DHCP, the DHCP information (for example, authentication, address assignment, and so on) is propagated in the associated bridge domain. This enables the DHCP server to configure client IP addresses residing within the bridge domain. IRB currently only supports static DHCP configurations. For additional information about how to configure IRB, see the *JUNOS® MX Series 3D Universal Edge Routers Solutions, Release 12.3*.

Options **exclude**—Exclude an interface or a range of interfaces from the group. This option and the **overrides** option are mutually exclusive.

interface-name—Name of the interface. You can repeat this option multiple times.

upto-interface-name—Upper end of the range of interfaces; the lower end of the range is the ***interface-name*** entry. The interface device name of the ***upto-interface-name*** must be the same as the device name of the ***interface-name***.

The remaining statements are explained separately.

Required Privilege Level	system—To view this statement in the configuration.
	system-control—To add this statement to the configuration.
Related Documentation	• Extended DHCP Local Server Overview on page 876
	• Grouping Interfaces with Common DHCP Configurations on page 926
	• Using External AAA Authentication Services with DHCP on page 923

interface-client-limit (DHCP Local Server)

Syntax	<code>interface-client-limit <i>number</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group group-name overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> overrides],</p> <p>[edit system services dhcp-local-server overrides],</p> <p>[edit system services dhcp-local-server dhcpv6 overrides],</p> <p>[edit system services dhcp-local-server dhcpv6 group group-name overrides],</p> <p>[edit system services dhcp-local-server dhcpv6 group interface <i>interface-name</i> <i>group-name</i> overrides],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> overrides],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i> overrides]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Set the maximum number of DHCP subscribers or DHCP clients per interface allowed for a specific group or for all groups. A group specification takes precedence over a global specification for the members of that group.
Default	No limit
Options	<i>number</i> —Maximum number of clients allowed.

Range: 1 through 500,000

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

Related Documentation

- [Specifying the Maximum Number of DHCP Clients Per Interface on page 929](#)
- [Overriding Default DHCP Local Server Configuration Settings on page 928](#)

interface-delete (Subscriber Management or DHCP Client Management)

Syntax interface-delete;

Hierarchy Level [edit system services subscriber-management maintain-subscriber]

Release Information Statement introduced in Junos OS Release 11.1.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description On router—Configure the router to maintain, rather than log out, subscribers when the subscriber interface is deleted. By default, the router logs out subscribers when the subscriber interface is deleted.

On switch—Configure the switch to maintain rather than log out DHCP clients when the client interface is deleted. By default, the switch logs out DHCP clients when the client interface is deleted.

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

Related Documentation

- [Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events](#)

interface-name (DHCP Local Server)

Syntax	interface-name;
Hierarchy Level	[edit system services dhcp-local-server authentication username-include], [edit system services dhcp-local-server dhcpv6 authentication username-include], [edit system services dhcp-local-server dhcpv6 group group-name authentication username-include], [edit system services dhcp-local-server group group-name authentication username-include] [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server ...], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...]
Release Information	Statement introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify that the interface name is concatenated with the username during the subscriber authentication or DHCP client authentication process. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Creating Unique Usernames for DHCP Clients on page 946

ip-address-first

Syntax	ip-address-first;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server pool-match-order],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server pool-match-order],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server pool-match-order],</p> <p>[edit system services dhcp-local-server pool-match-order]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Configure the extended DHCP local server to use the IP address method to determine which address-assignment pool to use. The local server uses the IP address in the gateway IP address if one is present in the DHCP client PDU. If no gateway IP address is present, the local server uses the IP address of the receiving interface to find the address-assignment pool. The DHCP local server uses this method by default when no method is explicitly specified.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 949 • Extended DHCP Local Server Overview on page 876 • Address-Assignment Pools Overview on page 890 • <i>Configuring a DHCP Server on Switches (CLI Procedure)</i>

liveness-detection

Syntax	<pre> liveness-detection { failure-action (clear-binding clear-binding-if-interface-up log-only); method { bfd { version (0 1 automatic); minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; multiplier <i>number</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } detection-time { threshold <i>milliseconds</i>; } session-mode (automatic multihop singlehop); holddown-interval <i>milliseconds</i>; } } } </pre>
Hierarchy Level	<pre> [edit system services dhcp-local-server], [edit system services dhcp-local-server dhcpv6], [edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit system services dhcp-local-server group group-name], [edit system services dhcp-local-server dhcpv6 group group-name], [edit forwarding-options dhcp-relay group group-name], [edit forwarding-options dhcp-relay dhcpv6 group group-name] </pre>
Release Information	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure bidirectional failure detection timers and authentication criteria for static routes.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • DHCP Liveness Detection Overview on page 980 • Configuring Detection of DHCP Local Server Client Connectivity on page 941 • Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 981 • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913

- *Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients*

logical-system-name (DHCP Local Server)

Syntax	logical-system-name;
Hierarchy Level	<p>[edit system services dhcp-local-server authentication username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit system services dhcp-local-server group group-name authentication username-include]</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...]</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server ...],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Specify that the logical system name be concatenated with the username during the subscriber authentication or DHCP client process. No logical system name is concatenated if the configuration is in the default logical system.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Using External AAA Authentication Services with DHCP on page 923

mac-address (DHCP Local Server)

Syntax	mac-address;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit system services dhcp-local-server authentication username-include],</p> <p>[edit system services dhcp-local-server group group-name authentication username-include]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Specify that the MAC address from the client PDU be concatenated with the username during the subscriber authentication or DHCP client authentication process.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Using External AAA Authentication Services with DHCP on page 923

method

Syntax	<pre> method { bfd { version (0 1 automatic); minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; multiplier <i>number</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } detection-time { threshold <i>milliseconds</i>; } session-mode (automatic multihop singlehop); holddown-interval <i>milliseconds</i>; } } </pre>
Hierarchy Level	<p> [edit system services dhcp-local-server liveness-detection], [edit system services dhcp-local-server dhcpv6 liveness-detection], [edit forwarding-options dhcp-relay liveness-detection], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection] </p>
Release Information	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure the liveness detection method.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913 • Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

minimum-interval

Syntax	<code>minimum-interval <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server liveness-detection method bfd transmit-interval], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd transmit-interval], [edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection method bfd transmit-interval], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd transmit-interval], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd transmit-interval], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd transmit-interval], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd transmit-interval], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd transmit-interval]</p>
Release Information	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure the minimum intervals at which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the transmit-interval minimal-interval and minimum-receive-interval statements.</p>
Options	<p><i>milliseconds</i> — Specify the minimum interval value for BFD liveliness detection.</p> <p>Range: 1 through 255,000</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913 • Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

minimum-receive-interval

Syntax	<code>minimum-receive-interval <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]</p>
Release Information	<p>Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Configure the minimum interval at which the local routing device (or switch) must receive a reply from a neighbor with which it has established a BFD session.
Options	<p><i>milliseconds</i> — Specify the minimum receive interval value. Range: 1 through 255,000</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913 • Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

multiplier

Syntax	<code>multiplier <i>number</i>;</code>
Hierarchy Level	<code>[edit system services dhcp-local-server liveness-detection method bfd],</code> <code>[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],</code> <code>[edit forwarding-options dhcp-relay liveness-detection method bfd],</code> <code>[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],</code> <code>[edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd],</code> <code>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd],</code> <code>[edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd],</code> <code>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]</code>
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the number of hello packets not received by the neighbor before Bidirectional Forwarding Detection (BFD) declares the neighbor down.
Options	number —Maximum allowable number of hello packets missed by the neighbor. Range: 1 through 255 Default: 3
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913• Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

no-adaptation

Syntax	no-adaptation;
Hierarchy Level	[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure Bidirectional Forwarding Detection (BFD) sessions to not adapt to changing network conditions.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913 • <i>Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</i>

no-arp (DHCP Local Server)

Syntax	no-arp;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> overrides],</p> <p>[edit system services dhcp-local-server overrides],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> overrides]</p> <p>[edit system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i> overrides]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.3.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	Turn off ARP table population in a distrusted environment.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring a DHCP Server on Switches (CLI Procedure)</i>

option-60 (DHCP Local Server)

Syntax	option-60;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> authentication username-include],</p> <p>[edit system services dhcp-local-server authentication username-include],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> authentication username-include]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Specify that the payload of Option 60 (Vendor Class Identifier) from the client PDU be concatenated with the username during the subscriber authentication or DHCP client authentication process.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Using External AAA Authentication Services with DHCP on page 923

option-82 (DHCP Local Server Authentication)

Syntax	<code>option-82 <circuit-id> <remote-id>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> authentication username-include],</p> <p>[edit system services dhcp-local-server authentication username-include],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> authentication username-include]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Specify the type of Option 82 information from the client PDU that is concatenated with the username during the subscriber authentication or DHCP client authentication process. You can specify either, both, or neither of the Agent Circuit ID and Agent Remote ID suboptions. If you specify both, the Agent Circuit ID is supplied first, followed by a delimiter, and then the Agent Remote ID. If you specify that neither suboption is supplied, the raw payload of Option 82 from the PDU is concatenated to the username.</p>
Options	<p>circuit-id—(Optional) Agent Circuit ID suboption (suboption 1).</p> <p>remote-id—(Optional) Agent Remote ID suboption (suboption 2).</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Using External AAA Authentication Services with DHCP on page 923

option-82 (DHCP Local Server Pool Matching)

Syntax	option-82;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server pool-match-order],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server pool-match-order],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server pool-match-order],</p> <p>[edit system services dhcp-local-server pool-match-order]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure the extended DHCP local server to use the option 82 value in the DHCP client DHCP PDU together with the ip-address-first method to determine which address-assignment pool to use. You must configure the ip-address-first statement before configuring the option-82 statement. The DHCP local server first determines which address-assignment pool to use based on the ip-address-first method. Then, the local server matches the option 82 value in the client PDU with the option 82 configuration in the address-assignment pool. This statement is supported for IPv4 address-assignment pools only.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 949 • Extended DHCP Local Server Overview on page 876 • Address-Assignment Pools Overview on page 890

overrides (DHCP Local Server)

Syntax	<pre> overrides { client-discover-match; delegated-pool; interface-client-limit <i>number</i>; no-arp; process-inform { pool <i>pool-name</i>; } rapid-commit; } </pre>
Hierarchy Level	<pre> [edit system services dhcp-local-server], [edit system services dhcp-local-server dhcpv6], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i>], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> interface <i>interface-name</i>], [edit system services dhcp-local-server group <i>group-name</i>], [edit system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server ...], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...] </pre>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Override the default configuration settings for the extended DHCP local server. Specifying the overrides statement with no subordinate statements removes all DHCP local server overrides at that hierarchy level.</p> <ul style="list-style-type: none"> To override global DHCP local server configuration options, include the overrides statement and its subordinate statements at the [edit system services dhcp-local-server] hierarchy level. To override configuration options for a named group of interfaces, include the statements at the [edit system services dhcp-local-server group <i>group-name</i>] hierarchy level. To override configuration options for a specific interface within a named group of interfaces, include the statements at the [edit system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i>] hierarchy level. Use the [edit system services dhcp-local-server dhcpv6] hierarchy level to override DHCPv6 configuration options. <p>The remaining statements are explained separately.</p> <p>The interface-client-limit and no-arp statements are not supported in the [edit system services dhcp-local-server dhcpv6] hierarchy level.</p> <p>The delegated-pool and the rapid-commit statements are supported in the [edit system services dhcp-local-server dhcpv6 ...] hierarchy level only.</p>

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Extended DHCP Local Server Overview on page 876• Overriding Default DHCP Local Server Configuration Settings on page 928• Deleting DHCP Local Server and DHCP Relay Override Settings on page 935• <i>Configuring a DHCP Server on Switches (CLI Procedure)</i>

password (DHCP Local Server)

Syntax	<code>password password-string;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group group-name authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group group-name authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication],</p> <p>[edit system services dhcp-local-server authentication],</p> <p>[edit system services dhcp-local-server dhcpv6],</p> <p>[edit system services dhcp-local-server dhcpv6 group group-name authentication],</p> <p>[edit system services dhcp-local-server group group-name authentication]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Configure the password that is sent to the external AAA authentication server for subscriber authentication or DHCP client authentication.
Options	<i>password-string</i> —Authentication password.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Using External AAA Authentication Services with DHCP on page 923

pool (DHCP Local Server Overrides)

Syntax `pool pool-name;`

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server overrides [process-inform](#)],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 overrides [process-inform](#)],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* overrides [process-inform](#)],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* interface *interface-name* overrides [process-inform](#)],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* overrides [process-inform](#)],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* interface *interface-name* overrides [process-inform](#)],
 [edit logical-systems *logical-system-name* system services dhcp-local-server overrides [process-inform](#)],
 [edit logical-systems *logical-system-name* system services dhcp-local-server dhcpv6 overrides [process-inform](#)],
 [edit logical-systems *logical-system-name* system services dhcp-local-server dhcpv6 group *group-name* overrides [process-inform](#)],
 [edit logical-systems *logical-system-name* system services dhcp-local-server dhcpv6 group *group-name* interface *interface-name* overrides [process-inform](#)],
 [edit logical-systems *logical-system-name* system services dhcp-local-server group *group-name* overrides [process-inform](#)],
 [edit logical-systems *logical-system-name* system services dhcp-local-server group *group-name* interface *interface-name* overrides [process-inform](#)],
 [edit routing-instances *routing-instance-name* system services dhcp-local-server overrides [process-inform](#)],
 [edit routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 overrides [process-inform](#)],
 [edit routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* overrides [process-inform](#)],
 [edit routing-instances *routing-instance-name* system services dhcp-local-server dhcpv6 group *group-name* interface *interface-name* overrides [process-inform](#)],
 [edit routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* overrides [process-inform](#)],
 [edit routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* interface *interface-name* overrides [process-inform](#)],
 [edit system services dhcp-local-server overrides [process-inform](#)],
 [edit system services dhcp-local-server dhcpv6 overrides [process-inform](#)],
 [edit system services dhcp-local-server dhcpv6 group *group-name* overrides [process-inform](#)],
 [edit system services dhcp-local-server dhcpv6 group *group-name* interface *interface-name* overrides [process-inform](#)],
 [edit system services dhcp-local-server group *group-name* overrides [process-inform](#)],
 [edit system services dhcp-local-server group *group-name* interface *interface-name* overrides [process-inform](#)]

Release Information Statement introduced in Junos OS Release 11.4.
 Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description	Configure DHCP or DHCPv6 local server to reply to DHCP information request messages (DHCPINFORM for DHCPv4 and INFORMATION-REQUEST for DHCPv6) with information taken from the specified pool without interacting with AAA.
Options	pool-name —Name of the address pool, which must be configured within family inet for DHCP local server and within family inet6 for DHCPv6 local server.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling Processing of Client Information Requests on page 933• Overriding Default DHCP Local Server Configuration Settings on page 928

pool-match-order

Syntax	<pre>pool-match-order { external-authority; ip-address-first; option-82; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server], [edit system services dhcp-local-server]
Release Information	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 12.1.
Description	Configure the order in which the DHCP local server uses information in the DHCP client PDU to determine how to obtain an address for the client. The remaining statements are explained separately.
Default	DHCP local server uses the ip-address-first method to determine which address pool to use.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring How the Extended DHCP Local Server Determines Which Address-Assignment Pool to Use on page 949• Extended DHCP Local Server Overview on page 876• Configuring a DHCP Server on Switches (CLI Procedure)

process-inform

Syntax	<pre>process-inform { pool pool-name; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> interface <i>interface-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> interface <i>interface-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i> overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> interface <i>interface-name</i> overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i> overrides],</p> <p>[edit system services dhcp-local-server overrides],</p> <p>[edit system services dhcp-local-server dhcpv6 overrides],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> overrides],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> interface <i>interface-name</i> overrides],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> overrides],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i> overrides]</p>
Release Information	<p>Statement introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Enable the processing of DHCP information request messages (DHCPINFORM for DHCPv4 and INFORMATION-REQUEST for DHCPv6) sent from the client to request DHCP options. For DHCP local servers, the messages are also passed to the configured server list.</p>

The remaining statement is explained separately.

Default Information request messages are not processed.

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

Related Documentation

- [Enabling Processing of Client Information Requests on page 933](#)
- [Overriding Default DHCP Local Server Configuration Settings on page 928](#)
- [Configuring a DHCP Server on Switches \(CLI Procedure\)](#)

radius-disconnect (DHCP Local Server)

Syntax	radius-disconnect;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server reconfigure trigger],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 reconfigure trigger],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure trigger],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure trigger],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server reconfigure trigger],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 reconfigure trigger],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure trigger],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure trigger],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server reconfigure trigger],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 reconfigure trigger],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure trigger],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure trigger],</p> <p>[edit system services dhcp-local-server reconfigure trigger],</p> <p>[edit system services dhcp-local-server dhcpv6 reconfigure trigger],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> reconfigure trigger],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure trigger]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.0.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit ... dhcpv6 ...] hierarchy levels introduced in Junos OS Release 10.4.</p>
Description	Configure all DHCP clients or only the DHCP clients serviced by the specified group of interfaces to be reconfigured when a RADIUS-initiated disconnect is received by the DHCP client or group of clients. A group configuration takes precedence over a DHCP local server configuration.
Default	The client is deleted when a RADIUS-initiated disconnect is received.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 936 • Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect on page 938

rapid-commit (DHCPv6 Local Server)

Syntax	rapid-commit;
Hierarchy Level	[edit system services dhcp-local-server dhcpv6 overrides], [edit system services dhcp-local-server dhcpv6 group group-name overrides], [edit system services dhcp-local-server dhcpv6 group group-name interface interface-name overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 ...], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 ...], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 ...]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure DHCPv6 local server to recognize the Rapid Commit option (DHCPv6 option 14) in DHCPv6 solicit messages sent from the DHCPv6 client. When rapid commit is enabled for both DHCPv6 local server and the DHCPv6 client, a two-message handshake is used instead of the standard four-message handshake. You can enable rapid commit support on DHCPv6 local server globally, for a named group, or for a specific interface.
Default	Rapid commit support is not enabled.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling DHCPv6 Rapid Commit Support on page 935• Overriding Default DHCP Local Server Configuration Settings on page 928

reconfigure (DHCP Local Server)

Syntax	<pre> reconfigure { attempts <i>attempt-count</i>; clear-on-abort; strict; timeout <i>timeout-value</i>; token <i>token-value</i>; trigger { radius-disconnect; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i>],</p> <p>[edit system services dhcp-local-server],</p> <p>[edit system services dhcp-local-server dhcpv6],</p> <p>[edit system services dhcp-local-server group <i>group-name</i>],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.0.</p> <p>Support at the [edit ... dhcpv6 ...] hierarchy levels introduced in Junos OS Release 10.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Enable dynamic reconfiguration triggered by the DHCP local server of all DHCP clients or only the DHCP clients serviced by the specified group of interfaces. A group configuration takes precedence over a DHCP local server configuration. The strict statement is available only for DHCPv6.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 936](#)
 - [Configuring a DHCP Server on Switches \(CLI Procedure\)](#)

relay-agent-interface-id (DHCP Local Server)

Syntax	relay-agent-interface-id;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication username-include]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Specify that the DHCPv6 Relay Agent Interface-ID option (option 18) in the client PDU name is concatenated with the username during the subscriber authentication or DHCP client authentication process.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Creating Unique Usernames for DHCP Clients on page 946

relay-agent-remote-id (DHCP Local Server)

Syntax	relay-agent-remote-id [enterprise-id remote-id];
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> authentication username-include]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>For MX Series routers only, enterprise-id and remote-id options at the [edit system-services dhcp-local-server dhcpv6 authentication username-include] hierarchy level introduced in Junos OS Release 12.3R3.</p>
Description	<p>Specify that the DHCPv6 Relay Agent Remote-ID option (option 37) in the client PDU name is concatenated with the username during the subscriber authentication or DHCP client authentication process.</p> <p>If you enter relay-agent-remote-id without options, all of the DHCPv6 Relay Agent Remote-ID option (option 37) contents in the client PDU name are concatenated with the username.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Creating Unique Usernames for DHCP Clients on page 946

routing-instance-name (DHCP Local Server)

Syntax	routing-instance-name;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit system services dhcp-local-server authentication username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit system services dhcp-local-server group group-name authentication username-include]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Specify that the routing instance name be concatenated with the username during the subscriber authentication or DHCP client authentication process. No routing instance name is concatenated if the configuration is in the default routing instance.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Using External AAA Authentication Services with DHCP on page 923](#)

service-profile (DHCP Local Server)

Syntax	<code>service-profile <i>dynamic-profile-name</i>;</code>
Hierarchy Level	<p>[edit system services dhcp-local-server], [edit system services dhcp-local-server dhcpv6], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i>], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> interface <i>interface-name</i>], [edit system services dhcp-local-server group <i>group-name</i>], [edit system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...]</p>
Release Information	<p>Statement introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Specify the default subscriber service or DHCP client management service, which is activated when the subscriber or client logs in and no other service is activated by a RADIUS server or a provisioning server.</p> <ul style="list-style-type: none"> • To specify the default service for all DHCP local server clients, include the service-profile statement at the [edit system services dhcp-local-server] hierarchy level. • To specify the default service for a named group of interfaces, include the service-profile statement at the [edit system services dhcp-local-server group <i>group-name</i>] hierarchy level. • To specify the default service for a particular interface within a named group of interfaces, include the service-profile statement at the [edit system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i>] hierarchy level. • For DHCPv6 clients, use the service-profile statement at the [edit system services dhcp-local-server dhcpv6] hierarchy level.
Options	<i>dynamic-profile-name</i> —Name of the dynamic profile that defines the service.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Extended DHCP Local Server Overview on page 876 • Default Subscriber Service Overview • Configuring a Default Subscriber Service


session-mode

Syntax	session-mode (automatic multihop singlehop);
Hierarchy Level	[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the session mode.
Options	automatic —Configure single-hop BFD sessions if the peer is directly connected to the router interface and multihop BFD sessions if the peer is not directly connected to the router interface. multihop —Configure multihop BFD sessions. single-hop —Configure single hop BFD sessions.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913• <i>Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</i>


strict (DHCP Local Server)

Syntax	strict;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure],</p> <p>[edit system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Specify whether the server denies a client to bind when the client does not indicate that it accepts reconfigure messages. This feature is available only for DHCPv6.
Default	Accept solicit messages from clients that do not support reconfiguration and permit them to bind.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 936 • Preventing Binding of Clients That Do Not Support Reconfigure Messages on page 939

threshold (detection-time)

Syntax	<code>threshold <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit system services dhcp-local-server liveness-detection method bfd detection-time], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd detection-time], [edit forwarding-options dhcp-relay liveness-detection method bfd detection-time], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd detection-time], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd detection-time], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd detection-time], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd detection-time], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd detection-time]</p>
Release Information	<p>Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Specify the threshold for the adaptation of the detection time. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.
<div style="display: flex; align-items: center;">  <div> <p>NOTE: The threshold time must be greater than or equal to the <code>minimum-interval</code> or the <code>minimum-receive-interval</code>.</p> </div> </div>	
Options	<p><i>milliseconds</i>— Value for the detection time adaptation threshold. Range: 1 through 255,000</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913 • Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

threshold (transmit-interval)

Syntax	<code>threshold <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit system services dhcp-local-server liveness-detection method bfd transmit-interval],</p> <p>[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd transmit-interval],</p> <p>[edit forwarding-options dhcp-relay liveness-detection method bfd transmit-interval],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd transmit-interval],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd transmit-interval],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd transmit-interval],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd transmit-interval],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd transmit-interval]</p>
Release Information	Statement introduced in Junos OS Release 12.1.
Description	Specify the threshold for detecting the adaptation of the transmit interval. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.
Options	<p><i>milliseconds</i> — Threshold value.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p>
<div>  <p>NOTE: The threshold value specified in the <code>threshold</code> statement must be greater than the value specified in the <code>minimum-interval</code> statement for the <code>transmit-interval</code> statement.</p> </div>	
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913 • Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

timeout (DHCP Local Server)

Syntax	<code>timeout <i>timeout-value</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure],</p> <p>[edit system services dhcp-local-server reconfigure],</p> <p>[edit system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.0.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit ... dhcpv6 ...] hierarchy levels introduced in Junos OS Release 10.4.</p>
Description	Configure the initial value in seconds between attempts to reconfigure all DHCP clients or only the DHCP clients serviced by the specified group of interfaces.
Options	<p><i>timeout-value</i>—Initial retry timeout value.</p> <p>Range: 1 through 10 seconds</p> <p>Default: 2 seconds</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 936 • Configuring Dynamic Reconfiguration Attempts for DHCP Clients on page 937

token (DHCP Local Server)

Syntax	<code>token <i>token-value</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure],</p> <p>[edit system services dhcp-local-server reconfigure],</p> <p>[edit system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.0.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit ... dhcpv6 ...] hierarchy levels introduced in Junos OS Release 10.4.</p>
Description	<p>Configure a plain-text token for all DHCP clients or only the DHCP clients serviced by the specified group of interfaces. The token enables rudimentary entity authentication to protect against inadvertently instantiated DHCP servers. A null token (empty string) indicates that the configuration token functionality is not enabled. A group configuration takes precedence over a DHCP local server configuration. For more information about tokens, see RFC 3118, <i>Authentication for DHCP Messages</i>, section 4.</p>
Options	<p><i>token-value</i>—Plain-text alphanumeric string.</p> <p>Default: null (empty string)</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 936 • Configuring a Token for DHCP Local Server Authentication on page 939

transmit-interval

Syntax	<pre>transmit-interval { threshold milliseconds; minimum-interval milliseconds; }</pre>
Hierarchy Level	[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the Bidirectional Forwarding Detection (BFD) transmit interval. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913• Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

trigger (DHCP Local Server)

Syntax	<pre>trigger { radius-disconnect; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure],</p> <p>[edit system services dhcp-local-server reconfigure],</p> <p>[edit system services dhcp-local-server dhcpv6 reconfigure],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> reconfigure],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> reconfigure]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.0.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit ... dhcpv6 ...] hierarchy levels introduced in Junos OS Release 10.4.</p>
Description	<p>Configure behavior in response to a trigger for all DHCP clients or only the DHCP clients serviced by the specified group of interfaces.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 936 • Configuring Reconfiguration of the Client on Receipt of RADIUS-Initiated Disconnect on page 938 • radius-disconnect on page 1051

use-primary (DHCP Local Server)

Syntax	<code>use-primary <i>primary-profile-name</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server <i>dynamic-profile profile-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <i>dynamic-profile profile-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server <i>dynamic-profile profile-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> <i>dynamic-profile profile-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server <i>dynamic-profile profile-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> <i>dynamic-profile profile-name</i>],</code> <code>[edit system services dhcp-local-server <i>dynamic-profile profile-name</i>]</code> <code>[edit system services dhcp-local-server group <i>group-name</i> <i>dynamic-profile profile-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify the dynamic profile to configure as the primary dynamic profile. The primary dynamic profile is instantiated when the first subscriber or DHCP client logs in. Subsequent subscribers (or clients) are not assigned the primary dynamic profile; instead, they are assigned the dynamic profile specified for the interface. When the first subscriber (or client) logs out, the next subscriber (or client) that logs in is assigned the primary dynamic profile.
Options	<i>primary-profile-name</i> —Name of the dynamic profile to configure as the primary dynamic profile
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 942

user-prefix (DHCP Local Server)

Syntax	<code>user-prefix <i>user-prefix-string</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group group-name authentication username-include],</p> <p>[edit system services dhcp-local-server authentication username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 authentication username-include],</p> <p>[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include],</p> <p>[edit system services dhcp-local-server group group-name authentication username-include]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Specify the user prefix that is concatenated with the username during the subscriber authentication or DHCP client authentication process.
Options	<i>user-prefix-string</i> —User prefix string.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Using External AAA Authentication Services with DHCP on page 923](#)

username-include (DHCP Local Server)

Syntax	<pre>username-include { circuit-type; client-id; delimiter <i>delimiter-character</i>; domain-name <i>domain-name-string</i>; interface-name; logical-system-name; mac-address; option-60; option-82 <circuit-id> <remote-id>; relay-agent-interface-id; relay-agent-remote-id; relay-agent-subscriber-id; routing-instance-name; user-prefix <i>user-prefix-string</i>; }</pre>
Hierarchy Level	<pre>[edit system services dhcp-local-server authentication], [edit system services dhcp-local-server dhcpv6 authentication], [edit system services dhcp-local-server dhcpv6 group group-name authentication], [edit system services dhcp-local-server group group-name authentication], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...], [edit logical-systems <i>logical-system-name</i> system services dhcp-local-server ...], [edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server ...]</pre>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure the username that the router or switch passes to the external AAA server. You must include at least one of the optional statements for the username to be valid. If you do not configure a username, the router (or switch) accesses the local authentication service only and does not use external authentication services, such as RADIUS.</p> <p>The statements are explained separately. The option-60 and option-82 statements are not supported in the DHCPv6 hierarchy levels. The client-id, relay-agent-interface-id, relay-agent-remote-id and relay-agent-subscriber-id statements are supported in the DHCPv6 hierarchy levels only.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Using External AAA Authentication Services with DHCP on page 923 • Creating Unique Usernames for DHCP Clients on page 946

version (BFD)

Syntax	version (0 1 automatic);
Hierarchy Level	[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the BFD protocol version to detect.
Options	0—Use BFD protocol version 0. 1—Use BFD protocol version 1. automatic —Autodetect the BFD protocol version. Default: automatic
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913 • Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

DHCP Relay Agent Configuration Statements

- [\[edit forwarding-options dhcp-relay\] Hierarchy Level on page 1069](#)

[\[edit forwarding-options dhcp-relay\] Hierarchy Level](#)

```
forwarding-options {
  dhcp-relay {
    active-server-group server-group-name;
  }
  arp-inspection;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
    }
  }
}
```

```
        interface-name;
        logical-system-name;
        mac-address;
        option-60;
        option-82 <circuit-id> <remote-id>;
        routing-instance-name;
        user-prefix user-prefix-string;
    }
}
dhcpv6 {
    active-server-group group-name;
}
authentication {
    password password-string;
}
username-include {
    circuit-type;
    client-id;
    delimiter delimiter-character;
    domain-name domain-name-string;
    interface-name;
    logical-system-name;
    relay-agent-interface-id;
    relay-agent-remote-id;
    relay-agent-subscriber-id;
    routing-instance-name;
    user-prefix user-prefix-string;
}
dynamic-profile profile-name {
    aggregate-clients (merge |replace);
    use-primary profile-name;
}
group group-name {
    ... the group subhierarchy appears after the main [edit forwarding-options
        dhcp-relay] hierarchy ...
}
liveness-detecton {
    ... the liveness-detection subhierarchy appears after the main [edit
        forwarding-options dhcp-relay] hierarchy ...
}
overrides {
    ... the overrides subhierarchy appears after the main [edit forwarding-options
        dhcp-relay] hierarchy ...
}
relay-agent-interface-id {
    prefix;
    user-interface-description;
}
relay-option {
    default-action;
    equals;
    option-number;
    starts-with;
}
server-group;
service-profile;
)
```

```

duplicate-clients-on-interface;
dynamic-profile profile-name {
    aggregate-clients (merge |replace);
    use-primary profile-name;
}
forward-snooped-clients (all-interfaces | configured-interfaces |
non-configured-interfaces);
group group-name {
    active-server-group server-group-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            logical-system-name;
            mac-address;
            option-60;
            option-82 <circuit-id> <remote-id>;
            routing-instance-name;
            user-prefix user-prefix-string;
        }
        dynamic-profile profile-name{
            aggregate-clients (merge |replace);
            use-primary profile-name;
        }
        interface;
        liveness-detection;
        overrides;
        relay-option;
        relay-option-82 ;
        service-profile;
    }
}
liveness-detection {
    failure-action (clear-binding |clear-binding-if-interface-up |log-only);
}
method {
    bfd {
        detection-time {
            threshold milliseconds;
        }
        holddown-interval;
        minimum--interval;
        minimum-receive-interval;
        multiplier;
        no-adaptation;
        session-mode;
    }
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version;
}
overrides {

```

```
(allow-snooped-clients | no-allow-snooped-clients);
always-write-giaddr;
always-write-option-82;
client-discover-match <option60-and-option82>;
disable-relay;
interface-client-limit number;
layer2-unicast-replies;
no-allow-snooped-clients;
no-arp;
no-bind-on-request;
no-unicast-replies;
proxy-mode;
replace-ip-source-with giaddr;
send-release-on-delete;
trust-option-82;
}
relay-option {
  default-action {
    drop;
    forward-only;
    local-server-group group-name;
    relay-server-group group-name;
  }
  equals {
    ascii string;
    hexadecimal string;
  }
}
option-number (60 | 77);
}
starts-with {
  ascii string;
  hexadecimal string;
}
relay-option-82 {
  circuit-id (value | ... the following prefix statement ...) {
    prefix {
      host-name;
      logical-system-name;
      routing-instance-name;
    }
    use-interface-description (device | logical);
  }
}
server-group {
  server-group-name {
    ip-address;
  }
}
service-profile name:
}
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - [\[edit forwarding-options\] Hierarchy Level on page 333](#)

access (Dynamic Access Routes)

Syntax

```
access {
  route prefix {
    next-hop next-hop;
    metric route-cost;
    preference route-distance;
    tag route-tag;
  }
}
```

Hierarchy Level [\[edit dynamic-profiles routing-options\]](#)

Release Information Statement introduced in Junos OS Release 9.5.

Description Dynamically configure access routes.



BEST PRACTICE: We recommend that you always include the **access-internal** stanza in the dynamic-profile when the access stanza is present for framed-route support.

Options The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

- Related Documentation**
- *Configuring Dynamic Access Routes for Subscriber Management*

access-internal (Dynamic Access-Internal Routes)

Syntax `access-internal {
 route subscriber-ip-address {
 qualified-next-hop underlying-interface {
 mac-address address;
 }
 }
 }`

Hierarchy Level [edit dynamic-profiles routing-options]

Release Information Statement introduced in Junos OS Release 9.5.

Description Dynamically configure access-internal routes. Access-internal routes are optional, but are used instead of access routes if the next-hop address is not specified in the Framed-Route Attribute [22] for IPv4 or the Framed-IPv6-Route attribute [99] for IPv6.



.....
BEST PRACTICE: We recommend that you always include the `access-internal` stanza in the dynamic-profile when the `access` stanza is present for framed-route support.
.....

The remaining statements are explained separately.

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


Related Documentation

- *Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management*
- *Configuring Dynamic Access-Internal Routes for PPP Subscriber Management*

active-server-group

Syntax	<code>active-server-group <i>server-group-name</i>;</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit forwarding-options dhcp-relay group <i>group-name</i>], [edit forwarding-options dhcp-relay group <i>group-name</i> dhcpv6], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6], [edit logical-systems <i>logical-system-name</i> forwarding-options group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options group <i>group-name</i> dhcpv6], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> dhcpv6], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay] [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Apply a DHCP relay agent configuration to the named group of DHCP server addresses. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.</p> <p>A group-specific configuration overrides a global option.</p>
Options	server-group-name —Name of the group of DHCP or DHCPv6 server addresses to which the DHCP or DHCPv6 relay agent configuration applies.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Extended DHCP Relay Agent Overview on page 900 • Configuring Active Server Groups on page 977 • Group-Specific DHCP Relay Options on page 906 • dhcp-relay on page 1087

allow-snooped-clients

Syntax	allow-snooped-clients;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> interface <i>interface-name</i> overrides],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> overrides],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 overrides],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> overrides],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit forwarding-options dhcp-relay overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.2.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 12.1.</p>
Description	<p>Explicitly enable DHCP snooping support on the router.</p> <p>Use the statement at the [edit ... dhcpv6] hierarchy levels to explicitly enable snooping support on the router for DHCPv6 relay agent.</p>
<div>  <p>NOTE: In Junos OS Release 10.0 and earlier, DHCP snooping is <i>enabled</i> by default. In Release 10.1 and later, DHCP snooping is <i>disabled</i> by default.</p> </div>	
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Extended DHCP Relay Agent Overview on page 900 Overriding the Default DHCP Relay Configuration Settings on page 955 DHCP Snooping Support on page 887

always-write-giaddr

Syntax	<code>always-write-giaddr;</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay overrides],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> overrides]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	Overwrite the gateway IP address (giaddr) of every DHCP packet with the giaddr of the DHCP relay agent before forwarding the packet to the DHCP server.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Extended DHCP Relay Agent Overview on page 900 • dhcp-relay on page 1087

always-write-option-82

Syntax	<code>always-write-option-82;</code>
Hierarchy Level	<code>[edit forwarding-options dhcp-relay overrides],</code> <code>[edit forwarding-options dhcp-relay group <i>group-name</i> overrides],</code> <code>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay overrides],</code> <code>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides],</code> <code>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides],</code> <code>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides],</code> <code>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> overrides]</code>
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>Override the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server. The use of this option causes the DHCP relay agent to perform one of the following actions, depending on how it is configured:</p> <ul style="list-style-type: none">• If the DHCP relay agent is configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the DHCP packets and inserts the new values before forwarding the packets to the DHCP server.• If the DHCP relay agent is not configured to add option 82 information to DHCP packets, it clears the existing option 82 values from the packets, but does not add any new values before forwarding the packets to the DHCP server.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Extended DHCP Relay Agent Overview on page 900

authentication (DHCP Relay Agent)

Syntax	<pre> authentication { password <i>password-string</i>; username-include { circuit-type; client-id; delimiter <i>delimiter-character</i>; domain-name <i>domain-name-string</i>; interface-name; logical-system-name; mac-address; option-60; option-82 [circuit-id] [remote-id]; relay-agent-interface-id; relay-agent-remote-id; relay-agent-subscriber-id; routing-instance-name; user-prefix <i>user-prefix-string</i>; } }</pre>
Hierarchy Level	<pre> [edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i>], [edit forwarding-options dhcp-relay group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</pre>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p>
Description	<p>Configure the parameters the router sends to the external AAA server. A group configuration takes precedence over a global DHCP relay configuration. Use the statement at the [edit...dhcpv6] hierarchy levels to configure DHCPv6 support.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • dhcp-relay on page 1087 • Using External AAA Authentication Services with DHCP on page 923

bfd

Syntax	<pre>bfd { version (0 1 automatic); minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; multiplier <i>number</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } detection-time { threshold <i>milliseconds</i>; } session-mode (automatic multihop singlehop); holddown-interval <i>milliseconds</i>; }</pre>
Hierarchy Level	<pre>[edit system services dhcp-local-server liveness-detection <i>method</i>], [edit system services dhcp-local-server dhcpv6 liveness-detection <i>method</i>], [edit forwarding-options dhcp-relay liveness-detection <i>method</i>], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection <i>method</i>], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection <i>method</i>], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection <i>method</i>], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection <i>method</i>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection <i>method</i>]</pre>
Release Information	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure Bidirectional Forwarding Detection (BFD) as the liveness detection method.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913• Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

circuit-id (DHCP Relay Agent)

Syntax	<pre>circuit-id { prefix <i>prefix</i>; use-interface-description (logical device); }</pre>
Hierarchy Level	<pre>[edit forwarding-options dhcp-relay relay-option-82], [edit forwarding-options dhcp-relay group <i>group-name</i> relay-option-82], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay relay-option-82], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-82], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay relay-option-82], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-82], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay relay-option-82], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-82]</pre>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 12.3 for EX Series switches.</p>
Description	<p>Specify the Agent Circuit ID suboption (suboption 1) of the DHCP relay agent information option (option 82) to include in DHCP packets destined for a DHCP server. Optionally specify that the suboption includes a prefix or textual description, or both, instead of circuit-id.</p> <p>The format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces that do not use virtual LANs (VLANs) or stacked VLANs (S-VLANs) is as follows:</p> <pre>(fe ge)-fpc/pic/port</pre> <p>The format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces that use VLANs is as follows:</p> <pre>(fe ge)-fpc/pic/port:vlan-id</pre> <p>The format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces that use S-VLANs is as follows:</p> <pre>(fe ge)-fpc/pic/port:svlan-id-vlan-id</pre> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Enabling and Disabling Insertion of Option 82 Information on page 973 • Configuring Agent Circuit ID Information on page 974

circuit-type (DHCP Relay Agent)

Syntax	circuit-type;
Hierarchy Level	[edit forwarding-options dhcp-relay authentication username-include], [edit forwarding-options dhcp-relay dhcpv6 authentication username-include], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include], [edit forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]
Release Information	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches. Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.
Description	Specify that the circuit type is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Using External AAA Authentication Services with DHCP on page 923• Creating Unique Usernames for DHCP Clients on page 946

client-discover-match (DHCP Relay Agent)

Syntax	client-discover-match <option60-and-option82>;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay overrides],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> overrides]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Configure DHCP relay to use option 60 and option 82 information to uniquely identify DHCP subscribers or clients when primary subscriber or client identification fails. The statement always uses the option60-and-option82 option. Specifying the option has no effect.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Extended DHCP Relay Agent Overview on page 900 • Overriding the Default DHCP Relay Configuration Settings on page 955

client-id (DHCP Relay Agent)

Syntax	client-id;
Hierarchy Level	[edit forwarding-options dhcp-relay dhcpv6 authentication username-include], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 ...]
Release Information	Statement introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify that the client ID is concatenated with the username during the subscriber authentication or client authentication process.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Using External AAA Authentication Services with DHCP on page 923• Creating Unique Usernames for DHCP Clients on page 946

delimiter (DHCP Relay Agent)

Syntax	<code>delimiter <i>delimiter-character</i>;</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay authentication username-include],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p>
Description	Specify the character used as the delimiter between the concatenated components of the username. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.
Options	<i>delimiter-character</i> —Character that separates components that make up the concatenated username. You cannot use the semicolon (;) as a delimiter.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Using External AAA Authentication Services with DHCP on page 923 • Creating Unique Usernames for DHCP Clients on page 946

detection-time

Syntax	<pre>detection-time { threshold milliseconds; }</pre>
Hierarchy Level	[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>Enable failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the one configured.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913• Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

dhcp-relay

```
Syntax  dhcp-relay {
    active-server-group server-group-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-name;
            logical-system-name;
            mac-address;
            option-60;
            option-82 <circuit-id> <remote-id>;
            routing-instance-name;
            user-prefix user-prefix-string;
        }
    }
    dhcpv6 {
        active-server-group server-group-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                logical-system-name;
                relay-agent-interface-id;
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
    }
    group group-name {
        active-server-group server-group-name;
        authentication {
            ...
        }
        dynamic-profile profile-name {
            ...
        }
    }
    interface interface-name {
        exclude;
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {

```

```
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
      session-mode (automatic | multihop | singlehop);
      holddown-interval milliseconds;
    }
  }
}
overrides {
  ...
}
relay-option {
  ...
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
service-profile dynamic-profile-name;
}
overrides {
  ...
}
relay-agent-interface-id {
  ...
}
relay-option {
  ...
}
}
service-profile dynamic-profile-name;
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
    }
  }
}
```

```

        session-mode(automatic | multihop | singlehop);
        holddown-interval milliseconds;
    }
}
overrides {
    allow-snooped-clients;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    send-release-on-delete;
}
relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device);
}
server-group {
    server-group-name {
        server-ip-address;
    }
}
duplicate-clients-on-interface;
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}
forward-snooped-clients (all-interfaces | configured-interfaces |
    non-configured-interfaces);
group group-name {
    active-server-group server-group-name;
    authentication {
        ...
    }
    dynamic-profile profile-name {
        ...
    }
}
interface interface-name {
    exclude;
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
            }
            detection-time {
                threshold milliseconds;
            }
        }
        session-mode(automatic | multihop | singlehop);
        holddown-interval milliseconds;
    }
}

```

```
    }
  }
}
overrides {
  ...
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
overrides {
  ...
}
relay-option {
  ...
}
relay-option-82 {
  ...
}
service-profile dynamic-profile-name;
}
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
      session-mode (automatic | multihop | singlehop);
      holddown-interval milliseconds;
    }
  }
}
}
overrides {
  allow-snooped-clients;
  always-write-giaddr;
  always-write-option-82;
  client-discover-match <option60-and-option82>;
  disable-relay;
  interface-client-limit number;
  layer2-unicast-replies;
  no-allow-snooped-clients;
  no-arp;
  no-bind-on-request;
  proxy-mode;
  replace-ip-source-with;
  send-release-on-delete;
```



```

    trust-option-82;
}
relay-option {
    option-number option-number;
    default-action {
        drop;
        forward-only;
        relay-server-group group-name;
    }
    equals (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        relay-server-group relay-server-group;
    }
    starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
        drop;
        forward-only;
        local-server-group local-server-group;
        relay-server-group relay-server-group;
    }
}
relay-option-82 {
    circuit-id {
        prefix prefix;
        use-interface-description (logical | device);
    }
}
server-group {
    server-group-name {
        server-ip-address;
    }
}
service-profile dynamic-profile-name;
}

```

Hierarchy Level [edit forwarding-options],
 [edit logical-systems *logical-system-name* forwarding-options],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
 forwarding-options],
 [edit routing-instances *routing-instance-name* forwarding-options]

Release Information Statement introduced in Junos OS Release 8.3.
 Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the router or switch and enable the router (or switch) to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.

DHCP relay supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication or client authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.

The extended DHCP and DHCPv6 relay agent options configured with the **dhcp-relay** and **dhcpv6** statements are incompatible with the DHCP/BOOTP relay agent options configured with the **bootp** statement. As a result, the extended DHCP or DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot both be enabled on the router (or switch) at the same time.

The remaining statements are explained separately.

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

Related Documentation

- [Extended DHCP Relay Agent Overview on page 900](#)
- [DHCPv6 Relay Agent Overview on page 905](#)
- [DHCP Relay Proxy Overview on page 903](#)
- [Using External AAA Authentication Services with DHCP on page 923](#)

dhcpx6 (DHCP Relay Agent)

```
Syntax  dhcpx6 {
    active-server-group server-group-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name-string;
            logical-system-name;
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix-string;
        }
    }
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    group group-name {
        active-server-group server-group-name;
        authentication {
            ...
        }
        dynamic-profile profile-name {
            ...
        }
    }
    interface interface-name {
        exclude;
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                    detection-time {
                        threshold milliseconds;
                    }
                }
                session-mode (automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
        }
    }
}
```

```
    overrides {
      ...
    }
    service-profile dynamic-profile-name;
    trace;
    upto upto-interface-name;
  }
}
overrides {
  ...
}
relay-agent-interface-id {
  ...
}
relay-option {
  ...
}
service-profile dynamic-profile-name;
}
liveness-detection {
  ...
}
overrides {
  allow-snooped-clients;
  interface-client-limit number;
  no-allow-snooped-clients;
  no-bind-on-request;
  send-release-on-delete;
}
relay-agent-interface-id {
  prefix prefix;
  use-interface-description (logical | device);
}
relay-option {
  option-number option-number;
  default-action {
    drop;
    forward-only;
    relay-server-group relay-server-group;
  }
  equals (ascii ascii-string | hexadecimal hexadecimal-string) {
    drop;
    forward-only;
    relay-server-group relay-server-group;
  }
  starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
    drop;
    forward-only;
    relay-server-group relay-server-group;
  }
}
}
server-group {
  server-group-name {
    server-ip-address;
  }
}
```

```

    service-profile dynamic-profile-name;
}

```

Hierarchy Level	[edit forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay]
Release Information	Statement introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3 for EX Series switches.
Description	<p>Configure DHCPv6 relay options on the router or switch and enable the router or switch to function as a DHCPv6 relay agent. A DHCPv6 relay agent forwards DHCPv6 request and reply packets between a DHCPv6 client and a DHCPv6 server.</p> <p>The DHCPv6 relay agent server is fully compatible with the extended DHCP local server and DHCP relay agent. However, the options configured with the dhcpv6 statement are incompatible with the DHCP/BOOTP relay agent options configured with the bootp statement. As a result, the DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot be enabled on the router or switch at the same time.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • dhcp-relay on page 1087 • DHCPv6 Relay Agent Overview on page 905 • Using External AAA Authentication Services with DHCP on page 923

disable-relay

Syntax	disable-relay;
Hierarchy Level	[edit forwarding-options dhcp-relay overrides], [edit forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> overrides]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Disable DHCP relay on specific interfaces in a group.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Extended DHCP Relay Agent Overview on page 900

domain-name (DHCP Relay Agent)

Syntax	<code>domain-name <i>domain-name-string</i>;</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay authentication username-include],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p>
Description	Specify the domain name that is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.
Options	<i>domain-name-string</i> —Domain name formatted string.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Using External AAA Authentication Services with DHCP on page 923 • Creating Unique Usernames for DHCP Clients on page 946

drop (DHCP Relay Agent Option)

Syntax	drop;
Hierarchy Level	[edit forwarding-options dhcp-relay relay-option (default-action equals starts-with)], [edit forwarding-options dhcp-relay dhcpv6 relay-option (default-action equals starts-with)], [edit forwarding-options dhcp-relay group <i>group-name</i> relay-option (default-action equals starts-with)], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> relay-option (default-action equals starts-with)], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]
Release Information	Statement introduced in Junos OS Release 12.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Drop (discard) specified DHCP client packets when you use DHCP relay agent selective processing. You can configure the drop operation globally or for a group of interfaces, and for either DHCP or DHCPv6 relay agent.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Using DHCP Option Information to Selectively Process DHCP Client Traffic</i>

duplicate-clients-on-interface (DHCP Relay Agent)

Syntax	duplicate-clients-on-interface;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Configure DHCP relay agent to include the client subinterface when distinguishing between duplicate DHCP clients (clients with the same MAC address or client ID) in the same subnet. By default, DHCP relay distinguishes clients by subnet. This feature is supported on DHCPv4 only.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring DHCP Duplicate Client Support on page 925 • Enabling and Disabling Insertion of Option 82 Information on page 973

dynamic-profile (DHCP Relay Agent)

Syntax	<pre>dynamic-profile <i>profile-name</i> { aggregate-clients (merge replace); use-primary <i>primary-profile-name</i>; }</pre>
Hierarchy Level	<pre>[edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> interface <i>interface-name</i>], [edit forwarding-options dhcp-relay group <i>group-name</i>], [edit forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</pre>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Specify the dynamic profile that is attached to all interfaces, to a named group of interfaces, or to a specific interface.</p> <p>M120 and M320 routers do not support DHCPv6.</p>
Options	<p><i>profile-name</i>—Name of the dynamic profile.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• dhcp-relay on page 1087• Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 942• Grouping Interfaces with Common DHCP Configurations on page 926• Configuring a Default Subscriber Service

failure-action

Syntax	failure-action (clear-binding clear-binding-if-interface-up log-only);
Hierarchy Level	[edit system services dhcp-local-server liveness-detection], [edit system services dhcp-local-server dhcpv6 liveness-detection], [edit forwarding-options dhcp-relay liveness-detection], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the action the router (or switch) takes when a liveness detection failure occurs.
Options	<p>clear-binding—The client session is cleared when a liveness detection failure occurs.</p> <p>clear-binding-if-interface-up—The client session is cleared only when a liveness detection failure occurs and the local interface is detected as being up.</p> <p>log-only—A message is logged to indicate the event; no action is taken and DHCP is left to manage the failure.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • DHCP Liveness Detection Overview on page 980 • Configuring Detection of DHCP Local Server Client Connectivity on page 941 • Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 981 • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913 • Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

forward-snooped-clients (DHCP Relay Agent)

Syntax	forward-snooped-clients (all-interfaces configured-interfaces non-configured-interfaces);
Hierarchy Level	[edit forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay]
Release Information	Statement introduced in Junos OS Release 10.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>Configure how DHCP relay agent handles DHCP snooped packets on specific interfaces. The router or switch determines the DHCP snooping action to perform based on a combination of the forward-snooped-clients configuration and the configuration of either the allow-snooped-clients statement or the no-allow-snooped-clients statement.</p> <p>The router (or switch) also uses this statement to determine how to handle snooped BOOTREPLY packets received on nonconfigured interfaces.</p>
Options	<p>all-interfaces—Perform the action on all interfaces.</p> <p>configured-interfaces—Perform the action only on configured interfaces.</p> <p>non-configured-interfaces—Perform the action only on nonconfigured interfaces.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• DHCP Snooping Support on page 887• Configuring DHCP Snooping for DHCP Relay Agent on page 963

group (DHCP Relay Agent)

```
Syntax  group group-name {
    active-server-group server-group-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            client-id;
            delimiter delimiter-character;
            domain-name domain-name-string;
            logical-system-name;
            mac-address;
            option-60;
            option-82 [circuit-id] [remote-id];
            relay-agent-interface-id;
            relay-agent-remote-id;
            relay-agent-subscriber-id;
            routing-instance-name;
            user-prefix user-prefix-string;
        }
    }
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    interface interface-name {
        exclude;
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                    detection-time {
                        threshold milliseconds;
                    }
                }
                session-mode (automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
        }
    }
    overrides {
        ...
    }
    service-profile dynamic-profile-name;
    trace;
```

```
    upto upto-interface-name;
  }
  overrides {
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    client-discover-match <option60-and-option82>;
    disable-relay;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-arp;
    no-bind-on-request;
    proxy-mode;
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
  }
  relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device):
  }
  relay-option {
    option-number option-number;
    default-action {
      drop;
      forward-only;
      local-server-group local-server-group;
      relay-server-group relay-server-group;
    }
    equals (ascii ascii-string | hexadecimal hexadecimal-string) {
      drop;
      forward-only;
      local-server-group local-server-group;
      relay-server-group relay-server-group;
    }
    starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
      drop;
      forward-only;
      local-server-group local-server-group;
      relay-server-group relay-server-group;
    }
  }
  relay-option-82 {
    circuit-id {
      prefix prefix;
      use-interface-description (logical | device);
    }
  }
  service-profile dynamic-profile-name;
}
```

Hierarchy Level	<p>[edit forwarding-options dhcp-relay],</p> <p>[edit forwarding-options dhcp-relay dhcpv6],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Specify the name of a group of interfaces that have a common DHCP or DHCPv6 relay agent configuration. A group must contain at least one interface. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.</p>
Options	<p>group-name—Name of a group of interfaces that have a common DHCP or DHCPv6 relay agent configuration.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • dhcp-relay on page 1087 • Extended DHCP Relay Agent Overview on page 900 • Group-Specific DHCP Relay Options on page 906 • Grouping Interfaces with Common DHCP Configurations on page 926 • Using External AAA Authentication Services with DHCP on page 923 • Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 942

holddown-interval

Syntax	<code>holddown-interval <i>milliseconds</i>;</code>
Hierarchy Level	<code>[edit system services dhcp-local-server liveness-detection method bfd],</code> <code>[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],</code> <code>[edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options</code> <code> dhcp-relay dhcpv6 liveness-detection method bfd],</code> <code>[edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd],</code> <code>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method</code> <code> bfd],</code> <code>[edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd],</code> <code>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method</code> <code> bfd]</code>
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the time (in milliseconds) for which Bidirectional Forwarding Detection (BFD) holds a session up notification.
Options	<i>milliseconds</i> —Interval specifying how long a BFD session must remain up before a state change notification is sent. Range: 0 through 255,000 Default: 0
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913• Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

interface (DHCP Relay Agent)

Syntax

```
interface interface-name {
  exclude;
  overrides {
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    client-discover-match <option60-and-option82>;
    disable-relay;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-arp;
    proxy-mode;
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
  }
  service-profile dynamic-profile-name;
  trace;
  upto upto-interface-name;
}
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay dhcpv6 group group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit logical-systems logical-system-name forwarding-options dhcp-relay ...],
[edit logical-systems logical-system-name routing-instances routing-instance-name
 forwarding-options dhcp-relay ...],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay ...]
```

Release Information

Statement introduced in Junos OS Release 8.3.

Options **upto** and **exclude** introduced in Junos OS Release 9.1.

Support at the **[edit ... dhcpv6]** hierarchy levels introduced in Junos OS Release 11.4.

Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description

Specify one or more interfaces, or a range of interfaces, that are within a specified group on which the DHCP or DHCPv6 relay agent is enabled. You can repeat the **interface *interface-name*** statement to specify multiple interfaces within a group, but you cannot specify the same interface in more than one group. Also, you cannot use an interface that is being used by the DHCP local server. Use the statement at the **[edit ... dhcpv6]** hierarchy levels to configure DHCPv6 support.



NOTE: DHCP values are supported in Integrated Routing and Bridging (IRB) configurations. When you configure an IRB interface in a network that is using DHCP, the DHCP information (for example, authentication, address assignment, and so on) is propagated in the associated bridge domain. This enables the DHCP server to configure client IP addresses residing within the bridge domain. IRB currently only supports static DHCP configurations. For additional information about how to configure IRB, see the *JUNOS® MX Series 3D Universal Edge Routers Solutions, Release 12.3*.

Options **exclude**—Exclude an interface or a range of interfaces from the group. This option and the **overrides** option are mutually exclusive.

interface-name—Name of the interface. You can repeat this option multiple times.

overrides—Override the specified default configuration settings for the interface. The **overrides** statement is described separately.

upto-interface-name—Upper end of the range of interfaces; the lower end of the range is the interface-name entry. The interface device name of the **upto-interface-name** must be the same as the device name of the **interface-name**.

The remaining statements are explained separately.

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

Related Documentation

- [Extended DHCP Relay Agent Overview on page 900](#)
- [Grouping Interfaces with Common DHCP Configurations on page 926](#)
- [Using External AAA Authentication Services with DHCP on page 923](#)

interface-client-limit (DHCP Relay Agent)

Syntax	<code>interface-client-limit <i>number</i>;</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay dhcpv6 overrides], [edit forwarding-options dhcp-relay overrides], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> overrides], [edit forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> overrides]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.2. Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Set the maximum number of DHCP (or DHCPv6) subscribers or clients per interface allowed for a specific group or for all groups. A group specification takes precedence over a global specification for the members of that group. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.</p> <p>M120 and M320 routers do not support DHCPv6.</p>
Default	No limit
Options	<p><i>number</i>—Maximum number of clients allowed. Range: 1 through 500,000</p>
Required Privilege Level	<p>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</p>

- Related Documentation**
- [dhcp-relay on page 1087](#)
 - [Extended DHCP Relay Agent Overview on page 900](#)
 - [Group-Specific DHCP Relay Options on page 906](#)
 - [Overriding the Default DHCP Relay Configuration Settings on page 955](#)

interface-delete (Subscriber Management or DHCP Client Management)

Syntax	interface-delete;
Hierarchy Level	[edit system services subscriber-management maintain-subscriber]
Release Information	Statement introduced in Junos OS Release 11.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>On router—Configure the router to maintain, rather than log out, subscribers when the subscriber interface is deleted. By default, the router logs out subscribers when the subscriber interface is deleted.</p> <p>On switch—Configure the switch to maintain rather than log out DHCP clients when the client interface is deleted. By default, the switch logs out DHCP clients when the client interface is deleted.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events

interface-name (DHCP Relay Agent)

Syntax	interface-name;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay authentication username-include], [edit forwarding-options dhcp-relay dhcpv6 authentication username-include], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include], [edit forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</p>
Release Information	<p>Statement introduced in Junos OS Release 11.4 Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Specify that the interface name is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.
Required Privilege Level	<p>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Creating Unique Usernames for DHCP Clients on page 946

layer2-unicast-replies

Syntax	layer2-unicast-replies;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay overrides], [edit forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> overrides]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	Override the setting of the broadcast bit in DHCP request packets and instead use the Layer 2 unicast transmission method to transmit DHCP Offer reply packets and DHCP ACK reply packets from the DHCP server to DHCP clients during the discovery process.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Extended DHCP Relay Agent Overview on page 900• dhcp-relay on page 1087

liveness-detection

Syntax	<pre> liveness-detection { failure-action (clear-binding clear-binding-if-interface-up log-only); method { bfd { version (0 1 automatic); minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; multiplier <i>number</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } detection-time { threshold <i>milliseconds</i>; } session-mode (automatic multihop singlehop); holddown-interval <i>milliseconds</i>; } } } </pre>
Hierarchy Level	<pre> [edit system services dhcp-local-server], [edit system services dhcp-local-server dhcpv6], [edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit system services dhcp-local-server group <i>group-name</i>], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i>], [edit forwarding-options dhcp-relay group <i>group-name</i>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i>] </pre>
Release Information	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure bidirectional failure detection timers and authentication criteria for static routes.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • DHCP Liveness Detection Overview on page 980 • Configuring Detection of DHCP Local Server Client Connectivity on page 941 • Configuring Detection of DHCP Relay or DHCP Relay Proxy Client Connectivity on page 981 • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913

- *Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients*

local-server-group (DHCP Relay Agent Option)

Syntax	<code>local-server-group <i>local-server-group</i>;</code>
Hierarchy Level	<code>[edit forwarding-options dhcp-relay relay-option (default-action equals starts-with)],</code> <code>[edit forwarding-options dhcp-relay group <i>group-name</i> relay-option (default-action equals starts-with)],</code> <code>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i></code> <code>forwarding-options dhcp-relay ...],</code> <code>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</code>
Release Information	Statement introduced in Junos OS Release 12.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Forward DHCP client packets to the specified group of DHCP local servers when you use the DHCP relay selective processing feature. You can configure the forwarding operation globally or for a group of interfaces. The local-server-group option is not supported for DHCPv6 relay agent.
Options	<i>local-server-group</i> —Name of DHCP local server group.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Using DHCP Option Information to Selectively Process DHCP Client Traffic</i>

logical-system-name (DHCP Relay Agent)

Syntax	logical-system-name;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay authentication username-include],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p>
Description	<p>Specify that the logical system name is concatenated with the username during the subscriber authentication or client authentication process. No logical system name is concatenated if the configuration is in the default logical system. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Using External AAA Authentication Services with DHCP on page 923 • Creating Unique Usernames for DHCP Clients on page 946

mac-address (DHCP Relay Agent)

Syntax	mac-address;
Hierarchy Level	[edit forwarding-options dhcp-relay authentication username-include], [edit forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include]
Release Information	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify that the MAC address from the client PDU be concatenated with the username during the subscriber authentication or client authentication process.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Using External AAA Authentication Services with DHCP on page 923

method

Syntax	<pre> method { bfd { version (0 1 automatic); minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; multiplier <i>number</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } detection-time { threshold <i>milliseconds</i>; } session-mode (automatic multihop singlehop); holddown-interval <i>milliseconds</i>; } } </pre>
Hierarchy Level	<p> [edit system services dhcp-local-server liveness-detection], [edit system services dhcp-local-server dhcpv6 liveness-detection], [edit forwarding-options dhcp-relay liveness-detection], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection] </p>
Release Information	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure the liveness detection method.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913 • Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

minimum-interval

Syntax	<code>minimum-interval <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server liveness-detection method bfd transmit-interval], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd transmit-interval], [edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection method bfd transmit-interval], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd transmit-interval], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd transmit-interval], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd transmit-interval], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd transmit-interval], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd transmit-interval]</p>
Release Information	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure the minimum intervals at which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the transmit-interval minimal-interval and minimum-receive-interval statements.</p>
Options	<p><i>milliseconds</i> — Specify the minimum interval value for BFD liveliness detection.</p> <p>Range: 1 through 255,000</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913 • Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

minimum-receive-interval

Syntax	<code>minimum-receive-interval <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]</p>
Release Information	<p>Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Configure the minimum interval at which the local routing device (or switch) must receive a reply from a neighbor with which it has established a BFD session.
Options	<p><i>milliseconds</i> — Specify the minimum receive interval value. Range: 1 through 255,000</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913 • Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

multiplier

Syntax	<code>multiplier <i>number</i>;</code>
Hierarchy Level	[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the number of hello packets not received by the neighbor before Bidirectional Forwarding Detection (BFD) declares the neighbor down.
Options	number —Maximum allowable number of hello packets missed by the neighbor. Range: 1 through 255 Default: 3
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913• Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients


next-hop (Dynamic Access-Internal Routes)

Syntax	<code>next-hop <i>next-hop</i>;</code>
Hierarchy Level	[edit dynamic-profiles routing-options access route <i>prefix</i>]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Dynamically configure the next-hop address for an access route. Access routes are typically unnumbered interfaces.
Options	<p><i>next-hop</i>—Either the specific next-hop address you want to assign to the access route or one of the following next-hop address predefined variables.</p> <ul style="list-style-type: none"> For IPv4 access routes, use the variable, \$junos-framed-route-nexthop. The route prefix variable is dynamically replaced with the value in Framed-Route RADIUS attribute [22]. For IPv6 access routes, use the variable, \$junos-framed-route-ipv6-nexthop. The variable is dynamically replaced with the value in Framed-IPv6-Route RADIUS attribute [99].
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Configuring Dynamic Access Routes for Subscriber Management</i>

no-adaptation

Syntax	no-adaptation;
Hierarchy Level	[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure Bidirectional Forwarding Detection (BFD) sessions to not adapt to changing network conditions.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913• <i>Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</i>

no-allow-snooped-clients

Syntax	no-allow-snooped-clients;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> interface <i>interface-name</i> overrides],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> overrides],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 overrides],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> overrides],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit forwarding-options dhcp-relay overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.2.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3 for EX Series switches.</p>
Description	<p>Explicitly disable DHCP snooping support on the router or switch.</p> <p>Use the statement at the [edit ... dhcpv6] hierarchy levels to explicitly disable snooping support on the router or switch for DHCPv6 relay agent.</p>
<div>  <p>NOTE: In Junos OS Release 10.0 and earlier, DHCP snooping is <i>enabled</i> by default. In Release 10.1 and later, DHCP snooping is <i>disabled</i> by default.</p> </div>	
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Extended DHCP Relay Agent Overview on page 900 • Overriding the Default DHCP Relay Configuration Settings on page 955 • DHCP Snooping Support on page 887

no-bind-on-request (DHCP Relay Agent)

Syntax no-bind-on-request;

Hierarchy Level [edit forwarding-options dhcp-relay dhcpv6 [overrides](#)],
 [edit forwarding-options dhcp-relay [overrides](#)],
 [edit forwarding-options dhcp-relay dhcpv6 group *group-name* [overrides](#)],
 [edit forwarding-options dhcp-relay group *group-name* [overrides](#)],
 [edit logical-systems *logical-system-name* forwarding-options dhcp-relay dhcpv6 [overrides](#)],
 [edit logical-systems *logical-system-name* forwarding-options dhcp-relay [overrides](#)],
 [edit logical-systems *logical-system-name* forwarding-options dhcp-relay dhcpv6 group
group-name [overrides](#)],
 [edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name*
[overrides](#)],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
 forwarding-options dhcp-relay dhcpv6 [overrides](#)],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
 forwarding-options dhcp-relay [overrides](#)],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
 forwarding-options dhcp-relay dhcpv6 group *group-name* [overrides](#)],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
 forwarding-options dhcp-relay group *group-name* [overrides](#)],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay dhcpv6
[overrides](#)],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay [overrides](#)],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay dhcpv6 group
group-name [overrides](#)],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group
group-name [overrides](#)],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group
group-name interface *interface-name* [overrides](#)]

Release Information Statement introduced in Junos OS Release 10.4.
 Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.
 Statement introduced in Junos OS Release 12.3 for EX Series switches.

Description Explicitly disable automatic binding of received DHCP request messages that have no entry in the database (*stray* requests). Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.

M120 and M320 routers do not support DHCPv6.



NOTE: Beginning with Junos OS Release 10.4, automatic binding of stray requests is enabled by default. In Junos OS Release 10.3 and earlier releases, automatic binding of stray requests is disabled by default.

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

- Related Documentation**
- [Extended DHCP Relay Agent Overview on page 900](#)
 - [Overriding the Default DHCP Relay Configuration Settings on page 955](#)
 - [Disabling Automatic Binding of Stray DHCP Requests on page 972](#)


no-arp (DHCP Relay Agent)

Syntax	no-arp;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay overrides], [edit forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> overrides]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.3.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	Turn off ARP table population in a distrusted environment.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Extended DHCP Relay Agent Overview on page 900 • Overriding the Default DHCP Relay Configuration Settings on page 955

option-60 (DHCP Relay Agent)

Syntax	option-60;
Hierarchy Level	[edit forwarding-options dhcp-relay authentication username-include], [edit forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include]
Release Information	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify that the payload of the Option 60 (Vendor Class Identifier) from the client PDU is concatenated with the username during the subscriber authentication or client authentication process.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Using External AAA Authentication Services with DHCP on page 923

option-82 (DHCP Relay Agent)

Syntax	<code>option-82 <circuit-id> <remote-id>;</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay authentication username-include],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Specify the option 82 that is concatenated with the username during the subscriber authentication or client authentication process. You can specify either, both, or neither the Agent Circuit ID and the Agent Remote ID suboptions. If you specify both, the Agent Circuit ID is supplied first, followed by a delimiter, and then the Agent Remote ID. If neither suboption is supplied, the raw payload of option 82 is concatenated to the username.
<div>  <p>NOTE: The option 82 value used in creating the username is based on the option 82 value that is encoded in the outgoing (relayed) PDU.</p> </div>	
Options	<p>circuit-id—(Optional) The string for the Agent Circuit ID suboption (suboption 1).</p> <p>remote-id—(Optional) The string for the Agent Remote ID suboption (suboption 2).</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Using External AAA Authentication Services with DHCP on page 923

option-number (DHCP Relay Agent Option)

Syntax	<code>option-number option-number;</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay <i>relay-option</i>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 <i>relay-option</i>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> <i>relay-option</i>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <i>relay-option</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</p>
Release Information	<p>Statement introduced in Junos OS Release 12.3.</p> <p>Statement introduced in Junos OS Release 12.3 for EX Series switches.</p>
Description	<p>Specify the DHCP option DHCP relay agent uses for selective processing of client traffic. You can configure support globally or for a named group of interfaces. You can also configure support for the extended DHCP relay agent on a per logical system and per routing instance basis.</p> <p>Use the [edit forwarding-options dhcp-relay dhcpv6] hierarchy level to configure the DHCPv6 relay agent support.</p>
Options	<i>option-number</i> —The DHCP or DHCPv6 option in the incoming traffic.
<div>  <p>NOTE: EX Series switches do not support the User Class Options.</p> </div>	
<ul style="list-style-type: none"> 15 (DHCPv6 only)—Use DHCPv6 option 15 (User Class Option) in packets 16 (DHCPv6 only)—(MX Series routers and EX Series switches only) Use DHCPv6 option 16 (Vendor Class Option) in packets 60 (DHCPv4 only)—(MX Series routers and EX Series switches only) Use DHCP option 60 (Vendor Class Identifier) in DHCP packets 77 (DHCPv4 only)—Use DHCP option 77 (User Class Identifier) in packets 	
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Using DHCP Option Information to Selectively Process DHCP Client Traffic</i> <i>Configuring an Extended DHCP Relay Server on EX Series Switches (CLI Procedure)</i>

overrides (DHCP Relay Agent)

Syntax	<pre> overrides { allow-snooped-clients; always-write-giaddr; always-write-option-82; client-discover-match <option60-and-option82>; disable-relay; interface-client-limit <i>number</i>; layer2-unicast-replies; no-allow-snooped-clients; no-arp; no-bind-on-request; proxy-mode; replace-ip-source-with; send-release-on-delete; trust-option-82; } </pre>
Hierarchy Level	<pre> [edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit forwarding-options dhcp-relay group <i>group-name</i>], [edit forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...] </pre>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Override the default configuration settings for the extended DHCP relay agent. Specifying the overrides statement with no subordinate statements removes all DHCP relay agent overrides at that hierarchy level. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.</p> <p>M120 and M320 routers do not support DHCPv6.</p> <p>The following statements are supported at both the [edit ... dhcp-relay] and [edit ... dhcpv6] hierarchy levels. All other statements are supported at the dhcp-relay hierarchy levels only.</p> <ul style="list-style-type: none"> • allow-snooped-clients • interface-client-limit • no-allow-snooped-clients • no-bind-on-request • send-release-on-delete

The remaining statements are explained separately.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Extended DHCP Relay Agent Overview on page 900• Overriding the Default DHCP Relay Configuration Settings on page 955• Deleting DHCP Local Server and DHCP Relay Override Settings on page 935• dhcp-relay on page 1087

password (DHCP Relay Agent)

Syntax	<code>password password-string;</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay authentication], [edit forwarding-options dhcp-relay dhcpv6 authentication], [edit forwarding-options dhcp-relay group <i>group-name</i> authentication], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 authentication], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p>
Description	Configure the password that is sent to the external AAA authentication server for subscriber authentication or client authentication. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.
Options	<i>password-string</i> —Authentication password.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Using External AAA Authentication Services with DHCP on page 923 • Configuring Passwords for Usernames on page 945

preference (Subscriber Management)

Syntax	<code>preference route-distance;</code>
Hierarchy Level	[edit dynamic-profiles routing-options access route <i>prefix</i>]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Dynamically configure the distance for an access route.
Options	<i>route-distance</i> —Either the specific distance you want to assign to the access route or the distance variable (<code>\$junos-framed-route-distance</code>). The distance variable is dynamically replaced with the value in Framed-Route Attribute [22].
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Dynamic Access Routes for Subscriber Management</i>

prefix (DHCP Relay Agent)

Syntax	<code>prefix <i>prefix</i>;</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay dhcpv6 <i>relay-agent-interface-id</i>],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <i>relay-agent-interface-id</i>],</p> <p>[edit forwarding-options dhcp-relay relay-option-82 <i>circuit-id</i>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> relay-option-82 <i>circuit-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 <i>relay-agent-interface-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <i>relay-agent-interface-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay relay-option-82 <i>circuit-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-82 <i>circuit-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 <i>relay-agent-interface-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <i>relay-agent-interface-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay relay-option-82 <i>circuit-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-82 <i>circuit-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 <i>relay-agent-interface-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> <i>relay-agent-interface-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay relay-option-82 <i>circuit-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> relay-option-82 <i>circuit-id</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.3 for EX Series switches.</p>
Description	<p>Add a prefix to the base option 82 Agent Circuit ID information in DHCP packets destined for a DHCP server. The prefix can consist of any combination of the hostname, logical system name, and routing instance name. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.</p> <p>If you include only the hostname, only the logical system name, or only the routing instance name in the prefix, the format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces with stacked virtual LANs (S-VLANs) is one of the following:</p> <pre> host-name:(fe ge)-fpc/pic/port:svlan-id-vlan-id logical-system-name:(fe ge)-fpc/pic/port:svlan-id-vlan-id routing-instance-name:(fe ge)-fpc/pic/port:svlan-id-vlan-id </pre> <p>If you include both the logical system name and the routing instance name in the prefix, the format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs is as follows:</p>

logical-system-name;routing-instance-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id

If you include the hostname, logical system name, and routing instance name in the prefix, the format of the Agent Circuit ID information for Fast Ethernet or Gigabit Ethernet interfaces with S-VLANs is as follows:

host-name/logical-system-name;routing-instance-name:(fe | ge)-fpc/pic/port:svlan-id-vlan-id

For Fast Ethernet or Gigabit Ethernet interfaces that use virtual LANs (VLANs) but not S-VLANs, only the **vlan-id** value appears in the Agent Circuit ID format. For Fast Ethernet or Gigabit Ethernet interfaces that do not use VLANs or S-VLANs, neither the **vlan-id** value nor the **svlan-id** value appears.

Options *prefix*—Any of the following:

- **host-name**—Prepend the hostname of the router configured with the **host-name** statement at the **[edit system]** hierarchy level to the Agent Circuit ID information.
- **logical-system-name**—Prepend the name of the logical system to the Agent Circuit ID information.
- **routing-instance-name**—Prepend the name of the routing instance to the Agent Circuit ID information.

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

Related Documentation

- [Enabling and Disabling Insertion of Option 82 Information on page 973](#)
- [Configuring an Option 82 Prefix on page 974](#)
- [Inserting DHCPv6 Interface-ID Option \(Option 18\) In DHCPv6 Packets on page 979](#)

proxy-mode

Syntax	proxy-mode;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay overrides],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides]</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> overrides]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Enable DHCP relay proxy mode on the extended DHCP relay. Proxy mode supports all extended DHCP relay functionality.</p> <p>The extended DHCP relay proxy is not supported for the J Series routers DHCP server. Also, you cannot configure both the DHCP relay proxy and the extended DHCP local server on the same interface.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • DHCP Relay Proxy Overview on page 903 • Extended DHCP Relay Agent Overview on page 900 • Enabling DHCP Relay Proxy Mode on page 977

relay-agent-interface-id (DHCPv6 Relay Agent)

Syntax	<pre>relay-agent-interface-id { <i>prefix</i> <i>prefix</i>; <i>use-interface-description</i> (logical device): }</pre>
Hierarchy Level	[edit forwarding-options dhcp-relay dhcpv6], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 ...]
Release Information	Statement introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3 for EX Series switches.
Description	<p>Insert the DHCPv6 Relay Agent Interface-ID option (option 18) in DHCPv6 packets destined for the DHCPv6 server.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• dhcp-relay on page 1087• Extended DHCP Relay Agent Overview on page 900• DHCPv6 Relay Agent Overview on page 905• Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets on page 979

relay-agent-remote-id (DHCPv6 Relay Agent)

Syntax	<code>relay-agent-remote-id [enterprise-id remote-id];</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay dhcpv6 authentication username-include], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include]</p>
Release Information	<p>Statement introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>For MX Series routers only, enterprise-id and remote-id options at the [edit forwarding-options dhcp-relay dhcpv6 authentication username-include] and the [edit system services dhcp-local-server dhcpv6 authentication username-include] hierarchy levels introduced in Junos OS Release 12.3R3.</p>
Description	<p>Specify that the DHCPv6 Relay Agent Remote-ID option (option 37) in the client PDU name is concatenated with the username during the subscriber authentication or client authentication process.</p> <p>If you enter relay-agent-remote-id without options, all of the DHCPv6 Relay Agent Remote-ID option (option 37) contents in the client PDU name are concatenated with the username.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Creating Unique Usernames for DHCP Clients on page 946 • DHCPv6 Relay Agent Overview on page 905

relay-option (DHCP Relay Agent)

Syntax	<pre> relay-option { option-number option-number; default-action { drop; forward-only; local-server-group local-server-group; relay-server-group relay-server-group; } equals (ascii <i>ascii-string</i> hexadecimal <i>hexadecimal-string</i>) { drop; forward-only; local-server-group local-server-group; relay-server-group relay-server-group; } starts-with (ascii <i>ascii-string</i> hexadecimal <i>hexadecimal-string</i>) { drop; forward-only; local-server-group local-server-group; relay-server-group relay-server-group; } } </pre>
Hierarchy Level	<pre> [edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit forwarding-options dhcp-relay group group-name], [edit forwarding-options dhcp-relay dhcpv6 group group-name], [edit logical-systems logical-system-name forwarding-options dhcp-relay ...], [edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-options dhcp-relay ...], [edit routing-instances routing-instance-name forwarding-options dhcp-relay ...] </pre>
Release Information	<p>Statement introduced in Junos OS Release 12.3.</p> <p>Statement introduced in Junos OS Release 12.3 for EX Series switches.</p>
Description	<p>Configure the extended DHCP relay agent selective processing that is based on DHCP options in DHCP client packets and specify the action to perform on client traffic. You can configure support globally or for a named group of interfaces, and for either DHCP or DHCPv6 relay agent.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Using DHCP Option Information to Selectively Process DHCP Client Traffic

relay-option-82

Syntax	<pre> relay-option-82 { circuit-id { prefix <i>prefix</i>; use-interface-description (logical device); } } </pre>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.3 for EX Series switches.</p>
Description	<p>Enable or disable the insertion of the DHCP relay agent information option (option 82) in DHCP packets destined for a DHCP server.</p> <p>If you enable insertion of option 82 information in DHCP packets, you must specify at least the circuit-id statement to include the Agent Circuit ID suboption (suboption 1) of the DHCP relay agent information option.</p> <p>You can use the relay-option-82 statement and its subordinate statements at the [edit forwarding-options dhcp-relay] hierarchy level to control insertion of option 82 information globally, or at the [edit forwarding-options dhcp-relay group group-name] hierarchy level to control insertion of option 82 information for a named group of interfaces.</p> <p>To restore the default behavior (option 82 information is not inserted into DHCP packets), use the delete relay-option-82 statement.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Enabling and Disabling Insertion of Option 82 Information on page 973 • dhcp-relay on page 1087

relay-server-group (DHCP Relay Agent Option)

Syntax	<code>relay-server-group <i>relay-server-group</i>;</code>
Hierarchy Level	<code>[edit forwarding-options dhcp-relay relay-option (default-action equals starts-with),</code> <code>[edit forwarding-options dhcp-relay dhcpv6 relay-option (default-action equals </code> <code>starts-with),</code> <code>[edit forwarding-options dhcp-relay group <i>group-name</i> relay-option (default-action equals</code> <code> starts-with),</code> <code>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> relay-option (default-action</code> <code> equals starts-with),</code> <code>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i></code> <code>forwarding-options dhcp-relay ...],</code> <code>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</code>
Release Information	Statement introduced in Junos OS Release 12.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Relay DHCP client packets to the specified group of DHCP servers when you use the DHCP relay selective processing feature. You can configure the relay operation globally or for a group of interfaces, and for either DHCP or DHCPv6 relay agent.
Options	<i>relay-server-group</i> —Name of DHCP server group.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Using DHCP Option Information to Selectively Process DHCP Client Traffic</i>

replace-ip-source-with

Syntax	replace-ip-source-with giaddr;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay overrides],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides]</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> overrides]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Replace the IP source address in DHCP relay request and release packets with the gateway IP address (giaddr).
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Extended DHCP Relay Agent Overview on page 900 • Replacing the DHCP Relay Request and Release Packet Source Address on page 957

routing-instance-name (DHCP Relay Agent)

Syntax	routing-instance-name;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay authentication username-include],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p>
Description	Specify that the routing instance name is concatenated with the username during the subscriber authentication or client authentication process. No routing instance name is concatenated if the configuration is in the default routing instance. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Using External AAA Authentication Services with DHCP on page 923 • Creating Unique Usernames for DHCP Clients on page 946

send-release-on-delete (DHCP Relay Agent)

Syntax	send-release-on-delete;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay dhcpv6 overrides], [edit forwarding-options dhcp-relay overrides], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> overrides], [edit forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> overrides]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.2.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.3 for EX Series switches.</p>
Description	<p>Send a release message to the DHCP (or DHCPv6) server whenever DHCP relay or relay proxy deletes a client. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.</p> <p>M120 and M320 routers do not support DHCPv6.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Extended DHCP Relay Agent Overview on page 900 • Overriding the Default DHCP Relay Configuration Settings on page 955 • Sending Release Messages When Clients Are Deleted on page 971

server-group

Syntax	<pre>server-group { server-group-name { server-ip-address; } }</pre>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Specify the name of a group of DHCP server addresses for use by the extended DHCP relay agent. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.</p>
Options	<p>server-group-name—Name of the group of DHCP or DHCPv6 server addresses.</p> <p>server-ip-address—IP address of the DHCP server belonging to this named server group. Use IPv6 addresses when configuring DHCPv6 support. You can configure a maximum of five IP addresses in each named server group.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• dhcp-relay on page 1087• Extended DHCP Relay Agent Overview on page 900• Configuring Server Groups on page 976


service-profile (DHCP Relay Agent)

Syntax	<code>service-profile <i>dynamic-profile-name</i>;</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit forwarding-options dhcp-relay group <i>group-name</i>], [edit forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay ...], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]</p>
Release Information	<p>Statement introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 12.3R2 for EX Series switches. Support at the [edit ... dhcpv6 ...] hierarchy levels introduced in Junos OS Release 11.4.</p>
Description	<p>Specify the default subscriber service (or the default DHCP client management service), which is activated when the subscriber (or client) logs in and no other service is activated by a RADIUS server or a provisioning server.</p> <ul style="list-style-type: none"> To specify the default service for all DHCP relay agent clients, include the service-profile statement at the [edit forwarding-options dhcp relay] hierarchy level. To specify the default service for a named group of interfaces, include the service-profile statement at the [edit forwarding-options dhcp relay group <i>group-name</i>] hierarchy level. To specify the default service for a particular interface within a named group of interfaces, include the service-profile statement at the [edit forwarding-options dhcp relay group <i>group-name</i> interface <i>interface-name</i>] hierarchy level.
Options	<i>dynamic-profile-name</i> —Name of the dynamic profile.
Required Privilege Level	<p>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> dhcp-relay on page 1087 Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 942 Grouping Interfaces with Common DHCP Configurations on page 926 Default Subscriber Service Overview Configuring a Default Subscriber Service


session-mode

Syntax	session-mode (automatic multihop singlehop);
Hierarchy Level	[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the session mode.
Options	automatic —Configure single-hop BFD sessions if the peer is directly connected to the router interface and multihop BFD sessions if the peer is not directly connected to the router interface. multihop —Configure multihop BFD sessions. single-hop —Configure single hop BFD sessions.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913• <i>Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</i>

threshold (detection-time)

Syntax	<code>threshold <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit system services dhcp-local-server liveness-detection method bfd detection-time],</p> <p>[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd detection-time],</p> <p>[edit forwarding-options dhcp-relay liveness-detection method bfd detection-time],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd detection-time],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd detection-time],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd detection-time],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd detection-time],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd detection-time]</p>
Release Information	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Specify the threshold for the adaptation of the detection time. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.
<div>  <p>NOTE: The threshold time must be greater than or equal to the <code>minimum-interval</code> or the <code>minimum-receive-interval</code>.</p> </div>	
Options	<p><i>milliseconds</i>— Value for the detection time adaptation threshold.</p> <p>Range: 1 through 255,000</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913 • Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

threshold (transmit-interval)

Syntax	<code>threshold <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit system services dhcp-local-server liveness-detection method bfd transmit-interval],</p> <p>[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd transmit-interval],</p> <p>[edit forwarding-options dhcp-relay liveness-detection method bfd transmit-interval],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd transmit-interval],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd transmit-interval],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd transmit-interval],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd transmit-interval],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd transmit-interval]</p>
Release Information	Statement introduced in Junos OS Release 12.1.
Description	Specify the threshold for detecting the adaptation of the transmit interval. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.
Options	<p><i>milliseconds</i> — Threshold value.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p>
<div>  <p>NOTE: The threshold value specified in the <code>threshold</code> statement must be greater than the value specified in the <code>minimum-interval</code> statement for the <code>transmit-interval</code> statement.</p> </div>	
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913 • Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

transmit-interval

Syntax	<pre>transmit-interval { threshold milliseconds; minimum-interval milliseconds; }</pre>
Hierarchy Level	<p>[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]</p>
Release Information	<p>Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure the Bidirectional Forwarding Detection (BFD) transmit interval.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913 • Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

trust-option-82

Syntax	trust-option-82;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay overrides], [edit forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> overrides], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i> overrides]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Enable processing of DHCP client packets that have a gateway IP address (giaddr) of 0 (zero) and contain option 82 information. By default, the DHCP relay agent treats such packets as if they originated at an untrusted source, and drops them without further processing.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Trusting Option 82 Information on page 958• Overriding the Default DHCP Relay Configuration Settings on page 955

use-interface-description

Syntax use-interface-description (logical | device);

Hierarchy Level [edit forwarding-options dhcp-relay **dhcpv6** relay-agent-interface-id],
 [edit forwarding-options dhcp-relay **dhcpv6** group *group-name* relay-agent-interface-id],
 [edit forwarding-options dhcp-relay relay-option-82 **circuit-id**],
 [edit forwarding-options dhcp-relay group *group-name* relay-option-82 **circuit-id**],
 [edit logical-systems *logical-system-name* forwarding-options dhcp-relay **dhcpv6** relay-agent-interface-id],
 [edit logical-systems *logical-system-name* forwarding-options dhcp-relay **dhcpv6** group *group-name* relay-agent-interface-id],
 [edit logical-systems *logical-system-name* forwarding-options dhcp-relay relay-option-82 **circuit-id**],
 [edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name* relay-option-82 **circuit-id**],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay **dhcpv6** relay-agent-interface-id],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay **dhcpv6** group *group-name* relay-agent-interface-id],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay relay-option-82 **circuit-id**],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* relay-option-82 **circuit-id**],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay **dhcpv6** relay-agent-interface-id],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay **dhcpv6** group *group-name* relay-agent-interface-id],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay relay-option-82 **circuit-id**],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* relay-option-82 **circuit-id**]

Release Information Statement introduced in Junos OS Release 9.6.
 Support at the [edit ... **dhcpv6**] hierarchy levels introduced in Junos OS Release 11.4.
 Statement introduced in Junos OS Release 12.3 for EX Series switches.

Description Use the textual interface description instead of the interface identifier when creating the agent-circuit-id suboption of the DHCP relay agent option 82. Use the statement at the [edit ... **dhcpv6**] hierarchy levels to configure DHCPv6 support.

If you specify that the textual description is used and no description is configured for the interface, DHCP relay defaults to using the interface identifier. The textual description is configured using the **description** statement at the [edit interfaces *interface-name*] hierarchy level.



NOTE: By default, DHCP relay accepts a maximum of 253 ASCII characters. If the textual interface description is longer than 253 characters, DHCP relay drops the packet, which results in the DHCP client failing to bind.

Options	logical —Use the textual description that is configured for the logical interface.
	device —Use the textual description that is configured for the device interface.
Required Privilege Level	interface—To view this statement in the configuration.
	interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling and Disabling Insertion of Option 82 Information on page 973• Using a Textual Description in Option 82 on page 976• Inserting DHCPv6 Interface-ID Option (Option 18) In DHCPv6 Packets on page 979

use-primary (DHCP Relay Agent)

Syntax	<code>use-primary <i>primary-profile-name</i>;</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay dhcpv6 dynamic-profile <i>profile-name</i>], [edit forwarding-options dhcp-relay dynamic-profile <i>profile-name</i>], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> dynamic-profile <i>profile-name</i>], [edit forwarding-options dhcp-relay group <i>group-name</i> dynamic-profile <i>profile-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 dynamic-profile <i>profile-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dynamic-profile <i>profile-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> dynamic-profile <i>profile-name</i>], [edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> dynamic-profile <i>profile-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 dynamic-profile <i>profile-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dynamic-profile <i>profile-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> dynamic-profile <i>profile-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> dynamic-profile <i>profile-name</i>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 dynamic-profile <i>profile-name</i>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dynamic-profile <i>profile-name</i>], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> dynamic-profile <i>profile-name</i>] [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> dynamic-profile <i>profile-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.3.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Specify the dynamic profile to configure as the primary dynamic profile. The primary dynamic profile is instantiated when the first subscriber logs in. Subsequent subscribers are not assigned the primary dynamic profile; instead, they are assigned the dynamic profile specified for the interface. When the first subscriber logs out, the next subscriber that logs in is assigned the primary dynamic profile.</p> <p>Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.</p> <p>EX Series switches do not support DHCPv6.</p>
Options	<i>primary-profile-name</i> —Name of the dynamic profile to configure as the primary dynamic profile
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Attaching Dynamic Profiles to DHCP Subscriber Interfaces or DHCP Client Interfaces on page 942](#)

user-prefix (DHCP Relay Agent)

Syntax	<code>user-prefix <i>user-prefix-string</i>;</code>
Hierarchy Level	<p>[edit forwarding-options dhcp-relay authentication username-include],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> authentication username-include],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> authentication username-include]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p>
Description	Specify the user prefix that is concatenated with the username during the subscriber authentication or client authentication process. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.
Options	<i>user-prefix-string</i> —User prefix string.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Using External AAA Authentication Services with DHCP on page 923

username-include (DHCP Relay Agent)

Syntax username-include {
 circuit-type;
 client-id;
 delimiter *delimiter-character*;
 domain-name *domain-name-string*;
 interface-name;
 logical-system-name;
 mac-address;
 option-60;
 option-82 <circuit-id> <remote-id>;
 relay-agent-interface-id;
 relay-agent-remote-id;
 relay-agent-subscriber-id;
 routing-instance-name;
 user-prefix *user-prefix-string*;
 }

Hierarchy Level [edit forwarding-options dhcp-relay [authentication](#)],
 [edit forwarding-options dhcp-relay dhcpv6 [authentication](#)],
 [edit forwarding-options dhcp-relay dhcpv6 group *group-name* [authentication](#)],
 [edit forwarding-options dhcp-relay group *group-name* [authentication](#)],
 [edit logical-systems *logical-system-name* forwarding-options dhcp-relay ...],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
 forwarding-options dhcp-relay ...],
 [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay ...]

Release Information Statement introduced in Junos OS Release 9.1.
 Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
 Support at the [\[edit ... dhcpv6\]](#) hierarchy levels introduced in Junos OS Release 11.4.

Description Configure the username that the router (or switch) passes to the external AAA server. You must include at least one of the optional statements for the username to be valid. If you do not configure a username, the router (or switch) accesses the local authentication service only and does not use external authentication services, such as RADIUS. Use the statement at the [\[edit...dhcpv6\]](#) hierarchy levels to configure DHCPv6 support.

The following statements are not supported in the DHCPv6 hierarchy levels:

- mac-address
- option-60
- option-82

The following statements are supported in the DHCPv6 hierarchy levels only:

- relay-agent-interface-id
- relay-agent-remote-id
- relay-agent-subscriber-id

The remaining statements are explained separately.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Creating Unique Usernames for DHCP Clients on page 946 • Using External AAA Authentication Services with DHCP on page 923

version (BFD)

Syntax	version (0 1 automatic);
Hierarchy Level	[edit system services dhcp-local-server liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 liveness-detection method bfd], [edit forwarding-options dhcp-relay liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd], [edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd], [edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the BFD protocol version to detect.
Options	0—Use BFD protocol version 0. 1—Use BFD protocol version 1. automatic—Autodetect the BFD protocol version. Default: automatic
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Group Liveness Detection for DHCP Local Server Clients on page 913 • Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients

Administration

- [Verifying and Managing DHCP Local Server Configurations on page 1158](#)
- [Verifying and Managing DHCP Relay Agent Configurations on page 1159](#)

- [DHCP Local Server Monitoring Commands on page 1159](#)
- [DHCP Relay Agent Monitoring Commands on page 1192](#)

Verifying and Managing DHCP Local Server Configurations

- [Verifying and Managing DHCP Local Server Configuration on page 1158](#)
- [Verifying and Managing DHCPv6 Local Server Configuration on page 1158](#)

Verifying and Managing DHCP Local Server Configuration

- Purpose** View or clear information about client address bindings and statistics for the extended DHCP local server.
- Action**
- To display the address bindings in the client table on the extended DHCP local server:
user@host> [show dhcp server binding](#)
 - To display extended DHCP local server statistics:
user@host> [show dhcp server statistics](#)
 - To clear the binding state of a DHCP client from the client table on the extended DHCP local server:
user@host> [clear dhcp server binding](#)
 - To clear all extended DHCP local server statistics:
user@host> [clear dhcp server statistics](#)
- Related Documentation**
- [Junos OS Operational Mode Commands](#)

Verifying and Managing DHCPv6 Local Server Configuration

- Purpose** View or clear information about client address bindings and statistics for the DHCPv6 local server.
- Action**
- To display the address bindings in the client table on the DHCPv6 local server:
user@host> [show dhcpv6 server binding](#)
 - To display DHCPv6 local server statistics:
user@host> [show dhcpv6 server statistics](#)
 - To clear all DHCPv6 local server statistics:
user@host> [clear dhcpv6 server binding](#)
 - To clear all DHCPv6 local server statistics:
user@host> [clear dhcpv6 server statistics](#)
- Related Documentation**
- [Junos OS Operational Mode Commands](#)

Verifying and Managing DHCP Relay Agent Configurations

- [Verifying and Managing DHCP Relay Configuration on page 1159](#)
- [Verifying and Managing DHCPv6 Relay Configuration on page 1159](#)

Verifying and Managing DHCP Relay Configuration

Purpose View or clear address bindings or statistics for extended DHCP relay agent clients:

Action • To display the address bindings for extended DHCP relay agent clients:

user@host> [show dhcp relay binding](#)

• To display extended DHCP relay agent statistics:

user@host> [show dhcp relay statistics](#)

• To clear the binding state of DHCP relay agent clients:

user@host> [clear dhcp relay binding](#)

• To clear all extended DHCP relay agent statistics:

user@host> [clear dhcp relay statistics](#)

Related Documentation • [Junos OS Operational Mode Commands](#)

Verifying and Managing DHCPv6 Relay Configuration

Purpose View or clear address bindings or statistics for extended DHCPv6 relay agent clients:

Action • To display the address bindings for extended DHCPv6 relay agent clients:

user@host> [show dhcpv6 relay binding](#)

• To display extended DHCPv6 relay agent statistics:

user@host> [show dhcpv6 relay statistics](#)

• To clear the binding state of DHCPv6 relay agent clients:

user@host> [clear dhcpv6 relay binding](#)

• To clear all extended DHCPv6 relay agent statistics:

user@host> [clear dhcpv6 relay statistics](#)

Related Documentation • [Junos OS Operational Mode Commands](#)

DHCP Local Server Monitoring Commands

clear dhcp server binding

Syntax `clear dhcp server binding`
 `<address>`
 `<all>`
 `<interface interface-name>`
 `<interfaces-vlan>`
 `<interfaces-wildcard>`
 `<logical-system logical-system-name>`
 `<routing-instance routing-instance-name>`

Release Information Command introduced in Junos OS Release 9.0.
 Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

Description Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table on the extended DHCP local server.

Options *address*—(Optional) Clear the binding state for the DHCP client, using one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

all—(Optional) Clear the binding state for all DHCP clients.

interface interface-name—(Optional) Clear the binding state for DHCP clients on the specified interface.



.....

NOTE: This option clears all bindings whose initial login requests were received over the specified interface. Dynamic demux login requests are not received over the dynamic demux interface, but rather the underlying interface of the dynamic demux interface. To clear a specific dynamic demux interface, use the *ip-address* or *mac-address* options.

.....

interfaces-vlan—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.

interfaces-wildcard—(Optional) Clear bindings on a set of interfaces. This option supports the use of the wildcard character (*).

logical-system logical-system-name—(Optional) Clear the binding state for DHCP clients on the specified logical system.

routing-instance routing-instance-name—(Optional) Clear the binding state for DHCP clients on the specified routing instance.

Required Privilege Level view

Related Documentation

- *Clearing DHCP Bindings for Subscriber Access*
- [show dhcp server binding on page 1177](#)

List of Sample Output

- [clear dhcp server binding <ip-address> on page 1161](#)
- [clear dhcp server binding all on page 1161](#)
- [clear dhcp server binding interface on page 1161](#)
- [clear dhcp server binding <interfaces-vlan> on page 1162](#)
- [clear dhcp server binding <interfaces-wildcard> on page 1162](#)

Output Fields See [show dhcp server binding](#) for an explanation of output fields.

Sample Output

clear dhcp server binding <ip-address>

The following sample output displays the address bindings in the DHCP client table on the extended DHCP local server before and after the **clear dhcp server binding** command is issued.

```
user@host> show dhcp server binding
```

```
2 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)
```

IP address	Hardware address	Type	Lease expires at
100.20.32.1	90:00:00:01:00:01	active	2007-01-17 11:38:47 PST
100.20.32.3	90:00:00:02:00:01	active	2007-01-17 11:38:41 PST

```
user@host> clear dhcp server binding 10.20.32.1
```

```
user@host> show dhcp server binding
```

```
1 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)
```

IP address	Hardware address	Type	Lease expires at
100.20.32.3	90:00:00:02:00:01	active	2007-01-17 11:38:41 PST

clear dhcp server binding all

The following command clears all DHCP local server bindings:

```
user@host> clear dhcp server binding all
```

clear dhcp server binding interface

The following command clears DHCP local server bindings on a specific interface:

```
user@host> clear dhcp server binding interface fe-0/0/2
```

clear dhcp server binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCP local server bindings on top of the underlying interface **ae0**, which clears DHCP bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcp server binding ae0
```

clear dhcp server binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCP local server bindings over a specific interface:

```
user@host> clear dhcp server binding ge-1/0/0.*
```


clear dhcp server statistics

Syntax	<code>clear dhcp server statistics</code> <code><interface <i>interface-name</i>></code> <code><logical-system <i>logical-system-name</i>></code> <code><routing-instance <i>routing-instance-name</i>></code>
Release Information	Command introduced in Junos OS Release 9.0.
Description	Clear all extended Dynamic Host Configuration Protocol (DHCP) local server statistics.
Options	<p>logical-system <i>logical-system-name</i>—(Optional) Clear the statistics for DHCP clients on the specified logical system. If you do not specify a logical system, statistics are cleared for the default logical system.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) Clear the statistics for DHCP clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.</p>
Required Privilege Level	view
List of Sample Output	clear dhcp server statistics on page 1163
Output Fields	See show dhcp server statistics for an explanation of output fields.

Sample Output

clear dhcp server statistics

The following sample output displays the extended DHCP local server statistics before and after the **clear dhcp server statistics** command is issued.

```

user@host> show dhcp server statistics
Packets dropped:
  Total                               0

Messages received:
  BOOTREQUEST                        89163
  DHCPDECLINE                        0
  DHCPDISCOVER                       8110
  DHCPINFORM                         0
  DHCPRELEASE                        0
  DHCPREQUEST                        81053

Messages sent:
  BOOTREPLY                          32420
  DHCPOFFER                          8110
  DHCPACK                            8110
  DHCPNAK                            8100

user@host> clear dhcp server statistics
user@host> show dhcp server statistics

```

Packets dropped:	
Total	0
Messages received:	
BOOTREQUEST	0
DHCPCDECLINE	0
DHCPDISCOVER	0
DHCPINFORM	0
DHCPRELEASE	0
DHCPREQUEST	0
Messages sent:	
BOOTREPLY	0
DHCPPOFFER	0
DHCPACK	0
DHCPNAK	0

clear dhcpv6 server binding

Syntax	<pre>clear dhcpv6 server binding <address> <all> <interface interface-name> <interfaces-vlan> <interfaces-wildcard> <logical-system logical-system-name> <routing-instance routing-instance-name></pre>
Release Information	<p>Command introduced in Junos OS Release 9.6.</p> <p>Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.</p>
Description	Clear the binding state of a Dynamic Host Configuration Protocol for IPv6 (DHCPv6) client from the client table on the extended DHCPv6 local server.
Options	<p>address—(Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:</p> <ul style="list-style-type: none"> • <i>CID</i>—The specified Client ID (CID). • <i>ipv6-prefix</i>—The specified IPv6 prefix. • <i>session-id</i>—The specified session ID. <p>all—(Optional) Clear the binding state for all DHCPv6 clients.</p> <p>interface interface-name—(Optional) Clear the binding state for DHCPv6 clients on the specified interface.</p> <p>interfaces-vlan—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.</p> <p>interfaces-wildcard—(Optional) Clear bindings on a set of interfaces. This option supports the use of the wildcard character (*).</p> <p>logical-system logical-system-name—(Optional) Clear the binding state for DHCPv6 clients on the specified logical system.</p> <p>routing-instance routing-instance-name—(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • Clearing DHCP Bindings for Subscriber Access • show dhcpv6 server binding on page 1185
List of Sample Output	<p>clear dhcpv6 server binding all on page 1166</p> <p>clear dhcpv6 server binding <ipv6-prefix> on page 1166</p>

[clear dhcpv6 server binding interface on page 1166](#)
[clear dhcpv6 server binding <interfaces-vlan> on page 1166](#)
[clear dhcpv6 server binding <interfaces-wildcard> on page 1166](#)

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

Sample Output

clear dhcpv6 server binding all

The following command clears all DHCPv6 local server bindings:

```
user@host> clear dhcpv6 server binding all
```

clear dhcpv6 server binding <ipv6-prefix>

The following command clears DHCPv6 local server bindings for a specific IPv6 prefix:

```
user@host> clear dhcpv6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005
```

clear dhcpv6 server binding interface

The following command clears DHCPv6 local server bindings on a specific interface:

```
user@host> clear dhcpv6 server binding interface fe-0/0/2
```

clear dhcpv6 server binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCPv6 local server bindings on top of the underlying interface **ae0**, which clears DHCPv6 bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcpv6 server binding interface ae0
```

clear dhcpv6 server binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCPv6 local server bindings over a specific interface:

```
user@host> clear dhcpv6 server binding ge-1/0/0.*
```

clear dhcpv6 server statistics


Syntax	<pre>clear dhcpv6 server statistics <interface <i>interface-name</i>> <logical-system <i>logical-system-name</i>> <routing-instance <i>routing-instance-name</i>></pre>
Release Information	Command introduced in Junos OS Release 9.6.
Description	Clear all extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.
Options	<p>logical-system <i>logical-system-name</i>—(Optional) Clear the statistics for DHCPv6 clients on the specified logical system. If you do not specify a logical system, statistics are cleared for the default logical system.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) Clear the statistics for DHCPv6 clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show dhcpv6 server statistics on page 1190
List of Sample Output	clear dhcpv6 server statistics on page 1167
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear dhcpv6 server statistics

```
user@host> clear dhcpv6 server statistics
```

request dhcp server reconfigure

Syntax	<code>request dhcp server reconfigure (all <i>address</i> interface <i>interface-name</i> logical-system <i>logical-system-name</i> routing-instance <i>routing-instance-name</i>)</code>
Release Information	Command introduced in Junos OS Release 10.0. Command introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>Initiate reconfiguration processing for the specified DHCP clients if they are in the bound state. If the clients are in the reconfiguring state, this command has no effect. If the clients are in any state other than bound or reconfiguring, this command has the same effect as the clear dhcp server binding command.</p> <p>When the local server state machine starts the reconfiguration process on a bound client, the client transitions to the reconfiguring state and the local server sends a forcerenew message to the client. Because the client was in the bound state before entering the reconfiguring state, all subscriber (or DHCP client) services, such as forwarding and statistics, continue to work. An exponential back-off timer determines the interval at which the forcerenew message is sent. If the final attempt is unsuccessful, the client is returned to its original state by default. You can optionally include the clear-on-abort statement to configure the client to be cleared when reconfiguration fails.</p>
Options	<p>all—Initiate reconfiguration for all DHCP clients.</p> <p><i>address</i>—Initiate reconfiguration for DHCP client with the specified IP address or MAC address.</p> <p>interface <i>interface-name</i>—Initiate reconfiguration for all DHCP clients on this logical interface (clients whose initial login requests were received over the specified interface).</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p> NOTE: You cannot use the interface <i>interface-name</i> option with the request dhcp server reconfigure command for DHCP passive clients (clients that are added as a result of DHCP snooped packets). For passive clients, the interface is not guaranteed to be the next-hop interface to the client, as is the case for active clients.</p> </div> <p>logical-system <i>logical-system-name</i>—Initiate reconfiguration for all DHCP clients on the specified logical system.</p> <p>routing-instance <i>routing-instance-name</i>—Initiate reconfiguration reconfigured for all DHCP clients in the specified routing instance.</p>
Required Privilege Level	view

Related Documentation • [Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 936](#)

List of Sample Output [request dhcp server reconfigure on page 1169](#)

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

Sample Output

[request dhcp server reconfigure](#)

```
user@host> request dhcp server reconfigure interface fe-0/0/0.100
```

request dhcpv6 server reconfigure

Syntax	request dhcpv6 server reconfigure (all address client-id interface <i>interface-name</i> logical-system <i>logical-system-name</i> routing-instance <i>routing-instance-name</i> session-id)
Release Information	Command introduced in Junos OS Release 10.4. Command introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>Initiate reconfiguration processing for the specified DHCPv6 clients if they are in the bound state. If the clients are in the reconfiguring state, this command has no effect. If the clients are in any state other than bound or reconfiguring, this command has the same effect as the clear dhcpv6 server binding command.</p> <p>When the local server state machine starts the reconfiguration process on a bound client, the client transitions to the reconfigure state and the local server sends a reconfigure message to the client. Because the client was in the bound state before entering the reconfiguring state, all subscriber (or DHCP client) services, such as forwarding and statistics, continue to work. An exponential back-off timer determines the interval at which the reconfigure message is sent. If the final attempt is unsuccessful, the client is returned to its original state by default. You can optionally include the clear-on-abort statement to configure the client to be cleared when reconfiguration fails.</p>
Options	<p>all—Initiate reconfiguration for all DHCPv6 clients.</p> <p>address—Initiate reconfiguration for DHCPv6 client with the specified IPv6 address.</p> <p>client-id—Initiate reconfiguration for DHCPv6 client with the specified client ID.</p> <p>interface <i>interface-name</i>—Initiate reconfiguration for all DHCPv6 clients on this logical interface (clients whose initial login requests were received over the specified interface).</p> <p>logical-system <i>logical-system-name</i>—Initiate reconfiguration for all DHCPv6 clients on the specified logical system.</p> <p>routing-instance <i>routing-instance-name</i>—Initiate reconfiguration reconfigured for all DHCPv6 clients in the specified routing instance.</p> <p>session-id—Initiate reconfiguration for DHCPv6 client with the specified session ID.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• Configuring Extended DHCP Local Server Dynamic Client Reconfiguration on page 936
List of Sample Output	request dhcpv6 server reconfigure on page 1171
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request dhcpv6 server reconfigure

```
user@host> request dhcpv6 server reconfigure 2001::2/16
```

request system reboot

Syntax	request system reboot <at <i>time</i> > <both-routing-engines> <in <i>minutes</i> > <media (compact-flash disk removable-compact-flash usb)> <message " <i>text</i> "> <other-routing-engine>
Syntax (EX Series Switches)	request system reboot <all-members> <at <i>time</i> > <both-routing-engines> <in <i>minutes</i> > <local> <media (external internal)> <member <i>member-id</i> > <message " <i>text</i> "> <other-routing-engine> <slice <i>slice</i> >
Syntax (TX Matrix Router)	request system reboot <all-chassis all-lcc lcc <i>number</i> scc> <at <i>time</i> > <both-routing-engines> <in <i>minutes</i> > <media (compact-flash disk)> <message " <i>text</i> "> <other-routing-engine>
Syntax (TX Matrix Plus Router)	request system reboot <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> > <at <i>time</i> > <both-routing-engines> <in <i>minutes</i> > <media (compact-flash disk)> <message " <i>text</i> "> <other-routing-engine> <partition (1 2 alternate)>
Syntax (MX Series Router)	request system reboot <all-members> <at <i>time</i> > <both-routing-engines> <in <i>minutes</i> > <local> <media (external internal)> <member <i>member-id</i> > <message " <i>text</i> "> <other-routing-engine>
Release Information	Command introduced before Junos OS Release 7.4.

Option **other-routing-engine** introduced in Junos OS Release 8.0.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Option **sfc** introduced for the TX Matrix Plus router in Junos OS Release 9.6.
 Option **both-routing-engines** introduced in Junos OS Release 12.1.

Description Reboot the software.

Options **none**—Reboot the software immediately.

all-chassis—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router or TX Matrix Plus router, reboot all routers connected to the TX Matrix or TX Matrix Plus router, respectively.

all-lcc—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router or TX Matrix Plus router, reboot all line card chassis connected to the TX Matrix or TX Matrix Plus router, respectively.

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

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

- **now**—Stop or reboot the software immediately. This is the default.
- **+minutes**—Number of minutes from now to reboot the software.
- **yymmddhhmm**—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.
- **hh:mm**—Absolute time on the current day at which to stop the software, specified in 24-hour time.

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

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

lcc number—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router or TX Matrix Plus router, the number of a specified line card chassis connected to the TX Matrix or TX Matrix Plus router, respectively. Replace **number** with a value from 0 through 3.

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

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

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

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

- **internal**—Reboot the internal flash device.

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

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

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

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

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

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

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

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

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

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

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



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



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

Required Privilege Level maintenance

Related Documentation

- *clear system reboot*
- *request system halt*
- *request system reboot*
- *Rebooting and Halting a QFX Series Product*

List of Sample Output

- [request system reboot on page 1175](#)
- [request system reboot \(at 2300\) on page 1175](#)
- [request system reboot \(in 2 Hours\) on page 1176](#)
- [request system reboot \(Immediately\) on page 1176](#)
- [request system reboot \(at 1:20 AM\) on page 1176](#)

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

Sample Output

`request system reboot`

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

`request system reboot (at 2300)`

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

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

request system reboot (in 2 Hours)

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

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

request system reboot (Immediately)

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

request system reboot (at 1:20 AM)

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

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

show dhcp server binding

Syntax	<pre>show dhcp server binding <address> <brief detail summary> <interface interface-name> <interfaces-vlan> <interfaces-wildcard> <logical-system logical-system-name> <routing-instance routing-instance-name></pre>
Release Information	<p>Command introduced in Junos OS Release 9.0.</p> <p>Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.</p>
Description	Display the address bindings in the client table on the extended Dynamic Host Configuration Protocol (DHCP) local server.
Options	<p>address—(Optional) Display DHCP binding information for a specific client identified by one of the following entries:</p> <ul style="list-style-type: none"> • <i>ip-address</i>—The specified IP address. • <i>mac-address</i>—The specified MAC address. • <i>session-id</i>—The specified session ID. <p>brief detail summary—(Optional) Display the specified level of output about active client bindings. The default is brief, which produces the same output as show dhcp server binding.</p> <p>interface interface-name—(Optional) Display information about active client bindings on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.</p> <p>interfaces-vlan—(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.</p> <p>interfaces-wildcard—(Optional) The set of interfaces on which to show the binding state information. This option supports the use of the wildcard character (*).</p> <p>logical-system logical-system-name—(Optional) Display information about active client bindings for DHCP clients on the specified logical system.</p> <p>routing-instance routing-instance-name—(Optional) Display information about active client bindings for DHCP clients on the specified routing instance.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Clearing DHCP Bindings for Subscriber Access</i> • <i>Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration</i> • clear dhcp server binding on page 1160

List of Sample Output

- [show dhcp server binding on page 1179](#)
- [show dhcp server binding detail on page 1179](#)
- [show dhcp server binding detail \(ACI Interface Set Configured\) on page 1180](#)
- [show dhcp server binding interface <vlan-id> on page 1180](#)
- [show dhcp server binding interface <svlan-id> on page 1180](#)
- [show dhcp server binding <ip-address> on page 1181](#)
- [show dhcp server binding <session-id> on page 1181](#)
- [show dhcp server binding summary on page 1181](#)
- [show dhcp server binding <interfaces-vlan> on page 1181](#)
- [show dhcp server binding <interfaces-wildcard> on page 1181](#)

Output Fields Table 80 on page 1178 lists the output fields for the **show dhcp server binding** command. Output fields are listed in the approximate order in which they appear.

Table 80: show dhcp server binding Output Fields

Field Name	Field Description	Level of Output
<i>number</i> clients, (<i>number</i> init, <i>number</i> bound, <i>number</i> selecting, <i>number</i> requesting, <i>number</i> renewing, <i>number</i> releasing)	Summary counts of the total number of DHCP clients and the number of DHCP clients in each state.	summary
IP address	IP address of the DHCP client.	brief detail
Session Id	Session ID of the subscriber session.	brief detail
Hardware address	Hardware address of the DHCP client.	brief detail
Expires	Number of seconds in which lease expires.	brief detail
State	State of the address binding table on the extended DHCP local server: <ul style="list-style-type: none"> • BOUND—Client has active IP address lease. • FORCERENEW—Client has received forcere Renew message from server. • INIT—Initial state. • RELEASE—Client is releasing IP address lease. • RENEWING—Client sending request to renew IP address lease. • REQUESTING—Client requesting a DHCP server. • SELECTING—Client receiving offers from DHCP servers. 	brief detail
Interface	Interface on which the request was received.	brief

Table 80: show dhcp server binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
Lease Expires	Date and time at which the client's IP address lease expires.	detail
Lease Expires in	Number of seconds in which lease expires.	detail
Lease Start	Date and time at which the client's IP address lease started.	detail
Last Packet Received	Date and time at which the router received the last packet.	detail
Incoming Client Interface	Client's incoming interface.	detail
Client Interface Svlan Id	S-VLAN ID of the client's incoming interface.	detail
Client Interface Vlan Id	VLAN ID of the client's incoming interface.	detail
Demux Interface	Name of the IP demultiplexing (demux) interface.	detail
Server IP Address or Server Identifier	IP address of DHCP server.	detail
Server Interface	Interface of DHCP server.	detail
Client Pool Name	Name of address pool used to assign client IP address lease.	detail
ACI Interface Set Name	Internally generated name of the dynamic agent circuit identifier (ACI) interface set.	detail
ACI Interface Set Index	Index number of the dynamic ACI interface set.	detail
ACI Interface Set Session ID	Identifier of the dynamic ACI interface set entry in the session database.	detail

Sample Output

show dhcp server binding

```

user@host> show dhcp server binding
IP address      Session Id  Hardware address  Expires  State  Interface
100.20.20.15    6          00:10:94:00:00:01  86180    BOUND  ge-1/0/0.0
100.20.20.16    7          00:10:94:00:00:02  86180    BOUND  ge-1/0/0.0
100.20.20.17    8          00:10:94:00:00:03  86180    BOUND  ge-1/0/0.0
100.20.20.18    9          00:10:94:00:00:04  86180    BOUND  ge-1/0/0.0
100.20.20.19    10         00:10:94:00:00:05  86180    BOUND  ge-1/0/0.0

```

show dhcp server binding detail

```

user@host> show dhcp server binding detail

```

```

Client IP Address: 100.20.20.15
  Hardware Address:      00:10:94:00:00:01
  State:                 BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)

  Lease Expires:         2009-07-21 10:10:25 PDT
  Lease Expires in:      86151 seconds
  Lease Start:           2009-07-20 10:10:25 PDT
  Incoming Client Interface: ge-1/0/0.0
  Server Ip Address:      100.20.20.9
  Server Interface:       none
  Session Id:             6
  Client Pool Name:       6
  Client IP Address:      100.20.20.16
  Hardware Address:       00:10:94:00:00:02
  State:                 BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)

  Lease Expires:         2009-07-21 10:10:25 PDT
  Lease Expires in:      86151 seconds
  Lease Start:           2009-07-20 10:10:25 PDT
  Incoming Client Interface: ge-1/0/0.0
  Server Ip Address:      100.20.20.9
  Server Interface:       none
  Session Id:             7
  Client Pool Name:       7

```

show dhcp server binding detail (ACI Interface Set Configured)

```

user@host> show dhcp server binding detail
Client IP Address: 100.20.22.14
  Hardware Address:      00:00:64:34:01:02
  State:                 BOUND(LOCAL_SERVER_STATE_BOUND)
  Lease Expires:         2012-03-13 09:53:32 PDT
  Lease Expires in:      82660 seconds
  Lease Start:           2012-03-12 10:23:32 PDT
  Last Packet Received:  2012-03-12 10:23:32 PDT
  Incoming Client Interface: demux0.1073741827
  Client Interface Svlan Id: 1802
  Client Interface Vlan Id: 302
  Demux Interface:       demux0.1073741832
  Server Identifier:      100.20.200.202
  Session Id:             11
  Client Pool Name:       poolA
  Client Profile Name:    DEMUXprofile
  ACI Interface Set Name:  aci-1002-demux0.1073741827
  ACI Interface Set Index: 2
  ACI Interface Set Session ID: 6

```

show dhcp server binding interface <vlan-id>

```

user@host> show dhcp server binding interface ge-1/1/0:100
IP address      Session Id  Hardware address  Expires  State  Interface
200.20.20.15    6          00:10:94:00:00:01  86124    BOUND
ge-1/1/0:100

```

show dhcp server binding interface <svlan-id>

```

user@host> show dhcp server binding interface ge-1/1/0:10-100
IP address      Session Id  Hardware address  Expires  State  Interface
200.20.20.16    7          00:10:94:00:00:02  86124    BOUND
ge-1/1/0:10-100

```

show dhcp server binding <ip-address>

```

user@host> show dhcp server binding 100.20.20.19
IP address      Session Id  Hardware address  Expires   State   Interface
100.20.20.19    10         00:10:94:00:00:05 86081     BOUND   ge-1/0/0.0

```

show dhcp server binding <session-id>

```

user@host> show dhcp server binding 6
IP address      Session Id  Hardware address  Expires   State   Interface
200.20.20.15    6          00:10:94:00:00:01 86124     BOUND   ge-1/0/0.0

```

show dhcp server binding summary

```

user@host> show dhcp server binding summary
3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)

```

show dhcp server binding <interfaces-vlan>

```

user@host> show dhcp server binding ge-1/0/0:100-200
IP address      Session Id  Hardware address  Expires   State   Interface
192.168.0.17    42         00:10:94:00:00:02 86346     BOUND   ge-1/0/0.1073741827
192.168.0.16    41         00:10:94:00:00:01 86346     BOUND   ge-1/0/0.1073741827

```

show dhcp server binding <interfaces-wildcard>

```

user@host> show dhcp server binding ge-1/3/*
IP address      Session Id  Hardware address  Expires   State   Interface
192.168.0.9     24         00:10:94:00:00:04 86361     BOUND   ge-1/3/0.110
192.168.0.8     23         00:10:94:00:00:03 86361     BOUND   ge-1/3/0.110
192.168.0.7     22         00:10:94:00:00:02 86361     BOUND   ge-1/3/0.110

```

show dhcp server statistics

Syntax	show dhcp server statistics <logical-system <i>logical-system-name</i>> <routing-instance <i>routing-instance-name</i>>
Release Information	Command introduced in Junos OS Release 9.0.
Description	Display extended Dynamic Host Configuration Protocol (DHCP) local server statistics.
Options	logical-system <i>logical-system-name</i> —(Optional) Display information about extended DHCP local server statistics on the specified logical system. If you do not specify a logical system, statistics are displayed for the default logical system. routing-instance <i>routing-instance-name</i> —(Optional) Display information about extended DHCP local server statistics on the specified routing instance. If you do not specify a routing instance, statistics are displayed for the default routing instance.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear dhcp server statistics on page 1163
List of Sample Output	show dhcp server statistics on page 1183
Output Fields	Table 81 on page 1183 lists the output fields for the show dhcp server statistics command. Output fields are listed in the approximate order in which they appear.

Table 81: show dhcp server statistics Output Fields

Field Name	Field Description
Packets dropped	<p>Number of packets discarded by the extended DHCP local server because of errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.</p> <ul style="list-style-type: none"> • Total—Total number of packets discarded by the extended DHCP local server • Authentication—Number of packets discarded because they could not be authenticated • Bad hardware address—Number of packets discarded because an invalid hardware address was specified • Bad opcode—Number of packets discarded because an invalid operation code was specified • Bad options—Number of packets discarded because invalid options were specified • Dynamic profile—Number of packets discarded due to dynamic profile information • Invalid server address—Number of packets discarded because an invalid server address was specified • No available addresses—Number of packets discarded because there were no addresses available for assignment • No interface match—Number of packets discarded because they did not belong to a configured interface • No routing instance match—Number of packets discarded because they did not belong to a configured routing instance • No valid local address—Number of packets discarded because there was no valid local address • Packet too short—Number of packets discarded because they were too short • Read error—Number of packets discarded because of a system read error • Send error—Number of packets that the extended DHCP local server could not send
Messages received	<p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> • BOOTREQUEST—Number of BOOTP protocol data units (PDUs) received • DHCPDECLINE—Number of DHCP PDUs of type DECLINE received • DHCPDISCOVER—Number of DHCP PDUs of type DISCOVER received • DHCPINFORM—Number of DHCP PDUs of type INFORM received • DHCPRELEASE—Number of DHCP PDUs of type RELEASE received • DHCPREQUEST—Number of DHCP PDUs of type REQUEST received
Messages sent	<p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> • BOOTREPLY—Number of BOOTP PDUs transmitted • DHCPOFFER—Number of DHCP OFFER PDUs transmitted • DHCPACK—Number of DHCP ACK PDUs transmitted • DHCPNACK—Number of DHCP NACK PDUs transmitted • DHCPFORCERENEW—Number of DHCP FORCERENEW PDUs transmitted

Sample Output

show dhcp server statistics

```

user@host> show dhcp server statistics
Packets dropped:
    Total                0

Messages received:

```

BOOTREQUEST	25
DHCPDECLINE	0
DHCPDISCOVER	10
DHCPINFORM	0
DHCPRELEASE	4
DHCPREQUEST	10

Messages sent:

BOOTREPLY	20
DHCPOFFER	10
DHCPACK	10
DHCPNAK	0
DHCPFORCERENEW	0

show dhcpv6 server binding

Syntax	<pre>show dhcpv6 server binding <address> <brief detail summary> <interface interface-name> <interfaces-vlan> <interfaces-wildcard> <logical-system logical-system-name> <routing-instance routing-instance-name></pre>
Release Information	<p>Command introduced in Junos OS Release 9.6.</p> <p>Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.</p>
Description	Display the address bindings in the client table on the extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server.
Options	<p>address—(Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:</p> <ul style="list-style-type: none"> • <i>CID</i>—The specified Client ID (CID). • <i>ipv6-prefix</i>—The specified IPv6 prefix. • <i>session-id</i>—The specified session ID. <p>brief detail summary—(Optional) Display the specified level of output about active client bindings. The default is brief, which produces the same output as show dhcpv6 server binding.</p> <p>interface interface-name—(Optional) Display information about active client bindings on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.</p> <p>interfaces-vlan—(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.</p> <p>interfaces-wildcard—(Optional) The set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).</p> <p>logical-system logical-system-name—(Optional) Display information about active client bindings for DHCPv6 clients on the specified logical system.</p> <p>routing-instance routing-instance-name—(Optional) Display information about active client bindings for DHCPv6 clients on the specified routing instance.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Clearing DHCP Bindings for Subscriber Access</i> • clear dhcpv6 server binding on page 1165

List of Sample Output

- [show dhcpv6 server binding on page 1187](#)
- [show dhcpv6 server binding detail on page 1187](#)
- [show dhcpv6 server binding interface on page 1188](#)
- [show dhcpv6 server binding interface detail on page 1188](#)
- [show dhcpv6 server binding <prefix> on page 1188](#)
- [show dhcpv6 server binding <session-id> on page 1189](#)
- [show dhcpv6 server binding <interfaces-vlan> on page 1189](#)
- [show dhcpv6 server binding <interfaces-wildcard> on page 1189](#)
- [show dhcpv6 server binding <interfaces-wildcard> on page 1189](#)
- [show dhcpv6 server binding summary on page 1189](#)

Output Fields [Table 82 on page 1186](#) lists the output fields for the **show dhcpv6 server binding** command. Output fields are listed in the approximate order in which they appear.

Table 82: show dhcpv6 server binding Output Fields

Field Name	Field Description	Level of Output
<i>number clients</i> , (<i>number init</i> , <i>number bound</i> , <i>number selecting</i> , <i>number requesting</i> , <i>number renewing</i> , <i>number releasing</i>)	Summary counts of the total number of DHCPv6 clients and the number of DHCPv6 clients in each state.	summary
Prefix	Client's DHCPv6 prefix, or prefix used to support multiple address assignment.	brief detail
Session Id	Session ID of the subscriber session.	brief detail
Expires	Number of seconds in which lease expires.	brief detail
State	State of the address binding table on the extended DHCPv6 local server: <ul style="list-style-type: none"> • BOUND—Client has active IP address lease. • INIT—Initial state. • RECONFIGURE—Server has sent reconfigure message to client. • RELEASE—Client is releasing IP address lease. • RENEWING—Client sending request to renew IP address lease. • REQUESTING—Client requesting a DHCPv6 server. • SELECTING—Client receiving offers from DHCPv6 servers. 	brief detail
Interface	Interface on which the DHCPv6 request was received.	brief
Client IPv6 Address	Client's IPv6 address.	detail
Client IPv6 Prefix	Client's IPv6 prefix.	detail
Client DUID	Client's DHCP Unique Identifier (DUID).	brief detail

Table 82: show dhcpv6 server binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
Lease expires	Date and time at which the client's IP address lease expires.	detail
Lease expires in	Number of seconds in which lease expires.	detail
Lease Start	Date and time at which the client's address lease was obtained.	detail
Incoming Client Interface	Client's incoming interface.	detail
Server IP Address	IP address of DHCPv6 server.	detail
Server Interface	Interface of DHCPv6 server.	detail
Client Pool Name	Address pool used to assign IPv6 address.	detail
Client Prefix Pool Name	Address pool used to assign IPv6 prefix.	detail
Client Id length	Length of the DHCPv6 client ID, in bytes.	detail
Client Id	ID of the DHCPv6 client.	detail

Sample Output

show dhcpv6 server binding

```

user@host> show dhcpv6 server binding
Prefix          Session Id Expires State Interface Client DUID
2001:bd8:1111:2222::/64 6      86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:01
2001:bd8:1111:2222::/64 7      86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
2001:bd8:1111:2222::/64 8      86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03
2001:bd8:1111:2222::/64 9      86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:04
2001:bd8:1111:2222::/64 10     86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:05
2002::1/74      11      86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:06

```

show dhcpv6 server binding detail

```

user@host> show dhcpv6 server binding detail
Session Id: 6
  Client IPv6 Prefix:      2001:bd8:1111:2222::/64
  Client DUID:             LL_TIME0x1-0x2e159c0-00:10:94:00:00:01

  State:
  BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
  Lease Expires:           2009-07-21 10:41:15 PDT
  Lease Expires in:        86308 seconds

```

```

Lease Start:                2009-07-20 10:41:15 PDT
Incoming Client Interface:  ge-1/0/0.0
Server Ip Address:         0.0.0.0
Server Interface:         none
Client Id Length:         14
Client Id:
/0x00010001/0x02e159c0/0x00109400/0x0001

Session Id:  7
Client IPv6 Address:      2002::1/128
Client IPv6 Prefix:      2001:bd8:1111:2222::/64
Client DUID:              LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

State:
BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
Lease Expires:            2009-07-21 10:41:15 PDT
Lease Expires in:        86308 seconds
Lease Start:             2009-07-20 10:41:15 PDT
Incoming Client Interface: ge-1/0/0.0
Server Ip Address:       0.0.0.0
Client Pool Name:        bos-v6-pool
Client Prefix Pool Name: bos-v6-prefix-pool
Client Id Length:        14
Client Id:
/0x00010001/0x02e159c0/0x00109400/0x0002

```

show dhcpv6 server binding interface

```

user@host> show dhcpv6 server binding interface ge-1/0/0:10-101
Prefix          Session Id Expires State Interface Client DUID
2001:bd8:1111:2222::/64 1      86055   BOUND   ge-1/0/0.100
LL_TIME0x1-0x4b0a53b9-00:10:94:00:00:01

```

show dhcpv6 server binding interface detail

```

user@host> show dhcpv6 server binding interface ge-1/0/0:10-101 detail
Session Id:  7
Client IPv6 Prefix:      2001:bd8:1111:2222::/64
Client DUID:              LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

State:
BOUND(bound)
Lease Expires:            2009-07-21 10:41:15 PDT
Lease Expires in:        86136 seconds
Lease Start:             2009-07-20 10:41:15 PDT
Incoming Client Interface: ge-1/0/0.0
Server Ip Address:       0.0.0.0
Server Interface:       none
Client Id Length:        14
Client Id:
/0x00010001/0x02e159c0/0x00109400/0x0002

```

show dhcpv6 server binding <prefix>

```

user@host> show dhcpv6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005
detail
Session Id:  7
Client IPv6 Prefix:      2001:bd8:1111:2222::/64
Client DUID:              LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

State:
BOUND(bound)
Lease Expires:            2009-07-21 10:41:15 PDT
Lease Expires in:        86136 seconds

```

```

Lease Start:                2009-07-20 10:41:15 PDT
Incoming Client Interface:  ge-1/0/0.0
Server Ip Address:          0.0.0.0
Server Interface:           none
Client Id Length:          14
Client Id:
/0x00010001/0x02e159c0/0x00109400/0x0002

```

show dhcpv6 server binding <session-id>

```

user@host> show dhcpv6 server binding 8
Prefix          Session Id Expires State Interface Client DUID
2001:DB8::/32   8          86235  BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03

```

show dhcpv6 server binding <interfaces-vlan>

```

user@host> show dhcpv6 server binding ge-1/0/0:100-200
Prefix          Session Id Expires State Interface Client DUID
2001:DB8::/32   11          87583  BOUND ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32   12          87583  BOUND ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

show dhcpv6 server binding <interfaces-wildcard>

```

user@host> show dhcpv6 server binding demux0
Prefix          Session Id Expires State Interface Client DUID
2001:DB8::/32   30          79681  BOUND demux0.1073741824
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32   31          79681  BOUND demux0.1073741825
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:CB9::/32   32          79681  BOUND demux0.1073741826
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

show dhcpv6 server binding <interfaces-wildcard>

```

user@host> show dhcpv6 server binding ge-1/3/*
Prefix          Session Id Expires State Interface Client DUID
2001:DB8::/32   22          79681  BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32   33          79681  BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:CB9::/32   24          79681  BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

show dhcpv6 server binding summary

```

user@host> show dhcpv6 server binding summary
5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)

```

show dhcpv6 server statistics

Syntax	show dhcpv6 server statistics <logical-system <i>logical-system-name</i>> <routing-instance <i>routing-instance-name</i>>
Release Information	Command introduced in Junos OS Release 9.6.
Description	Display extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.
Options	logical-system <i>logical-system-name</i> —(Optional) Display information about extended DHCPv6 local server statistics on the specified logical system. If you do not specify a logical system, statistics are displayed for the default logical system. routing-instance <i>routing-instance-name</i> —(Optional) Display information about extended DHCPv6 local server statistics on the specified routing instance. If you do not specify a routing instance, statistics are displayed for the default routing instance.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear dhcpv6 server statistics on page 1167
List of Sample Output	show dhcpv6 server statistics on page 1191
Output Fields	Table 83 on page 1191 lists the output fields for the show dhcpv6 server statistics command. Output fields are listed in the approximate order in which they appear.

Table 83: show dhcpv6 server statistics Output Fields

Field Name	Field Description
Packets dropped	<p>Number of packets discarded by the extended DHCPv6 local server because of errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.</p> <ul style="list-style-type: none"> • Total—Total number of packets discarded by the extended DHCPv6 local server • Strict Reconfigure—Number of solicit messages discarded because the client does not support reconfiguration • Bad hardware address—Number of packets discarded because an invalid hardware address was specified • Bad opcode—Number of packets discarded because an invalid operation code was specified • Bad options—Number of packets discarded because invalid options were specified • Invalid server address—Number of packets discarded because an invalid server address was specified • No available addresses—Number of packets discarded because there were no addresses available for assignment • No interface match—Number of packets discarded because they did not belong to a configured interface • No routing instance match—Number of packets discarded because they did not belong to a configured routing instance • No valid local address—Number of packets discarded because there was no valid local address • Packet too short—Number of packets discarded because they were too short • Read error—Number of packets discarded because of a system read error • Send error—Number of packets that the extended DHCPv6 local server could not send
Messages received	<p>Number of DHCPv6 messages received.</p> <ul style="list-style-type: none"> • DHCPV6_CONFIRM—Number of DHCPv6 CONFIRM PDUs received. • DHCPV6_DECLINE—Number of DHCPv6 DECLINE PDUs received. • DHCPV6_INFORMATION_REQUEST—Number of DHCPv6 INFORMATION-REQUEST PDUs received. • DHCPV6_REBIND—Number of DHCPv6 REBIND PDUs received. • DHCPV6_RELAY_FORW—Number of DHCPv6 RELAY-FORW PDUs received. • DHCPV6_RELAY_REPL—Number of DHCPv6 RELAY-REPL PDUs received. • DHCPV6_RELEASE—Number of DHCPv6 RELEASE PDUs received. • DHCPV6_RENEW—Number of DHCPv6 RENEW PDUs received. • DHCPV6_REQUEST—Number of DHCPv6 REQUEST PDUs received. • DHCPV6_SOLICIT—Number of DHCPv6 SOLICIT PDUs received.
Messages sent	<p>Number of DHCPv6 messages sent.</p> <ul style="list-style-type: none"> • DHCPV6_ADVERTISE—Number of DHCPv6 ADVERTISE PDUs transmitted. • DHCPV6_REPLY—Number of DHCPv6 ADVERTISE PDUs transmitted. • DHC6_RECONFIGURE—Number of DHCPv6 RECONFIGURE PDUs transmitted.

Sample Output

show dhcpv6 server statistics

```
user@host> show dhcpv6 server statistics
```

```
Dhcpv6 Packets dropped:
  Total                0

Messages received:
  DHCPV6_DECLINE      0
  DHCPV6_SOLICIT      9
  DHCPV6_INFORMATION_REQUEST 0
  DHCPV6_RELEASE      0
  DHCPV6_REQUEST      5
  DHCPV6_CONFIRM      0
  DHCPV6_RENEW        0
  DHCPV6_REBIND       0
  DHCPV6_RELAY_FORW   0
  DHCPV6_RELAY_REPL   0

Messages sent:
  DHCPV6_ADVERTISE    9
  DHCPV6_REPLY        5
  DHCPV6_RECONFIGURE  0
```

DHCP Relay Agent Monitoring Commands

clear dhcp relay binding

Syntax	<pre>clear dhcp relay binding <address> <all> <interface interface-name> <interfaces-vlan> <interfaces-wildcard> <logical-system logical-system-name> <routing-instance routing-instance-name></pre>
Release Information	<p>Command introduced in Junos OS Release 8.3.</p> <p>Options all and interface added in Junos OS Release 8.4.</p> <p>Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches.</p>
Description	Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table.
Options	<p>address—(Optional) Clear the binding state for the DHCP client, using one of the following entries:</p> <ul style="list-style-type: none"> <i>ip-address</i>—The specified IP address. <i>mac-address</i>—The specified MAC address. <i>session-id</i>—The specified session ID. <p>all—(Optional) Clear the binding state for all DHCP clients.</p> <p>interface interface-name—(Optional) Clear the binding state for DHCP clients on the specified interface.</p> <p>interfaces-vlan—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.</p> <p>interfaces-wildcard—(Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).</p> <p>logical-system logical-system-name—(Optional) Clear the binding state for DHCP clients on the specified logical system.</p> <p>routing-instance routing-instance-name—(Optional) Clear the binding state for DHCP clients on the specified routing instance.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> <i>Clearing DHCP Bindings for Subscriber Access</i> show dhcp relay binding on page 1203

List of Sample Output [clear dhcp relay binding on page 1194](#)
[clear dhcp relay binding all on page 1194](#)
[clear dhcp relay binding interface on page 1194](#)
[clear dhcp relay binding <interfaces-vlan> on page 1194](#)
[clear dhcp relay binding <interfaces-wildcard> on page 1194](#)

Output Fields See [show dhcp relay binding](#) for an explanation of output fields.

Sample Output

clear dhcp relay binding

The following sample output displays the address bindings in the DHCP client table before and after the **clear dhcp relay binding** command is issued.

```
user@host> show dhcp relay binding
IP address      Hardware address  Type    Lease expires at
100.20.32.1     90:00:00:01:00:01 active    2007-02-08 16:41:17 EST
192.168.14.8    90:00:01:01:02:01 active    2007-02-10 10:01:06 EST
```

```
user@host> clear dhcp relay binding 100.20.32.1
```

```
user@host> show dhcp relay binding
IP address      Hardware address  Type    Lease expires at
192.168.14.8    90:00:01:01:02:01 active    2007-02-10 10:01:06 EST
```

clear dhcp relay binding all

The following command clears all DHCP relay agent bindings:

```
user@host> clear dhcp relay binding all
```

clear dhcp relay binding interface

The following command clears DHCP relay agent bindings on a specific interface:

```
user@host> clear dhcp relay binding interface fe-0/0/3
```

clear dhcp relay binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCP relay agent bindings on top of the underlying interface **ae0**, which clears DHCP bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcp relay binding interface ae0
```

clear dhcp relay binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCP relay agent bindings over a specific interface:

```
user@host> clear dhcp relay binding ge-1/0/0.*
```


clear dhcp relay statistics

Syntax	<pre>clear dhcp relay statistics <logical-system <i>logical-system-name</i>> <routing-instance <i>routing-instance-name</i>></pre>
Syntax	<p>Syntax for EX Series switches:</p> <pre>show dhcp relay statistics <routing-instance <i>routing-instance-name</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches.</p>
Description	Clear all Dynamic Host Configuration Protocol (DHCP) relay statistics.
Options	<p>logical-system <i>logical-system-name</i>—(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are cleared for the default logical system.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show dhcp relay statistics on page 1208
List of Sample Output	clear dhcp relay statistics on page 1196
Output Fields	Table 84 on page 1196 lists the output fields for the clear dhcp relay statistics command.

Table 84: clear dhcp relay statistics Output Fields

Field Name	Field Description
Packets dropped	<p>Number of packets discarded by the extended DHCP relay agent application due to errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.</p> <ul style="list-style-type: none"> • Total—Total number of packets discarded by the extended DHCP relay agent application. • Bad hardware address—Number of packets discarded because an invalid hardware address was specified. • Bad opcode—Number of packets discarded because an invalid operation code was specified. • Bad options—Number of packets discarded because invalid options were specified. • Invalid server address—Number of packets discarded because an invalid server address was specified. • No available addresses—Number of packets discarded because there were no addresses available for assignment. • No interface match—Number of packets discarded because they did not belong to a configured interface. • No routing instance match—Number of packets discarded because they did not belong to a configured routing instance. • No valid local address—Number of packets discarded because there was no valid local address. • Packet too short—Number of packets discarded because they were too short. • Read error—Number of packets discarded because of a system read error. • Send error—Number of packets that the extended DHCP relay application could not send. • Option 60—Number of packets discarded containing DHCP option 60 vendor-specific information. • Option 82—Number of packets discarded because DHCP option 82 information could not be added.
Messages received	<p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> • BOOTREQUEST—Number of BOOTP protocol data units (PDUs) received • DHCPDECLINE—Number of DHCP PDUs of type DECLINE received • DHCPDISCOVER—Number of DHCP PDUs of type DISCOVER received • DHCPINFORM—Number of DHCP PDUs of type INFORM received • DHCPRELEASE—Number of DHCP PDUs of type RELEASE received • DHCPREQUEST—Number of DHCP PDUs of type REQUEST received
Messages sent	<p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> • BOOTREPLY—Number of BOOTP PDUs transmitted • DHCPOFFER—Number of DHCP OFFER PDUs transmitted • DHCPACK—Number of DHCP ACK PDUs transmitted • DHCPNACK—Number of DHCP NACK PDUs transmitted

Sample Output

clear dhcp relay statistics

The following sample output displays the DHCP relay statistics before and after the **clear dhcp relay statistics** command is issued.

```
user@host> show dhcp relay statistics
```

```
Packets dropped:
  Total                0

Messages received:
  BOOTREQUEST          116
  DHCPDECLINE           0
  DHCPDISCOVER          11
  DHCPINFORM            0
  DHCPRELEASE           0
  DHCPREQUEST          105

Messages sent:
  BOOTREPLY             44
  DHCPOFFER             11
  DHCPACK               11
  DHCPNAK               11
```

```
user@host> clear dhcp relay statistics
```

```
user@host> show dhcp relay statistics
```

```
Packets dropped:
  Total                0

Messages received:
  BOOTREQUEST          0
  DHCPDECLINE           0
  DHCPDISCOVER          0
  DHCPINFORM            0
  DHCPRELEASE           0
  DHCPREQUEST          0

Messages sent:
  BOOTREPLY             0
  DHCPOFFER             0
  DHCPACK               0
  DHCPNAK               0
```

clear dhcpv6 relay binding

Syntax	<code>clear dhcpv6 relay binding</code> <code><address></code> <code><all></code> <code><interface <i>interface-name</i>></code> <code><interfaces-vlan></code> <code><interfaces-wildcard></code> <code><logical-system <i>logical-system-name</i>></code> <code><routing-instance <i>routing-instance-name</i>></code>
Release Information	Command introduced in Junos OS Release 11.4. Command introduced in Junos OS Release 12.3R2 for EX Series switches. Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1. Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches.
Description	Clear the binding state of Dynamic Host Configuration Protocol for IPv6 (DHCPv6) clients from the client table.
Options	<p><i>address</i>—(Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:</p> <ul style="list-style-type: none">• <i>CID</i>—The specified Client ID (CID).• <i>ipv6-prefix</i>—The specified IPv6 prefix.• <i>session-id</i>—The specified session ID. <p><i>all</i>—(Optional) Clear the binding state for all DHCPv6 clients.</p> <p><i>interfaces-vlan</i>—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.</p> <p><i>interfaces-wildcard</i>—(Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).</p> <p><i>interface interface-name</i>—(Optional) Clear the binding state for DHCPv6 clients on the specified interface.</p> <p><i>logical-system logical-system-name</i>—(Optional) Clear the binding state for DHCPv6 clients on the specified logical system.</p> <p><i>routing-instance routing-instance-name</i>—(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• <i>Clearing DHCP Bindings for Subscriber Access</i>• show dhcpv6 relay binding on page 1211

- List of Sample Output** [clear dhcpv6 relay binding on page 1199](#)
[clear dhcpv6 relay binding <prefix> on page 1199](#)
[clear dhcpv6 relay binding all on page 1199](#)
[clear dhc6p relay binding interface on page 1199](#)
[clear dhcpv6 relay binding <interfaces-vlan> on page 1200](#)
[clear dhcpv6 relay binding <interfaces-wildcard> on page 1200](#)
- Output Fields** See [show dhcpv6 relay binding](#) for an explanation of output fields.

Sample Output

clear dhcpv6 relay binding

The following sample output displays the DHCPv6 bindings before and after the **clear dhcpv6 relay binding** command is issued.

```
user@host> show dhcpv6 relay binding
```

Prefix	Session Id	Expires	State	Interface	Client DUID
2001:bd8:3c4d:15::/64	1	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01					
2001:bd8:3c4d:16::/64	2	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02					
2001:bd8:3c4d:17::/64	3	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03					
2001:bd8:3c4d:18::/64	4	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04					
2001:bd8:3c4d:19::/64	5	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05					
2001:bd8:3c4d:20::/64	6	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06					

clear dhcpv6 relay binding <prefix>

```
user@host> clear dhcpv6 relay binding 2001:bd8:3c4d:15::/64
```

```
user@host> show dhcpv6 relay binding
```

Prefix	Session Id	Expires	State	Interface	Client DUID
2001:bd8:3c4d:16::/64	2	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02					
2001:bd8:3c4d:17::/64	3	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03					
2001:bd8:3c4d:18::/64	4	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04					
2001:bd8:3c4d:19::/64	5	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05					
2001:bd8:3c4d:20::/64	6	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06					

clear dhcpv6 relay binding all

The following command clears all DHCP relay agent bindings:

```
user@host> clear dhcpv6 relay binding all
```

clear dhc6p relay binding interface

The following command clears DHCPv6 relay agent bindings on a specific interface:

```
user@host> clear dhcpv6 relay binding interface fe-0/0/2
```

clear dhcpv6 relay binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCPv6 relay agent bindings on top of the underlying interface **ae0**, which clears DHCPv6 bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcpv6 relay binding interface ae0
```

clear dhcpv6 relay binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCPv6 relay agent bindings over a specific interface:

```
user@host> clear dhcpv6 relay binding ge-1/0/0.*
```

clear dhcpv6 relay statistics

Syntax	<code>clear dhcpv6 relay statistics</code> <code><logical-system <i>logical-system-name</i>></code> <code><routing-instance <i>routing-instance-name</i>></code>
Release Information	Command introduced in Junos OS Release 11.4. Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches. Command introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Clear all Dynamic Host Configuration Protocol for IPv6 (DHCPv6) relay statistics.
Options	logical-system <i>logical-system-name</i> —(Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are cleared for the default logical system. routing-instance <i>routing-instance-name</i> —(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance.
Required Privilege Level	view
List of Sample Output	clear dhcpv6 relay statistics on page 1201
Output Fields	See show dhcpv6 relay statistics for an explanation of output fields.

Sample Output

clear dhcpv6 relay statistics

The following sample output displays the DHCPv6 relay statistics before and after the **clear dhcpv6 relay statistics** command is issued.

```

user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
    Total                               0

Messages received:
    DHCPV6_DECLINE                      0
    DHCPV6_SOLICIT                      10
    DHCPV6_INFORMATION_REQUEST          0
    DHCPV6_RELEASE                      0
    DHCPV6_REQUEST                      10
    DHCPV6_CONFIRM                      0
    DHCPV6_RENEW                        0
    DHCPV6_REBIND                       0
    DHCPV6_RELAY_REPL                   0

Messages sent:
    DHCPV6_ADVERTISE                    0
    DHCPV6_REPLY                        0
    DHCPV6_RECONFIGURE                  0
    DHCPV6_RELAY_FORW                   0

```

```
user@host> clear dhcpv6 relay statistics
user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
    Total                                0

Messages received:
    DHCPV6_DECLINE                      0
    DHCPV6_SOLICIT                      0
    DHCPV6_INFORMATION_REQUEST          0
    DHCPV6_RELEASE                      0
    DHCPV6_REQUEST                      0
    DHCPV6_CONFIRM                      0
    DHCPV6_RENEW                        0
    DHCPV6_REBIND                       0
    DHCPV6_RELAY_REPL                  0

Messages sent:
    DHCPV6_ADVERTISE                    0
    DHCPV6_REPLY                        0
    DHCPV6_RECONFIGURE                  0
    DHCPV6_RELAY_FORW                   0
```


show dhcp relay binding

Syntax	<pre> show dhcp relay binding <address> <brief> <detail> <interface interface-name> <interfaces-vlan> <interfaces-wildcard> <ip-address mac-address> <logical-system logical-system-name> <routing-instance routing-instance-name> <summary> </pre>
Release Information	<p>Command introduced in Junos OS Release 8.3.</p> <p>Options interface and mac-address added in Junos OS Release 8.4.</p> <p>Options interfaces-vlan and interfaces-wildcard added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches.</p>
Description	Display the address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.
Options	<p>address—(Optional) Display DHCP binding information for a specific client identified by one of the following entries:</p> <ul style="list-style-type: none"> • ip-address—The specified IP address. • mac-address—The specified MAC address. • session-id—The specified session ID. <p>brief—(Optional) Display brief information about the active client bindings. This is the default, and produces the same output as show dhcp relay binding.</p> <p>detail—(Optional) Display detailed client binding information.</p> <p>interface interface-name—(Optional) Perform this operation on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.</p> <p>interfaces-vlan—(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.</p> <p>interfaces-wildcard—(Optional) The set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).</p> <p>logical-system logical-system-name—(Optional) Perform this operation on the specified logical system.</p> <p>routing-instance routing-instance-name—(Optional) Perform this operation on the specified routing instance.</p> <p>summary—(Optional) Display a summary of DHCP client information.</p>

Required Privilege Level view

Related Documentation

- *Clearing DHCP Bindings for Subscriber Access*
- [clear dhcp relay binding on page 1193](#)

List of Sample Output

- [show dhcp relay binding on page 1205](#)
- [show dhcp relay binding detail on page 1205](#)
- [show dhcp relay binding interface on page 1206](#)
- [show dhcp relay binding interface vlan-id on page 1206](#)
- [show dhcp relay binding interface svlan-id on page 1206](#)
- [show dhcp relay binding ip-address on page 1206](#)
- [show dhcp relay binding mac-address on page 1206](#)
- [show dhcp relay binding session-id on page 1206](#)
- [show dhcp relay binding <interfaces-vlan> on page 1207](#)
- [show dhcp relay binding <interfaces-wildcard> on page 1207](#)
- [show dhcp relay binding summary on page 1207](#)

Output Fields Table 85 on page 1204 lists the output fields for the **show dhcp relay binding** command. Output fields are listed in the approximate order in which they appear.

Table 85: show dhcp relay binding Output Fields

Field Name	Field Description	Level of Output
<i>number</i> clients, (<i>number</i> init, <i>number</i> bound, <i>number</i> selecting, <i>number</i> requesting, <i>number</i> renewing, <i>number</i> rebinding, <i>number</i> releasing)	Summary counts of the total number of DHCP clients and the number of DHCP clients in each state.	summary
IP address	IP address of the DHCP client.	briefdetail
Session Id	Session ID of the subscriber session.	briefdetail
Hardware address	Hardware address of the DHCP client.	briefdetail
Expires	Number of seconds in which the lease expires.	briefdetail
State	State of the DHCP relay address binding table on the DHCP client: <ul style="list-style-type: none"> • BOUND—Client has an active IP address lease. • INIT—Initial state. • REBINDING—Client is broadcasting a request to renew the IP address lease. • RELEASE—Client is releasing the IP address lease. • RENEWING—Client is sending a request to renew the IP address lease. • REQUESTING—Client is requesting a DHCP server. • SELECTING—Client is receiving offers from DHCP servers. 	briefdetail

Table 85: show dhcp relay binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
Interface	Incoming client interface.	brief
Lease Expires	Date and time at which the client's IP address lease expires.	detail
Lease Expires in	Number of seconds in which the lease expires.	detail
Lease Start	Date and time at which the client's IP address lease started.	detail
Incoming Client Interface	Client's incoming interface.	detail
Server IP Address	IP address of the DHCP server.	detail
Server Interface	Interface of the DHCP server.	detail
Bootp Relay Address	IP address of BOOTP relay.	detail
Type	Type of DHCP packet processing performed on the router: <ul style="list-style-type: none"> active—Router actively processes and relays DHCP packets. passive—Router passively snoops DHCP packets passing through the router. 	All levels
Lease expires at	Date and time at which the client's IP address lease expires.	All levels

Sample Output

show dhcp relay binding

```

user@host> show dhcp relay binding
IP address      Session Id  Hardware address  Expires    State    Interface
100.20.32.11    41          00:10:94:00:00:01 86371      BOUND    ge-1/0/0.0
100.20.32.12    42          00:10:94:00:00:02 86371      BOUND    ge-1/0/0.0
100.20.32.13    43          00:10:94:00:00:03 86371      BOUND    ge-1/0/0.0
100.20.32.14    44          00:10:94:00:00:04 86371      BOUND    ge-1/0/0.0
100.20.32.15    45          00:10:94:00:00:05 86371      BOUND    ge-1/0/0.0

```

show dhcp relay binding detail

```

user@host> show dhcp relay binding detail

Client IP Address: 100.20.32.11
Hardware Address:   00:10:94:00:00:01
State:              BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
Lease Expires:      2009-07-21 11:00:06 PDT
Lease Expires in:   86361 seconds

```

```

Lease Start:                2009-07-20 11:00:06 PDT
Last Packet Received:       2009-07-20 11:00:06 PDT
Incoming Client Interface:   ge-1/0/0.0
Server Ip Address:          100.20.22.2
Server Interface:           none
Bootp Relay Address:        100.20.32.2
Session Id:                 41

```

```

Client IP Address: 100.20.32.12
Hardware Address:   00:10:94:00:00:02
State:              BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
Lease Expires:      2009-07-21 11:00:06 PDT
Lease Expires in:   86361 seconds
Lease Start:        2009-07-20 11:00:06 PDT
Last Packet Received: 2009-07-20 11:00:06 PDT
Incoming Client Interface: ge-1/0/0.0
Server Ip Address:   100.20.22.2
Server Interface:    none
Bootp Relay Address: 100.20.32.2
Session Id:          42

```

show dhcp relay binding interface

```
user@host> show dhcp relay binding interface fe-0/0/2
```

IP address	Hardware address	Type	Lease expires at
100.20.32.1	90:00:00:01:00:01	active	2007-03-27 15:06:20 EDT

show dhcp relay binding interface vlan-id

```
user@host> show dhcp relay binding interface ge-1/1/0:100
```

IP address	Session Id	Hardware address	Expires	State	Interface
200.20.20.15	6	00:10:94:00:00:01	86124	BOUND	ge-1/1/0:100

show dhcp relay binding interface svlan-id

```
user@host> show dhcp relay binding interface ge-1/1/0:10-100
```

IP address	Session Id	Hardware address	Expires	State	Interface
200.20.20.16	7	00:10:94:00:00:02	86124	BOUND	ge-1/1/0:10-100

show dhcp relay binding ip-address

```
user@host> show dhcp relay binding 100.20.32.13
```

IP address	Session Id	Hardware address	Expires	State	Interface
100.20.32.13	43	00:10:94:00:00:03	86293	BOUND	ge-1/0/0.0

show dhcp relay binding mac-address

```
user@host> show dhcp relay binding 00:10:94:00:00:05
```

IP address	Session Id	Hardware address	Expires	State	Interface
100.20.32.15	45	00:10:94:00:00:05	86279	BOUND	ge-1/0/0.0

show dhcp relay binding session-id

```
user@host> show dhcp relay binding 41
```

IP address	Session Id	Hardware address	Expires	State	Interface
100.20.32.11	41	00:10:94:00:00:01	86305	BOUND	ge-1/0/0.0

show dhcp relay binding <interfaces-vlan>

```
user@host> show dhcp relay binding ge-1/0/0:100-200
```

IP address	Session Id	Hardware address	Expires	State	Interface
192.168.0.17	42	00:10:94:00:00:02	86346	BOUND	ge-1/0/0.1073741827
192.168.0.16	41	00:10:94:00:00:01	86346	BOUND	ge-1/0/0.1073741827

show dhcp relay binding <interfaces-wildcard>

```
user@host> show dhcp relay binding ge-1/3/*
```

IP address	Session Id	Hardware address	Expires	State	Interface
192.168.0.9	24	00:10:94:00:00:04	86361	BOUND	ge-1/3/0.110
192.168.0.8	23	00:10:94:00:00:03	86361	BOUND	ge-1/3/0.110
192.168.0.7	22	00:10:94:00:00:02	86361	BOUND	ge-1/3/0.110

show dhcp relay binding summary

```
user@host> show dhcp relay binding summary
```

3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 rebinding, 0 releasing)

show dhcp relay statistics

Syntax	<code>show dhcp relay statistics</code> <code><logical-system <i>logical-system-name</i>></code> <code><routing-instance <i>routing-instance-name</i>></code>
Syntax	Syntax for EX Series switches: <code>show dhcp relay statistics</code> <code><routing-instance <i>routing-instance-name</i>></code>
Release Information	Command introduced in Junos OS Release 8.3. Command introduced in Junos OS Release 12.1 for EX Series switches. Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches.
Description	Display Dynamic Host Configuration Protocol (DHCP) relay statistics.
Options	<code>logical-system <i>logical-system-name</i></code> —(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system. <code>routing-instance <i>routing-instance-name</i></code> —(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear dhcp relay statistics on page 1195
List of Sample Output	show dhcp relay statistics on page 1210
Output Fields	Table 86 on page 1209 lists the output fields for the <code>show dhcp relay statistics</code> command. Output fields are listed in the approximate order in which they appear.

Table 86: show dhcp relay statistics Output Fields

Field Name	Field Description
Packets dropped	<p>Number of packets discarded by the extended DHCP relay agent application due to errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.</p> <ul style="list-style-type: none"> • Total—Total number of packets discarded by the extended DHCP relay agent application. • Bad hardware address—Number of packets discarded because an invalid hardware address was specified. • Bad opcode—Number of packets discarded because an invalid operation code was specified. • Bad options—Number of packets discarded because invalid options were specified. • Invalid server address—Number of packets discarded because an invalid server address was specified. • No available addresses—Number of packets discarded because there were no addresses available for assignment. • No interface match—Number of packets discarded because they did not belong to a configured interface. • No routing instance match—Number of packets discarded because they did not belong to a configured routing instance. • No valid local address—Number of packets discarded because there was no valid local address. • Packet too short—Number of packets discarded because they were too short. • Read error—Number of packets discarded because of a system read error. • Send error—Number of packets that the extended DHCP relay application could not send. • Option 60—Number of packets discarded containing DHCP option 60 vendor-specific information. • Option 82—Number of packets discarded because DHCP option 82 information could not be added.
Messages received	<p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> • BOOTREQUEST—Number of BOOTP protocol data units (PDUs) received • DHCPDECLINE—Number of DHCP PDUs of type DECLINE received • DHCPDISCOVER—Number of DHCP PDUs of type DISCOVER received • DHCPINFORM—Number of DHCP PDUs of type INFORM received • DHCPRELEASE—Number of DHCP PDUs of type RELEASE received • DHCPREQUEST—Number of DHCP PDUs of type REQUEST received
Messages sent	<p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> • BOOTREPLY—Number of BOOTP PDUs transmitted • DHCPOFFER—Number of DHCP OFFER PDUs transmitted • DHCPACK—Number of DHCP ACK PDUs transmitted • DHCPNACK—Number of DHCP NACK PDUs transmitted • DHCPFORCERENEW—Number of DHCP FORCERENEW PDUs transmitted
Packets forwarded	<p>Number of packets forwarded.</p> <ul style="list-style-type: none"> • BOOTREQUEST—Number of BOOTREQUEST protocol data units (PDUs) forwarded • BOOTREPLY—Number of BOOTREPLY protocol data units (PDUs) forwarded

Sample Output

show dhcp relay statistics

```
user@host> show dhcp relay statistics
Packets dropped:
  Total                30
  Bad hardware address 1
  Bad opcode            1
  Bad options           3
  Invalid server address 5
  No available addresses 1
  No interface match    2
  No routing instance match 9
  No valid local address 4
  Packet too short      2
  Read error            1
  Send error            1
  Option 60             1
  Option 82             2

Messages received:
  BOOTREQUEST          116
  DHCPDECLINE           0
  DHCPDISCOVER          11
  DHCPINFORM            0
  DHCPRELEASE           0
  DHCPREQUEST          105

Messages sent:
  BOOTREPLY             0
  DHCPOFFER             2
  DHCPACK               1
  DHCPNAK               0
  DHCPFORCERENEW        0

Packets forwarded:
  Total                4
  BOOTREQUEST           2
  BOOTREPLY             2
```


show dhcpv6 relay binding

Syntax	<pre>show dhcpv6 relay binding <address client-id session-id> <brief> <detail> <interface interface-name> <interfaces-vlan> <interfaces-wildcard> <logical-system logical-system-name> <routing-instance routing-instance-name> <summary></pre>
Release Information	<p>Command introduced in Junos OS Release 11.4.</p> <p>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches.</p> <p>Command introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p><i>interfaces-vlan</i> and <i>interfaces-wildcard</i> options introduced in Junos OS Release 12.1.</p>
Description	Display the DHCPv6 address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.
Options	<p>address—(Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:</p> <ul style="list-style-type: none"> • <i>CID</i>—The specified Client ID (CID). • <i>ipv6-prefix</i>—The specified IPv6 prefix. • <i>session-id</i>—The specified session ID. <p>brief—(Optional) Display brief information about the active client bindings. This is the default, and produces the same output as show dhcpv6 relay binding.</p> <p>detail—(Optional) Display detailed client binding information.</p> <p>interface interface-name—(Optional) Perform this operation on the specified interface. You can optionally filter on VLAN ID and S-VLAN ID.</p> <p>interfaces-vlan—(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.</p> <p>interfaces-wildcard—(Optional) The set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).</p> <p>logical-system logical-system-name—(Optional) Perform this operation on the specified logical system.</p> <p>routing-instance routing-instance-name—(Optional) Perform this operation on the specified routing instance.</p> <p>summary—(Optional) Display a summary of DHCPv6 client information.</p>

Required Privilege Level view

Related Documentation

- [Clearing DHCP Bindings for Subscriber Access](#)
- [clear dhcpv6 relay binding on page 1198](#)

List of Sample Output

- [show dhcpv6 relay binding on page 1213](#)
- [show dhcpv6 relay binding \(address\) on page 1213](#)
- [show dhcpv6 relay binding \(client-id\) on page 1214](#)
- [show dhcpv6 relay binding detail on page 1214](#)
- [show dhcpv6 relay binding detail \(Multi-Relay Topology\) on page 1215](#)
- [show dhcpv6 relay binding \(session-id\) on page 1215](#)
- [show dhcpv6 relay binding \(interfaces-vlan\) on page 1215](#)
- [show dhcpv6 relay binding \(interfaces-wildcard\) on page 1215](#)
- [show dhcpv6 relay binding \(interfaces-wildcard\) on page 1215](#)
- [show dhcpv6 relay binding summary on page 1215](#)

Output Fields Table 87 on page 1212 lists the output fields for the **show dhcpv6 relay binding** command. Output fields are listed in the approximate order in which they appear.

Table 87: show dhcpv6 relay binding Output Fields

Field Name	Field Description	Level of Output
<i>number</i> clients, (<i>number</i> init, <i>number</i> bound, <i>number</i> selecting, <i>number</i> requesting, <i>number</i> renewing, <i>number</i> rebinding, <i>number</i> releasing)	Summary counts of the total number of DHCPv6 clients and the number of DHCPv6 clients in each state.	summary
Client IPv6 Prefix	Prefix of the DHCPv6 client.	briefdetail
Client DUID	DHCP for IPv6 Unique Identifier (DUID) of the client.	briefdetail
Session Id	Session ID of the subscriber (DHCP client) session.	briefdetail
Expires	Number of seconds in which the lease expires.	briefdetail
State	State of the DHCPv6 relay address binding table on the DHCPv6 client: <ul style="list-style-type: none"> • BOUND—Client has an active IP address lease. • INIT—Initial state. • REBINDING—Client is broadcasting a request to renew the IP address lease. • RELEASE—Client is releasing the IP address lease. • RENEWING—Client is sending a request to renew the IP address lease. • REQUESTING—Client is requesting a DHCPv6 server. • SELECTING—Client is receiving offers from DHCPv6 servers. 	briefdetail
Interface	Incoming client interface.	brief

Table 87: show dhcpv6 relay binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
Lease Expires	Date and time at which the client's IP address lease expires.	detail
Lease Expires in	Number of seconds in which the lease expires.	detail
Lease Start	Date and time at which the client's IP address lease started.	detail
Incoming Client Interface	Client's incoming interface.	detail
Server Address	IP address of the DHCPv6 server. Displays unknown for a DHCPv6 relay agent in a multi-relay topology that is not directly adjacent to the DHCPv6 server and does not detect the IP address of the server. In that case, the output instead displays the Next Hop Server Facing Relay field.	detail
Next Hop Server Facing Relay	Next-hop address in the direction of the DHCPv6 server.	detail
Server Interface	Interface of the DHCPv6 server.	detail
Relay Address	IP address of the relay.	detail
Client Pool Name	Address pool that granted the client lease.	
Client ID Length	Length of client ID.	All levels
Client Id	Client ID.	All levels

Sample Output

show dhcpv6 relay binding

```

user@host> show dhcpv6 relay binding
Prefix                Session Id Expires State Interface Client DUID
2001:bd8:3c4d:15::/64 1 83720 BOUND ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
2001:bd8:3c4d:16::/64 2 83720 BOUND ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02
2001:bd8:3c4d:17::/64 3 83720 BOUND ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03
2001:bd8:3c4d:18::/64 4 83720 BOUND ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04
2001:bd8:3c4d:19::/64 5 83720 BOUND ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05
2001:bd8:3c4d:20::/64 6 83720 BOUND ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06

```

show dhcpv6 relay binding (address)

```

user@host> show dhcpv6 relay binding 2001:bd8:1111:2222::/64 detail

```

```

Session Id: 1
  Client IPv6 Prefix:      2001:bd8:3c4d:15::/64
  Client DUID:             LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

  State:                   BOUND(RELAY_STATE_BOUND)
  Lease Expires:           2011-05-25 07:12:09 PDT
  Lease Expires in:       77115 seconds
  Lease Start:            2011-05-24 07:12:09 PDT
  Incoming Client Interface: ge-1/0/0.0
  Server Address:          2008:aaaa:bbbb::1
  Server Interface:       none
  Relay Address:           2001:bd8:1111:2222::
  Client Pool Name:       pool-25
  Client Id Length:       14
  Client Id:
/0x00010001/0x4bfa26af/0x00109400/0x0001

```

show dhcpv6 relay binding (client-id)

```

user@host> show dhcpv6 relay binding 14/0x00010001/0x4bfa26af/0x00109400/0x0001
detail
Session Id: 1
  Client IPv6 Prefix:      2001:bd8:3c4d:15::/64
  Client DUID:             LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

  State:                   BOUND(RELAY_STATE_BOUND)
  Lease Expires:           2011-05-25 07:12:09 PDT
  Lease Expires in:       77115 seconds
  Lease Start:            2011-05-24 07:12:09 PDT
  Incoming Client Interface: ge-1/0/0.0
  Server Address:          2008:aaaa:bbbb::1
  Server Interface:       none
  Relay Address:           2001:bd8:1111:2222::
  Client Pool Name:       pool-25
  Client Id Length:       14
  Client Id:
/0x00010001/0x4bfa26af/0x00109400/0x0001

```

show dhcpv6 relay binding detail

```

user@host> show dhcpv6 relay binding detail
Session Id: 1
  Client IPv6 Prefix:      2001:bd8:3c4d:15::/64
  Client DUID:             LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

  State:                   BOUND(RELAY_STATE_BOUND)
  Lease Expires:           2011-05-25 07:12:09 PDT
  Lease Expires in:       77115 seconds
  Lease Start:            2011-05-24 07:12:09 PDT
  Incoming Client Interface: ge-1/0/0.0
  Server Address:          2008:aaaa:bbbb::1
  Server Interface:       none
  Relay Address:           2001:bd8:1111:2222::
  Client Pool Name:       pool-25
  Client Id Length:       14
  Client Id:
/0x00010001/0x4bfa26af/0x00109400/0x0001

```

show dhcpv6 relay binding detail (Multi-Relay Topology)

```

user@host > show dhcpv6 relay binding detail
Session Id: 13
  Client IPv6 Prefix:          3000:0:0:8001::5/128
  Client DUID:                 LL0x1-00:00:65:03:01:02
  State:                       BOUND(DHCPV6_RELAY_STATE_BOUND)
  Lease Expires:               2011-11-21 06:14:50 PST
  Lease Expires in:            293 seconds
  Lease Start:                 2011-11-21 06:09:50 PST
  Incoming Client Interface:   ge-1/0/0.0
  Server Address:              unknown
  Next Hop Server Facing Relay: 4000::2
  Server Interface:            none
  Client Id Length:            10
  Client Id:                   /0x00030001/0x00006503/0x0102

```

show dhcpv6 relay binding (session-id)

```

user@host> show dhcpv6 relay binding 41
Prefix          Session Id Expires   State   Interface   Client DUID
2001:bd8::/32  41          78837    BOUND   ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

```

show dhcpv6 relay binding (interfaces-vlan)

```

user@host> show dhcpv6 relay binding ge-1/0/0:100-200
Prefix          Session Id Expires   State   Interface   Client DUID
2001:DB8::/32   11          87583    BOUND   ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32   12          87583    BOUND   ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

show dhcpv6 relay binding (interfaces-wildcard)

```

user@host> show dhcpv6 relay binding demux0
Prefix          Session Id Expires   State   Interface   Client DUID
2001:DB8::/32   30          79681    BOUND   demux0.1073741824
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32   31          79681    BOUND   demux0.1073741825
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:CB9::/32   32          79681    BOUND   demux0.1073741826
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

show dhcpv6 relay binding (interfaces-wildcard)

```

user@host> show dhcpv6 relay binding ge-1/3/*
Prefix          Session Id Expires   State   Interface   Client DUID
2001:DB8::/32   22          79681    BOUND   ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32   33          79681    BOUND   ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:CB9::/32   24          79681    BOUND   ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

show dhcpv6 relay binding summary

```

user@host> show dhcpv6 relay binding summary
5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)

```

show dhcpv6 relay statistics

Syntax	show dhcpv6 relay statistics <logical-system <i>logical-system-name</i>> <routing-instance <i>routing-instance-name</i>>
Release Information	Command introduced in Junos OS Release 11.4. Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches. Command introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Display Dynamic Host Configuration Protocol for IPv6 (DHCPv6) relay statistics.
Options	logical-system <i>logical-system-name</i> —(Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system. routing-instance <i>routing-instance-name</i> —(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear dhcpv6 relay statistics on page 1201
List of Sample Output	show dhcpv6 relay statistics on page 1217
Output Fields	Table 88 on page 1216 lists the output fields for the show dhcpv6 relay statistics command. Output fields are listed in the approximate order in which they appear.

Table 88: show dhcpv6 relay statistics Output Fields

Field Name	Field Description
DHCPv6 Packets dropped	<p>Number of packets discarded by the extended DHCPv6 relay agent application due to errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.</p> <ul style="list-style-type: none"> • Total—Total number of packets discarded by the DHCPv6 relay agent application. • Bad options—Number of packets discarded because invalid options were specified. • Bad send—Number of packets that the extended DHCP relay application could not send. • Bad src address—Number of packets discarded because the family type was not AF_INET6. • No client id—Number of packets discarded because they could not be matched to a client. • No safd—Number of packets discarded because they arrived on an unconfigured interface. • Short packet—Number of packets discarded because they were too short. • Relay hop count—Number of packets discarded because the hop count in the packet exceeded 32.

Table 88: show dhcpv6 relay statistics Output Fields (*continued*)

Field Name	Field Description
Messages received	<p>Number of DHCPv6 messages received.</p> <ul style="list-style-type: none"> DHCPV6_DECLINE—Number of DHCPv6 PDUs of type DECLINE received DHCPV6_SOLICIT—Number of DHCPv6 PDUs of type SOLICIT received DHCPV6_INFORMATION_REQUEST—Number of DHCPv6 PDUs of type INFORMATION-REQUEST received DHCPV6_RELEASE—Number of DHCPv6 PDUs of type RELEASE received DHCPV6_REQUEST—Number of DHCPv6 PDUs of type REQUEST received DHCPV6_CONFIRM—Number of DHCPv6 PDUs of type CONFIRM received DHCPV6_RENEW—Number of DHCPv6 PDUs of type RENEW received DHCPV6_REBIND—Number of DHCPv6 PDUs of type REBIND received DHCPV6_RELAY_REPL—Number of DHCPv6 PDUs of type RELAY-REPL received
Messages sent	<p>Number of DHCPv6 messages sent.</p> <ul style="list-style-type: none"> DHCPV6_ADVERTISE—Number of DHCPv6 ADVERTISE PDUs transmitted DHCP_REPLY—Number of DHCPv6 REPLY PDUs transmitted DHCP_RECONFIGURE—Number of DHCPv6 RECONFIGURE PDUs transmitted DHCP_RELAY_FORW—Number of DHCPv6 RELAY-FORW PDUs transmitted
Packets forwarded	<p>Number of packets forwarded by the extended DHCPv6 relay agent application.</p> <ul style="list-style-type: none"> FWD REQUEST—Number of DHCPv6 REQUEST packets forwarded FWD REPLY—Number of DHCPv6 REPLY packets forwarded

Sample Output

show dhcpv6 relay statistics

```

user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
  Total                               0

Messages received:
  DHCPV6_DECLINE                      0
  DHCPV6_SOLICIT                      10
  DHCPV6_INFORMATION_REQUEST          0
  DHCPV6_RELEASE                      0
  DHCPV6_REQUEST                      10
  DHCPV6_CONFIRM                      0
  DHCPV6_RENEW                        0
  DHCPV6_REBIND                       0
  DHCPV6_RELAY_REPL                   0

Messages sent:
  DHCPV6_ADVERTISE                    0
  DHCPV6_REPLY                        0
  DHCPV6_RECONFIGURE                   0
  DHCPV6_RELAY_FORW                   0

Packets forwarded:
  Total                               4

```

FWD REQUEST	2
FWD REPLY	2

show route extensive

Syntax	show route extensive <destination-prefix> <logical-system (all logical-system-name)>
Syntax (EX Series Switches)	show route extensive <destination-prefix>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display extensive information about the active entries in the routing tables.
Options	<p>none—Display all active entries in the routing table.</p> <p>destination-prefix—(Optional) Display active entries for the specified address or range of addresses.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route extensive on page 1224 show route extensive (Access Route) on page 1230 show route extensive (Route Reflector) on page 1231 show route extensive (FRR and LFA) on page 1231 show route extensive (FRR and LFA) on page 1232
Output Fields	Table 89 on page 1219 describes the output fields for the show route extensive command. Output fields are listed in the approximate order in which they appear.

Table 89: show route extensive Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> active (routes that are active). holddown (routes that are in the pending state before being declared inactive). hidden (routes that are not used because of a routing policy).

Table 89: show route extensive Output Fields (*continued*)

Field Name	Field Description
<i>route-destination</i> (entry, announced)	<p>Route destination (for example: 10.0.0.1/24). The entry value is the number of route for this destination, and the announced value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote.
TSI	Protocol header information.
label stacking	<p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> • S=0 route indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed). • If there is no S= information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[<i>protocol, preference</i>]	<p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>
Level	<p>(IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.</p>
Route Distinguisher	IP subnet augmented with a 64-bit prefix.

Table 89: show route extensive Output Fields (*continued*)

Field Name	Field Description
Next-hop type	Type of next hop. For a description of possible values for this field, see the Output Field table in the show route detail command.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	<p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word Selected. This field can also contain the following information:</p> <ul style="list-style-type: none"> • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path <i>lsp-path-name</i>	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).
Offset	Whether the metric has been increased or decreased by an offset value.
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop.
<i>label-operation</i>	MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).
Indirect next hops	When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved.
State	State of the route (a route can be in more than one state). See the Output Field table in the show route detail command.

Table 89: show route extensive Output Fields (*continued*)

Field Name	Field Description
Session ID	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).
Inactive reason	<p>If the route is inactive, the reason for its current state is indicated. Typical reasons include:</p> <ul style="list-style-type: none"> • Active preferred—Currently active route was selected over this route. • Always compare MED—Path with a lower multiple exit discriminator (MED) is available. • AS path—Shorter AS path is available. • Cisco Non-deterministic MED selection—Cisco nondeterministic MED is enabled and a path with a lower MED is available. • Cluster list length—Path with a shorter cluster list length is available. • Forwarding use only—Path is only available for forwarding purposes. • IGP metric—Path through the next hop with a lower IGP metric is available. • IGP metric type—Path with a lower OSPF link-state advertisement type is available. • Interior > Exterior > Exterior via Interior—Direct, static, IGP, or EBGP path is available. • Local preference—Path with a higher local preference value is available. • Next hop address—Path with a lower metric next hop is available. • No difference—Path from a neighbor with a lower IP address is available. • Not Best in its group—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed). • Number of gateways—Path with a higher number of next hops is available. • Origin—Path with a lower origin code is available. • OSPF version—Path does not support the indicated OSPF version. • RIB preference—Route from a higher-numbered routing table is available. • Route distinguisher—64-bit prefix added to IP subnets to make them unique. • Route metric or MED comparison—Route with a lower metric or MED is available. • Route preference—Route with a lower preference value is available. • Router ID—Path through a neighbor with a lower ID is available. • Unusable path—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved. • Update source—Last tiebreaker is the lowest IP address value.
Local AS	Autonomous system (AS) number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.

Table 89: show route extensive Output Fields (*continued*)

Field Name	Field Description
TTL-Action	<p>For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.</p> <p>For sample output, see show route table.</p>
Task	Name of the protocol that has added the route.
Announcement bits	List of protocols that announce this route. n-Resolve inet indicates that the route is used for route resolution for next hops found in the routing table. n is an index used by Juniper Networks customer support only.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> I—IGP. E—EGP. Recorded—The AS path is recorded by the sample process (sampled). ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> []—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. ()—Parentheses enclose a confederation. ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>
AS path: I <Originator>	(For route reflected output only) Originator ID attribute set by the route reflector.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.
Originator ID	(For route reflected output only) Address of router that originally sent the route to the route reflector.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field.

Table 89: show route extensive Output Fields (*continued*)

Field Name	Field Description
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: none or Site Down.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix.
Node path count	Number of nodes in the path.
Forwarding nexthops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Sample Output

show route extensive

```

user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kerne1 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:34:06
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

10.31.1.0/30 (2 entries, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 2

```

```

Next hop: via so-0/3/0.0, selected
State: <Active Int>
Local AS: 69
Age: 1:32:40
Task: IF
Announcement bits (1): 3-Resolve tree 2
AS path: I
OSPF Preference: 10
Next-hop reference count: 1
Next hop: via so-0/3/0.0, selected
State: <Int>
Inactive reason: Route Preference
Local AS: 69
Age: 1:32:40 Metric: 1
Area: 0.0.0.0
Task: OSPF
AS path: I

10.31.1.1/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 7
Interface: so-0/3/0.0
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:32:43
Task: IF
Announcement bits (1): 3-Resolve tree 2
AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
TSI:
KRT in-kerne 10.31.2.0/30 -> {10.31.1.6}
*OSPF Preference: 10
Next-hop reference count: 9
Next hop: via so-0/3/0.0
Next hop: 10.31.1.6 via ge-3/1/0.0, selected
State: <Active Int>
Local AS: 69
Age: 1:32:19 Metric: 2
Area: 0.0.0.0
Task: OSPF
Announcement bits (2): 0-KRT 3-Resolve tree 2
AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
TSI:
KRT in-kerne 224.0.0.2/32 -> {}
*PIM Preference: 0
Next-hop reference count: 18
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:34:08
Task: PIM Recv
Announcement bits (2): 0-KRT 3-Resolve tree 2
AS path: I

```

```
...

224.0.0.22/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 224.0.0.22/32 -> {}
    *IGMP Preference: 0
        Next-hop reference count: 18
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:34:06
        Task: IGMP
        Announcement bits (2): 0-KRT 3-Resolve tree 2
        AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
State: <FlashAll>
    *RSVP Preference: 7
        Next-hop reference count: 6
        Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
        Label-switched-path green-r1-r3
        Label operation: Push 100096
        State: <Active Int>
        Local AS: 69
        Age: 1:28:12 Metric: 2
        Task: RSVP
        Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
        AS path: I

10.255.71.238/32 (1 entry, 1 announced)
State: <FlashAll>
    *RSVP Preference: 7
        Next-hop reference count: 6
        Next hop: via so-0/3/0.0 weight 0x1, selected
        Label-switched-path green-r1-r2
        State: <Active Int>
        Local AS: 69
        Age: 1:28:12 Metric: 1
        Task: RSVP
        Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
        AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 1
        Next hop: via lo0.0, selected
        State: <Active Int>
        Local AS: 69
        Age: 1:34:07
        Task: IF
        AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
```



```

0 (1 entry, 1 announced)
TSI:
KRT in-kernel 0 /36 -> {}
    *MPLS Preference: 0
        Next hop type: Receive
        Next-hop reference count: 6
        State: <Active Int>
        Local AS: 69
        Age: 1:34:08 Metric: 1
        Task: MPLS
        Announcement bits (1): 0-KRT
        AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299776 (1 entry, 1 announced)
TSI:
KRT in-kernel 299776 /52 -> {Flood}
    *RSVP Preference: 7
        Next hop type: Flood
        Next-hop reference count: 130
        Flood nexthop branches exceed maximum
        Address: 0x8ea65d0

...

800010 (1 entry, 1 announced)
TSI:
KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
    *VPLS Preference: 7
        Next-hop reference count: 2
        Next hop: via vt-3/2/0.32769, selected
        Label operation: Pop
        State: <Active Int>
        Age: 1:31:53
        Task: Common L2 VC
        Announcement bits (1): 0-KRT
        AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
TSI:
KRT in-kernel vt-3/2/0.32769.0 /16 -> {indirect(1048574)}
    *VPLS Preference: 7
        Next-hop reference count: 2
        Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
        Label-switched-path green-r1-r3
        Label operation: Push 800012, Push 100096(top)
        Protocol next hop: 10.255.70.103
        Push 800012
        Indirect next hop: 87272e4 1048574
        State: <Active Int>
        Age: 1:31:53 Metric2: 2
        Task: Common L2 VC
        Announcement bits (2): 0-KRT 1-Common L2 VC
        AS path: I
        Communities: target:11111:1 Layer2-info: encaps:VPLS,
        control flags:, mtu: 0
        Indirect next hops: 1
            Protocol next hop: 10.255.70.103 Metric: 2

```

```

        Push 800012
        Indirect next hop: 87272e4 1048574
        Indirect path forwarding next hops: 1
            Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1
        10.255.70.103/32 Originating RIB: inet.3
            Metric: 2                      Node path count: 1
            Forwarding nexthops: 1
                Nexthop: 10.31.1.6 via ge-3/1/0.0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 1
        Next hop: via lo0.0, selected
        State: <Active Int>
        Local AS: 69
        Age: 1:34:07
        Task: IF
        AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 1
        Next hop: via lo0.0, selected
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:34:07
        Task: IF
        AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
    *PIM Preference: 0
        Next-hop reference count: 18
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:34:08
        Task: PIM Recv6
        Announcement bits (1): 0-KRT
        AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
    *PIM Preference: 0
        Next-hop reference count: 18
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:34:08
        Task: PIM Recv6
        Announcement bits (1): 0-KRT
        AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
    *MLD Preference: 0
```

```

Next-hop reference count: 18
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:34:06
Task: MLD
Announcement bits (1): 0-KRT
AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 1
Next hop: via lo0.16385, selected
State: <Active NoReadvrt Int>
Age: 1:34:07
Task: IF
AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.70.103:1
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Secondary Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 1:28:12 Metric2: 1
Task: BGP_69.10.255.70.103+179
Announcement bits (1): 0-green-l2vpn
AS path: I
Communities: target:11111:1 Layer2-info: encaps:VPLS,
control flags:, mtu: 0
Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699540
*L2VPN Preference: 170/-1
Next-hop reference count: 5
Protocol next hop: 10.255.71.52
Indirect next hop: 0 -
State: <Active Int Ext>
Age: 1:34:03 Metric2: 1
Task: green-l2vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
mtu: 0
Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699528

```

```
*L2VPN Preference: 170/-101
Next-hop reference count: 5
Protocol next hop: 10.255.71.52
Indirect next hop: 0 -
State: <Active Int Ext>
Age: 1:34:03 Metric2: 1
Task: green-l2vpn
Announcement bits (1): 1-BGP.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

TSI:

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
*L2CKT Preference: 7
Next hop: via so-1/1/2.0 weight 1, selected
Label-switched-path my-lsp
Label operation: Push 100000[0]
Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
State: <Active Int>
Local AS: 99
Age: 10:21
Task: l2 circuit
Announcement bits (1): 0-LDP
AS path: I
VC Label 100000, MTU 1500, VLAN ID 512

55.0.0.0/24 (1 entry, 1 announced)
TSI:
KRT queued (pending) add
55.0.0.0/24 -> {Push 300112}
*BGP Preference: 170/-101
Next hop type: Router
Address: 0x925c208
Next-hop reference count: 2
Source: 10.0.0.9
Next hop: 10.0.0.9 via ge-1/2/0.15, selected
Label operation: Push 300112
Label TTL action: prop-ttl
State: <Active Ext>
Local AS: 7019 Peer AS: 13979
Age: 1w0d 23:06:56
AIGP: 25
Task: BGP_13979.10.0.0.9+56732
Announcement bits (1): 0-KRT
AS path: 13979 7018 I
Accepted
Route Label: 300112
Localpref: 100
Router ID: 10.9.9.1
```

show route extensive (Access Route)

```
user@host> show route 13.160.0.102 extensive
```

```

inet.0: 39256 destinations, 39258 routes (39255 active, 0 holddown, 1 hidden)
13.160.0.102/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.102/32 -> {13.160.0.2}
OSPF area : 0.0.0.0, LSA ID : 13.160.0.102, LSA type : Extern
    *Access Preference: 13
        Next-hop reference count: 78472
        Next hop: 13.160.0.2 via fe-0/0/0.0, selected
        State: <Active Int>
    Age: 12
        Task: RPD Unix Domain Server./var/run/rpd_serv.local
        Announcement bits (2): 0-KRT 1-OSPFv2
        AS path: I

```

show route extensive (Route Reflector)

```

user@host> show route extensive
1.0.0.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 1.0.0.0/8 -> {indirect(40)}
    *BGP    Preference: 170/-101
        Source: 192.168.4.214
        Protocol next hop: 207.17.136.192 Indirect next hop: 84ac908 40
        State: <Active Int Ext>
        Local AS: 10458 Peer AS: 10458
        Age: 3:09      Metric: 0      Metric2: 0
        Task: BGP_10458.192.168.4.214+1033
        Announcement bits (2): 0-KRT 4-Resolve inet.0
        AS path: 3944 7777 I <Originator>
        Cluster list: 1.1.1.1
        Originator ID: 10.255.245.88
        Communities: 7777:7777
        Localpref: 100
        Router ID: 4.4.4.4
        Indirect next hops: 1
            Protocol next hop: 207.17.136.192 Metric: 0
            Indirect next hop: 84ac908 40
            Indirect path forwarding next hops: 0
            Next hop type: Discard

```

show route extensive (FRR and LFA)

```

user@host> show route 20:31:2:0 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
20.31.2.0/24 (2 entries, 1 announced)
    State: FlashAll

TSI:
KRT in-kernel 20.31.2.0/24 -> {Push 299776, Push 299792}
    *RSVP    Preference: 7/1
        Next hop type: Router, Next hop index: 1048574
        Address: 0xbbbc010
        Next-hop reference count: 5
        Next hop: 10.31.1.2 via ge-2/1/8.0 weight 0x1, selected
        Label-switched-path europa-d-to-europa-e
        Label operation: Push 299776
        Label TTL action: prop-ttl
        Session Id: 0x201
        Next hop: 10.31.2.2 via ge-2/1/4.0 weight 0x4001
        Label-switched-path europa-d-to-europa-e
        Label operation: Push 299792

```

```
Label TTL action: prop-ttl
Session Id: 0x202
State: Active Int
Local AS: 100
Age: 5:31 Metric: 2
Task: RSVP
Announcement bits (1): 0-KRT
AS path: I
OSPF Preference: 10
Next hop type: Router, Next hop index: 615
Address: 0xb9d78c4
Next-hop reference count: 7
Next hop: 10.31.1.2 via ge-2/1/8.0, selected
Session Id: 0x201
State: Int
Inactive reason: Route Preference
Local AS: 100
Age: 5:35 Metric: 3
Area: 0.0.0.0
Task: OSPF
AS path: I
```

show route extensive (FRR and LFA)

```
user@host> show route 20:31:2:0 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
20.31.2.0/24 (2 entries, 1 announced)
State: FlashAll
TSI:
KRT in-kernel 20.31.2.0/24 -> {Push 299776, Push 299792}
*RSVP Preference: 7/1
Next hop type: Router, Next hop index: 1048574
Address: 0xbbbc010
Next-hop reference count: 5
Next hop: 10.31.1.2 via ge-2/1/8.0 weight 0x1, selected
Label-switched-path europa-d-to-europa-e
Label operation: Push 299776
Label TTL action: prop-ttl
Session Id: 0x201
Next hop: 10.31.2.2 via ge-2/1/4.0 weight 0x4001
Label-switched-path europa-d-to-europa-e
Label operation: Push 299792
Label TTL action: prop-ttl
Session Id: 0x202
State: Active Int
Local AS: 100
Age: 5:31 Metric: 2
Task: RSVP
Announcement bits (1): 0-KRT
AS path: I
OSPF Preference: 10
Next hop type: Router, Next hop index: 615
Address: 0xb9d78c4
Next-hop reference count: 7
Next hop: 10.31.1.2 via ge-2/1/8.0, selected
Session Id: 0x201
State: Int
Inactive reason: Route Preference
Local AS: 100
Age: 5:35 Metric: 3
Area: 0.0.0.0
```

Task: OSPF
AS path: I

show route protocol

Syntax	<code>show route protocol <i>protocol</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route protocol <i>protocol</i></code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. ospf2 and ospf3 options introduced in Junos OS Release 9.2. ospf2 and ospf3 options introduced in Junos OS Release 9.2 for EX Series switches. flow option introduced in Junos OS Release 10.0. flow option introduced in Junos OS Release 10.0 for EX Series switches.
Description	Display the route entries in the routing table that were learned from a particular protocol.
Options	brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>protocol</i> —Protocol from which the route was learned: <ul style="list-style-type: none">• access—Access route for use by DHCP application• access-internal—Access-internal route for use by DHCP application• aggregate—Locally generated aggregate route• arp—Route learned through the Address Resolution Protocol• atmvpn—Asynchronous Transfer Mode virtual private network• bgp—Border Gateway Protocol• ccc—Circuit cross-connect• direct—Directly connected route• dvmrp—Distance Vector Multicast Routing Protocol• esis—End System-to-Intermediate System• flow—Locally defined flow-specification route• frr—Precomputed protection route or backup route used when a link goes down• isis—Intermediate System-to-Intermediate System• ldp—Label Distribution Protocol• l2circuit—Layer 2 circuit• l2vpn—Layer 2 virtual private network

- **local**—Local address
- **mpls**—Multiprotocol Label Switching
- **msdp**—Multicast Source Discovery Protocol
- **ospf**—Open Shortest Path First versions 2 and 3
- **ospf2**—Open Shortest Path First versions 2 only
- **ospf3**—Open Shortest Path First version 3 only
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation
- **rsvp**—Resource Reservation Protocol
- **rtarget**—Local route target virtual private network
- **static**—Statically defined route
- **tunnel**—Dynamic tunnel
- **vpn**—Virtual private network



NOTE: EX Series switches run a subset of these protocols. See the switch CLI for details.

Required Privilege Level	view
List of Sample Output	show route protocol access on page 1236 show route protocol access-internal extensive on page 1236 show route protocol arp on page 1236 show route protocol bgp on page 1237 show route protocol bgp detail on page 1237 show route protocol bgp extensive on page 1237 show route protocol bgp terse on page 1238 show route protocol direct on page 1238 show route protocol frr on page 1239 show route protocol l2circuit detail on page 1239 show route protocol l2vpn extensive on page 1240 show route protocol ldp on page 1241 show route protocol ldp extensive on page 1241 show route protocol ospf (Layer 3 VPN) on page 1242 show route protocol ospf detail on page 1243 show route protocol rip on page 1243 show route protocol rip detail on page 1243 show route protocol ripng table inet6 on page 1244 show route protocol static detail on page 1244

Output Fields For information about output fields, see the output field tables for the [show route](#) command, the [show route detail](#) command, the [show route extensive](#) command, or the [show route terse](#) command.

Sample Output

show route protocol access

```
user@host> show route protocol access
inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
```

show route protocol access-internal extensive

```
user@host> show route protocol access-internal 13.160.0.19 extensive
inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.19/32 -> {13.160.0.2}
    *Access-internal Preference: 12
        Next-hop reference count: 200000
        Next hop: 13.160.0.2 via fe-0/0/0.0, selected
        State: <Active Int>
    Age: 36
        Task: RPD Unix Domain Server./var/run/rpd_serv.local
        Announcement bits (1): 0-KRT
        AS path: I
```

show route protocol arp

```
user@host> show route protocol arp
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.4/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.5/32      [ARP/4294967293] 00:04:32, from 20.20.1.1
                  Unusable
20.20.1.6/32      [ARP/4294967293] 00:04:34, from 20.20.1.1
                  Unusable
20.20.1.7/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.8/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.9/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.10/32     [ARP/4294967293] 00:04:35, from 20.20.1.1
```

```

Unusable
20.20.1.11/32      [ARP/4294967293] 00:04:33, from 20.20.1.1
Unusable
20.20.1.12/32      [ARP/4294967293] 00:04:33, from 20.20.1.1
Unusable
20.20.1.13/32      [ARP/4294967293] 00:04:33, from 20.20.1.1
Unusable
...

```

show route protocol bgp

```

user@host> show route protocol bgp 192.168.64.0/21
inet.0: 335832 destinations, 335833 routes (335383 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21      *[BGP/170] 6d 10:41:16, localpref 100, from 192.168.69.71
AS path: 10458 14203 2914 4788 4788 I
> to 192.168.167.254 via fxp0.0

```

show route protocol bgp detail

```

user@host> show route protocol bgp 66.117.63.0/24 detail
inet.0: 335805 destinations, 335806 routes (335356 active, 0 holddown, 450 hidden)
66.117.63.0/24      (1 entry, 1 announced)
    *BGP      Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 1006436
                Source: 192.168.69.71
                Next hop type: Router, Next hop index: 324
                Next hop: 192.168.167.254 via fxp0.0, selected
                Protocol next hop: 192.168.69.71
                Indirect next hop: 8e166c0 342
                State: <Active Ext>
                Local AS: 69 Peer AS: 10458
                Age: 6d 10:42:42 Metric2: 0
                Task: BGP_10458.192.168.69.71+179
                Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree

1
    AS path: 10458 14203 2914 4788 4788 I
    Communities: 2914:410 2914:2403 2914:3400
    Accepted
    Localpref: 100
    Router ID: 207.17.136.192

```

show route protocol bgp extensive

```

user@host> show route protocol bgp 192.168.64.0/21 extensive

inet.0: 335827 destinations, 335828 routes (335378 active, 0 holddown, 450 hidden)
192.168.64.0/21 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.9.0.0/16 -> {indirect(342)}
Page 0 idx 1 Type 1 val db31a80
  Nexthop: Self
  AS path: [69] 10458 14203 2914 4788 4788 I
  Communities: 2914:410 2914:2403 2914:3400
Path 1.9.0.0 from 192.168.69.71 Vector len 4. Val: 1
    *BGP      Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 1006502
                Source: 192.168.69.71
                Next hop type: Router, Next hop index: 324

```

```

Next hop: 192.168.167.254 via fxp0.0, selected
Protocol next hop: 192.168.69.71
Indirect next hop: 8e166c0 342
State: <Active Ext>
Local AS: 69 Peer AS: 10458
Age: 6d 10:44:45 Metric2: 0
Task: BGP_10458.192.168.69.71+179
Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree
1
AS path: 10458 14203 2914 4788 4788 I
Communities: 2914:410 2914:2403 2914:3400
Accepted
Localpref: 100
Router ID: 207.17.136.192
Indirect next hops: 1
  Protocol next hop: 192.168.69.71
  Indirect next hop: 8e166c0 342
  Indirect path forwarding next hops: 1
    Next hop type: Router
    Next hop: 192.168.167.254 via fxp0.0
  192.168.0.0/16 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
    Nexthop: 192.168.167.254 via fxp0.0

```

show route protocol bgp terse

```
user@host> show route protocol bgp 192.168.64.0/21 terse
```

```
inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
192.168.64.0/21	B 170	100		>100.1.3.2	10023 21 I

show route protocol direct

```
user@host> show route protocol direct
```

```
inet.0: 335843 destinations, 335844 routes (335394 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both
```

```

8.8.8.0/24      *[Direct/0] 17w0d 10:31:49
                 > via fe-1/3/1.0
10.255.165.1/32 *[Direct/0] 25w4d 04:13:18
                 > via lo0.0
30.30.30.0/24   *[Direct/0] 17w0d 23:06:26
                 > via fe-1/3/2.0
192.168.164.0/22 *[Direct/0] 25w4d 04:13:20
                 > via fxp0.0

```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```

47.0005.80ff.f800.0000.0108.0001.0102.5516.5001/152
    *[Direct/0] 25w4d 04:13:21
    > via lo0.0

```

```
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```

abcd::10:255:165:1/128
    *[Direct/0] 25w4d 04:13:21
    > via lo0.0
fe80::2a0:a5ff:fe12:ad7/128
    *[Direct/0] 25w4d 04:13:21
    > via lo0.0

```

show route protocol frr

```

user@host> show route protocol frr
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.3 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.4/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.4 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.5/32      *[FRR/200] 00:05:35, from 20.20.1.1
                  > to 20.20.1.5 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.6/32      *[FRR/200] 00:05:37, from 20.20.1.1
                  > to 20.20.1.6 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.7/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.7 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.8/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.8 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.9/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.9 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.10/32     *[FRR/200] 00:05:38, from 20.20.1.1
...

```

show route protocol l2circuit detail

```

user@host> show route protocol l2circuit detail

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)
    *L2CKT Preference: 7
        Next hop: via ge-2/0/0.0, selected
        Label operation: Pop      Offset: 4
        State: <Active Int>
        Local AS: 99
        Age: 9:52
        Task: Common L2 VC
        Announcement bits (1): 0-KRT
        AS path: I

ge-2/0/0.0 (1 entry, 1 announced)
    *L2CKT Preference: 7
        Next hop: via so-1/1/2.0 weight 1, selected
        Label-switched-path my-lsp

```

```

Label operation: Push 100000, Push 100000(top)[0] Offset: -4
Protocol next hop: 10.245.255.63
Push 100000 Offset: -4
  Indirect next hop: 86af0c0 298
State: <Active Int>
Local AS: 99
Age: 9:52
Task: Common L2 VC
Announcement bits (2): 0-KRT 1-Common L2 VC
AS path: I

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
    State: <Active Int>
    Local AS: 99
    Age: 10:21
    Task: l2 circuit
    Announcement bits (1): 0-LDP
    AS path: I
    VC Label 100000, MTU 1500, VLAN ID 512

```

show route protocol l2vpn extensive

```

user@host> show route protocol l2vpn extensive

inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)
TSI:
KRT in-kernel 800001 /36 -> {so-0/0/0.0}
  *L2VPN Preference: 7
    Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
    Label operation: Pop Offset: 4
    State: <Active Int>
    Local AS: 69
    Age: 7:48
    Task: Common L2 VC
    Announcement bits (1): 0-KRT
    AS path: I

so-0/0/0.0 (1 entry, 1 announced)
TSI:
KRT in-kernel so-0/0/0.0 /16 -> {indirect(288)}
  *L2VPN Preference: 7
    Next hop: via so-0/0/1.0, selected
    Label operation: Push 800000 Offset: -4
    Protocol next hop: 10.255.14.220
    Push 800000 Offset: -4
    Indirect next hop: 85142a0 288
    State: <Active Int>

```

```

Local AS:    69
Age: 7:48
Task: Common L2 VC
Announcement bits (2): 0-KRT 1-Common L2 VC
AS path: I
Communities: target:69:1 Layer2-info: encaps:PPP,
control flags:2, mtu: 0

```

show route protocol ldp

```

user@host> show route protocol ldp
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.16.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Push 100000
192.168.17.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0

private1___.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100064            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100064(S=0)        *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100080            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Swap 100000

```

show route protocol ldp extensive

```

user@host> show route protocol ldp extensive
192.168.16.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP    Preference: 9
          Next-hop reference count: 3
          Next hop: via t1-4/0/0.0, selected
          Label operation: Push 100000
          State: <Active Int>
          Local AS: 65500
          Age: 1d 23:03:58      Metric: 1
          Task: LDP
          Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
          AS path: I

192.168.17.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP    Preference: 9
          Next-hop reference count: 3
          Next hop: via t1-4/0/0.0, selected
          State: <Active Int>
          Local AS: 65500
          Age: 1d 23:03:58      Metric: 1
          Task: LDP
          Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
          AS path: I

```

```
private1__inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

```
mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
```

```
100064 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 100064 /36 -> {t1-4/0/0.0}
```

```
*LDP      Preference: 9
           Next-hop reference count: 2
           Next hop: via t1-4/0/0.0, selected
           State: <Active Int>
           Local AS: 65500
           Age: 1d 23:03:58      Metric: 1
           Task: LDP
           Announcement bits (1): 0-KRT
           AS path: I
           Prefixes bound to route: 192.168.17.1/32
```

```
100064(S=0) (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 100064 /40 -> {t1-4/0/0.0}
```

```
*LDP      Preference: 9
           Next-hop reference count: 2
           Next hop: via t1-4/0/0.0, selected
           Label operation: Pop
           State: <Active Int>
           Local AS: 65500
           Age: 1d 23:03:58      Metric: 1
           Task: LDP
           Announcement bits (1): 0-KRT
           AS path: I
```

```
100080 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 100080 /36 -> {t1-4/0/0.0}
```

```
*LDP      Preference: 9
           Next-hop reference count: 2
           Next hop: via t1-4/0/0.0, selected
           Label operation: Swap 100000
           State: <Active Int>
           Local AS: 65500
           Age: 1d 23:03:58      Metric: 1
           Task: LDP
           Announcement bits (1): 0-KRT
           AS path: I
           Prefixes bound to route: 192.168.16.1/32
```

show route protocol ospf (Layer 3 VPN)

```
user@host> show route protocol ospf
```

```
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

```
10.39.1.4/30      *[OSPF/10] 00:05:18, metric 4
                  > via t3-3/2/0.0
10.39.1.8/30      [OSPF/10] 00:05:18, metric 2
                  > via t3-3/2/0.0
10.255.14.171/32  *[OSPF/10] 00:05:18, metric 4
                  > via t3-3/2/0.0
10.255.14.179/32 *[OSPF/10] 00:05:18, metric 2
                  > via t3-3/2/0.0
```



```

224.0.0.5/32      *[OSPF/10] 20:25:55, metric 1

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30      [OSPF/10] 00:05:43, metric 1
> via so-0/2/2.0
10.255.14.173/32   *[OSPF/10] 00:05:43, metric 1
> via so-0/2/2.0
224.0.0.5/32      *[OSPF/10] 20:26:20, metric 1

```

show route protocol ospf detail

```

user@host> show route protocol ospf detail
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
    OSPF    Preference: 10
            Nexthop: via so-0/2/2.0, selected
            State: <Int>
            Inactive reason: Route Preference
            Age: 6:25      Metric: 1
            Area: 0.0.0.0
            Task: VPN-AB-OSPF
            AS path: I
            Communities: Route-Type:0.0.0.0:1:0

...

```

show route protocol rip

```

user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32   *[RIP/100] 20:24:34, metric 2
> to 10.39.1.22 via t3-0/2/2.0
224.0.0.9/32      *[RIP/100] 00:03:59, metric 1

```

show route protocol rip detail

```

user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
    *RIP    Preference: 100
            Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
            State: <Active Int>
            Age: 20:25:02  Metric: 2
            Task: VPN-AB-RIPv2
            Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179
            AS path: I
            Route learned from 10.39.1.22 expires in 96 seconds

```

show route protocol ripng table inet6

```

user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::4/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::6/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0

```

show route protocol static detail

```

user@host> show route protocol static detail
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
10.5.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
        State: <Active NoReadvrt Int Ext>
        Age: 7w3d 21:24:25
        Validation State: unverified
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.10.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
        State: <Active NoReadvrt Int Ext>
        Age: 7w3d 21:24:25
        Validation State: unverified
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.13.10.0/23 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
        State: <Active NoReadvrt Int Ext>
        Age: 7w3d 21:24:25
        Validation State: unverified

```

Task: RT
Announcement bits (1): 0-KRT
AS path: I

Troubleshooting

- [Acquiring Troubleshooting Information on page 1245](#)
- [Troubleshooting Configuration Statements on page 1253](#)

Acquiring Troubleshooting Information

- [Tracing Extended DHCP Operations on page 1245](#)
- [Tracing Extended DHCP Operations for Specific Interfaces on page 1251](#)

Tracing Extended DHCP Operations

Both the extended DHCP local server and the extended DHCP relay agent support tracing operations. DHCP tracing operations track extended DHCP operations and record them in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.

You can configure DHCP trace operations at the global level and at the interface level. Global DHCP tracing logs all DHCP-related events, whereas interface-level tracing logs only interface-specific DHCP events. If you configure interface-level trace operations, you can specify tracing for a range of interfaces or an individual interface. However, only a single interface-level log file is supported. That is, you cannot specify different interface-level log files for different interfaces or groups of interfaces.

By default, nothing is traced. When you enable the tracing operation, the default tracing behavior is as follows:

- Important events for both global and per-interface tracing are logged in a file located in the `/var/log` directory. By default, the router uses the filename, `jdhcpd`. You can specify a different filename, but you cannot change the directory in which trace files are located.
- When the trace log file *filename* reaches 128 kilobytes (KB), it is compressed and renamed *filename.0.gz*. Subsequent events are logged in a new file called *filename*, until it reaches capacity again. At this point, *filename.0.gz* is renamed *filename.1.gz* and *filename* is compressed and renamed *filename.0.gz*. This process repeats until the number of archived files reaches the maximum file number. Then the oldest trace file—the one with the highest number—is overwritten.

You can optionally specify the number of trace files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB). (For more information about how log files are created, see the *Junos OS System Log Messages Reference*.)

- By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure global DHCP tracing operations.

- Specify tracing operations for DHCP local server and DHCP relay:

```
[edit system processes dhcp-service]
user@host# edit traceoptions
```

The tracing configuration is applied globally to all DHCP applications in every LS:RI. Configuration of event tracing on a per-LS:RI basis is not supported. DHCP tracing is configurable only in the default LS:RI. However, DHCP applications (local server or relay) do not have to be configured in the default LS:RI. This behavior was different in software releases before Junos OS Release 11.4, where you had to configure a DHCP application in the default LS:RI in order to configure DHCP tracing, even when you wanted to run DHCP and trace its operations only in a nondefault LS:RI.

In the earlier software releases, you configured tracing statements at the **[edit system services dhcp-local-server]** and **[edit forwarding-options dhcp-relay]** hierarchy levels. These statements have been deprecated and hidden in favor of the statements at the **[edit system processes dhcp-service]** hierarchy level.



NOTE: The deprecated statements may be removed from a future release; we recommend that you transition to the new statements.

Because you can configure DHCP tracing at three different hierarchy levels (one new and recommended, two old and deprecated), the following rules apply to manage the interaction:

- When you configure a filename or any other options for the trace log file, the configuration at the **[edit system processes dhcp-service]** hierarchy level has the highest precedence, followed by the configuration at the **[edit system services dhcp-local-server]** hierarchy level, and finally with the lowest precedence, the configuration at the **[edit forwarding-options dhcp-relay]** hierarchy level.
- The flag configurations for multiple hierarchy levels are merged and applied to all trace log events.
- The deprecated statements do not support filtering the generation of DHCP trace log events by severity level. If you use these statements, trace logging operates with an implicit severity of **all**, regardless of the severity level configured at the **[edit system processes dhcp-service]** hierarchy level.

For information about configuring per-interface tracing options, see [“Tracing Extended DHCP Operations for Specific Interfaces” on page 1250](#).

The extended DHCP traceoptions operations are described in the following sections:

- [Configuring the Extended DHCP Log Filename on page 1247](#)
- [Configuring the Number and Size of Extended DHCP Log Files on page 1247](#)
- [Configuring Access to the Extended DHCP Log File on page 1248](#)
- [Configuring a Regular Expression for Extended DHCP Messages to Be Logged on page 1248](#)

- [Configuring the Extended DHCP Tracing Flags on page 1249](#)
- [Configuring the Severity Level to Filter Which Extended DHCP Messages Are Logged on page 1249](#)
- [Tracing Extended DHCP Operations for Specific Interfaces on page 1250](#)

Configuring the Extended DHCP Log Filename

By default, the name of the file that records trace output is **jdhcpd**. You can specify a different name by including the **file** option. DHCP local server and DHCP relay agent both support the **file** option for the **traceoptions** statement and the **interface-traceoptions** statement.

To change the filename:

- Specify a filename for global tracing operations.

```
[edit system processes dhcp-service traceoptions]
user@host# set file filename
```

- Specify a filename for per-interface tracing operations.

```
[edit system processes dhcp-service interface-traceoptions]
user@host# set file filename
```

Configuring the Number and Size of Extended DHCP Log Files

You can optionally specify the number of compressed, archived trace log files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB); the default size is 128 kilobytes (KB).

The archived files are differentiated by a suffix in the format **.number.gz**. The newest archived file is **.0.gz** and the oldest archived file is **.(maximum number)-1.gz**. When the current trace log file reaches the maximum size, it is compressed and renamed, and any existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation, **filename**, reaches 2 MB, **filename** is compressed and renamed **filename.0.gz**, and a new file called **filename** is created. When the new **filename** reaches 2 MB, **filename.0.gz** is renamed **filename.1.gz** and **filename** is compressed and renamed **filename.0.gz**. This process repeats until there are 20 trace files. Then the oldest file, **filename.19.gz**, is simply overwritten when the next oldest file, **filename.18.gz** is compressed and renamed to **filename.19.gz**.

DHCP local server and DHCP relay agent both support the **files** and **size** options for the **traceoptions** statement and the **interface-traceoptions** statement. To configure the number and size of trace files:

- Specify the name, number, and size of the file used for the trace output for global tracing operations.

```
[edit system processes dhcp-service traceoptions]
user@host# set file filename files number size maximum-file-size
```

- Specify the name, number, and size of the file used for the trace output for per-interface tracing operations.

```
[edit system processes dhcp-service interface-traceoptions]  
user@host# set file filename files number size maximum-file-size
```

Configuring Access to the Extended DHCP Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

DHCP local server and DHCP relay agent both support the **world-readable** option and the **no-world-readable** option for the **traceoptions** statement and the **interface-traceoptions** statement. To specify that all users can read the log file:

- Configure the log file to be world-readable for global tracing operations.

```
[edit system processes dhcp-service traceoptions]  
user@host# set file filename world-readable
```

- Configure the log file to be world-readable for per-interface tracing operations.

```
[edit system processes dhcp-service interface-traceoptions]  
user@host# set file filename world-readable
```

To explicitly set the default behavior, in which the log file can only be read by the user who configured tracing:

- Configure the log file to be no-world-readable for global tracing operations.

```
[edit system processes dhcp-service traceoptions]  
user@host# set file filename no-world-readable
```

- Configure the log file to be no-world-readable for per-interface tracing operations.

```
[edit system processes dhcp-service interface-traceoptions]  
user@host# set file filename no-world-readable
```

Configuring a Regular Expression for Extended DHCP Messages to Be Logged

By default, the trace operation output includes all messages relevant to the logged events. You can refine the output by including regular expressions to be matched.

DHCP local server and DHCP relay agent both support the **match** option for the **traceoptions** statement and the **interface-traceoptions** statement. To configure regular expressions to be matched:

- Specify the regular expression for global tracing operations.

```
[edit system processes dhcp-service traceoptions]  
user@host# set file filename match regular-expression
```

- Specify the regular expression for per-interface tracing operations.

```
[edit system processes dhcp-service interface-traceoptions]  
user@host# set file filename match regular-expression
```

Configuring the Extended DHCP Tracing Flags

By default, only important events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

DHCP local server and DHCP relay agent both support the **flag** option for the **traceoptions** statement and the **interface-traceoptions** statement. A smaller set of flags is supported for interface-level tracing than for global tracing. To configure the flags for the events to be logged:

- Specify the flags for global tracing operations.

```
[edit system processes dhcp-service traceoptions]  
user@host# set flag flag
```

- Specify the flags for per-interface tracing operations.

```
[edit system processes dhcp-service interface-traceoptions]  
user@host# set flag flag
```

Configuring the Severity Level to Filter Which Extended DHCP Messages Are Logged

The messages associated with a logged event are categorized according to severity level. You can use the severity level to determine which messages are logged for the event type. A low severity level is less restrictive—filters out fewer messages—than a higher level. When you configure a severity level, all messages at that level and all higher (more restrictive) levels are logged.

The following list presents severity levels in order from lowest (least restrictive) to highest (most restrictive). This order also represents the significance of the messages; for example, **error** messages are of greater concern than **info** messages.

- **verbose**
- **info**
- **notice**
- **warning**
- **error**

The severity level that you configure depends on the issue that you are trying to resolve. In some cases you might be interested in seeing all messages relevant to the logged event, so you specify **all**. You can also specify **verbose** with the same result, because **verbose** is the lowest (least restrictive) severity level; it has nothing to do with the terseness or verbosity of the messages. Either choice generates a large amount of output. You can specify a more restrictive severity level, such as **notice** or **info** to filter the messages. By default, the trace operation output includes only messages with a severity level of **error**.

DHCP local server and DHCP relay agent both support the **level** option for the **traceoptions** statement and the **interface-traceoptions** statement. To configure the flags for the events to be logged:

- Specify the severity level for global tracing operations.

```
[edit system processes dhcp-service traceoptions]  
user@host# set level severity
```

- Specify the severity level for per-interface tracing operations.

```
[edit system processes dhcp-service interface-traceoptions]  
user@host# set level severity
```

Tracing Extended DHCP Operations for Specific Interfaces

In addition to the global DHCP tracing operations, subscriber management enables you to trace extended DHCP operations for a specific interface or for a range of interfaces.

Configuring per-interface tracing is a two-step procedure. In the first step, you specify the tracing options that you want to use, such as file information and flags. In the second step, you enable the tracing operation on the specific interfaces.

To configure per-interface tracing operations:

1. Specify the tracing options you want to use.



NOTE: Per-interface tracing uses the same default tracing behavior as the global extended DHCP tracing operation. The default behavior is described in [“Tracing Extended DHCP Operations” on page 1245](#).

- a. Specify that you want to configure per-interface tracing options.

- For DHCP local server, DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent:

```
[edit system processes dhcp-service]  
user@host# edit interface-traceoptions
```

- b. (Optional) Specify the tracing file options.

- Configure the name for the file used for the trace output.

See [“Configuring the Extended DHCP Log Filename” on page 1247](#).

- Configure the number and size of the log files.

See [“Configuring the Number and Size of Extended DHCP Log Files” on page 1247](#).

- Configure access to the log file.

See [“Configuring Access to the Extended DHCP Log File” on page 1248](#).

- Configure a regular expression to filter logging events.

See [“Configuring a Regular Expression for Extended DHCP Messages to Be Logged” on page 1248](#).

- c. (Optional) Specify tracing flag options.

See [“Configuring the Extended DHCP Tracing Flags” on page 1249](#).

- d. (Optional) Configure a severity level for messages to specify which event messages are logged.

See [“Configuring the Severity Level to Filter Which Extended DHCP Messages Are Logged” on page 1249](#).

2. Enable tracing on an interface or interface range.

The following examples show a DHCP local server configuration. You can also use the **trace** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level and at the **[edit system services dhcp-local-server dhcpv6]** hierarchy level.

- Enable tracing on a specific interface.

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

- Enable tracing on a range of interfaces.

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

Tracing Extended DHCP Operations for Specific Interfaces

In addition to the global DHCP tracing operations, subscriber management enables you to trace extended DHCP operations for a specific interface or for a range of interfaces.

Configuring per-interface tracing is a two-step procedure. In the first step, you specify the tracing options that you want to use, such as file information and flags. In the second step, you enable the tracing operation on the specific interfaces.

To configure per-interface tracing operations:

1. Specify the tracing options you want to use.



NOTE: Per-interface tracing uses the same default tracing behavior as the global extended DHCP tracing operation. The default behavior is described in [“Tracing Extended DHCP Operations” on page 1245](#).

- a. Specify that you want to configure per-interface tracing options.

- For DHCP local server, DHCPv6 local server, DHCP relay agent, and DHCPv6 relay agent:

```
[edit system processes dhcp-service]
user@host# edit interface-traceoptions
```

- b. (Optional) Specify the tracing file options.

- Configure the name for the file used for the trace output.

See [“Configuring the Extended DHCP Log Filename”](#) on page 1247.

- Configure the number and size of the log files.

See [“Configuring the Number and Size of Extended DHCP Log Files”](#) on page 1247.

- Configure access to the log file.

See [“Configuring Access to the Extended DHCP Log File”](#) on page 1248.

- Configure a regular expression to filter logging events.

See [“Configuring a Regular Expression for Extended DHCP Messages to Be Logged”](#) on page 1248.

- c. (Optional) Specify tracing flag options.

See [“Configuring the Extended DHCP Tracing Flags”](#) on page 1249.

- d. (Optional) Configure a severity level for messages to specify which event messages are logged.

See [“Configuring the Severity Level to Filter Which Extended DHCP Messages Are Logged”](#) on page 1249.

- 2. Enable tracing on an interface or interface range.

The following examples show a DHCP local server configuration. You can also use the **trace** statement at the **[edit forwarding-options dhcp-relay]** hierarchy level and at the **[edit system services dhcp-local-server dhcpv6]** hierarchy level.

- Enable tracing on a specific interface.

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

- Enable tracing on a range of interfaces.

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

**Related
Documentation**

- [Tracing Extended DHCP Operations](#) on page 1245

Troubleshooting Configuration Statements

interface-traceoptions (DHCP)

Syntax	<pre>interface-traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i> > <size <i>maximum-file-size</i>> <world-readable no-world-readable>; flag <i>flag</i>; level (all error info notice verbose warning); no-remote-trace; }</pre>
Hierarchy Level	[edit system processes dhcp-service]
Release Information	Statement introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>Configure extended DHCP tracing operations that can be enabled on a specific interface or group of interfaces.</p> <p>Replaces deprecated interface-traceoptions statements at the [edit forwarding-options dhcp-relay] and [edit system services dhcp-local-server] hierarchy levels.</p> <p>To enable the tracing operation on the specific interfaces, you use the interface <i>interface-name</i> trace statement.</p>
Options	<p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000</p> <p>Default: 3 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements</p> <ul style="list-style-type: none">• all—Trace all events• packet—Trace packet and option decoding operations• state—Trace changes in state <p>level—Level of tracing to perform; also known as severity level. The option you configure enables tracing of events at that level and all higher (more restrictive) levels. You can specify any of the following levels:</p> <ul style="list-style-type: none">• all—Match messages of all levels.• error—Match error messages.• info—Match informational messages.• notice—Match notice messages about conditions requiring special handling.

- **verbose**—Match verbose messages. This is the lowest (least restrictive) severity level; when you configure **verbose**, messages at all higher levels are traced. Therefore, the result is the same as when you configure **all**.

- **warning**—Match warning messages.

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access.

size *maximum-file-size*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: *sizek* to specify KB, *sizem* to specify MB, or *sizeg* to specify GB

Range: 10240 through 1073741824

world-readable—(Optional) Enable unrestricted file access.

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

Related Documentation	<ul style="list-style-type: none"> • Tracing Extended DHCP Operations for Specific Interfaces on page 1250
------------------------------	---

trace (DHCP Local Server)

Syntax	trace;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server dhcpv6 group <i>group-name</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i>],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> interface <i>interface-name</i>],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Enable trace operations for a group of interfaces or for a specific interface within a group.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Tracing Extended DHCP Operations on page 1245• Tracing Extended DHCP Operations for Specific Interfaces on page 1250

trace (DHCP Relay Agent)

Syntax	trace;
Hierarchy Level	<p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> interface <i>interface-name</i>],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay group <i>group-name</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.4.</p> <p>Support at the [edit ... dhcpv6] hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
Description	<p>Enable trace operations for a group of interfaces or for a specific interface within a group. Use the statement at the [edit ... dhcpv6] hierarchy levels to configure DHCPv6 support.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Tracing Extended DHCP Operations on page 1245 • Tracing Extended DHCP Operations for Specific Interfaces on page 1250

traceoptions (DHCP)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i> > <size <i>maximum-file-size</i>> <world-readable no-world-readable>; flag <i>flag</i>; level (all error info notice verbose warning); no-remote-trace; }</pre>
Hierarchy Level	[edit system processes dhcp-service]
Release Information	Statement introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.1 for EX Series switches.
Description	<p>Define global tracing operations for extended DHCP local server and extended DHCP relay agent processes.</p> <p>Replaces deprecated traceoptions statements at the [edit forwarding-options dhcp-relay] and [edit system services dhcp-local-server] hierarchy levels.</p>
Options	<p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000</p> <p>Default: 3 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements</p> <ul style="list-style-type: none">• all—Trace all events.• auth—Trace authentication events.• database—Trace database events.• fwd—Trace firewall process events.• general—Trace miscellaneous events.• ha—Trace high availability-related events.• interface—Trace interface operations.• io—Trace I/O operations.• packet—Trace packet and option decoding operations.• performance—Trace performance measurement operations.• profile—Trace profile operations.• rpd—Trace routing protocol process events.

- **rtsock**—Trace routing socket operations.
- **session-db**—Trace session database events.
- **state**—Trace changes in state.
- **statistics**—Trace baseline statistics.
- **ui**—Trace user interface operations.

level—Level of tracing to perform; also known as severity level. You can specify any of the following levels:

- **all**—Match all levels.
- **error**—Match error conditions.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access, allowing only the user **root** and users who have the Junos OS **maintenance** permission to access the trace files.

size *maximum-file-size*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: *sizek* to specify KB, *sizem* to specify MB, or *sizeg* to specify GB

Range: 10240 through 1073741824

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

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

Related Documentation	<ul style="list-style-type: none"> • Tracing Extended DHCP Operations on page 1245
------------------------------	---

PART 3

Features

- [Class of Service on page 1263](#)
- [Device Security on page 1803](#)
- [Ethernet Switching on page 1821](#)
- [High Availability on page 2061](#)
- [Interfaces on page 2279](#)
- [Layer 3 Protocols on page 2505](#)
- [LLDP on page 3243](#)
- [Logical Systems on page 3259](#)
- [MPLS on page 3535](#)
- [Multicast on page 3637](#)
- [Network Management and Monitoring on page 4073](#)
- [Port Mirroring on page 4191](#)
- [Routing Policy and Packet Filtering on page 4463](#)
- [Spanning-Tree Protocols on page 5031](#)
- [VPNs on page 5151](#)

CHAPTER 6

Class of Service

- [Class of Service Overview and Examples on page 1263](#)
- [Class of Service on Ethernet Interfaces on page 1311](#)
- [Class of Service for MPLS on page 1317](#)
- [Classifying Packets by Behavior Aggregate on page 1319](#)
- [Defining Code-Point Aliases on page 1371](#)
- [Classifying Packets Based on Various Packet Header Fields on page 1386](#)
- [Tricolor Marking Policers on page 1440](#)
- [Forwarding Classes on page 1525](#)
- [Forwarding Policy Options on page 1556](#)
- [Schedulers on page 1573](#)
- [Queue-Level Bandwidth Sharing on page 1666](#)
- [RED Drop Profiles on page 1673](#)
- [Rewriting Packet Header Information on page 1689](#)
- [Routing Engine Protocol Queue Assignments on page 1745](#)
- [CoS for Tunnels on page 1765](#)

Class of Service Overview and Examples

- [Overview on page 1263](#)
- [Configuration on page 1275](#)

Overview

- [CoS Overview on page 1263](#)
- [CoS Input and Output Configuration on page 1272](#)
- [Packet Flow Through the CoS Process on page 1273](#)

CoS Overview

- [CoS Overview on page 1264](#)
- [CoS Standards on page 1264](#)
- [Understanding Packet Flow Across a Network on page 1265](#)

- [Junos CoS Components on page 1266](#)
- [Default CoS Settings on page 1267](#)
- [CoS Applications Overview on page 1269](#)
- [Interface Types That Do Not Support CoS on page 1270](#)
- [VPLS and Default CoS Classification on page 1271](#)

CoS Overview

When a network experiences congestion and delay, some packets must be dropped. The Juniper Networks® Junos® operating system (Junos OS) class of service (CoS) enables you to divide traffic into classes and offer various levels of throughput and packet loss when congestion occurs. This allows packet loss to happen according to rules that you configure.

For interfaces that carry IPv4, IPv6, and MPLS traffic, you can configure the Junos OS CoS features to provide multiple classes of service for different applications. On the routing device, you can configure multiple forwarding classes for transmitting packets, define which packets are placed into each output queue, schedule the transmission service level for each queue, and manage congestion using a random early detection (RED) algorithm.

The Junos OS CoS features provide a set of mechanisms that you can use to provide differentiated services when best-effort traffic delivery is insufficient. In designing CoS applications, you must give careful consideration to your service needs, and you must thoroughly plan and design your CoS configuration to ensure consistency across all routing devices in a CoS domain. You must also consider all the routing devices and other networking equipment in the CoS domain to ensure interoperability among all equipment.

Because Juniper Networks routing devices implement CoS in hardware rather than in software, you can experiment with and deploy CoS features without adversely affecting packet forwarding and routing performance.

Related Documentation

- *Hardware Capabilities and Limitations*

CoS Standards

The standards for Juniper Networks® Junos® operating system (Junos OS) class of service (CoS) capabilities are defined in the following RFCs:

- RFC 2474, *Definition of the Differentiated Services Field in the IPv4 and IPv6 Headers*
- RFC 2597, *Assured Forwarding PHB Group*
- RFC 2598, *An Expedited Forwarding PHB*
- RFC 2698, *A Two Rate Three Color Marker*


Understanding Packet Flow Across a Network

CoS works by examining traffic entering at the edge of your network. The edge routing devices classify traffic into defined service groups, to provide the special treatment of traffic across the network. For example, voice traffic can be sent across certain links, and data traffic can use other links. In addition, the data traffic streams can be serviced differently along the network path to ensure that higher-paying customers receive better service. As the traffic leaves the network at the far edge, you can reclassify the traffic.

To support CoS, you must configure each routing device in the network. Generally, each routing device examines the packets that enter it to determine their CoS settings. These settings then dictate which packets are first transmitted to the next downstream routing device. In addition, the routing devices at the edges of the network might be required to alter the CoS settings of the packets that enter the network from the customer or peer networks.

In [Figure 5 on page 1265](#), Router A is receiving traffic from a customer network. As each packet enters, Router A examines the packet's current CoS settings and classifies the traffic into one of the groupings defined by the Internet service provider (ISP). This definition allows Router A to prioritize its resources for servicing the traffic streams it is receiving. In addition, Router A might alter the CoS settings (forwarding class and loss priority) of the packets to better match the ISP's traffic groups. When Router B receives the packets, it examines the CoS settings, determines the appropriate traffic group, and processes the packet according to those settings. It then transmits the packets to Router C, which performs the same actions. Router D also examines the packets and determines the appropriate group. Because Router D sits at the far end of the network, the ISP might decide once again to alter the CoS settings of the packets before Router D transmits them to the neighboring network.

Figure 5: Packet Flow Across the Network



Junos CoS Components

The Juniper Networks® Junos® operating system (Junos OS) CoS consists of many components that you can combine and tune to provide the level of services required by customers.

The Junos OS CoS components include:

- **Code-point aliases**—A *code-point alias* assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components, such as classifiers, drop-profile maps, and rewrite rules.
- **Classifiers**—*Packet classification* refers to the examination of an incoming packet. This function associates the packet with a particular CoS servicing level. In the Junos OS, classifiers associate incoming packets with a forwarding class and loss priority and, based on the associated forwarding class, assign packets to output queues. Two general types of classifiers are supported:
 - **Behavior aggregate or CoS value traffic classifiers**—A *behavior aggregate* (BA) is a method of classification that operates on a packet as it enters the routing device. The CoS value in the packet header is examined, and this single field determines the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point (DSCP) value, DSCP IPv6 value, IP precedence value, MPLS EXP bits, and IEEE 802.1p value. The default classifier is based on the IP precedence value.
 - **Multifield traffic classifiers**—A *multifield* classifier is a second method for classifying traffic flows. Unlike a behavior aggregate, a multifield classifier can examine multiple fields in the packet. Examples of some fields that a multifield classifier can examine include the source and destination address of the packet as well as the source and destination port numbers of the packet. With multifield classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.
- **Forwarding classes**—The *forwarding classes* affect the forwarding, scheduling, and marking policies applied to packets as they transit a routing device. The forwarding class plus the loss priority define the per-hop behavior. Four categories of forwarding classes are supported: best effort, assured forwarding, expedited forwarding, and network control. For Juniper Networks M Series Multiservice Edge Routers, four forwarding classes are supported. You can configure up to one each of the four types of forwarding classes. For M120 and M320 Multiservice Edge Routers, Juniper Networks MX Series 3D Universal Edge Routers, Juniper Networks T Series Core Routers and EX Series switches, 16 forwarding classes are supported, so you can classify packets more granularly. For example, you can configure multiple classes of expedited forwarding (EF) traffic: EF, EF1, and EF2.
- **Loss priorities**—*Loss priorities* allow you to set the priority of dropping a packet. Loss priority affects the scheduling of a packet without affecting the packet's relative ordering. You can use the packet loss priority (PLP) bit as part of a congestion control strategy. You can use the loss priority setting to identify packets that have experienced congestion. Typically you mark packets exceeding some service level with a high loss priority. You set loss priority by configuring a classifier or a policer. The loss priority is used later in the workflow to select one of the drop profiles used by RED.

- Forwarding policy options—These options allow you to associate forwarding classes with next hops. Forwarding policy also allows you to create classification overrides, which assign forwarding classes to sets of prefixes.
- Transmission scheduling and rate control—These parameters provide you with a variety of tools to manage traffic flows:
 - Queuing—After a packet is sent to the outgoing interface on a routing device, it is queued for transmission on the physical media. The amount of time a packet is queued on the routing device is determined by the availability of the outgoing physical media as well as the amount of traffic using the interface.
 - Schedulers—An individual routing device interface has multiple queues assigned to store packets. The routing device determines which queue to service based on a particular method of scheduling. This process often involves a determination of which type of packet should be transmitted before another. The Junos OS schedulers allow you to define the priority, bandwidth, delay buffer size, rate control status, and RED drop profiles to be applied to a particular queue for packet transmission.
 - Fabric schedulers—For M320 and T Series routers only, fabric schedulers allow you to identify a packet as high or low priority based on its forwarding class, and to associate schedulers with the fabric priorities.
 - Policers for traffic classes—*Policers* allow you to limit traffic of a certain class to a specified bandwidth and burst size. Packets exceeding the policer limits can be discarded, or can be assigned to a different forwarding class, a different loss priority, or both. You define policers with filters that can be associated with input or output interfaces.
- Rewrite rules—A *rewrite rule* sets the appropriate CoS bits in the outgoing packet. This allows the next downstream routing device to classify the packet into the appropriate service group. Rewriting, or marking, outbound packets is useful when the routing device is at the border of a network and must alter the CoS values to meet the policies of the targeted peer.

Default CoS Settings

If you do not configure any CoS settings on your routing device, the software performs some CoS functions to ensure that user traffic and protocol packets are forwarded with minimum delay when the network is experiencing congestion. Some default mappings are automatically applied to each logical interface that you configure. Other default mappings, such as explicit default classifiers and rewrite rules, are in operation only if you explicitly associate them with an interface.

You can display default CoS settings by issuing the **show class-of-service** operational mode command. This section includes sample output displaying the default CoS settings. The sample output is truncated for brevity.

show class-of-service

```
user@host> show class-of-service
```

Default Forwarding Classes

Forwarding class	Queue
best-effort	0
expedited-forwarding	1
assured-forwarding	2
network-control	3

Default Code-Point Aliases

```
Code point type: dscp
  Alias      Bit pattern
  af11      001010
  af12      001100
...
Code point type: dscp-ipv6
...
Code point type: exp
...
Code point type: ieee-802.1
...
Code point type: inet-precedence
...
```

Default Classifiers

```
Classifier: dscp-default, Code point type: dscp, Index: 7
...

Classifier: dscp-ipv6-default, Code point type: dscp-ipv6, Index: 8
...

Classifier: exp-default, Code point type: exp, Index: 9
...

Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 10
...

Classifier: ipprec-default, Code point type: inet-precedence, Index: 11
...

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 12
...
```

Default Frame Relay Loss Priority Map

```
Loss-priority-map: frame-relay-de-default, Code point type: frame-relay-de, Index:
13
  Code point      Loss priority
  0               low
  1               high
```

Default Rewrite Rules

```
Rewrite rule: dscp-default, Code point type: dscp, Index: 24
  Forwarding class      Loss priority      Code point
  best-effort           low           000000
  best-effort           high          000000
...

Rewrite rule: dscp-ipv6-default, Code point type: dscp-ipv6, Index: 25
```

```

...

Rewrite rule: exp-default, Code point type: exp, Index: 26
...

Rewrite rule: ieee8021p-default, Code point type: ieee-802.1, Index: 27
...

Rewrite rule: ipprec-default, Code point type: inet-precedence, Index: 28
...

```

Default Drop Profile

```

Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100              100

```

Default Schedulers

```

Scheduler map: <default>, Index: 2

  Scheduler: <default-be>, Forwarding class: best-effort, Index: 17
    Transmit rate: 95 percent, Rate Limit: none, Buffer size: 95 percent, Priority:
    low
    Drop profiles:
      Loss priority  Protocol    Index    Name
      Low           Any         1        <default-drop-profile>
      High          Any         1        <default-drop-profile>
...

```

Related Documentation

- [Default Forwarding Classes](#)
- [Default Behavior Aggregate Classification Overview on page 1322](#)
- [Default Drop Profile on page 1676](#)
- [Default Schedulers Overview on page 1575](#)
- [Default Fabric Priority Queuing on page 1578](#)

CoS Applications Overview

You can configure CoS features to meet your application needs. Because the components are generic, you can use a single CoS configuration syntax across multiple routing devices. CoS mechanisms are useful for two broad classes of applications. These applications can be referred to as *in the box* and *across the network*.

In-the-box applications use CoS mechanisms to provide special treatment for packets passing through a single node on the network. You can monitor the incoming traffic on each interface, using CoS to provide preferred service to some interfaces (that is, to some customers) while limiting the service provided to other interfaces. You can also filter outgoing traffic by the packet's destination, thus providing preferred service to some destinations.

Across-the-network applications use CoS mechanisms to provide differentiated treatment to different classes of packets across a set of nodes in a network. In these types of applications, you typically control the ingress and egress routing devices to a routing domain and all the routing devices within the domain. You can use the Junos OS CoS

features to modify packets traveling through the domain to indicate the packet's priority across the domain.

Specifically, you modify the CoS code points in packet headers, remapping these bits to values that correspond to levels of service. When all routing devices in the domain are configured to associate the precedence bits with specific service levels, packets traveling across the domain receive the same level of service from the ingress point to the egress point. For CoS to work in this case, the mapping between the precedence bits and service levels must be identical across all routing devices in the domain.

The Junos OS CoS applications support the following range of mechanisms:

- **Differentiated Services (DiffServ)**—The CoS application supports DiffServ, which uses 6-bit IPv4 and IPv6 header type-of-service (ToS) byte settings. The configuration uses CoS values in the IP and IPv6 ToS fields to determine the forwarding class associated with each packet.
- **Layer 2 to Layer 3 CoS mapping**—The CoS application supports mapping of Layer 2 (IEEE 802.1p) packet headers to routing device forwarding class and loss-priority values. Layer 2 to Layer 3 CoS mapping involves setting the forwarding class and loss priority based on information in the Layer 2 header. Output involves mapping the forwarding class and loss priority to a Layer 2-specific marking. You can mark the Layer 2 and Layer 3 headers simultaneously.
- **MPLS EXP**—Supports configuration of mapping of MPLS experimental (EXP) bit settings to routing device forwarding classes and vice versa.
- **VPN outer-label marking**—Supports setting of outer-label EXP bits, also known as CoS bits, based on MPLS EXP mapping.

Interface Types That Do Not Support CoS

For original Channelized OC12 PICs, limited CoS functionality is supported. For more information, contact Juniper Networks customer support.



NOTE: Transmission scheduling is not supported on 8-port, 12-port, and 48-port Fast Ethernet PICs.

You can configure CoS on all interfaces, except the following:

- **cau4**—Channelized STM1 IQ interface (configured on the Channelized STM1 IQ PIC).
- **coc1**—Channelized OC1 IQ interface (configured on the Channelized OC12 IQ PIC).
- **coc12**—Channelized OC12 IQ interface (configured on the Channelized OC12 IQ PIC).
- **cstm-1**—Channelized STM1 IQ interface (configured on the Channelized STM1 IQ PIC).
- **ct1**—Channelized T1 IQ interface (configured on the Channelized DS3 IQ PIC or Channelized OC12 IQ PIC).
- **ct3**—Channelized T3 IQ interface (configured on the Channelized DS3 IQ PIC or Channelized OC12 IQ PIC).

- **ce1**—Channelized E1 IQ interface (configured on the Channelized E1 IQ PIC or Channelized STM1 IQ PIC).
- **dsc**—Discard interface.
- **fxp**—Management and internal Ethernet interfaces.
- **lo**—Loopback interface. This interface is internally generated.
- **pe**—Encapsulates packets destined for the rendezvous point routing device. This interface is present on the first-hop routing device.
- **pd**—De-encapsulates packets at the rendezvous point. This interface is present on the rendezvous point.
- **vt**—Virtual loopback tunnel interface.



NOTE: For channelized interfaces, you can configure CoS on channels, but not at the controller level. For a complex configuration example, see the *Junos OS Feature Guides*.

VPLS and Default CoS Classification

A VPLS routing instance with the **no-tunnel-services** option configured has a default classifier applied to the label-switched interface for all VPLS packets coming from the remote VPLS PE. This default classifier is modifiable only on MX Series routers. On T Series, when **no-tunnel-services** option is configured, the custom classifier for VPLS instances is not supported.



NOTE: With **no-tunnel-services** configured, custom classifier for VPLS routing instances on T Series and LMNR based FPC for M320 is not supported. When a wild card configuration or an explicit routing instances are configured for VPLS on CoS CLI, the custom classifier binding results in default classifier binding on Packet Forwarding Engine (PFE).

For example, on routing devices with eight queues (Juniper Networks M120 and M320 Multiservice Edge Routers, MX Series Ethernet Services Routers, and T Series Core Routers), the default classification applied to **no-tunnel-services** VPLS packets are shown in [Table 90 on page 1271](#).

Table 90: Default VPLS Classifiers

MPLS Label EXP Bits	Forwarding Class/Queue
000	0
001	1
010	2

Table 90: Default VPLS Classifiers (*continued*)

MPLS Label EXP Bits	Forwarding Class/Queue
011	3
100	4
101	5
110	6
111	7



NOTE: Forwarding class to queue number mapping is not always one-to-one. Forwarding classes and queues are only the same when default forwarding-class-to-queue mapping is in effect. For more information about configuring forwarding class and queues, see [“Configuring Forwarding Classes” on page 1530](#).

On MX Series routers, VPLS filters and policers act on a Layer 2 frame that includes the media access control (MAC) header (after any VLAN rewrite or other rules are applied), but does not include the cyclical redundancy check (CRC) field.



NOTE: On MX Series routers, if you apply a counter in a firewall for egress MPLS or VPLS packets with the EXP bits set to 0, the counter will not tally these packets.

CoS Input and Output Configuration

- [CoS Inputs and Outputs Overview on page 1272](#)

CoS Inputs and Outputs Overview

Some CoS components map one set of values to another set of values. Each mapping contains one or more inputs and one or more outputs.

Some CoS components map one set of values to another set of values. Each mapping contains one or more inputs and one or more outputs. When you configure a mapping, you set the outputs for a given set of inputs, as shown in [Table 91 on page 1272](#).

Table 91: CoS Mappings—Inputs and Outputs

CoS Mappings	Inputs	Outputs	Comments
classifiers	code-points	forwarding-class loss-priority	The map sets the forwarding class and PLP for a specific set of code points.
drop-profile-map	loss-priority protocol	drop-profile	The map sets the drop profile for a specific PLP and protocol type.

Table 91: CoS Mappings—Inputs and Outputs (*continued*)

CoS Mappings	Inputs	Outputs	Comments
rewrite-rules	loss-priority	code-points	The map sets the code points for a specific forwarding class and PLP.

Related Documentation

- [Default Behavior Aggregate Classification Overview on page 1322](#)
- [Configuring Drop Profile Maps for Schedulers on page 1590](#)
- [Applying Default Rewrite Rules](#)
- [CoS Inputs and Outputs Examples](#)

Packet Flow Through the CoS Process

- [Packet Flow Through the CoS Process Overview on page 1273](#)

Packet Flow Through the CoS Process Overview

Perhaps the best way to understand Junos CoS is to examine how a packet is treated on its way through the CoS process. This topic includes a description of each step and figures illustrating the process.

The following steps describe the CoS process:

1. A logical interface has one or more classifiers of different types applied to it (at the **[edit class-of-service interfaces]** hierarchy level). The types of classifiers are based on which part of the incoming packet the classifier examines (for example, EXP bits, IEEE 802.1p bits, or DSCP bits). You can use a translation table to rewrite the values of these bits on ingress.



NOTE: You can only rewrite the values of these bits on ingress on the Juniper Networks M40e, M120, M320 Multiservice Edge Routers, and T Series Core Routers with IQE PICs. For more information about rewriting the values of these bits on ingress, see *Configuring ToS Translation Tables*.

2. The classifier assigns the packet to a forwarding class and a loss priority (at the **[edit class-of-service classifiers]** hierarchy level).
3. Each forwarding class is assigned to a queue (at the **[edit class-of-service forwarding-classes]** hierarchy level).
4. Input (and output) policers meter traffic and might change the forwarding class and loss priority if a traffic flow exceeds its service level.
5. The physical or logical interface has a scheduler map applied to it (at the **[edit class-of-service interfaces]** hierarchy level).

At the **[edit class-of-service interfaces]** hierarchy level, the **scheduler-map** and **rewrite-rules** statements affect the outgoing packets, and the **classifiers** statement affects the incoming packets.

6. The scheduler defines how traffic is treated in the output queue—for example, the transmit rate, buffer size, priority, and drop profile (at the **[edit class-of-service schedulers]** hierarchy level).
7. The scheduler map assigns a scheduler to each forwarding class (at the **[edit class-of-service scheduler-maps]** hierarchy level).
8. The drop-profile defines how aggressively to drop packets that are using a particular scheduler (at the **[edit class-of-service drop-profiles]** hierarchy level).
9. The rewrite rule takes effect as the packet leaves a logical interface that has a rewrite rule configured (at the **[edit class-of-service rewrite-rules]** hierarchy level). The rewrite rule writes information to the packet (for example, EXP or DSCP bits) according to the forwarding class and loss priority of the packet.

Figure 6 on page 1274 and Figure 7 on page 1274 show the components of the Junos CoS features, illustrating the sequence in which they interact.

Figure 6: CoS Classifier, Queues, and Scheduler


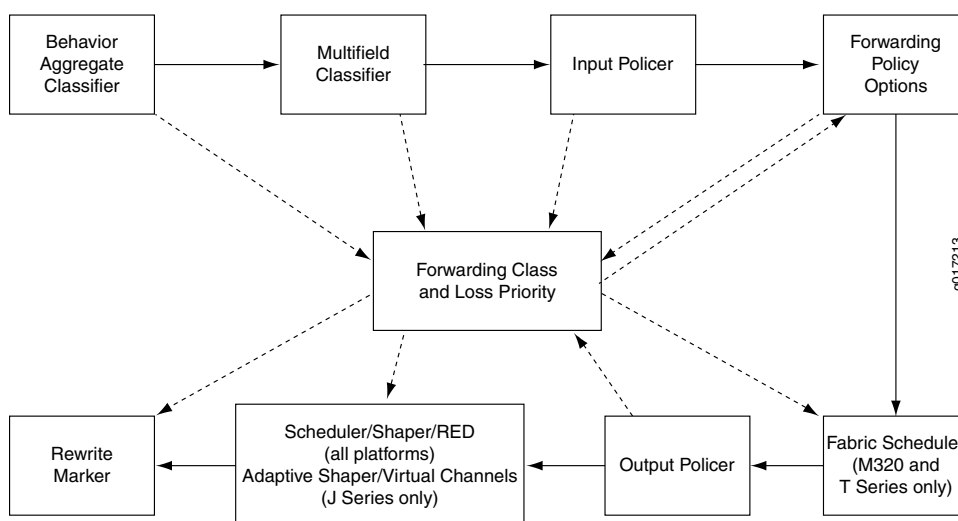


Figure 7: Packet Flow Through CoS Configurable Components



Each outer box in [Figure 7 on page 1274](#) represents a process component. The components in the upper row apply to inbound packets, and the components in the lower row apply to outbound packets. The arrows with the solid lines point in the direction of packet flow.

The middle box (forwarding class and loss priority) represents two data values that can either be inputs to or outputs of the process components. The arrows with the dotted lines indicate inputs and outputs (or settings and actions based on settings). For example, the multifield classifier sets the forwarding class and loss priority of incoming packets. This means that the forwarding class and loss priority are outputs of the classifier; thus, the arrow points away from the classifier. The scheduler receives the forwarding class and loss priority settings, and queues the outgoing packet based on those settings. This means that the forwarding class and loss priority are inputs to the scheduler; thus, the arrow points to the scheduler.

Typically, only a combination of some components (not all) is used to define a CoS service offering.

Related Documentation

- [Packet Flow Through the CoS Process Configuration Example](#)

Configuration

- [Configuration Statements on page 1275](#)

Configuration Statements

- [\[edit chassis\] Hierarchy Level on page 1275](#)
- [\[edit class-of-service\] Hierarchy Level on page 1283](#)
- [\[edit firewall\] Hierarchy Level on page 1287](#)
- [\[edit interfaces\] Hierarchy Level on page 1300](#)

[edit chassis] Hierarchy Level

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
      lacp {
        link-protection {
          non-revertive;
        }
        system-priority;
      }
    }
    sonet {
      device-count number;
    }
    maximum-links maximum-links-limit;
  }
  alarm {
    ds1 {
      ais (ignore | red | yellow);
      ylw (ignore | red | yellow);
    }
  }
}
```

```
ethernet {
    link-down (ignore | red | yellow);
}
integrated-services {
    failure (ignore | red | yellow);
}
management-ethernet {
    link-down (ignore | red | yellow);
}
relay
input {
    port port-number {
        mode (close | open);
        trigger (ignore | red | yellow);
    }
}
output {
    port port-number {
        input-relay input-relay;
        mode (close | open);
        temperature;
    }
}
serial {
    cts-absent (ignore | red | yellow);
    dcd-absent (ignore | red | yellow);
    dsr-absent (ignore | red | yellow);
    loss-of-rx-clock (ignore | red | yellow);
    loss-of-tx-clock (ignore | red | yellow);
    tm-absent (ignore | red | yellow);
}
services {
    hw-down (ignore | red | yellow);
    linkdown (ignore | red | yellow);
    pic-hold-reset (ignore | red | yellow);
    pic-reset (ignore | red | yellow);
    rx-errors (ignore | red | yellow);
    sw-down (ignore | red | yellow);
    tx-errors (ignore | red | yellow);
}
sonet {
    (ais-l | ais-p | ber-sd | ber-sf | locd | lol | lop-p | los | pll | plm-p | rfi-l | rfl-p | uneq-p)
    (ignore | red | yellow);
}
t3 {
    (ais | exz | ferf | idle | lcv | lof | los | pll | ylw) (ignore | red | yellow);
}
}
cluster {
    control-link-recovery;
    control-ports {
        fpc slot-number port port-number;
    }
    heartbeat-interval milliseconds;
    heartbeat-threshold number;
    redundancy-group {
```

```

... the redundancy-group subhierarchy appears at the end of the [edit chassis cluster]
    hierarchy ...
}
reth-count number;
traceoptions {
  file <filename> <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag;
  level severity;
  no-remote-trace;
}
redundancy-group group-number {
  gratuitous-arp-count number;
  hold-down-interval seconds;
  interface-monitor {
    interface-name weight number;
  }
  ip-monitoring {
    family {
      inet {
        ipv4-address {
          interface rethindex.logical-unit-number secondary-ip-address ipv4-address;
          weight number;
        }
      }
    }
    global-threshold number;
    global-weight number;
    retry-count count;
    retry-interval interval;
  }
  node node-number priority priority-number;
  preempt;
}
config-button {
  no-clear;
  no-rescue;
}
container-devices {
  device-count number;
}
craft-lockout;
disable-power-management;
disk-partition partition-name (/config | /var) {
  level (full | high) {
    free-space threshold-value (mb | percent);
  }
}
enhanced-policer;
extended-statistics;
fabric {
  degraded {
    action-fpc-restart-disable;
    degraded-fabric-detection-enable
    degraded-fpc-bad-plane-threshold number-bad-planes;
  }
}

```

```
    redundancy-mode (increased-bandwidth | redundant);
  }
  filter;
  fpc slot-number {
    ... the fpc subhierarchy appears after the main [edit chassis] hierarchy ...
  }
  fpc-feb-connectivity {
    fpc slot-number feb (slot-number | none);
  }
  fpc-resync;
  fru-poweron-sequence sequence;
  lcc index {
    ... the lcc subhierarchy appears after the main [edit chassis] hierarchy ...
  }
  maximum-ecmp value;
  memory-enhanced {
    filter;
    route;
    vpn-label;
  }
  network-services (ethernet | enhanced-ethernet | ip | enhanced-ip);
  (packet-scheduling | no-packet-scheduling);
  pem {
    minimum number;
  }
  policer-drop-probability-low;
  ppp-subscriber-services (disable | enable);
  redundancy {
    cfeb slot (always | preferred);
    failover {
      on-disk-failure;
      on-loss-of-keepalives;
    }
    feb {
      redundancy-group group-name {
        description description;
        feb slot-number <backup | primary>;
        no-auto-failover;
      }
    }
    graceful-switchover;
    keepalive-time seconds;
    routing-engine slot-number (backup | disabled | master);
    sfm slot-number (always | preferred);
    ssb slot-number (always | preferred);
  }
  route-memory-enhanced;
  route-localization {
    inet (chassis);
    inet6;
  }
  routing-engine {
    bios {
      no-auto-upgrade;
    }
    on-disk-failure disk-failure-action (halt | reboot);
```

```

}
sfm slot-number {
    power off;
}
sib {
    minimum number;
}
(source-route | no-source-route);
state [
    cb-upgrade [on | off];
]
synchronization { # for M Series and T Series routers
    primary (external-a | external-b);
    secondary (external-a | external-b);
    signal-type (e1 | t1);
    switching-mode (non-revertive | revertive);
    transmitter-enable;
    validation-interval seconds;
    y-cable-line-termination;
}
synchronization { # for MX80 and MX240 routers
    clock-mode (auto-select | free-run);
    esmc-transmit {
        interfaces (all | interface-name);
    }
    hold-interval {
        configuration-change seconds;
        restart seconds;
        switchover seconds;
    }
    network-option (option-1 | option-2);
    quality-mode-enable;
    selection-mode (configured-quality|received-quality);
    source {
        (external-a | external-b) {
            priority number;
            quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
            request (force-switch | lockout);
        }
        interfaces interface-name {
            priority number;
            quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
            request (force-switch | lockout);
            wait-to-restore minutes;
        }
    }
    switchover-mode (revertive | non-revertive);
}
synchronization { # for ACX Series routers
    clock-mode (auto-select | free-run);
    esmc-transmit {
        interfaces (all | interface-name);
    }
    hold-interval {
        configuration-change seconds;
        restart seconds;

```

```

    switchover seconds;
}
network-option (option-1 | option-2);
quality-mode-enable;
selection-mode (configured-quality | received-quality);
source {
    (bits | gps) {
        priority number;
        quality-level (prc | prs | sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
        request (force-switch | lockout);
    }
    interfaces interface-name {
        priority number;
        quality-level (prc | prs | sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
        request (force-switch | lockout);
        wait-to-restore minutes;
    }
}
switchover-mode(non-revertive | revertive);
}
system-domains {
    protected-system-domains psdnumerical-index {
        control-plane-bandwidth-percent percent;
        control-slot-numbers [ slot-numbers ];
        control-system-id control-system-id;
        description description;
        fpcs [ slot-numbers ];
    }
    root-domain-id root-domain-id;
}
vrf-mtu-check;
}

chassis {
    fpc slot-number {
        number-of-ports active-ports;
        offline;
        pic slot-number {
            ... the pic subhierarchy appears after the main [edit chassis fpc slot-number] hierarchy
            ...
        }
        port-mirror-instance port-mirror-instance-name;
        power (off | on);
        sampling-instance instance-name;
    }
}

fpc slot-number {
    pic slot-number {
        adaptive-services {
            service-package (layer-2 | layer-3 | ...the following extension-provider subhierarchy
                ...);
            extension-provider {
                control-cores number;
                data-cores number;
                data-flow-affinity {
                    hash-key (layer-3 | layer-4);

```

```

    }
    channelization;
    forwarding-db-size megabytes;
    object-cache-size megabytes;
    package package-name;
    policy-db-size megabytes;
    syslog {
        facility {
            severity;
            destination (pic-console | routing-engine);
        }
    }
    wired-process-mem-size megabytes;
}
aggregated-devices {
    ima {
        device-count number;
    }
}
aggregate-ports;
atm-cell-relay-accumulation;
atm-l2circuit-mode (aal5 | cell | trunk trunk);
cel {
    e1 port-number {
        channel-group group-number timeslots slot-number;
    }
}
ct3 {
    port port-number {
        t1 link-number {
            channel-group group-number timeslots slot-number;
        }
    }
}
ethernet {
    pic-mode (enhanced-switching | routing | switching);
}
fibre-channel {
    port port-number;
    port-range port-range-low port-range-high
}
egress-policer-overhead bytes;
forwarding-mode {
    sa-multicast;
    vlan-steering {
        vlan-rule (high-low | odd-even);
    }
}
framing (e1 | e3 | sdh | sonet | t1 | t3);
idle-cell-format {
    itu-t;
    payload-pattern payload-pattern-byte;
}
ingress-policer-overhead bytes;
inline-services {

```

```
        bandwidth (1g | 10g);
    }
    linerate-mode;
    max-queues-per-interface (4 | 8);
    mlfr-uni-nni-bundles number;
    no-concatenate;
    no-multi-rate;
    port port-number {
        framing (e1 | e3 | sdh | sonet | t1 | t3);
        forwarding-mode {
            sa-multicast;
        }
        speed ( oc3-stm1 | oc12-stm4 | oc48-stm16);
    }
    port-mirror-instance port-mirror-instance-name;
    q-pic-large-buffer {
        (large-scale | small-scale);
    }
    red-buffer-occupancy {
        weighted-averaged <instant-usage-weight-exponent weight-value>;
    }
    shdsl {
        pic-mode (1-port-atm | 2-port-atm);
    }
    sparse-dlcis;
    traffic-manager {
        egress-shaping-overhead number;
        ingress-shaping-overhead number;
        mode {
            egress-only;
            ingress-and-egress;
            session-shaping;
        }
    }
    tunnel-queuing;
    tunnel-services {
        bandwidth (1g | 10g | 20g | 40g);
        tunnel-only;
    }
    vtmapping (itu-t | klm);
}
}

chassis {
    lcc index {
        fpc slot-number {
            ... the fpc subhierarchy appears after the main [edit chassis lcc index] hierarchy ...
        }
        offline;
        online-expected;
    }
}

lcc index {
    fpc slot-number {
```



```

pic slot-number {
  ... the pic subhierarchy appears after the main [edit chassis lcc index fpc slot-number]
  hierarchy ...
}
power (off | on);
sampling-instance instance-name;
}

fpc slot-number {
  pic slot-number {
    aggregate-ports;
    atm-cell-relay-accumulation;
    atm-l2circuit-mode (aal5 | cell | trunk trunk);
    framing (e1 | e3 | sdh | sonet | t1 | t3);
    idle-cell-format {
      itu-t;
      payload-pattern payload-pattern-byte;
    }
    linerate-mode;
    max-queues-per-interface (4 | 8);
    no-concatenate;
    no-mcast-replication;
    no-pre-classifier;
    port port-number {
      framing (e1 | e3 | sdh | sonet | t1 | t3);
    }
    q-pic-large-buffer {
      (large-scale | small-scale);
    }
    red-buffer-occupancy {
      weighted-averaged <instant-usage-weight-exponent weight-value>;
    }
    shdsl {
      pic-mode (1-port-atm | 2-port-atm);
    }
    traffic-manager {
      egress-shaping-overhead bytes;
      ingress-shaping-overhead bytes;
      mode {
        egress-only;
        ingress-and-egress;
      }
    }
  }
}
}

```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

[edit class-of-service] Hierarchy Level

```

class-of-service {
  classifiers {
    type classifier-name {
      forwarding-class class-name {

```

```

        loss-priority (high | low | medium-high | medium-low) code-points [ aliases bits ];
    }
    import ( classifier-name | default );
}
}
code-point-aliases {
    ( dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence ) {
        alias-name bits;
    }
}
drop-profiles {
    profile-name {
        fill-level percentage drop-probability percentage;
        interpolate {
            drop-probability value;
            fill-level value;
        }
    }
}
fabric {
    scheduler-map {
        priority (high | low) scheduler scheduler-name;
    }
}
forwarding-class-map {
    map-name {
        class class-name queue-num queue-number <restricted-queue queue-number>;
    }
}
forwarding-classes {
    class class-name policing-priority (normal | premium) queue-num queue-number
        priority (high | low);
    queue queue-number class-name policing-priority (normal | premium) priority (high |
        low);
}
forwarding-policy {
    class class-name {
        classification-override {
            forwarding-class class-name;
        }
    }
    next-hop-map map-name {
        forwarding-class class-name {
            discard;
            lsp-next-hop [ lsp-regular-expressions ];
            next-hop [ next-hop-names ];
            non-lsp-next-hop;
        }
    }
}
fragmentation-maps {
    map-name {
        forwarding-class class-name {
            drop-timeout milliseconds;
            fragment-threshold bytes;
            multilink-class number;
        }
    }
}

```

```

        no-fragmentation;
    }
}
host-outbound-traffic {
    dscp-code-point value;
    forwarding-class class-name;
    ieee-802.1 {
        default value;
        rewrite-rules;
    }
    tcp {
        raise-internet-control-priority;
    }
}
interfaces {
    ... the interfaces subhierarchy appears after the main [edit class-of-service] hierarchy
    ...
}
restricted-queues {
    forwarding-class class-name queue-number;
}
rewrite-rules {
    (dscp | dscp-ipv6 | exp | frame-relay-de | ieee-802.1 | ieee-802.1ad | inet-precedence)
    rewrite-rule {
        forwarding-class class-name {
            loss-priority level code-point (alias | bits);
        }
        import (rewrite-rule | default);
    }
}
routing-instances routing-instance-name {
    classifiers {
        dscp (classifier-name | default);
        dscp-ipv6 (classifier-name | default);
        exp (classifier-name | default);
        ieee-208.1 (classifier-name | default | encapsulated | vlan-tag (inner | outer));
    }
}
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
schedulers {
    scheduler-name {
        adjust-minimum value;
        adjust-percent value;
        buffer-size (exact | percent percentage | remainder);
        drop-profile-map loss-priority (any | high | low | medium-high | medium-low)
            protocol any;
        excess-priority (high | low | medium-high | medium-low);
        excess-rate (percent percentage | proportion proportion);
        priority (high | low | medium-high | medium-low | strict-high);
        shaping-rate (bps | percent percentage | burst-size size);
    }
}

```

```

        transmit-rate (bps | percent percentage | remainder) <exact | rate-limit>;
    }
}
traceoptions {
    file <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
traffic-control-profiles {
    profile-name {
        adjust-minimum rate;
        delay-buffer-rate (bps | cps cps | percent percentage);
        excess-rate (percent percentage | proportion value);
        guaranteed-rate (bps | percent percentage) <burst-size bytes>;
        overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
        scheduler-map map-name;
        shaping-rate (bps | percent percentage) <burst-size bytes>;
    }
}
tri-color;
}

class-of-service {
    interfaces {
        interface-name {
            excess-bandwidth-share (equal | proportional value);
            input-excess-bandwidth-share (equal | proportional value);
            input-scheduler-map map-name;
            input-shaping-rate bps;
            input-traffic-control-profile profile-name;
            output-forwarding-class-map map-name;
            output-traffic-control-profile profile-name;
            scheduler-map map-name;
            scheduler-map-chassis (map-name | derived);
            shaping-rate bps;
            unit (logical-unit-number | *) {
                classifiers {
                    dscp (classifier-name | default) {
                        family [ inet mpls ];
                    }
                    dscp-ipv6 (classifier-name | default) {
                        family [ inet mpls ];
                    }
                    exp (classifier-name | default);
                    ieee-208.1 (classifier-name | default) <vlan-tag (inner | outer)>;
                    ieee-208.1ad (classifier-name | default);
                    inet-precedence (classifier-name | default);
                }
                forwarding-class class-name;
                input-scheduler-map map-name;
                input-shaping-rate bps;
                input-traffic-control-profile profile-name shared-instance instance-name;
                loss-priority-maps {
                    (map-name | default);
                }
            }
        }
    }
}

```

```

    loss-priority-rewrites {
        (map-name | default);
    }
    output-forwarding-class-map map-name;
    output-traffic-control-profile profile-name shared-instance instance-name;
    rewrite-rules {
        dscp (rule-name | default) <protocol mpls>;
        dscp-ipv6 (rule-name | default);
        exp (rule-name | default) <protocol [ mpls-any | mpls-inet-both |
            mpls-inet-both-non-vpn ]>;
        exp-push-push-push default;
        exp-swap-push-push default;
        ieee-802.1 (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
        ieee-802.1ad (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
        inet-precedence (rewrite-name | default) <protocol mpls>;
    }
    scheduler-map map-name;
    shaping-rate bps;
    translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 |
        to-exp-from-exp | to-inet-precedence-from-inet-precedence) table-name;
    }
}
interface-set interface-set-name {
    excess-bandwidth-share (equal | proportional value);
    input-excess-bandwidth-share (equal | proportional value);
    input-traffic-control-profile profile-name;
    input-traffic-control-profile-remaining profile-name;
    internal-node;
    output-traffic-control-profile profile-name;
    output-traffic-control-profile-remaining profile-name;
}
}
}

```

Related Documentation

- [Notational Conventions Used in Junos OS Configuration Hierarchies](#)

[edit firewall] Hierarchy Level

Several statements in the **[edit firewall]** hierarchy are valid at numerous locations within the hierarchy. To make the complete hierarchy easier to read, the repeated statements are listed in the following sections, which are referenced at the appropriate locations in [“Complete \[edit firewall\] Hierarchy” on page 325](#).

- [Common Firewall Actions on page 1288](#)
- [Common IP Firewall Actions on page 1288](#)
- [Common IPv4 Firewall Actions on page 1288](#)
- [Common IP Firewall Match Conditions on page 1289](#)
- [Common IPv4 Firewall Match Conditions on page 1290](#)
- [Common Layer 2 Firewall Match Conditions on page 1290](#)
- [Complete \[edit firewall\] Hierarchy on page 1292](#)

Common Firewall Actions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- [edit firewall family (any | ccc | ethernet-switching | inet | inet6 | mpls | vpls) filter *filter-name* term *term-name* then]
- [edit firewall filter *filter-name* term *term-name* then]

The common firewall actions are as follows:

```
count counter-name;  
forwarding-class class-name;  
loss-priority (high | low | medium-high | medium-low);  
next term;  
policer policer-name;  
three-color-policer policer-name {  
    (single-rate single-rate-policer-name | two-rate two-rate-policer-name);  
}
```

Common IP Firewall Actions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- [edit firewall family inet filter *filter-name* term *term-name* then]
- [edit firewall family inet6 filter *filter-name* term *term-name* then]
- [edit firewall filter *filter-name* term *term-name* then]

The common IP firewall actions are as follows:

```
log;  
logical-system logical-system-name <routing-instance routing-instance-name>  
    <topology topology-name>;  
port-mirror;  
port-mirror-instance instance-name;  
routing-instance routing-instance-name <topology topology-name>;  
sample;  
service-filter-hit;  
syslog;  
topology topology-name;
```

Common IPv4 Firewall Actions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- [edit firewall family inet filter *filter-name* term *term-name* then]
- [edit firewall filter *filter-name* term *term-name* then]

The common IP version 4 (IPv4) firewall actions are as follows:

```
(accept | discard <accounting collector-name> | reject <administratively-prohibited |
  bad-host-tos | bad-network-tos | fragmentation-needed | host-prohibited |
  host-unknown | host-unreachable | network-prohibited | network-unknown |
  network-unreachable | port-unreachable | precedence-cutoff | precedence-violation |
  protocol-unreachable | source-host-isolated | source-route-failed | tcp-reset>);
ipsec-sa sa-name;
load-balance sa-name;
next-hop-group group-name;
prefix-action action-name;
```

Common IP Firewall Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit firewall\] Hierarchy](#)” on page 325 instead of the statements being repeated.

- **[edit firewall family inet dialer-filter *filter-name* term *term-name* from]** (with the exceptions noted at this level in “[Complete \[edit firewall\] Hierarchy](#)” on page 325)
- **[edit firewall family inet filter *filter-name* term *term-name* from]**
- **[edit firewall family inet6 dialer-filter *filter-name* term *term-name* from]** (with the exceptions noted at this level in “[Complete \[edit firewall\] Hierarchy](#)” on page 325)
- **[edit firewall family inet6 filter *filter-name* term *term-name* from]**
- **[edit firewall filter *filter-name* term *term-name* from]**

The common IP firewall match conditions are as follows:

```
address {
  ip-prefix </prefix-length> <except>;
}
destination-address {
  ip-prefix </prefix-length> <except>;
}
destination-class [ class-names ] | destination-class-except [ class-names ];
(destination-port [ port-names ] | destination-port-except [ port-names ]);
destination-prefix-list {
  list-name <except>;
}
(forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
 icmp-code [ codes ] | icmp-code-except [ codes ];
 icmp-type [ types ] | icmp-type-except [ types ];
interface interface-name;
(interface-group [ group-names ] | interface-group-except [ group-names ]);
interface-set set-name;
(loss-priority [ priorities ] | loss-priority-except [ priorities ]);
(packet-length [ values ] | packet-length-except [ values ]);
(port [ port-names ] | port-except [ port-names ]);
prefix-list {
  list-name <except>;
}
service-filter-hit;
source-address {
```

```
    ip-prefix </prefix-length> <except>;  
  }  
  (source-class [ class-names ] | source-class-except [ class-names ] );  
  (source-port [ port-names ] | source-port-except [ port-names ] );  
  source-prefix-list {  
    list-name <except>;  
  }  
  tcp-established;  
  tcp-flags flag;  
  tcp-initial;
```

Common IPv4 Firewall Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- [edit firewall family inet dialer-filter *filter-name* term *term-name* from] (with the exceptions noted at this level in [“Complete \[edit firewall\] Hierarchy” on page 325](#))
- [edit firewall family inet filter *filter-name* term *term-name* from]
- [edit firewall filter *filter-name* term *term-name* from]

The common IPv4 firewall match conditions are as follows:

```
(ah-spi [ values ] | ah-spi-except [ values ] );  
(dscp [ code-point-values ] | dscp-except [ code-point-values ] );  
(esp-spi [ values ] | esp-spi-except [ values ] );  
first-fragment;  
fragment-flags flag;  
(fragment-offset [ offsets ] | fragment-offset-except [ offsets ] );  
(ip-options [ option-names ] | ip-options-except [ option-names ] );  
is-fragment;  
(precedence [ precedence-names ] | precedence-except [ precedence-names ] );  
(protocol [ protocol-names ] | protocol-except [ protocol-names ] );  
(ttl [ ttl-values ] | ttl-except [ ttl-values ] );
```

Common Layer 2 Firewall Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- [edit firewall family ethernet-switching filter *filter-name* term *term-name* from]
- [edit firewall family vpls filter *filter-name* term *term-name* from]

The common Layer 2 firewall match conditions are as follows:

```
destination-mac-address {  
  mac-address <except>;  
}  
(destination-port [ port-names ] | destination-port-except [ port-names ] );  
(dscp [ code-point-values ] | dscp-except [ code-point-values ] );  
(ether-type [ protocol-types ] | ether-type-except [ protocol-types ] );  
(forwarding-class [ class-names ] | forwarding-class-except [ class-names ] );
```



```

icmp-code [ codes ] | icmp-code-except [ codes ];
icmp-type [ types ] | icmp-type-except [ types ];
(interface-group [ group-names ] | interface-group-except [ group-names ]);
ip-address {
    ip-prefix </prefix-length> <except>;
}
ip-destination-address {
    ip-prefix </prefix-length> <except>;
}
(ip-precedence [ precedence-names ] | ip-precedence-except [ precedence-names ]);
(ip-protocol [ protocol-names ] | ip-protocol-except [ protocol-names ]);
ip-source-address ip-prefix </prefix-length>;
(learn-vlan-lp-priority [ priorities ] | learn-vlan-lp-priority [ priorities ]);
(learn-vlan-id [ vlan-ids ] | learn-vlan-id-except [ vlan-ids ]);
(loss-priority [ priorities ] | loss-priority-except [ priorities ]);
(port [ port-names ] | port-except [ port-names ]);
source-mac-address {
    mac-address <except>;
}
(source-port [ port-names ] | source-port-except [ port-names ]);
tcp-flags flag;
(traffic-type [ broadcast known-unicast multicast unknown-unicast ] |
    traffic-type-except [ broadcast known-unicast multicast unknown-unicast ]);
(user-vlan-lp-priority [ priorities ] | user-vlan-lp-priority [ priorities ]);
(user-vlan-id [ vlan-ids ] | user-vlan-id-except [ vlan-ids ]);
(vlan-ether-type [ protocol-types ] | vlan-ether-type-except [ protocol-types ]);

```

Complete [edit firewall] Hierarchy

```
firewall {
  family (any | ccc | ethernet-switching | inet | inet6 | mpls | vpls) {
    ... the family subhierarchies appear after the main [edit firewall] hierarchy ...
  }
  filter filter-name {
    accounting-profile [ profile-names ];
    enhanced-mode;
    interface-shared-with;
    interface-specific;
    physical-interface-policer;
    term term-name {
      filter filter-name;
      from {
        ... statements in Common IP Firewall Match Conditions on page 322 AND
        statements in Common IPv4 Firewall Match Conditions on page 323 ...
      }
      then {
        ... statements in Common Firewall Actions on page 320 AND
        statements in Common IP Firewall Actions on page 321 AND
        statements in Common IPv4 Firewall Actions on page 321 ...
      }
    }
  }
  hierarchical-policer policer-name {
    aggregate {
      if-exceeding {
        bandwidth-limit bps;
        burst-size-limit bytes;
      }
      then {
        discard;
        forwarding-class class-name;
        loss-priority (high | low | medium-high | medium-low);
      }
    }
    logical-interface-policer;
    physical-interface-policer;
    premium {
      if-exceeding {
        bandwidth-limit bps;
        burst-size-limit bytes;
      }
      then {
        discard;
      }
    }
  }
  shared-bandwidth-policer;
  interface-set interface-set-name {
    interface-name;
  }
  load-balance-group group-name {
    next-hop-group [ group-names ];
  }
}
```

```

}
policer policer-name {
  filter-specific;
  if-exceeding {
    (bandwidth-limit bps | bandwidth-percent percentage);
    burst-size-limit bytes;
  }
  logical-bandwidth-policer;
  logical-interface-policer;
  physical-interface-policer;
  then {
    discard;
    forwarding-class class-name;
    loss-priority (high | low | medium-high | medium-low);
  }
}
three-color-policer policer-name {
  action {
    loss-priority high then discard;
  }
  filter-specific;
  logical-interface-policer;
  physical-interface-policer;
  shared-bandwidth-policer;
  single-rate {
    (color-aware | color-blind);
    committed-burst-size bytes;
    committed-information-rate bps;
    excess-burst-size bytes;
  }
  two-rate {
    (color-aware | color-blind);
    committed-burst-size bytes;
    committed-information-rate bps;
    peak-burst-size bytes;
    peak-information-rate bps;
  }
}
}

firewall {
  family any {
    filter filter-name {
      interface-shared;
      term term-name {
        from {
          (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
          interface interface-name;
          interface-set set-name;
          (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
          (packet-length [ values ] | packet-length-except [ values ]);
        }
        then {
          ... statements in Common Firewall Actions on page 320 PLUS ...
          (accept | discard);
        }
      }
    }
  }
}

```

```
    }  
  }  
}
```

```
firewall {  
  family ccc {  
    filter filter-name {  
      accounting-profile [ profile-names ];  
      physical-interface-filter;  
      interface-specific;  
      term term-name {  
        filter filter-name;  
        from {  
          (forwarding-class [ class-names ] | forwarding-class-except [ class-names ] );  
          (interface-group [ group-names ] | interface-group-except [ group-names ] );  
          (learn-vlan-1p-priority [ priorities ] | learn-vlan-1p-priority [ priorities ] );  
          (loss-priority [ priorities ] | loss-priority-except [ priorities ] );  
          (user-vlan-1p-priority [ priorities ] | user-vlan-1p-priority [ priorities ] );  
        }  
        then {  
          ... statements in Common Firewall Actions on page 320 PLUS ...  
          (accept | discard);  
          port-mirror-instance instance-name;  
        }  
      }  
    }  
  }  
}
```

```
firewall {  
  family ethernet-switching {  
    filter filter-name {  
      interface-specific;  
      term term-name {  
        from {  
          destination-address {  
            ip-prefix</prefix-length>;  
          }  
          destination-mac-address {  
            mac-address;  
          }  
          destination-port [ port-names ];  
          destination-prefix-list {  
            list-name;  
          }  
          dot1q-tag [ tag-values ];  
          dot1q-user-priority [ priority-values ];  
          dscp [ code-point-values ];  
          ether-type [ protocol-names ];  
          fragment-flags flag;  
          icmp-code [ codes ];  
          icmp-type [ types ];  
          interface interface-name;  
          is-fragment;  
        }  
      }  
    }  
  }  
}
```

```

precedence [ precedence-names ];
protocol [ protocol-names ];
source-address {
    ip-prefix </prefix-length>;
}
source-mac-address {
    mac-address;
}
source-port [ port-names ];
source-prefix-list {
    list-name;
}
tcp-established;
tcp-flags flag;
tcp-initial;
vlan [ vlan-names ];
}
then {
    (accept | discard);
    analyzer analyzer-name;
    count counter-name;
    forwarding-class class-name;
    interface interface-name;
    log;
    loss-priority (high | low);
    policer policer-name;
    syslog;
    vlan vlan-name;
}
}
}
}
}

firewall {
    family inet {
        dialer-filter filter-name {
            accounting-profile [ profile-names ];
            term term-name {
                from {
                    ... statements in Common IP Firewall Match Conditions on page 322 AND
                    statements in Common IPv4 Firewall Match Conditions on page 323 EXCEPT
                    FOR ...
                    (ah-spi [ values ] | ah-spi-except [ values ]); # NOT valid at this level
                    (destination-class [ class-names ] |
                     destination-class-except [ class-names ]); # NOT valid at this level
                    interface interface-name; # NOT valid at this level
                    (loss-priority [ priorities ] | loss-priority-except [ priorities ]); # NOT valid at
                     this level
                    service-filter-hit; # NOT valid at this level
                    (source-class [ class-names ] | source-class-except [ class-names ]); # NOT
                     valid at this level
                }
            }
            then {
                (ignore | note);
                log;
            }
        }
    }
}

```

```

        sample;
        syslog;
    }
}
filter filter-name {
    accounting-profile [ profile-names ];
    interface-specific;
    term term-name {
        filter filter-name;
        from {
            ... statements in Common IP Firewall Match Conditions on page 322 AND
               statements in Common IPv4 Firewall Match Conditions on page 323 ...
        }
        then {
            ... statements in Common Firewall Actions on page 320 AND
               statements in Common IP Firewall Actions on page 321 AND
               statements in Common IPv4 Firewall Actions on page 321 ...
        }
    }
}
prefix-action name {
    count;
    destination-prefix-length prefix-length;
    filter-specific;
    policer policer-name;
    source-prefix-length prefix-length;
    subnet-prefix-length prefix-length;
}
service-filter filter-name {
    term term-name {
        from {
            address {
                ip-prefix</prefix-length>;
            }
            (ah-spi [ values ] | ah-spi-except [ values ]);
            destination-address {
                ip-prefix</prefix-length>;
            }
            (destination-port [ port-names ] | destination-port-except [ port-names ]);
            destination-prefix-list {
                list-name;
            }
            (esp-spi [ values ] | esp-spi-except [ values ]);
            first-fragment;
            fragment-flags flag;
            (fragment-offset [ offsets ] | fragment-offset-except [ offsets ]);
            (interface-group [ group-names ] | interface-group-except [ group-names ]);
            (ip-options [ option-names ] | ip-options-except [ option-names ]);
            is-fragment;
            (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
            (port [ port-names ] | port-except [ port-names ]);
            prefix-list {
                list-name;
            }
            (protocol [ protocol-names ] | protocol-except [ protocol-names ]);
        }
    }
}

```

```

        source-address {
            ip-prefix</prefix-length>;
        }
        (source-port [ port-names ] | source-port-except [ port-names ]);
        source-prefix-list {
            list-name;
        }
        tcp-flags flag-name;
    }
    then {
        count counter-name;
        log;
        port-mirror;
        sample;
        (service | skip);
    }
}
}
simple-filter filter-name {
    term term-name {
        from {
            destination-address ip-prefix</prefix-length>;
            destination-port port-name;
            forwarding-class [ class-names ];
            protocol protocol-name;
            source-address ip-prefix</prefix-length>;
            source-port port-name;
        }
        then {
            forwarding-class class-name;
            loss-priority (high | low | medium-high | medium-low);
            policer policer-name;
        }
    }
}
}
}
}
firewall {
    family inet6 {
        dialer-filter filter-name {
            accounting-profile [ profile-names ];
            term term-name {
                from {
                    ... statements in Common IP Firewall Match Conditions on page 322 PLUS ...
                    (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
                    ... BUT NOT ...
                    (destination-class [ class-names ] |
                     destination-class-except [ class-names ]); # NOT valid at this level
                    (forwarding-class [ class-names ] |
                     forwarding-class-except [ class-names ]); # NOT valid at this level
                    interface interface-name; # NOT valid at this level
                    (interface-group [ group-names ] | interface-group-except [ group-names ]); #
                     NOT valid at this level
                    (loss-priority [ priorities ] | loss-priority-except [ priorities ]); # NOT valid at
                     this level
                }
            }
        }
    }
}

```

```

        service-filter-hit; # NOT valid at this level
        (source-class [ class-names ] | source-class-except [ class-names ]); # NOT
            valid at this level
        tcp-established; # NOT valid at this level
        tcp-flags flag; # NOT valid at this level
        tcp-initial; # NOT valid at this level
    }
    then {
        (ignore | note);
        log;
        sample;
        syslog;
    }
}
}
filter filter-name {
    accounting-profile [ profile-names ];
    interface-specific;
    term term-name {
        filter filter-name;
        from {
            ... statements in Common IP Firewall Match Conditions on page 322 PLUS ...
            (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
            (traffic-class [ code-point-values ] | traffic-class-except [ code-point-values ]);
        }
        then {
            ... statements in Common Firewall Actions on page 320 AND
            statements in Common IP Firewall Actions on page 321 PLUS ...
            (accept | discard | reject <address-unreachable | administratively-prohibited |
                beyond-scope | fragmentation-needed | no-route | port-unreachable |
                tcp-reset>);
        }
    }
}
service-filter filter-name {
    term term-name {
        from {
            address {
                ip-prefix</prefix-length>;
            }
            (ah-spi [ values ] | ah-spi-except [ values ]);
            destination-address {
                ip-prefix</prefix-length>;
            }
            (destination-port [ port-names ] | destination-port-except [ port-names ]);
            destination-prefix-list {
                list-name;
            }
            (esp-spi [ values ] | esp-spi-except [ values ]);
            (interface-group [ group-names ] | interface-group-except [ group-names ]);
            (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
            (port [ port-names ] | port-except [ port-names ]);
            prefix-list {
                list-name;
            }
            source-address {

```



```

        ip-prefix </prefix-length>;
    }
    (source-port [ port-names ] | source-port-except [ port-names ]);
    source-prefix-list {
        list-name;
    }
    tcp-flags flag-name;
}
then {
    count counter-name;
    log;
    port-mirror;
    sample;
    (service | skip);
}
}
}
}
}

firewall {
    family mpls {
        filter filter-name {
            accounting-profile [ profile-names ];
            interface-specific;
            physical-interface-filter;
            term term-name {
                from {
                    (exp [ exp-bits ] | exp-except [ exp-bits ]);
                }
                then {
                    (ignore | note);
                    log;
                    sample;
                    syslog;
                }
            }
        }
    }
    filter filter-name {
        accounting-profile [ profile-names ];
        interface-specific;
        physical-interface-filter;
        term term-name {
            filter filter-name;
            from {
                (exp [ exp-bits ] | exp-except [ exp-bits ]);
                (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
                interface interface-name;
                interface-set set-name;
                (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
            }
            then {
                ... statements in Common Firewall Actions on page 320 PLUS ...
                (accept | discard);
                sample;
            }
        }
    }
}

```

```

    }
  }
}

firewall {
  family vpls {
    filter filter-name {
      accounting-profile [ profile-names ];
      interface-specific;
      term term-name {
        filter filter-name;
        from {
          ... statements in Common Layer 2 Firewall Match Conditions on page 323 ...
        }
        then {
          ... statements in Common Firewall Actions on page 320 PLUS ...
          (accept | discard);
          port-mirror;
          port-mirror-instance instance-name;
        }
      }
    }
  }
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

[edit interfaces] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```

interfaces {
  interface-name {
    ... the "interface-name" subhierarchy appears after the main [edit interfaces] hierarchy level ...
  }
  interface-set interface-set-name {
    interface interface-name {
      (unit unit-number | vlan-tags-outer vlan-tag);
    }
  }
  irb (Interfaces) {
    accounting-profile name;
    description text;
    disable;

    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;

    traceoptions {

```

```

    flag flag;
}
(traps | no-traps);
unit logical-unit-number {
    accounting-profile name;
    bandwidth rate;
    description text;
    disable;
    encapsulation type;
    family inet {
        accounting {
            destination-class-usage;
            source-class-usage {
                input;
                output;
            }
        }
    }
    address ipv4-address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        preferred;
        primary;
        vrrp-group group-id {
            (accept-data | no-accept-data);
            advertise-interval seconds;
            advertisements-threshold number;
            authentication-key key;
            authentication-type authentication;
            fast-interval milliseconds;
            (preempt | no-preempt) {
                hold-time seconds;
            }
            priority number;
            track {
                interface interface-name {
                    bandwidth-threshold bits-per-second priority-cost priority;
                    priority-cost priority;
                }
                priority-hold-time seconds;
                route prefix/prefix-length routing-instance instance-name priority-cost priority;
            }
            virtual-address [ addresses ];
            vrrp-inherit-from vrrp-group;
        }
    }
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
    fail-filter filter-name;
    mode {

```

```
        loose;
    }
}
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bandwidth priority-cost number;
                priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-inet6-address [addresses];
        virtual-link-local-address ipv6-address;
        vrrp-inherit-from {
            active-group group-number;
            active-interface interface-name;
        }
    }
}
}
(dad-disable | no-dad-disable);
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
```

```

no-redirects;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    filter {
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    policer {
        input policer-name;
        output policer-name;
    }
}
native-inner-vlan-id vlan-id;
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id-list [vlan-id's];
vlan-id-range [vlan-id-range];
}
}
traceoptions {
    file <filename> <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
    flag flag <disable>;
    no-remote-trace;
}
}

interfaces {
    interface-name {
        disable;
        accounting-profile name;
        aggregated-ether-options {
            ethernet-switch-profile {
                tag-protocol-id [ hexadecimal-identifiers ];
            }
        }
        (flow-control | no-flow-control);
        lacp {
            (active | passive);
            admin-key key;
            fast-failover;
            link-protection {

```

```

        disable;
        (revertive | non-revertive);
    }
    periodic (fast | slow);
    system-id mac-address;
    system-priority priority;
}
(link-protection | no-link-protection);
link-speed (100m | 1g | 8g | 10g | 40g | 50g | 80g | 100g | oc192);
logical-interface-fpc-redundancy;
(loopback | no-loopback);
mc-ae {
    chassis-id chassis-id;
    events {
        iccp-peer-down {
            force-icl-down;
            prefer-status-control-active;
        }
    }
    mc-ae-id mc-ae-id;
    mode (active-active | active-standby);
    redundancy-group group-id;
    status-control (active | standby);
}
minimum-links number;
rebalance-periodic {
    start-time time;
    interval number;
}
source-address-filter {
    mac-address;
}
(source-filtering | no-source-filtering);
}
auto-configure {
    remove-when-no-subscribers;
    stacked-vlan-ranges {
        access-profile profile-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                mac-address;
                option-82 ( circuit-id | remote-id);
                radius-realm radius-realm-string;
                user-prefix user-prefix-string;
            }
        }
    }
    dynamic-profile profile-name {
        accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
        ranges (any | low-tag-high-tag), (any | low-tag-high-tag);
    }
}
}

```

```

vlan-ranges {
  access-profile profile-name;
  authentication {
    password password-string;
    username-include {
      circuit-type;
      delimiter delimiter-character;
      domain-name domain-name-string;
      interface-name;
      mac-address;
      option-82;
      radius-realm radius-realm-string;
      user-prefix user-prefix-string;
    }
  }
  dynamic-profile profile-name {
    accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
    ranges (any | low-tag)–(any | high-tag);
  }
}
override tag vlan-tag dynamic-profile profile name;
}
encapsulation (ethernet-bridge | ethernet-vpls | extended-vlan-bridge |
  extended-vlan-vpls | flexible-ethernet-services | vlan-vpls);
ether-options {
  802.3ad {
    aex;
    (backup | primary);
    lacp {
      force-up;
      port-priority
    }
  }
}
asynchronous-notification;
(auto-negotiation | no-auto-negotiation);
ethernet-switch-profile {
  ethernet-policer-profile {
    input-priority-map {
      ieee802.1p premium [ values ];
    }
    output-priority-map {
      classifier {
        premium {
          forwarding-class class-name {
            loss-priority (high | low);
          }
        }
      }
    }
  }
}
policer cos-policer-name {
  aggregate {
    bandwidth-limit bps;
    burst-size-limit bytes;
  }
  premium {
    bandwidth-limit bps;

```

```

        burst-size-limit bytes;
    }
}
tag-protocol-id;
}
(mac-learn-enable | no-mac-learn-enable);
}
(flow-control | no-flow-control);
ignore-l3-incompletes;
link-mode (automatic | full-duplex | half-duplex);
(loopback | no-loopback);
keepalives <interval seconds> <down-count number> <up-count number>;
speed (1g | 10m | 100m | 10m-100m | auto-negotiation);
source-address-filter {
    mac-address;
}
source-filtering | no-source-filtering;
}
flexible-vlan-tagging;
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time (up milliseconds | down milliseconds);
interface-transmit-statistics;
(keepalives <down-count number> <interval seconds> <up-count number> |
no-keepalives);
layer2-policer {
    apply-groups [ group-names ];
    apply-groups-except [ group-names ];
}
link-mode (automatic | full-duplex);
mac mac-address;
mtu bytes;
multi-chassis-protection peer-ip-address {
    interface interface-name;
}
native-vlan-id number;
no-gratuitous-arp-request;
optics-options {
    alarm low-light-alarm {
        (link-down | syslog);
    }
    warning low-light-warning {
        (link-down | syslog);
    }
}
wavelength nm;
}
passive-monitor-mode;
per-unit-scheduler;
speed (10m | 100m | 1g | auto | oc3 | oc12 | oc48);
stacked-vlan-tagging;
traceoptions {
    flag flag;
}
transmit-bucket {
    overflow discard;
    rate percentage;
    threshold bytes;
}

```



```

}
(traps | no-traps);
unidirectional;
vlan-tagging;
}

interface-name {
  unit logical-unit-number {
    disable;
    accept-source-mac {
      mac-address mac-address {
        policer {
          input policer-name;
          output policer-name;
        }
      }
    }
  }
  account-layer2-overhead (Interface Level) {
    value;
    egress bytes;
    ingress bytes;
  }
  accounting-profile name;
  advisory-options {
    downstream-rate rate;
    upstream-rate rate;
  }
  arp-resp (restricted|unrestricted);
  bandwidth rate;
  clear-dont-fragment-bit;
  copy-tos-to-outer-ip-header;
  demux-destination family;
  encapsulation (vlan-bridge | vlan-vpls);
  epd-threshold cells plp1 cells;
  filter filter-name;
  inner-vlan-id-range start start-id end end-id;
  input-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
  }
  interface-shared-with psdnumerical-index;
  layer2-policer {
    input-hierarchical-policer policer-name;
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
  }
  multi-chassis-protection peer-ip-address {
    interface interface-name;
  }
  native-inner-vlan-id number;

```

```

output-vlan-map {
  (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
  inner-tag-protocol-id tpid;
  inner-vlan-id number;
  tag-protocol-id tpid;
  vlan-id number;
}
peer-interface interface-name;
peer-unit unit-number;
plp-to-clp;
proxy-arp <restricted | unrestricted>;
rpm {
  (client | server);
  twamp-server;
}
swap-by-poppush;
vlan-id number;
vlan-id-list [ vlan-id vlan-id-vlan-id ];
vlan-id-range number-number;
vlan-tags (inner <tpid.>vlan-id | inner-list [ vlan-id vlan-id-vlan-id ] |
  inner-range <tpid.>vlan-id-vlan-id) outer <tpid.>vlan-id;
}

unit logical-unit-number {
  family ethernet-switching {
    filter {
      group filter-group-number;
      (input filter-name | input-list [ filter-names ]);
      (output filter-name | output-list [ filter-names ]);
      (inner-vlan-id-list [ vlan-ids ] | vlan-id number | vlan-id-list [ number
        number-number ]);
      interface-mode (access | trunk);
    }
    policer {
      input policer-name;
      output policer-name;
    }
    vlan-rewrite {
      translate old-vlan-id new-vlan-id;
    }
    vlan {
      members [ all vlan-identifiers ];
    }
  }
}
family inet {
  filter {
    group filter-group-number;
    (input filter-name | input-list [ filter-names ]);
    (output filter-name | output-list [ filter-names ]);
  }
  input-hierarchical-policer policer-name;
  mac-validate (loose | strict);
  mtu bytes;
  no-neighbor-learn;
  no-redirects;
  policer {
    arp policer-template-name;
  }
}

```

```

    input policer-name;
    output policer-name;
}
primary;
receive-options-packets;
receive-ttl-exceeded;
rpf-check {
    fail-filter filter-name;
    mode loose;
}
sampling {
    (input | output | input output);
}
simple-filter {
    input filter-name;
}
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
}
unnumbered-address interface-name <destination address>
    <destination-profile profile-name> <preferred-source-address address>;
}

family inet6 {
    address ipv6-address {
        destination destination-address;
        eui-64;
        ndp ipv6-address <l2-interface interface-name> <(mac mac-address |
            multicast-mac multicast-mac-address) <publish>>;
        preferred;
        primary;
        vrrp-inet6-group group-number {
            (accept-data | no-accept-data);
            fast-interval milliseconds;
            inet6-advertise-interval seconds;
            (no-preempt; | ... the following preempt statement ...)
            preempt {
                hold-time seconds;
            }
            priority number;
            track {
                interface interface-name {
                    bandwidth-threshold bits-per-second priority-cost priority;
                    priority-cost priority;
                }
                priority-hold-time seconds;
                route ip-address-prefix/prefix-length routing-instance instance-name
                    priority-cost priority;
            }
        }
        virtual-inet6-address [ addresses ];
        virtual-link-local-address ipv6-address;
        vrrp-inherit-from {
            active-group group-number;
            active-interface interface-name;

```

```
    }
  }
}
(dad-disable | no-dad-disable);
filter {
  group filter-group-number;
  (input filter-name | input-list [ filter-names ]);
  (output filter-name | output-list [ filter-names ]);
}
input-hierarchical-policer policer-name;
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
policer {
  input policer-name;
  output policer-name;
}
rpf-check {
  fail-filter filter-name;
  mode loose;
}
sampling {
  (input | output | input output);
}
unnumbered-address interface-name preferred-source-address address;
}

family iso {
  address iso-address;
  mtu bytes;
}

family mlfr-end-to-end {
  bundle logical-interface-name;
}

family mpls {
  filter {
    group filter-group-number;
    (input filter-name | input-list [ filter-names ]);
    (output filter-name | output-list [ filter-names ]);
  }
  input-hierarchical-policer policer-name;
  maximum-labels maximum-labels;
  mtu bytes;
  policer {
    input policer-name;
    output policer-name;
  }
}
```

```

family vpls {
  core-facing;
  filter {
    group filter-group-number;
    (input filter-name | input-list [ filter-names ]);
    (output filter-name | output-list [ filter-names ]);
  }
  policer {
    input policer-name;
    output policer-name;
  }
}
}
}

```

**Related
Documentation**

- [Notational Conventions Used in Junos OS Configuration Hierarchies](#)

Class of Service on Ethernet Interfaces

- [Overview on page 1311](#)
- [Configuration on page 1312](#)

Overview

- [Class of Service for Ethernet on page 1311](#)

Class of Service for Ethernet

- [CoS for L2TP Tunnels on Ethernet Interface Overview on page 1311](#)

CoS for L2TP Tunnels on Ethernet Interface Overview

For effective packet tunneling, CoS is implemented over L2TP tunnels. For Ethernet interfaces, CoS is supported for L2TP session traffic to a LAC on platforms configured as an LNS that include egress IQ2 or IQ2E PICs.

This feature is supported on the following platforms:

- EX Series switches
- M7i and M10i routers
- M120 routers

To enable session-aware CoS on an L2TP interface, include the **per-session-scheduler** statement at the **[edit interfaces unit logical-unit-number]** hierarchy level.

After CoS is configured on an L2TP tunnel, Junos OS dynamically creates a traffic shaper for the traffic-shaping-profile and the L2TP tunnel based on the tunnel identification number. This ensures that the packets are monitored at the LAC and classified to allow the traffic flow to be adjusted on congested networks.

This feature has the following limitations:

- Only 991 shapers are supported on each IQ2 or IQ2E PIC.
- For a 4-port IQ2E PIC, you can configure up to 1976 shapers for an 8-queue session and 3952 shapers for a 4-queue session.
- For an 8-port IQ2E PIC, you can configure up to 1912 shapers for an 8-queue session and up to 3824 shapers for a 4-queue session.
- Sessions in excess of the maximum supported values specified for the PICs cannot be shaped (but they can be policed).
- The overall traffic rate cannot exceed the L2TP traffic rate, or else random drops result.
- There is no support for logical interface scheduling and shaping at the ingress because all schedulers are now reserved for L2TP.
- There is no support for physical interface rate shaping at the ingress.
- You cannot delete or deactivate the primary Ethernet interface on which the tunnel is established.

You can provide policing support for sessions with more than the maximum supported value on each IQ2 or IQ2E PIC. Each session can have four or eight different classes of traffic (queues). Each class needs its own policer; for example, one for voice and one for data traffic.

Related Documentation

- [Configuring CoS for L2TP Tunnels on Ethernet Interfaces](#)
- [Configuring LNS CoS for Link Redundancy on page 1313](#)
- [Example: L2TP LNS CoS Support for Link Redundancy on page 1314](#)

Configuration

- [Configuration Task on page 1312](#)

Configuration Task

- [Configuring CoS for L2TP Tunnels on Ethernet Interfaces on page 1312](#)
- [Configuring LNS CoS for Link Redundancy on page 1313](#)
- [Example: L2TP LNS CoS Support for Link Redundancy on page 1314](#)

Configuring CoS for L2TP Tunnels on Ethernet Interfaces

The Layer 2 Tunneling Protocol (L2TP) is often used to carry traffic securely between an L2TP Network Server (LNS) to an L2TP Access Concentrator (LAC). CoS is supported for L2TP session traffic to a LAC on platforms configured as an LNS that include egress IQ2 and IQ2E Ethernet PICs.

This feature is supported on the following platforms:

- EX Series switches
- M7i and M10i routers

- M120 routers

To configure CoS for L2TP on Ethernet interfaces:

1. Configure L2TP services on the Ethernet interface.
2. On the Ethernet interface, enable session-aware CoS for L2TP sessions.

```
[[edit interfaces interface-name unit logical-unit-number]
user@host# set per-session-scheduler
```
3. Configure the traffic manager in the IQ2 or IQ2E PIC to enable per-session CoS support.

```
[edit chassis fpc slot-number pic pic-number]
user@host# set traffic-manager mode-session-shaping
```
4. (Optional) To fine tune the system, you may also set the traffic-manager mode to session-shaping and configure the value of ingress-shaping-overhead parameter from 50 through 130 depending on your network requirement.

```
[edit chassis fpc slot-number pic pic-number]
user@host# set traffic-manager ingress-shaping-overhead value mode-session-shaping
```



NOTE: If you deactivate or delete the primary Ethernet interface on which the L2TP tunnel is configured, the tunnel with sessions having CoS is torn down.

After CoS is enabled for L2TP tunnels on Ethernet interface, you can run the **show class-of-service l2tp-session** command to verify the mapping of CoS with the configured L2TP session.

Related Documentation

- [L2TP Minimum Configuration](#)
- [CoS for L2TP Tunnels on Ethernet Interface Overview](#)
- [Example: Configuring CoS for L2TP Tunnels on Ethernet Interfaces](#)
- [Configuring LNS CoS for Link Redundancy on page 1313](#)
- [Example: L2TP LNS CoS Support for Link Redundancy on page 1314](#)
- `show class-of-service l2tp-session`

Configuring LNS CoS for Link Redundancy

You can configure multiple ports on the same IQ2 and IQ2E PICs to support link redundancy for CoS on L2TP tunnels configured on an Ethernet interface. Link redundancy is useful when the active port is unavailable due to events such as:

- Disconnection of the cable
- Rebooting of the remote end system
- Traffic re-routing through a different port due to network conditions

When link redundancy is enabled in such scenarios, the L2TP tunnels and its session are maintained by switching traffic to another port configured on the same IQ2 or IQ2E PIC.

To configure multiple ports (IQ and IQ2PE PIC) on an Ethernet interface for redundancy with CoS, configure per-session-scheduler for all Ethernet ports:

```
user@host#edit interfaces ge-2/0/0 unit 0 per-session-scheduler
user@host#edit interfaces ge-2/0/1 unit 0 per-session-scheduler
```

You can similarly configure all the ports on the IQ2 or IQ2E PIC to support link redundancy for CoS on L2TP tunnels.



NOTE:

- If one or more redundancy ports is removed from the configuration, the tunnels established through those redundancy ports also go down.
 - You must configure per-session-scheduler for all the ports that are to be used for redundancy. If you do not do so, new tunnels or sessions with CoS do not get established.
-

**Related
Documentation**

- [per-session-scheduler](#)

Example: L2TP LNS CoS Support for Link Redundancy

This example shows how link redundancy is supported when CoS for L2TP is configured on Ethernet interfaces.



NOTE: In this example, support for link redundancy is demonstrated by manually disabling the interface. However, link redundancy is also supported when the interface goes down due to events such as disconnection of the cable or rebooting of the remote end system.

- [Requirements on page 1314](#)
- [Overview on page 1315](#)
- [Configuration on page 1315](#)
- [Verification on page 1316](#)

Requirements

Before you begin:

- Configure service and loopback interfaces.
- Configure CoS for L2TP.


This feature applies to M Series Multiservice Edge Router running Junos OS Release 12.1 or later and EX Series switches.

Overview

Junos OS now supports link redundancy for CoS configured on an L2TP LNS. In this example, we verify that an L2TP tunnel does not go down when the Ethernet interface, through which the tunnels and its sessions with CoS are established, goes down.

Figure 8 on page 1315 shows a sample scenario in which L2TP access concentrator (LAC) devices operate on one side of an L2TP tunnel. LAC devices are configured with the address range of 192.168.100.0 with a subnet mask of 24. The LAC devices are connected to two backbone routers, P1 and P2. These two routers, P1 and P2, are connected over two Gigabit Ethernet ports on a single Ethernet IQ2 PIC to an L2TP network server (LNS). The LNS device is a router running Junos OS that supports redundancy for terminating L2TP sessions configured with CoS parameters. The CoS settings are applied on the interfaces using a RADIUS server when the L2TP session is set up. One of the Gigabit Ethernet interfaces on the IQ2 PIC present on the LNS device, ge-0/3/1, is connected to P1, while the other interface, ge-0/3/3, is linked to P2. Such a method of connection enables the subscriber sessions that reach the LAC devices to be forwarded to one of the two ports of the IQ2 PIC on the LNS device.

Figure 8: Topology to Verify Link Redundancy Support for L2TP LNS CoS



Configuration

Step-by-Step Procedure

To configure Ethernet interfaces for redundancy:

1. Configure Gigabit Ethernet interfaces.

```
[edit interfaces]
```

```
user@host# set ge-0/3/1 unit 0 family inet address 192.168.1.1/30
```

```
user@host# set ge-0/3/3 unit 0 family inet address 192.168.1.5/30
```

```
user@host# set ge-0/3/1 unit 0 per-session-scheduler
```

```
user@host# set ge-0/3/3 unit 0 per-session-scheduler
```

2. Configure static routing options.

```
[edit routing-options]
```

```
user@host# set static route 192.168.100.0/24 next-hop [ 192.168.1.2 192.168.1.6 ]
```

Step-by-Step Procedure Verify that CoS is now implemented over L2TP on an Ethernet interface and the LAC is reachable.

1. Verify that LAC is reachable.

```
user@host> show route 192.168.100.1
inet.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.100.0/24    *[Static/5] 1d 02:09:09
                  to 192.168.1.2 via ge-0/3/1.0
                  > to 192.168.1.6 via ge-0/3/3.0
```

2. Bring up an L2TP session and verify that L2TP sessions come up.

```
user@host> show services l2tp session
Interface: sp-1/3/0, Tunnel group: GEN-TUN-GRP-BIO, Tunnel local ID: 44806
Local Remote Interface State      Bundle Username
ID   ID   unit
12491 33795      1 Established      - test1
```

3. Send a traffic stream towards the subscriber.
4. Verify that the shaping at the subscriber end is as per the shaping rate configured.

```
user@host# show class-of-service l2tp-session
L2TP Session Username: test1, Index: 12491
Physical interface: ge-0/3/3, Index: 131
Queues supported: 4, Queues in use: 4
Scheduler map: GEN-SCHED-MAP-EF-65%, Index: 5212
Shaping rate: 2162200 bps
Encapsulation Overhead: 6, Cell Overhead: Enabled
```

In the output of the **show class-of-service l2tp-session** command, ge-0/3/3, index 131 represents the port used to establish the L2TP tunnel to which the current L2TP session belongs. It does not represent the port that was active when the L2TP session came up.

Verification

Verify that, when CoS is configured on an L2TP tunnel, link redundancy works if one of the ports on which the L2TP tunnel is established goes down.

- [Bring Down ge-0/3/3 Interface Through Which the L2TP Tunnel Is Established on page 1316](#)
- [Verify LAC Reachability and the Status of L2TP Sessions on page 1317](#)

Bring Down ge-0/3/3 Interface Through Which the L2TP Tunnel Is Established

Purpose Bring down the interface through which the L2TP session and its tunnels are established.

Action [edit interfaces]
 user@host# set ge-0/3/3 disable
 user@host# commit

Verify LAC Reachability and the Status of L2TP Sessions

Purpose Verify that link redundancy works and the L2TP session does not go down when the active port on the IQ2 PIC is down. Verify that the traffic flow is unaffected after it is switched to another port configured on the same IQ2 or IQ2E PIC.

Action

```

user@host> show route 192.168.100.1
inet.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.100.0/24    *[Static/5] 1d 02:35:09
                  to 192.168.1.2 via ge-0/3/1.0

user@host> show services l2tp session
Interface: sp-1/3/0, Tunnel group: GEN-TUN-GRP-BIO, Tunnel local ID: 44806
Local Remote Interface State          Bundle Username
ID    ID    unit
12491 33795    1 Established          - test1

```

Related Documentation

- [Configuring LNS CoS for Link Redundancy on page 1313](#)

Class of Service for MPLS

- [Overview on page 1317](#)
- [Configuration on page 1318](#)

Overview

- [CoS for MPLS on page 1317](#)

CoS for MPLS

- [CoS for MPLS Overview on page 1317](#)

CoS for MPLS Overview

When IP traffic enters a label-switched path (LSP) tunnel, the ingress routing device marks all packets with a class-of-service (CoS) value, which is used to place the traffic into a transmission priority queue. On the routing device, each interface has up to eight transmit queues. The CoS value is encoded as part of the Multiprotocol Label Switching (MPLS) header and remains in the packets until the MPLS header is removed when the packets exit from the egress routing device. The routing devices within the LSP utilize the CoS value set at the ingress routing device. The CoS value is encoded by means of the CoS bits (also known as the EXP or experimental bits).

MPLS class of service works in conjunction with the routing device's general CoS functionality. If you do not configure any CoS features, the default general CoS settings are used. For MPLS class of service, you might want to prioritize how the transmit queues are serviced by configuring weighted round-robin, and to configure congestion avoidance using random early detection (RED).

The next-hop label-switching router (LSR) uses the default classification shown in [Table 92 on page 1318](#).

Table 92: LSR Default Classification

Code Point	Forwarding Class	Loss Priority
000	best-effort	low
001	best-effort	high
010	expedited-forwarding	low
011	expedited-forwarding	high
100	assured-forwarding	low
101	assured-forwarding	high
110	network-control	low
111	network-control	high

Configuration

- [Configuration Task on page 1318](#)

Configuration Task

- [Configuring CoS for MPLS Traffic on page 1318](#)

Configuring CoS for MPLS Traffic

To configure CoS for MPLS packets in an LSP, include the **class-of-service** statement with the appropriate CoS value:

```
class-of-service cos-value;
```

If you do not specify a CoS value, the IP precedence bits from the packet's IP header are used as the packet's CoS value.

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit protocols mpls interface *interface-name* label-map *label-value*]
- [edit protocols mpls label-switched-path *path-name*]
- [edit protocols mpls label-switched-path *path-name* primary *path-name*]
- [edit protocols mpls label-switched-path *path-name* secondary *path-name*]
- [edit protocols mpls static-path *prefix*]
- [edit protocols rsvp interface *interface-name* link-protection]

- [edit protocols rsvp interface *interface-name* link-protection bypass *destination*]
- [edit logical-systems *logical-system-name* protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *path-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *path-name* primary *path-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *path-name* secondary *path-name*]
- [edit logical-systems *logical-system-name* protocols mpls static-path *prefix*]
- [edit logical-systems *logical-system-name* protocols mpls interface *interface-name* label-map *label-value*]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *destination*]

The **class-of-service** statement at the [edit protocols mpls label-switched-path] hierarchy level assigns an initial EXP value for the MPLS shim header of packets in the LSP. This value is initialized at the ingress routing device only and overrides the rewrite configuration established for that forwarding class. However, the CoS processing (weighted round robin [WRR] and RED) of packets entering the ingress routing device is not changed by the **class-of-service** statement on an MPLS LSP. Classification is still based on the behavior aggregate (BA) classifier at the [edit class-of-service] hierarchy level or the multifield classifier at the [edit firewall] hierarchy level.

We recommend configuring all routing devices along the LSP to have the same input classifier for EXP, and, if a rewrite rule is configured, all routing devices should have the same rewrite configuration. Otherwise, traffic at the next LSR might be classified into a different forwarding class, resulting in a different EXP value being written to the EXP header.

For more information, see the *Junos OS MPLS Applications Configuration Guide*.

Classifying Packets by Behavior Aggregate

- [Overview on page 1319](#)
- [Configuration on page 1330](#)

Overview

- [Behavior Aggregate Classifier on page 1320](#)
- [BA Classifier Default Values on page 1324](#)

Behavior Aggregate Classifier

- [BA Classifier Overview on page 1320](#)
- [Overview of BA Classifier Types on page 1321](#)
- [Default Behavior Aggregate Classification Overview on page 1322](#)
- [Understanding DSCP Classification for VPLS on page 1323](#)
- [BA Classifiers and ToS Translation Tables on page 1323](#)
- [Classifiers and Rewrite Rules at the Global and Physical Interface Levels Overview on page 1324](#)

BA Classifier Overview

The behavior aggregate (BA) classifier maps a class-of-service (CoS) value to a forwarding class and loss priority. The forwarding class determines the output queue. The loss priority is used by schedulers in conjunction with the random early discard (RED) algorithm to control packet discard during periods of congestion.

The types of BA classifiers are based on which part of the incoming packet the classifier examines:

- Differentiated Services code point (DSCP) for IP DiffServ
- DSCP for IPv6 DiffServ
- IP precedence bits
- MPLS EXP bits
- IEEE 802.1p CoS bits
- IEEE 802.1ad drop eligible indicator (DEI) bit

Unlike multifield classifiers (which are discussed in [“Multifield Classifier Overview” on page 1387](#)), BA classifiers are based on fixed-length fields, which makes them computationally more efficient than multifield classifiers. For this reason, core devices are normally configured to perform BA classification, because of the higher traffic volumes they handle.

In most cases, you need to rewrite a given marker (IP precedence, DSCP, IEEE 802.1p, IEEE 802.1ad, or MPLS EXP settings) at the ingress node to accommodate BA classification by core and egress devices. For more information about rewrite markers, see [“Rewriting Packet Header Information Overview” on page 1690](#).

For Juniper Networks M Series Multiservice Edge Routers, four classes can forward traffic independently. For M320 Multiservice Edge Routers and Juniper Networks T Series Core Routers, eight classes can forward traffic independently. Therefore, you must configure additional classes to be aggregated into one of these classes. You use the BA classifier to configure class aggregation.

For Juniper Networks MX Series 3D Universal Edge Routers and Intelligent Queuing 2 (IQ2) PICs, the following restrictions apply:

- You cannot use BA classifiers for IPv4 DSCP bits for Layer 2 VPNs.

- You cannot use BA classifiers for IPv6 DSCP bits for VPLS.
- You cannot use BA classifiers for IPv6 DSCP bits for Layer 2 VPNs.

For the 10-port 10-Gigabit Oversubscribed Ethernet (OSE) PICs, the following restrictions on BA classifiers apply:

- Multiple classifiers can be configured to a single logical interface. However, there are some restrictions on which the classifiers can coexist.

For example, the DSCP and IP precedence classifiers cannot be configured on the same logical interface. The DSCP and IP precedence classifiers can coexist with the DSCP IPv6 classifier on the same logical interface. An IEEE 802.1 classifier can coexist with other classifiers and is applicable only if a packet does not match any of the configured classifiers. For information about the supported combinations, see [“Applying Classifiers to Logical Interfaces” on page 1331](#).

- If the classifiers are not defined explicitly, then the default classifiers are applied as follows:
 - All MPLS packets are classified using the MPLS (EXP) classifier. If there is no explicit MPLS (EXP) classifier, then the default MPLS (EXP) classifier is applied.
 - All IPv4 packets are classified using the IP precedence and DSCP classifiers. If there is no explicit IP precedence or DSCP classifier, then the default IP precedence classifier is applied.
 - All IPv6 packets are classified using a DSCP IPv6 classifier. If there is no explicit DSCP IPv6 classifier, then the default DSCP IPv6 classifier is applied.
 - If the IEEE 802.1p classifier is configured and a packet does not match any explicitly configured classifier, then the IEEE 802.1p classifier is applied.



NOTE: For a specified interface, you can configure both a multifield classifier and a BA classifier without conflicts. Because the classifiers are always applied in sequential order, the BA classifier followed by the multifield classifier, any BA classification result is overridden by a multifield classifier if they conflict. For more information about multifield classifiers, see [“Multifield Classifier Overview” on page 1387](#).

Overview of BA Classifier Types

The idea behind class of service (CoS) is that packets are not treated identically by the routers or switches on the network. In order to selectively apply service classes to specific packets, the packets of interest must be classified in some fashion.

The simplest way to classify a packet is to use behavior aggregate classification. The DSCP, DSCP IPv6, or IP precedence bits of the IP header convey the behavior aggregate class information. The information might also be found in the MPLS EXP bits, IEEE 802.1ad, or IEEE 802.1p CoS bits.

You can configure the following classifier types:

- DSCP, DSCP IPv6, or IP precedence—IP packet classification (Layer 3 headers)
- MPLS EXP—MPLS packet classification (Layer 2 headers)
- IEEE 802.1p—Packet classification (Layer 2 headers)
- IEEE 802.1ad—Packet classification for IEEE 802.1ad formats (including DEI bit)

If you apply an IEEE 802.1 classifier to a logical interface, this classifier takes precedence and is not compatible with any other classifier type. On Juniper Networks MX Series Ethernet Services Routers and EX Series switches using IEEE 802.1ad frame formats, you can apply classification on the basis of the IEEE 802.1p bits (three bits in either the inner virtual LAN (VLAN) tag or the outer VLAN tag) and the drop eligible indicator (DEI) bit. On routers with IQ2 PICs using IEEE 802.1ad frame format, you can apply classification based on the IEEE 802.1p bits and the DEI bit. Classifiers for IP (DSCP or IP precedence) and MPLS (EXP) can coexist on a logical interface if the hardware requirements are met. (See [“Applying Classifiers to Logical Interfaces”](#) on page 1331.)

The Enhanced Queuing DPC (EQ DPC) does not support BA classification for packets received from a Layer 3 routing interface or a virtual routing and forwarding (VRF) interface and routed to an integrated routing and bridging interface (IRB) to reach the remote end of a pseudowire connection. The EQ DPC also does not support BA classification for Layer 2 frames received from a Virtual Private LAN Service (VPLS) pseudowire connection from a remote site and routed to a Layer 3 routing interface through an IRB interface.

Default Behavior Aggregate Classification Overview

The software automatically assigns an implicit default IP precedence classifier to all logical interfaces.



NOTE: Only the IEEE 802.1p classifier is supported in Layer 2 interfaces. You must explicitly apply this classifier to the interface as shown in [“Default IEEE 802.1p Classifier”](#) on page 1327.

If you enable the MPLS protocol family on a logical interface, a default MPLS EXP classifier is automatically applied to that logical interface.

Other default classifiers (such as those for IEEE 802.1p bits and DSCP) require that you explicitly associate a default classification table with a logical interface. When you explicitly associate a default classifier with a logical interface, you are in effect overriding the implicit default classifier with an explicit default classifier.



NOTE: Although several code points map to the expedited-forwarding (ef) and assured-forwarding (af) classes, by default no resources are assigned to these forwarding classes. All af classes other than af1x are mapped to best-effort, because RFC 2597, *Assured Forwarding PHB Group*, prohibits a node from aggregating classes.

You can apply IEEE 802.1p classifiers to interfaces that are part of VPLS routing instances.

Related Documentation

- [Default IP Precedence Classifier \(ipprec-compatibility\) on page 1325](#)
- [Default MPLS EXP Classifier on page 1325](#)
- [Default DSCP and DSCP IPv6 Classifier on page 1326](#)
- [Default IEEE 802.1p Classifier on page 1327](#)
- [Default IEEE 802.1ad Classifier on page 1328](#)
- [Default IP Precedence Classifier \(ipprec-default\) on page 1329](#)

Understanding DSCP Classification for VPLS

You can perform Differentiated Services Code Point (DSCP) classification for IPv4 packets on Ethernet interfaces that are part of a virtual private LAN service (VPLS) routing instance on the ingress provider edge (PE) router. This is supported on the M320 router with Enhanced type III FPC and the M120 router. On the Intelligent Queuing 2 (IQ2) or Intelligent Queuing 2 Enhanced (IQ2E) PICs, the **vlan-vpls** encapsulation statement is required. DSCP for IPv6 and Internet precedence for IPv6 are not supported.

In order to perform DSCP classification for IPv4 packets on Ethernet interfaces that are part of a VPLS routing instance on the ingress PE router, you must make sure of the following:

- The correct encapsulation statement based on PIC type is configured for the interface.
- The DSCP classifier is defined (default is allowed) at the **[edit class-of-service classifiers]** hierarchy level.
- The defined DSCP classifier is applied to the interface.
- The interface is included in the VPLS routing instance on the ingress of the PE router.

Related Documentation

- [BA Classifier Overview on page 1320](#)

BA Classifiers and ToS Translation Tables

On some PICs, the behavior aggregate (BA) translation tables are included for every logical interface (unit) protocol family configured on the logical interface. The proper default translation table is active even if you do not include any explicit translation tables. You can display the current translation table values with the **show class-of-service classifiers** command.

On Juniper Networks M40e, M120, M320 Multiservice Edge Routers, and T Series Core Routers with Enhanced IQ (IQE) PICs, or on any router or switch with IQ2 or Enhanced IQ2 (IQ2E) PICs, you can replace the type-of-service (ToS) bit value on the incoming packet header on a logical interface with a user-defined value. The new ToS value is used for all class-of-service processing and is applied before any other class-of-service or firewall treatment of the packet. The PIC uses the **translation-table** statement to determine the new ToS bit values.

You can configure a physical interface (port) or logical interface (unit) with up to three translation tables. For example, you can configure a port or unit with BA classification

for IPv4 DSCP, IPv6 DSCP, and MPLS EXP. The number of frame relay data-link connection identifiers (DLCIs) (units) that you can configure on each PIC varies based on the number and type of BA classification tables configured on the interfaces.

For more information on configuring ToS translation tables, along with examples, see *Configuring ToS Translation Tables*.

Classifiers and Rewrite Rules at the Global and Physical Interface Levels Overview

On ACX Series Universal Access Routers and EX Series switches, CoS supports classification and rewrite at the global level and physical interface levels.

At a global level, you can define EXP classification.

At a physical interface level, you can define the following features:

- DSCP and inet-precedence classifiers
- DSCP and inet-precedence rewrites
- ieee-802.1 classifiers (inner and outer)
- ieee-802.1 rewrites (outer)

At a logical interface level, you can define the fixed classification and EXP rewrites.

To configure global EXP classifiers, include the **classifiers exp classifier-name** statement at the **[edit class-of-service] system-defaults** hierarchy level.

To configure classifiers or rewrite rules at the physical interface, include either the **classifiers** statement or the **rewrite-rules** statement at the **[edit class-of-service] interfaces interface-name]** hierarchy level.

To display classifiers configured under **system-defaults**, enter the **show class-of-service system-defaults** command.

To display classifiers and rewrite rules bound to physical interfaces, enter the **show class-of-service interfaces interface-name** command.

Related Documentation

- [Configuring Classifiers and Rewrite Rules at the Global and Physical Interface Levels on page 1339](#)

BA Classifier Default Values

- [Default IP Precedence Classifier \(ipprec-compatibility\) on page 1325](#)
- [Default MPLS EXP Classifier on page 1325](#)
- [Default DSCP and DSCP IPv6 Classifier on page 1326](#)
- [Default IEEE 802.1p Classifier on page 1327](#)
- [Default IEEE 802.1ad Classifier on page 1328](#)
- [Default IP Precedence Classifier \(ipprec-default\) on page 1329](#)

Default IP Precedence Classifier (*ipprec-compatibility*)

By default, all logical interfaces are automatically assigned an implicit IP precedence classifier called **ipprec-compatibility**. The **ipprec-compatibility** IP precedence classifier maps IP precedence bits to forwarding classes and loss priorities, as shown in [Table 93 on page 1325](#).

Table 93: Default IP Precedence Classifier

IP Precedence CoS Values	Forwarding Class	Loss Priority
000	best-effort	low
001	best-effort	high
010	best-effort	low
011	best-effort	high
100	best-effort	low
101	best-effort	high
110	network-control	low
111	network-control	high

Default MPLS EXP Classifier

For all PICs except PICs mounted on Juniper Networks M Series Multiservice Edge Router standard (nonenhanced) FPCs, if you enable the MPLS protocol family on a logical interface, the default MPLS EXP classifier is automatically applied to that logical interface. The default MPLS classifier maps EXP bits to forwarding classes and loss priorities, as shown in [Table 94 on page 1325](#).

Table 94: Default MPLS Classifier

Code Point	Forwarding Class	Loss Priority
000	best-effort	low
001	best-effort	high
010	expedited-forwarding	low
011	expedited-forwarding	high
100	assured-forwarding	low
101	assured-forwarding	high
110	network-control	low

Table 94: Default MPLS Classifier (*continued*)

Code Point	Forwarding Class	Loss Priority
111	network-control	high

Default DSCP and DSCP IPv6 Classifier

Table 95 on page 1326 shows the forwarding class and packet loss priority (PLP) that are assigned to each well-known DSCP when you apply the explicit default DSCP or DSCP IPv6 classifier. To do this, include the **default** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* classifiers (dscp | dscp-ipv6)]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number classifiers (dscp
| dscp-ipv6)]
default;
```



NOTE: If you deactivate or delete dscp-ipv6 statement from the configuration, the default IPv6 classifier is not activated on the on the M5, M10, M7i, M10i, M20, M40, M40e, and M160 routing platforms. As a workaround, explicitly specify the default option to the dscp-ipv6 statement.

Table 95: Default DSCP Classifier

DSCP and DSCP IPv6	Forwarding Class	PLP
ef	expedited-forwarding	low
af11	assured-forwarding	low
af12	assured-forwarding	high
af13	assured-forwarding	high
af21	best-effort	low
af22	best-effort	low
af23	best-effort	low
af31	best-effort	low
af32	best-effort	low
af33	best-effort	low
af41	best-effort	low
af42	best-effort	low

Table 95: Default DSCP Classifier (*continued*)

DSCP and DSCP IPv6	Forwarding Class	PLP
af43	best-effort	low
be	best-effort	low
cs1	best-effort	low
cs2	best-effort	low
cs3	best-effort	low
cs4	best-effort	low
cs5	best-effort	low
nc1/cs6	network-control	low
nc2/cs7	network-control	low
other	best-effort	low

Default IEEE 802.1p Classifier

Table 96 on page 1327 shows the forwarding class and PLP that are assigned to the IEEE 802.1p CoS bits when you apply the explicit default IEEE 802.1p classifier. To do this, include the **default** statement at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number* classifiers ieee-802.1] hierarchy level:



NOTE: Only the IEEE 802.1p classifier is supported in Layer 2 interfaces. You must explicitly apply this classifier as shown.

```
[edit class-of-service interfaces interface-name unit logical-unit-number classifiers
  ieee-802.1]
  default;
```

Table 96: Default IEEE 802.1p Classifier

Code Point	Forwarding Class	PLP
000	best-effort	low
001	best-effort	high
010	expedited-forwarding	low
011	expedited-forwarding	high

Table 96: Default IEEE 802.1p Classifier (*continued*)

Code Point	Forwarding Class	PLP
100	assured-forwarding	low
101	assured-forwarding	high
110	network-control	low
111	network-control	high

Default IEEE 802.1ad Classifier

Table 97 on page 1328 shows the code point, forwarding class alias, and PLP that are assigned to the IEEE 802.1ad bits when you apply the explicit default IEEE 802.1ad classifier. The table is very similar to the IEEE 802.1p default table, but the loss priority is determined by the DEI bit. To apply the default table, include the **default** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* classifiers ieee-802.1]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number classifiers
  ieee-802.1ad]
  default;
```

Table 97: Default IEEE 802.1ad Classifier

IEEE 802.1ad Code Point	Forwarding Class Alias	PLP
0000	be	low
0010	be1	low
0100	ef	low
0110	ef1	low
1000	af11	low
1010	af12	low
1100	nc1	low
1110	nc2	low
0001	be-dei	high
0011	be1-dei	high
0101	ef-dei	high
0111	ef1-dei	high

Table 97: Default IEEE 802.1ad Classifier (*continued*)

IEEE 802.1ad Code Point	Forwarding Class Alias	PLP
1001	af11-dei	high
1011	af12-dei	high
1101	nc1-dei	high
1111	nc2-dei	high

Default IP Precedence Classifier (*ipprec-default*)

There are two separate tables for default IP precedence classification. All logical interfaces are implicitly assigned the **ipprec-compatibility** classifier by default, as described in [Table 93 on page 1325](#).

The other default IP precedence classifier (called **ipprec-default**) overrides the **ipprec-compatibility** classifier when you explicitly associate it with a logical interface. To do this, include the **default** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* classifiers inet-precedence]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number classifiers
  inet-precedence]
  default;
```

[Table 98 on page 1329](#) shows the forwarding class and PLP that are assigned to the IP precedence CoS bits when you apply the default IP precedence classifier.

Table 98: Default IP Precedence (*ipprec-default*) Classifier

Code Point	Forwarding Class	PLP
000	best-effort	low
001	assured-forwarding	low
010	best-effort	low
011	best-effort	low
100	best-effort	low
101	expedited-forwarding	low
110	network-control	low
111	network-control	high

Configuration

- [Configuration Tasks for Classifiers on page 1330](#)
- [Configuration Tasks for BA Classifiers on page 1338](#)
- [Configuration Task for DSCP IPv6 Classifiers on page 1341](#)
- [Configuration Tasks for MPLS EXP Classifiers on page 1341](#)
- [Configuration Task for IEEE 802.1ad Classifiers on page 1346](#)
- [Configuration Statements on page 1347](#)

Configuration Tasks for Classifiers

- [Defining Classifiers on page 1330](#)
- [Applying Classifiers to Logical Interfaces on page 1331](#)
- [DSCP Classifier Configuration Examples on page 1335](#)
- [Setting Packet Loss Priority on page 1337](#)

Defining Classifiers

You can override the default IP precedence classifier by defining a classifier and applying it to a logical interface. To define new classifiers for all code point types, include the **classifiers** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
classifiers {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) classifier-name {
    import [classifier-name | default];
    forwarding-class class-name {
      loss-priority level code-points [ aliases ] [ bit-patterns ];
    }
  }
}
```

The map sets the forwarding class and PLP for a specific set of code-point aliases and bit patterns. The inputs of the map are code-point aliases and bit patterns. The outputs of the map are the forwarding class and the PLP. For more information about how CoS maps work, see “[CoS Inputs and Outputs Overview](#)” on page 1272.

The classifiers work as follows:

- **dscp**—Handles incoming IPv4 packets.
- **dscp-ipv6**—Handles incoming IPv6 packets. For more information, see “[Applying DSCP IPv6 Classifiers](#)” on page 1341.
- **exp**—Handles MPLS packets using Layer 2 headers.
- **ieee-802.1**—Handles Layer 2 CoS.
- **inet-precedence**—Handles incoming IPv4 packets. IP precedence mapping requires only the upper three bits of the DSCP field.

A classifier takes a specified bit pattern as either the literal pattern or as a defined alias and attempts to match it to the type of packet arriving on the interface. If the information in the packet's header matches the specified pattern, the packet is sent to the appropriate queue, defined by the forwarding class associated with the classifier.

The code-point aliases and bit patterns are the input for the map. The loss priority and forwarding class are outputs of the map. In other words, the map sets the PLP and forwarding class for a given set of code-point aliases and bit patterns.



NOTE: On M Series, MX Series, and T Series routers, and EX Series switches that do not have tricolor marking enabled, the loss priority can be configured only by setting the PLP within a multifield classifier. This setting can then be used by the appropriate drop profile map and rewrite rule. For more information, see [“Setting Packet Loss Priority” on page 1337](#).

Importing a Classifier

You can use any table, including the default, in the definition of a new classifier by including the **import** statement. The imported classifier is used as a template and is not modified. Whenever you commit a configuration that assigns a new **class-name** and **loss-priority** value to a code-point alias or set of bits, it replaces that entry in the imported classifier template. As a result, you must explicitly specify every CoS value in every designation that requires modification.

To do this, include the **import default** statement at the **[edit class-of-service classifiers type classifier-name]** hierarchy level:

```
[edit class-of-service classifiers type classifier-name]
import default;
```

For instance, to import the default DSCP classifier, include the **dscp default** statement at the **[edit class-of-service classifiers dscp classifier-name]** hierarchy level:

```
[edit class-of-service classifiers dscp classifier-name]
import default;
```

Applying Classifiers to Logical Interfaces

To apply the classification map to a logical interface:

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
user@host#set classifiers (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence)
(classifier-name | default);
```

You can use interface wildcards for **interface-name** and **logical-unit-number**.

For most PICs, if you apply an IEEE 802.1p classifier to a logical interface, you cannot apply non-IEEE classifiers to other logical interfaces on the same physical interface. This restriction does not apply to Gigabit Ethernet IQ2 PICs.

There are some restrictions on applying multiple BA classifiers to a single logical interface. [Table 99 on page 1332](#) shows the supported combinations. In this table, the OSE PICs refer to the 10-port 10-Gigabit OSE PICs.

Table 99: Logical Interface Classifier Combinations

Classifier Combinations	Gigabit Ethernet IQ2 PICs	OSE PICs	Other PICs on M320, MX Series, T Series routers and on EX Series Switches	Other M Series with Regular FPCs	Other M Series with Enhanced FPCs
dscp and inet-precedence	No	No	No	No	No
dscp-ipv6 and (dscp inet-precedence)	Yes	Yes	Yes	No	No
exp and ieee 802.1	Yes	Yes	No	No	No
ieee 802.1 and (dscp dscp-ipv6 exp inet-precedence)	Yes	Yes	No	No	Yes
exp and (dscp dscp-ipv6 inet-precedence)	Yes	Yes	Yes	No	Yes

For Gigabit Ethernet IQ2 and 10-port 10-Gigabit Oversubscribed Ethernet (OSE) interfaces, family-specific classifiers take precedence over IEEE 802.1p BA classifiers. For example, if you configure a logical interface to use both an MPLS EXP and an IEEE 802.1p classifier, the EXP classifier takes precedence. MPLS-labeled packets are evaluated by the EXP classifier, and all other packets are evaluated by the IEEE 802.1p classifier. The same is true about other classifiers when combined with IEEE 802.1p classifiers on the same logical interface.

In Junos OS Releases 9.6 and later, the DSCP and IPv6 DSCP classifiers are not compatible with older formats. You cannot directly replace the old classifier with the new one. You must first delete the old classifier and then apply the new one, although both steps can be done in one configuration session. Otherwise, the commit will fail.



NOTE: If an interface is mounted on an M Series router FPC, you can apply only the default **exp** classifier. If an interface is mounted on an enhanced FPC, you can create a new **exp** classifier and apply it to an interface.

On MX960, MX480, MX240, MX80, M120, and M320 routers and EX Series switches with Enhanced Type III FPCs only, you can configure user-defined DSCP-based BA classification for MPLS interfaces (this feature is not available for IQE PICs or on MX Series routers and EX Series switches when ingress queuing is used) or VPLS/L3VPN routing instances (LSI interfaces). The DSCP-based classification for MPLS packets for Layer 2 VPNs is not supported. To classify MPLS packets on the routing instance at the egress PE, include the **dscp** or **dscp-ipv6** statements at the **[edit class-of-service routing-instances routing-instance-name classifiers]** hierarchy level. To classify MPLS packets at the core-facing interface, apply the classifier at the **[edit class-of-service interface**

interface-name unit *unit-name* classifiers (dscp | dscp-ipv6) *classifier-name* family mpls] hierarchy level.



NOTE: If you do not apply a DSCP classifier, the default EXP classifier is applied to MPLS traffic.

You can apply DSCP classification for MPLS traffic in the following usage scenarios:

- In a Layer 3 VPN (L3VPN) using an LSI routing instance.
 - Supported on the M120, M320, MX960, MX480, MX240, and MX80 routers.
 - DSCP classifier configured under **[edit class-of-service routing-instances]** on the egress PE router.
- In VPLS using an LSI routing instance.
 - Supported on the M120, M320, MX960, MX480, MX240, and MX80 routers.
 - DSCP classifier configured under **[edit class-of-service routing-instances]** on the egress PE router.
- In a Layer 3 VPN (L3VPN) using a VT routing instance.
 - Supported on the M120, M320, MX960, MX480, MX240, and MX80 routers.
 - DSCP classifier configured under **[edit class-of-service interfaces]** on the core-facing interface on the egress PE router.
- In VPLS using the VT routing instance.
- MPLS forwarding.
 - Supported on the M120, M320, MX960, MX480, MX240, and MX80 routers (not supported on IQE and MX when ingress queuing is enabled).
 - DSCP classifier configured under **[edit class-of-service interfaces]** on the ingress core-facing interface on the P or egress PE router.

MPLS forwarding when the label stacking is greater than 2 is not supported:

The following example configures a DSCP classifier for IPv4 named **dscp-ipv4-classifier** for the **fc-af11-class** forwarding class and a corresponding IPv6 DSCP classifier:

```
class-of-service {
  routing-instances routing-instance-one {
    classifiers {
      dscp dscp-ipv4-classifier {
        loss-priority low code-points 000100;
      }
      dscp dscp-ipv6-classifier {
        forwarding-class fc-af11-class {
          loss-priority low {
            code-points af11;
          }
        }
      }
    }
  }
}
```

```

    }
  }
}

```



NOTE: This is not a complete configuration.

This example applies the IPv4 classifier to MPLS traffic and the IPv6 classifier to Internet traffic on interface **ge-2/0/3.0**:

```

class-of-service {
  interfaces ge-2/0/3 {
    unit 0 {
      classifiers {
        dscp dscp-ipv4-classifier {
          family mpls;
        }
        dscp-ipv6 dscp-ipv6-classifier {
          family inet; # This is the default if not present.
        }
      }
    }
  }
}

```



NOTE: This is not a complete configuration.

This example applies the same classifier to both MPLS and IP traffic on interface **ge-2/2/0**.

```

[edit class-of-services interface ge-2/2/0]
unit 0 {
  classifiers {
    dscp dscp-mpls {
      family [ mpls inet ];
    }
  }
}

```



NOTE: This is not a complete configuration.



NOTE: You can apply DSCP and DSCP IPv6 classifiers to explicit null MPLS packets. The **family mpls** statement works the same on both explicit null and non-null MPLS labels.

**Related
Documentation**

- [DSCP Classifier Configuration Examples on page 1335](#)

DSCP Classifier Configuration Examples

On MX960, MX480, MX240, MX80, M120, and M320 routers with Enhanced Type III FPCs and EX Series switches only, you can configure user-defined DSCP-based BA classification for MPLS interfaces (this feature is not available for IQE PICs or on MX Series routers or EX Series switches when ingress queuing is used) or VPLS/L3VPN routing instances (LSI interfaces). The following examples show how you can apply DSCP classifiers for MPLS traffic in these cases.

**Applying a DSCP
Classifier to MPLS
Packets on the
Core-facing Interface**

Configure the core-facing interface and associated logical interfaces:

```
interfaces ge-5/3/1 {
  unit 0 {
    family inet {
      address 1.1.1.1/24;
    }
    family iso;
    family inet6 {
      address 2000::1/64;
    }
    family mpls
  }
}
```

Configure the DSCP classifier.

```
class-of-service {
  classifiers {
    dscp dscp11 {
      forwarding-class expedited-forwarding {
        loss-priority low code-points [ ef cs5 ];
      }
      forwarding-class assured-forwarding {
        loss-priority low code-points [ af21 af31 af41 cs4 ];
        loss-priority high code-points [ af23 af33 af43 cs2 af22 af32 af42 cs3 ];
      }
      forwarding-class best-effort {
        loss-priority low code-points [ af11 cs1 af12 ];
        loss-priority high code-points af13;
      }
      forwarding-class network-control {
        loss-priority low code-points [ cs6 cs7 ];
      }
    }
  }
}
```

Attach the classifier to the logical interface for the mpls family. You cannot configure more than one classifier per family.

```
class-of-service {
  interfaces {
    ge-5/3/1 {
      unit 0 {
        classifiers {
          dscp dscp11 {
            family mpls;
          }
        }
      }
    }
  }
}
```

```

    }
  }
}

```

The above classifiers are applicable on egress PE routers for VPLS and L3VPN cases. For plain interfaces (not VPLS/L3VPN (LSI) interfaces), these classifiers are applicable on P and egress PE routers on core facing interfaces.

Applying a DSCP Classifier to MPLS Traffic for L3VPN/VPLS

Configure routing instances of type either vrf or vpls.

```

routing-instances {
  vpls1 {
    instance-type vpls;
    interface ge-2/2/2.0; #customer facing interface for VPLS
    route-distinguisher 10.255.245.51:1;
    vrf-target target:1234:1;
    protocols {
      vpls {
        site-range 10;
        no-tunnel-services;
        site vpls-1-site-1 {
          site-identifier 1;
        }
      }
    }
  }
}

```

Configure the DSCP classifier.

```

class-of-service {
  classifiers {
    dscp dscp11 {
      forwarding-class expedited-forwarding {
        loss-priority low code-points [ ef cs5 ];
      }
      forwarding-class assured-forwarding {
        loss-priority low code-points [ af21 af31 af41 cs4 ];
        loss-priority high code-points [ af23 af33 af43 cs2 af22 af32 af42 cs3 ];
      }
      forwarding-class best-effort {
        loss-priority low code-points [ af11 cs1 af12 ];
        loss-priority high code-points af13;
      }
      forwarding-class network-control {
        loss-priority low code-points [ cs6 cs7 ];
      }
    }
  }
}

```

Attach the classifier to a routing instance. You cannot configure more than one classifier per routing instance.

```

class-of-service {
  routing-instances {
    vpls1 {
      classifiers {
        dscp dscp11;
      }
    }
  }
}

```

Related Documentation

- [Applying Classifiers to Logical Interfaces on page 1331](#)

Setting Packet Loss Priority

By default, the least significant bit of the CoS value sets the packet loss priority (PLP) value. For example, CoS value 000 is associated with PLP **low**, and CoS value 001 is associated with PLP **high**. In general, you can change the PLP by configuring a behavior aggregate (BA) or multifield classifier, as discussed in “[Overview of BA Classifier Types](#)” on page 1321 and “[Multifield Classifier Overview](#)” on page 1387.

However, on Juniper Networks M320 Multiservice Edge Routers, MX Series 3D Universal Edge Routers, and T Series Core Routers and EX Series switches that do not have tricolor marking enabled, the loss priority can be configured by setting the PLP within a multifield classifier or by behavior aggregate (BA) classifier. This setting can then be used by the appropriate drop profile map and rewrite rule.

On M320 routers and T Series routers with Enhanced II Flexible PIC Concentrators (FPCs) and tricolor marking enabled, you can set the PLP with a BA or multifield classifier, as described in “[Using BA Classifiers to Set PLP](#)” on page 1460 and “[Using Multifield Classifiers to Set PLP](#)” on page 1461.

On T Series routers with different Packet Forwarding Engines (non-Enhanced Scaling and Enhanced Scaling FPCs), you can configure PLP bit copying for ingress and egress unicast and multicast traffic. To configure, include the **copy-plp-all** statement at the **[edit class-of-service]** hierarchy level.

Example: Overriding the Default PLP on M320 Routers

The following example shows a two-step procedure to override the default PLP settings on M320 routers:

1. The following example specifies that while the DSCP code points are 110, the loss priority is set to **high**; however, on M320 routers, overriding the default PLP this way has no effect.

```

class-of-service {
  classifiers {
    dscp ba-classifier {
      forwarding-class expedited-forwarding {
        loss-priority high code-points 110;
      }
    }
  }
}

```

2. For M320 routers, this multifield classifier sets the PLP.

```
firewall {
  filter ef-filter {
    term ef-multifield {
      from {
        precedence 6;
      }
      then {
        loss-priority high;
        forwarding-class expedited-forwarding;
      }
    }
  }
}
```

Configuration Tasks for BA Classifiers

- [Example: DSCP IPv6 Rewrites and Forwarding Class Maps on page 1338](#)
- [Tunneling and BA Classifiers on page 1339](#)
- [Configuring Classifiers and Rewrite Rules at the Global and Physical Interface Levels on page 1339](#)

Example: DSCP IPv6 Rewrites and Forwarding Class Maps

You cannot configure a DSCP IPv6 rewrite rule and output forwarding class map on the same logical interface (unit). These must be used on different logical interfaces. Although a warning is issued, there is nothing in the CLI that prevents this configuration. An error message appears when you attempt to commit the configuration.

This example shows the warning and error message that results when the default DSCP IPv6 rewrite rule is configured on logical interface **ge-1/0/4.0** with output forwarding class map **vg1**.

```
[edit class-of-service]
interfaces {
  ge-1/0/4 {
    unit 0 {
      ##
      ## Warning: DSCP-IPv6 rewrite and forwarding class map not allowed on same unit
      ##
      output-forwarding-class-map vg1;
      rewrite-rules {
        dscp-ipv6 default;
      }
    }
  }
}

user@router# commit
[edit class-of-service interfaces ge-1/0/4 unit 0 output-forwarding-class-map]
'output-forwarding-class-map vg1'
DSCP-IPv6 rewrite and forwarding class map not allowed on same unit
error: commit failed: (statements constraint check failed)
```


Related Documentation • [Applying Forwarding Classes to Interfaces on page 1531](#)

Tunneling and BA Classifiers

BA classifiers can be used with GRE and IP-IP tunnels on the following routers and switches:

- EX Series switches
- M7i and M10i routers
- M Series routers with E-FPC or EP-FPC
- M120 routers
- M320 routers
- T Series routers

When a GRE or IP-IP tunnel is configured on an incoming (core-facing) interface, the queue number and PLP information are carried through the tunnel. At the egress (customer-facing) interface, the packet is queued and the CoS bits rewritten based on the information carried through the tunnel.

If no BA classifier is configured in the incoming interface, the default classifier is applied. If no rewrite rule is configured, the default rewrite rule is applied.

Configuring Classifiers and Rewrite Rules at the Global and Physical Interface Levels

On ACX Series Universal Access Routers and EX Series switches, CoS supports classification and rewrite at the global and physical interface levels.

To configure the global EXP classifier, include the following statements at the **[edit class-of-service] system-defaults** hierarchy level.

```
[edit class-of-service]
{
  system-defaults
  {
    classifiers exp classifier-name
  }
}
```

CoS supports one global system default classifier of the EXP type, as shown in the following example:

```
[edit class-of-service]
{
  system-defaults {
    classifiers {
      exp exp-classf-core;
    }
  }
}
```

To configure classifiers and rewrite rules at the physical interface level, include the following statements at the **[edit class-of-service] interfaces** hierarchy level.

```
[edit class-of-service]
interfaces {
  interface-name
    classifiers dscp classifier-name
    classifiers inet-precedence classifier-name
    classifiers ieee-802.1 [vlan-tag (outer | inner)] classifier-name
    rewrite-rules dscp rewrite-name
    rewrite-rules inet-prec rewrite-name
    rewrite-rules ieee-802.1 rewrite-name
}
```

The following example shows classifiers and rewrite rules configured on physical interfaces:

```
ge-0/1/0 {
  unit 0 {
    rewrite-rules {
      exp custom-exp;
    }
  }
  classifiers {
    dscp d1;
    ieee-802.1 ci;
  }
  rewrite-rules {
    dscp default;
  }
}
ge-0/1/2 {
  classifiers {
    ieee-802.1 ci;
  }
  rewrite-rules {
    ieee-802.1 ri;
  }
}
ge-0/1/3 {
  unit 0 {
    rewrite-rules {
      exp custom-exp2;
    }
  }
}
ge-0/1/7 {
  classifiers {
    dscp d1;
  }
}
ge-0/1/8 {
  classifiers {
    dscp d1;
  }
}
```

**Related
Documentation**

- [Classifiers and Rewrite Rules at the Global and Physical Interface Levels Overview on page 1324](#)

Configuration Task for DSCP IPv6 Classifiers

- [Applying DSCP IPv6 Classifiers on page 1341](#)

Applying DSCP IPv6 Classifiers

For M320 and T Series routers and forEX Series switches, you can apply separate classifiers for IPv4 and IPv6 packets per logical interface by including the **classifiers** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level and specifying the **dscp** and **dscp-ipv6** classifier types:

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
  classifiers dscp (classifier-name | default) family (mpls | inet);
  classifiers dscp-ipv6 (classifier-name | default) family (mpls | inet));
```

For M Series router enhanced FPCs, you cannot apply separate classifiers for IPv4 and IPv6 packets on a single logical interface. Instead, classifier assignment works as follows:

- If you assign a DSCP classifier only, IPv4 and IPv6 packets are classified using the DSCP classifier.
- If you assign an IP precedence classifier only, IPv4 and IPv6 packets are classified using the IP precedence classifier. In this case, the lower three bits of the DSCP field are ignored because IP precedence mapping requires the upper three bits only.
- If you assign either the DSCP or the IP precedence classifier in conjunction with the DSCP IPv6 classifier, the commit fails.
- If you assign a DSCP IPv6 classifier only, IPv4 and IPv6 packets are classified using the DSCP IPv6 classifier, but the commit displays a warning message.

For more information, see “[Applying Classifiers to Logical Interfaces](#)” on page 1331. For a complex configuration example, see the *Junos OS Feature Guides*.

Configuration Tasks for MPLS EXP Classifiers

- [Applying MPLS EXP Classifiers to Routing Instances on page 1341](#)
- [Applying MPLS EXP Classifiers for Explicit-Null Labels on page 1345](#)

Applying MPLS EXP Classifiers to Routing Instances

When you enable VRF table labels and you do not explicitly apply a classifier configuration to the routing instance, the default MPLS EXP classifier is applied to the routing instance. For detailed information about VRF table labels, see the *Junos OS VPNs Configuration Guide*.

The default MPLS EXP classification table contents are shown in [Table 100 on page 1341](#).

Table 100: Default MPLS EXP Classification Table

Forwarding Class	Loss Priority	CoS Value
best-effort	low	000
best-effort	high	001

Table 100: Default MPLS EXP Classification Table (*continued*)

Forwarding Class	Loss Priority	CoS Value
expedited-forwarding	low	010
expedited-forwarding	high	011
assured-forwarding	low	100
assured-forwarding	high	101
network-control	low	110
network-control	high	111

For PICs that are installed on enhanced FPCs, you can override the default MPLS EXP classifier and apply a custom classifier to the routing instance. To do this, perform the following configuration tasks:

1. Filter traffic based on the IP header by including the **vrf-table-label** statement at the **[edit routing-instances *routing-instance-name*]** hierarchy level:

```
[edit routing-instances routing-instance-name]
vrf-table-label;
```

2. Configure a custom MPLS EXP classifier by including the following statements at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
classifiers {
  exp classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
      loss-priority level code-points [ aliases ] [ bit-patterns ];
    }
  }
}
forwarding-classes {
  queue queue-number class-name priority (high | low);
}
```

3. Configure the routing instance to use the custom MPLS EXP classifier by including the **exp** statement at the **[edit class-of-service routing-instances *routing-instance-name* classifiers]** hierarchy level:

```
[edit class-of-service routing-instances routing-instance-name classifiers]
exp classifier-name;
```

To display the MPLS EXP classifiers associated with all routing instances, issue the **show class-of-service routing-instances** command.



NOTE: The following caveats apply to custom MPLS EXP classifiers for routing instances:

- An enhanced FPC is required.
- Logical systems are not supported.

For more details, see the following sections:

- [Configuring Global Classifiers and Wildcard Routing Instances on page 1343](#)
- [Examples: Applying MPLS EXP Classifiers to Routing Instances on page 1344](#)

Configuring Global Classifiers and Wildcard Routing Instances

To configure a global routing instance classifier, include the **all** statement at the **[edit class-of-service routing-instances]** hierarchy level:

```
[edit class-of-service routing-instances]
all {
  classifiers {
    exp classifier-name;
  }
}
```

For routing instances associated with specific classifiers, the global configuration is ignored.

To use a wildcard in the routing instance classifier configuration, include an asterisk (*) in the name of the routing instance:

```
[edit class-of-service routing-instances]
routing-instance-name* {
  classifiers {
    exp classifier-name;
  }
}
```

The wildcard configuration follows the longest match. If there is a specific configuration, it is given precedence over the wildcard configuration.



NOTE: Wildcards and the all keyword are supported at the [edit class-of-service routing-instances] hierarchy level but not at the [edit routing-instances] hierarchy level.

If you configure a routing instance at the [edit routing-instances] hierarchy level with, for example, the name `vpn*`, the Junos OS treats `vpn*` as a valid and distinct routing instance name. If you then try to apply a classifier to the `vpn*` routing instance at the [edit class-of-service routing-instances] hierarchy level, the Junos OS treats the `vpn*` routing instance name as a wildcard, and all the routing instances that start with `vpn` and do not have a specific classifier applied receive the classifier associated with `vpn*`. This same behavior applies with the all keyword.

Examples: Applying MPLS EXP Classifiers to Routing Instances

Configure a global classifier for all routing instances and override the global classifier for a specific routing instance. In this example, there are three routing instances: `vpn1`, `vpn2`, and `vpn3`, each with VRF table label enabled. The classifier `exp-classifier-global` is applied to `vpn1` and `vpn2` (that is, all but `vpn3`, which is listed separately). The classifier `exp-classifier-3` is applied to `vpn3`.

Configuring a Global Classifier

```
[edit routing-instances]
vpn1 {
  vrf-table-label;
}
vpn2 {
  vrf-table-label;
}
vpn3 {
  vrf-table-label;
}

[edit class-of-service routing-instances]
all {
  classifiers {
    exp exp-classifier-global;
  }
}
vpn3 {
  classifiers {
    exp exp-classifier-3;
  }
}
```

Configure a wildcard routing instance and override the wildcard with a specific routing instance. In this example, there are three routing instances: `vpn-red`, `vpn-yellow`, and `vpn-green`, each with VRF table label enabled. The classifier `exp-class-wildcard` is applied to `vpn-yellow` and `vpn-green`. The classifier `exp-class-red` is applied to `vpn-red`.

Configuring a Wildcard Routing Instance

```
[edit routing-instances]
vpn-red {
  vrf-table-label;
```

```

}
vpn-yellow {
    vrf-table-label;
}
vpn-green {
    vrf-table-label;
}

[edit class-of-service routing-instances]
vpn* {
    classifiers {
        exp exp-class-wildcard;
    }
}
vpn-red {
    classifiers {
        exp exp-class-red;
    }
}

```

Display the MPLS EXP classifiers associated with two routing instances:

Monitoring a Configuration

```

[edit class-of-service routing-instances]
vpn1 {
    classifiers {
        exp default;
    }
}
vpn2 {
    classifiers {
        exp class2;
    }
}

```

```
user@host> show class-of-service routing-instances
```

```

Routing Instance : vpn1
  Object      Name      Type      Index
  Classifier   exp-default  exp       8

Routing Instance : vpn2
  Object      Name      Type      Index
  Classifier   class2     exp       57507

```

Applying MPLS EXP Classifiers for Explicit-Null Labels

When you configure MPLS explicit-null labels, label 0 is advertised to the egress router of an LSP. When label 0 is advertised, the egress router (instead of the penultimate router) removes the label. Ultimate-hop popping ensures that any packets traversing an MPLS network include a label. For more information about explicit-null labels and ultimate-hop popping, see the *Junos OS MPLS Applications Configuration Guide*.

On M320 and T Series routers, when you configure MPLS explicit-null labels with an MPLS EXP classifier, the MPLS EXP classifier can be different from an IPv4 or IPv6 classifier configured on the same logical interface. In other words, you can apply separate classifiers for MPLS EXP, IPv4, and IPv6 packets per logical interface. To combine an

EXP classifier with a distinct IPv6 classifier, the PIC must be mounted on an Enhanced FPC.



NOTE: For J Series routers and other M Series routers, MPLS explicit-null labels with MPLS EXP classification are supported if you set the same classifier for EXP and IPv4 traffic, or EXP and IPv6 traffic.

For more information about how IPv4 and IPv6 packet classification is handled, see [“Applying DSCP IPv6 Classifiers” on page 1341](#).

To configure an MPLS EXP classifiers for explicit-null labels, include the **exp** statement at the **[edit class-of-service classifiers]** and **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* classifiers]** hierarchy levels:

```
[edit class-of-service classifiers]
exp classifier-name {
  import (classifier-name | default);
  forwarding-class class-name {
    loss-priority level code-points [ aliases ] [ bit-patterns ];
  }
}
[edit class-of-service interfaces interface-name unit logical-unit-number classifiers]
exp (classifier-name | default);
```

Configuration Task for IEEE 802.1ad Classifiers

- [Configuring and Applying IEEE 802.1ad Classifiers on page 1346](#)

Configuring and Applying IEEE 802.1ad Classifiers

For Juniper Network MX Series 3D Universal Edge Router interfaces or IQ2 PICs with IEEE 802.1ad frame formats or EX Series switches, you can set the forwarding class and loss priority for traffic on the basis of the three IEEE 802.1p bits and the DEI bit. You can apply the default map or customize one or more of the default values.

You then apply the classifier to the interface on which you configure IEEE 802.1ad frame formats.

Defining Custom IEEE 802.1ad Maps

You can customize the default IEEE 802.1ad map by defining values for IEEE 802.1ad code points.

```
class-of-service {
  classifiers {
    ieee-802.1ad dot1p_dei_class {
      forwarding-class best-effort {
        loss-priority low code-points [ 0000 1101 ];
      }
    }
  }
}
```


Applying Custom IEEE 802.1ad Maps

You then apply the classifier map to the logical interface:

```

interfaces {
  ge-2/0/0 {
    unit 0 {
      classifiers {
        ieee-802.1ad dot1p_dei_class;
      }
    }
  }
}

```

Verifying Custom IEEE 802.1ad Map Configuration

To verify your configuration, you can issue the following operational mode commands:

- `show class-of-service forwarding-table loss-priority-map`
- `show class-of-service forwarding-table loss-priority-map mapping`
- `show chassis forwarding`
- `show pfe fwdd`

Configuration Statements

- [\[edit class-of-service\] Hierarchy Level on page 1347](#)

[edit class-of-service] Hierarchy Level

```

class-of-service {
  classifiers {
    type classifier-name {
      forwarding-class class-name {
        loss-priority (high | low | medium-high | medium-low) code-points [ aliases bits ];
      }
      import (classifier-name | default);
    }
  }
  code-point-aliases {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) {
      alias-name bits;
    }
  }
  drop-profiles {
    profile-name {
      fill-level percentage drop-probability percentage;
      interpolate {
        drop-probability value;
        fill-level value;
      }
    }
  }
  fabric {
    scheduler-map {

```

```

        priority (high | low) scheduler scheduler-name;
    }
}
forwarding-class-map {
    map-name {
        class class-name queue-num queue-number <restricted-queue queue-number>;
    }
}
forwarding-classes {
    class class-name policing-priority (normal | premium) queue-num queue-number
    priority (high | low);
    queue queue-number class-name policing-priority (normal | premium) priority (high |
    low);
}
forwarding-policy {
    class class-name {
        classification-override {
            forwarding-class class-name;
        }
    }
    next-hop-map map-name {
        forwarding-class class-name {
            discard;
            lsp-next-hop [ lsp-regular-expressions ];
            next-hop [ next-hop-names ];
            non-lsp-next-hop;
        }
    }
}
fragmentation-maps {
    map-name {
        forwarding-class class-name {
            drop-timeout milliseconds;
            fragment-threshold bytes;
            multilink-class number;
            no-fragmentation;
        }
    }
}
host-outbound-traffic {
    dscp-code-point value;
    forwarding-class class-name;
    ieee-802.1 {
        default value;
        rewrite-rules;
    }
    tcp {
        raise-internet-control-priority;
    }
}
interfaces {
    ... the interfaces subhierarchy appears after the main [edit class-of-service] hierarchy
    ...
}
restricted-queues {

```

```

    forwarding-class class-name queue-number;
  }
rewrite-rules {
  (dscp | dscp-ipv6 | exp | frame-relay-de | ieee-802.1 | ieee-802.1ad | inet-precedence)
  rewrite-rule {
    forwarding-class class-name {
      loss-priority level code-point (alias | bits);
    }
    import (rewrite-rule | default);
  }
}
routing-instances routing-instance-name {
  classifiers {
    dscp (classifier-name | default);
    dscp-ipv6 (classifier-name | default);
    exp (classifier-name | default);
    ieee-208.1 (classifier-name | default | encapsulated | vlan-tag (inner | outer));
  }
}
scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
schedulers {
  scheduler-name {
    adjust-minimum value;
    adjust-percent value;
    buffer-size (exact | percent percentage | remainder);
    drop-profile-map loss-priority (any | high | low | medium-high | medium-low)
      protocol any;
    excess-priority (high | low | medium-high | medium-low);
    excess-rate (percent percentage | proportion proportion);
    priority (high | low | medium-high | medium-low | strict-high);
    shaping-rate (bps | percent percentage | burst-size size);
    transmit-rate (bps | percent percentage | remainder) <exact | rate-limit>;
  }
}
traceoptions {
  file <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
traffic-control-profiles {
  profile-name {
    adjust-minimum rate;
    delay-buffer-rate (bps | cps cps | percent percentage);
    excess-rate (percent percentage | proportion value);
    guaranteed-rate (bps | percent percentage) <burst-size bytes>;
    overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
    scheduler-map map-name;
    shaping-rate (bps | percent percentage) <burst-size bytes>;
  }
}
tri-color;

```

```

}

class-of-service {
  interfaces {
    interface-name {
      excess-bandwidth-share (equal | proportional value);
      input-excess-bandwidth-share (equal | proportional value);
      input-scheduler-map map-name;
      input-shaping-rate bps;
      input-traffic-control-profile profile-name;
      output-forwarding-class-map map-name;
      output-traffic-control-profile profile-name;
      scheduler-map map-name;
      scheduler-map-chassis (map-name | derived);
      shaping-rate bps;
      unit (logical-unit-number | *) {
        classifiers {
          dscp (classifier-name | default) {
            family [ inet mpls ];
          }
          dscp-ipv6 (classifier-name | default) {
            family [ inet mpls ];
          }
          exp (classifier-name | default);
          ieee-208.1 (classifier-name | default) <vlan-tag (inner | outer)>;
          ieee-208.1ad (classifier-name | default);
          inet-precedence (classifier-name | default);
        }
        forwarding-class class-name;
        input-scheduler-map map-name;
        input-shaping-rate bps;
        input-traffic-control-profile profile-name shared-instance instance-name;
        loss-priority-maps {
          (map-name | default);
        }
        loss-priority-rewrites {
          (map-name | default);
        }
      }
      output-forwarding-class-map map-name;
      output-traffic-control-profile profile-name shared-instance instance-name;
      rewrite-rules {
        dscp (rule-name | default) <protocol mpls>;
        dscp-ipv6 (rule-name | default);
        exp (rule-name | default) <protocol [ mpls-any | mpls-inet-both |
          mpls-inet-both-non-vpn ]>;
        exp-push-push-push default;
        exp-swap-push-push default;
        ieee-802.1 (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
        ieee-802.1ad (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
        inet-precedence (rewrite-name | default) <protocol mpls>;
      }
      scheduler-map map-name;
      shaping-rate bps;
      translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 |
        to-exp-from-exp | to-inet-precedence-from-inet-precedence) table-name;
    }
  }
}

```


```

    }
    interface-set interface-set-name {
        excess-bandwidth-share (equal | proportional value);
        input-excess-bandwidth-share (equal | proportional value);
        input-traffic-control-profile profile-name;
        input-traffic-control-profile-remaining profile-name;
        internal-node;
        output-traffic-control-profile profile-name;
        output-traffic-control-profile-remaining profile-name;
    }
}

```

Related Documentation • [Notational Conventions Used in Junos OS Configuration Hierarchies](#)


classifiers (Application)

Syntax	<pre> classifiers { <i>type</i> (<i>classifier-name</i> default) family (mpls inet); } </pre>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Apply a CoS aggregate behavior classifier to a logical interface. You can apply a default classifier or one that is previously defined.
Options	<p><i>classifier-name</i>—Name of the aggregate behavior classifier.</p> <p><i>type</i>—Traffic type.</p> <p>Values: dscp, dscp-ipv6, exp, ieee-802.1, inet-precedence</p>
	<div>  <p>NOTE: You can only specify a family for the dscp and dscp-ipv6 types.</p> </div>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	• Applying Classifiers to Logical Interfaces on page 1331

classifiers (Application for Routing Instances)

Syntax	<pre>classifiers { exp (classifier-name default); dscp (classifier-name default); dscp-ipv6 (classifier-name default); }</pre>
Hierarchy Level	[edit class-of-service routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. dscp and dscp-ipv6 support introduced in Junos OS Release 9.6.
Description	For routing instances with VRF table labels enabled, apply a custom Multiprotocol Label Switching (MPLS) EXP classifier or DSCP classifier to the routing instance. You can apply the default classifier or one that is previously defined.
Options	classifier-name —Name of the behavior aggregate MPLS EXP or DSCP classifier.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Applying MPLS EXP Classifiers to Routing Instances on page 1341• Applying Classifiers to Logical Interfaces on page 1331

classifiers (Definition)

Syntax	<pre> classifiers { type classifier-name { import (classifier-name default); forwarding-class class-name { loss-priority level code-points [aliases] [bit-patterns]; } } } </pre>
Hierarchy Level	[edit class-of-service], [edit class-of-service routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. ieee-802.1ad option introduced in Junos OS Release 9.2.
Description	Define a CoS behavior aggregate (BA) classifier for classifying packets. You can associate the classifier with a forwarding class or code-point mapping, and import a default classifier or one that is previously defined.
<div style="display: flex; align-items: center;">  <div> <p>NOTE: The [edit class-of-service routing-instances <i>routing-instance-name</i>] hierarchy level and the dscp-ipv6 and ieee-802.1ad classifier types are not supported on ACX Series routers.</p> </div> </div>	
Options	<p>classifier-name—Name of the aggregate behavior classifier.</p> <p>type—Traffic type: dscp, dscp-ipv6, exp, ieee-802.1, ieee-802.1ad, inet-precedence.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Overview of BA Classifier Types on page 1321 • Example: Configuring CoS for a PBB Network • Configuring CoS on ACX Series Universal Access Routers

classifiers (Physical Interface)

Syntax	<code>classifiers { type (classifier-name default) ; }</code>
Hierarchy Level	<code>[edit class-of-service interfaces interface-name]</code>
Release Information	Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.
Description	Apply a CoS aggregate behavior classifier to a physical interface. You can apply a default classifier or one that is previously defined.
Options	classifier-name —Name of the aggregate behavior classifier. type —Traffic type. Values: <code>dscp</code> , <code>ieee-802.1</code> , and <code>inet-precedence</code>
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.

code-points

Syntax	<code>code-points ([aliases] [bit-patterns]);</code>
Hierarchy Level	<code>[edit class-of-service classifiers type classifier-name forwarding-class class-name loss-priority level]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 8.5 for J Series devices. Statement introduced in Junos OS Release 9.2 for SRX Series devices.
Description	Specify one or more DSCP code-point aliases or bit sets for association with a forwarding class.
Options	aliases —Name of the DSCP alias. bit-patterns —Value of the code-point bits, in six-bit binary form.
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Overview of BA Classifier Types on page 1321• Example: Configuring CoS for a PBB Network• Example: Configuring Behavior Aggregate Classifiers• Example: Configuring Forwarding Classes

copy-plp-all

Syntax	<code>copy-plp-all;</code>
Hierarchy Level	<code>[edit class-of-service]</code>
Release Information	Statement introduced in Junos OS Release 10.3 on T series platforms.
Description	Enable PLP bit copying for ingress and egress for unicast and multicast traffic when traffic is ingressing one FPC and egressing the other.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Classifying Packets by Behavior Aggregate</i> • Packet Loss Priority Configuration Overview on page 1676 • Setting Packet Loss Priority on page 1337

dscp (AS PIC Classifiers)

Syntax	<code>dscp (<i>alias</i> <i>bits</i>);</code>
Hierarchy Level	<code>[edit services cos application-profile <i>profile-name</i> (ftp sip) (data video voice)],</code> <code>[edit services cos rule <i>rule-name</i> term <i>term-name</i> then],</code> <code>[edit services cos rule <i>rule-name</i> term <i>term-name</i> then (reflexive reverse)]</code>
Release Information	Statement introduced in Junos OS Release 8.1.
Description	Define the Differentiated Services code point (DSCP) mapping that is applied to the packets.
Options	<i>alias</i> —Name assigned to a set of CoS markers. <i>bits</i> —Mapping value in the packet header.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Actions in a CoS Rule</i>

dscp (Classifier on Physical Interface)

Syntax	<code>dscp (<i>classifier-name</i> default);</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> classifiers]
Release Information	Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.
Description	For ACX Series Universal Access routers and EX Series switches, map the DSCP field of the incoming packet to the forwarding class and packet loss priority based on the specified DSCP classifier.
Options	<i>classifier-name</i> —Name of the previously defined DSCP behavior aggregate classifier. default —Default DSCP behavior aggregate classifier.
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.

dscp-ipv6 (Class-of-Service)

Syntax	<code>dscp-ipv6 (<i>rewrite-name</i> <default>) { protocol mpls }</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For IPv6 traffic, apply a DSCP rewrite rule.
Options	<p><i>rewrite-name</i>—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules dscp-ipv6] hierarchy level.</p> <p>default— Default mapping.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Rewrite Rules on page 1694 • dscp (Rewrite Rules) on page 1382 • exp on page 1358 • exp-push-push-push on page 1723 • exp-swap-push-push on page 1724 • ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360 • ieee-802.1ad on page 1362 • inet-precedence on page 1364 • rewrite-rules (Definition) on page 1481

exp

Syntax	<code>exp (rewrite-name default) protocol protocol-types;</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS Release 12.2. for ACX series
Description	Apply an MPLS experimental (EXP) rewrite rule.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules exp] hierarchy level.</p> <p>default—The default mapping.</p> <p>By default, IP precedence rewrite rules alter the first three bits on the type-of-service (ToS) byte while leaving the last three bits unchanged. This default behavior applies to rewrite rules you configure for MPLS packets with IPv4 payloads. You configure these types of rewrite rules by including the mpls-inet-both or mpls-inet-both-non-vpn option at the [edit class-of-service interfaces <i>interface</i> <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules exp <i>rewrite-rule-name</i> protocol] hierarchy level. The IP precedence rewrite rules explanation does not apply to ACX Series Universal Access routers.</p> <p>On interfaces configured on Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) on MX Series Ethernet Services Routers and EX Series switches, we highly recommend that you configure the default option when you configure a behavior aggregate (BA) classifier that does not include a specific rewrite rule for MPLS packets. Doing so ensures that MPLS exp value is rewritten according to the BA classifier rules configured for forwarding or packet loss priority. This does not apply to ACX Series Universal Access routers.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Rewriting the EXP Bits of All Three Labels of an Outgoing Packet</i>• dscp (Rewrite Rules) on page 1382• dscp-ipv6 (Class-of-Service) on page 1357• exp-push-push-push on page 1723• exp-swap-push-push on page 1724• ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360• ieee-802.1ad on page 1362• inet-precedence on page 1364

- [rewrite-rules \(Definition\)](#) on page 1481

forwarding-class (BA Classifiers)

Syntax	<code>forwarding-class <i>class-name</i> { loss-priority <i>level</i> code-points [<i>aliases</i>] [<i>bit-patterns</i>]; }</code>
Hierarchy Level	[edit class-of-service classifiers <i>type classifier-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define forwarding class name and option values.
Options	<p><i>class-name</i>—Name of the forwarding class.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Defining Classifiers on page 1330 • <i>Example: Configuring CoS for a PBB Network</i>

ieee-802.1 (Rewrite Rules on Logical Interface)

Syntax	ieee-802.1 (<i>rewrite-name</i> default) vlan-tag (outer outer-and-inner);
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4. vlan-tag statement introduced in Junos OS Release 8.1.
Description	Apply an IEEE-802.1 rewrite rule. For IQ PICs, you can only configure one IEEE 802.1 rewrite rule on a physical port. All logical ports (units) on that physical port should apply the same IEEE 802.1 rewrite rule.
Options	rewrite-name —Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules ieee-802.1] hierarchy level. default —The default mapping.
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694• <i>Example: Configuring CoS for a PBB Network</i>• dscp (Rewrite Rules) on page 1382• dscp-ipv6 (Class-of-Service) on page 1357• exp on page 1358• exp-push-push-push on page 1723• exp-swap-push-push on page 1724• ieee-802.1ad on page 1362• inet-precedence on page 1364• rewrite-rules (Definition) on page 1481

ieee-802.1 (Classifier on Physical Interface)

Syntax	ieee-802.1 (<i>classifier-name</i> default) vlan-tag (inner outer);
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> classifiers]
Release Information	Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.
Description	For ACX Series Universal Access routers and EX Series switches, map the ieee-802.1p field of the incoming packet to the forwarding class and packet loss priority based on the specified 802.1p classifier. In the case of double tagged packets, you can configure whether to use the 802.1p of the outer or inner VLAN tag.
Options	<p>vlan-tag inner—In the case of double tagged packets, classify based on the 802.1p of the inner VLAN tag.</p> <p>vlan-tag outer—Classify based on the 802.1p of the outermost VLAN tag.</p> <p>classifier-name—Name of the previously defined ieee-802.1p behavior aggregate classifier.</p> <p>default—Default ieee-802.1p behavior aggregate classifier.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

ieee-802.1ad

Syntax	ieee-802.1ad (<i>rewrite-name</i> default) vlan-tag (outer outer-and-inner);
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Apply a IEEE-802.1ad rewrite rule.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules ieee-802.1ad] hierarchy level.</p> <p>default—The default rewrite bit mapping.</p> <p>vlan-tag—The rewrite rule is applied to the outer or outer-and-inner VLAN tag.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694• <i>Example: Configuring CoS for a PBB Network</i>• dscp (Rewrite Rules) on page 1382• dscp-ipv6 (Class-of-Service) on page 1357• exp on page 1358• exp-push-push-push on page 1723• exp-swap-push-push on page 1724• ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360• inet-precedence on page 1364• rewrite-rules (Definition) on page 1481

import (Classifiers)

Syntax	<code>import (classifier-name default);</code>
Hierarchy Level	<code>[edit class-of-service classifiers type classifier-name]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify a default or previously defined classifier.
Options	<p>classifier-name—Name of the classifier mapping configured at the <code>[edit class-of-service classifiers]</code> hierarchy level.</p> <p>default—The default classifier mapping.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Overview of BA Classifier Types on page 1321

inet-precedence

Syntax	<code>inet-precedence (rewrite-name default);</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Apply a IPv4 precedence rewrite rule.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules inet-precedence] hierarchy level.</p> <p>default—The default mapping. By default, IP precedence rewrite rules alter the first three bits on the type of service (ToS) byte while leaving the last three bits unchanged.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694• dscp (Rewrite Rules) on page 1382• dscp-ipv6 (Class-of-Service) on page 1357• exp on page 1358• exp-push-push-push on page 1723• exp-swap-push-push on page 1724• ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360• ieee-802.1ad on page 1362• rewrite-rules (Definition) on page 1481

inet-precedence (Classifier on Physical Interface)

Syntax	<code>inet-precedence (<i>classifier-name</i> default);</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> classifiers]
Release Information	Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.
Description	On ACX Series Universal Access routers and EX Series switches, map the inet-precedence field of the incoming packet to the forwarding class and packet loss priority, based on the specified inet-precedence classifier. When no classifier is configured on the physical interface, the default ipprec-compatibility inet-precedence classifier is applied on the physical interface.
Options	<p><i>classifier-name</i>—Name of the previously defined inet-precedence behavior aggregate classifier.</p> <p>default—Default inet-precedence behavior aggregate classifier.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

interfaces

```
Syntax  interfaces {
    interface-name {
        classifiers{
            dscp(classifier-name | default) {
            }
            ieee-802.1 (classifier-name | default) vlan-tag (inner | outer | classifier-name);
            inet-precedence (rewrite-name | default);
        }
        input-scheduler-map map-name;
        input-shaping-rate rate;
        irb {
            unit logical-unit-number {
                classifiers {
                    type (classifier-name | default);
                }
                rewrite-rules {
                    dscp (rewrite-name | default);
                    dscp-ipv6 (rewrite-name | default);
                    exp (rewrite-name | default) protocol protocol-types;
                    ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                    inet-precedence (rewrite-name | default);
                }
            }
        }
        member-link-scheduler (replicate | scale);
        rewrite-rules {
            dscp (rewrite-name | default);
            ieee-802.1 (rewrite-name | default) vlan-tag (outer);
            inet-precedence (rewrite-name | default);
        }
        scheduler-map map-name;
        scheduler-map-chassis map-name;
        shaping-rate rate;
        unit logical-unit-number {
            classifiers {
                type (classifier-name | default) family (mpls | inet);
            }
            forwarding-class class-name;
            fragmentation-map map-name;
            input-shaping-rate (percent percentage | rate);
            input-traffic-control-profile profile-name shared-instance instance-name;
            output-traffic-control-profile profile-name shared-instance instance-name;
            per-session-scheduler;
            rewrite-rules {
                dscp (rewrite-name | default);
                dscp-ipv6 (rewrite-name | default);
                exp (rewrite-name | default) protocol protocol-types;
                exp-push-push-push default;
                exp-swap-push-push default;
                ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                inet-precedence (rewrite-name | default);
            }
        }
    }
}
```

```

    }
    scheduler-map map-name;
    shaping-rate rate;
    translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 | to-exp-from-exp
    | to-inet-precedence-from-inet-precedence) table-name;
  }
}
interface-set interface-set-name {
  excess-bandwidth-share;
  internal-node;
  output-traffic-control-profile profile-name;
  output-traffic-control-profile-remaining profile-name;
}
}

```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 7.4.
Interface-set level added in Junos OS Release 8.5.

Description Configure interface-specific CoS properties for incoming packets.



NOTE: The `dscp-ipv6` and `ieee-802.1ad` classifier types are not supported on ACX Series routers. For further information about support on ACX Series routers, see *Understanding CoS CLI Configuration Statements on ACX Series Universal Access Routers*.

Options The remaining statements are explained separately.

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

Related Documentation

- [Overview of BA Classifier Types on page 1321](#)
- [Configuring Rewrite Rules on page 1694](#)
- *Understanding CoS CLI Configuration Statements on ACX Series Universal Access Routers*

loss-priority (BA Classifiers)

Syntax	loss-priority <i>level</i> ;
Hierarchy Level	[edit class-of-service classifiers <i>type classifier-name</i> forwarding-class <i>class-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify packet loss priority value for a specific set of code-point aliases and bit patterns.
Options	<p><i>level</i> can be one of the following:</p> <ul style="list-style-type: none">• high—Packet has high loss priority.• medium-high—Packet has medium-high loss priority.• medium-low—Packet has medium-low loss priority.• low—Packet has low loss priority.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Overview of BA Classifier Types on page 1321• Example: Configuring CoS for a PBB Network• Configuring Tricolor Marking on page 1447

routing-instances (Class of Service)

Syntax	<pre> routing-instances <i>routing-instance-name</i> { classifiers { dscp (<i>classifier-name</i> default); dscp-ipv6 (<i>classifier-name</i> default); exp (<i>classifier-name</i> default); } }</pre>
Hierarchy Level	[edit class-of-service]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For routing instances with VRF table labels enabled, apply a custom MPLS EXP classifier or DSCP classifier to the routing instance. You can apply the default MPLS EXP classifier or one that is previously defined.
Default	If you do not include this statement, the default MPLS EXP classifier is applied to the routing instance. When no DSCP classifier is configured, the default MPLS EXP classifier is applied.
Options	<p><i>routing-instance-name</i>—Name of a routing instance.</p> <p><i>classifier-name</i>—Name of the behavior aggregate MPLS EXP classifier or DSCP classifier.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Forwarding Classes on page 1530 • Applying Custom MPLS EXP Classifiers to Routing Instances in Layer 3 VPNs

system-defaults

Syntax	<pre>system-defaults { classifiers{ type classifier-name; } }</pre>
Hierarchy Level	[edit class-of-service]
Release Information	Statement introduced in Junos OS release 12.2
Description	Define a CoS classifier to support global classifiers.
Options	classifier-name —Name of the behavior aggregate (BA) classifier. type —Traffic type: dscp, dscp-ipv6, or exp.



NOTE: The **dscp** and **dscp-ipv6** classifier types are not supported on ACX Series routers.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Configuring CoS on ACX Series Universal Access Routers

unit

Syntax	<pre> unit logical-unit-number { classifiers { type (classifier-name default) family (mpls all); } forwarding-class class-name; fragmentation-map map-name; input-traffic-control-profile profile-name shared-instance instance-name; output-traffic-control-profile profile-name shared-instance instance-name; per-session-scheduler; rewrite-rules { dscp (rewrite-name default); dscp-ipv6 (rewrite-name default); exp (rewrite-name default) protocol protocol-types; exp-push-push-push default; exp-swap-push-push default; ieee-802.1 (rewrite-name default) vlan-tag (outer outer-and-inner); inet-precedence (rewrite-name default); } scheduler-map map-name; shaping-rate rate; } </pre>
Hierarchy Level	[edit class-of-service interfaces interface-name]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.
Options	<p>logical-unit-number—Number of the logical unit.</p> <p>Range: 0 through 16,384</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Overview of BA Classifier Types on page 1321 • Configuring Rewrite Rules on page 1694

Defining Code-Point Aliases

- [Overview on page 1371](#)
- [Configuration on page 1375](#)

Overview

- [Code-Point Alias on page 1372](#)

Code-Point Alias

- [Default Code-Point Alias Overview on page 1372](#)
- [Default CoS Values on page 1372](#)

Default Code-Point Alias Overview

Behavior aggregate (BA) classifiers use class-of-service (CoS) values such as Differentiated Services code points (DSCPs), DSCP IPv6, IP precedence, IEEE 802.1 and MPLS experimental (EXP) bits to associate incoming packets with a particular CoS servicing level. On a Services Router or EX Series switch, you can assign a meaningful name or alias to the CoS values and use this alias instead of bits when configuring CoS components. These aliases are not part of the specifications but are well known through usage. For example, the alias for DSCP 101110 is widely accepted as **ef** (expedited forwarding).



NOTE: The code point aliases must begin with a letter and can be up to 64 characters long.

When you configure classes and define classifiers, you can refer to the markers by alias names. You can configure user-defined classifiers in terms of alias names. If the value of an alias changes, it alters the behavior of any classifier that references it.

To configure class-of-service (CoS) code point aliases, include the **code-point-aliases** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
code-point-aliases {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) {
    alias-name bits;
  }
}
```

Related
Documentation

- [code-point-aliases on page 1381](#)

Default CoS Values

[Table 101 on page 1372](#) shows the default mappings between the bit values and standard aliases. For example, it is widely accepted that the alias for DSCP 101110 is **ef** (expedited forwarding).

Table 101: Default CoS Values

CoS Value Types	Mapping
DSCP and DSCP IPv6 CoS Values	
ef	101110
af11	001010

Table 101: Default CoS Values (*continued*)

CoS Value Types	Mapping
af12	001100
af13	001110
af21	010010
af22	010100
af23	010110
af31	011010
af32	011100
af33	011110
af41	100010
af42	100100
af43	100110
be	000000
cs1	001000
cs2	010000
cs3	011000
cs4	100000
cs5	101000
nc1/cs6	110000
nc2/cs7	111000
MPLS EXP CoS Values	
be	000
be1	001
ef	010
ef1	011

Table 101: Default CoS Values (*continued*)

CoS Value Types	Mapping
af11	100
af12	101
nc1/cs6	110
nc2/cs7	111
IEEE 802.1 CoS Values	
be	000
be1	001
ef	010
ef1	011
af11	100
af12	101
nc1/cs6	110
nc2/cs7	111
Legacy IP Precedence CoS Values	
be	000
be1	001
ef	010
ef1	011
af11	100
af12	101
nc1/cs6	110
nc2/cs7	111

Configuration

- [Configuration Task on page 1375](#)
- [Configuration Statements on page 1376](#)

Configuration Task

- [Defining Code Point Aliases for Bit Patterns on page 1375](#)

Defining Code Point Aliases for Bit Patterns

To define a code-point alias, include the **code-point-aliases** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
code-point-aliases {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) {
    alias-name bit-pattern;
  }
}
```

The CoS marker types are as follows:

- **dscp**—Handles incoming IPv4 packets.
- **dscp-ipv6**—Handles incoming IPv6 packets.
- **exp**—Handles MPLS packets using Layer 2 headers.
- **ieee-802.1**—Handles Layer 2 CoS.
- **inet-precedence**—Handles incoming IPv4 packets. IP precedence mapping requires only the upper three bits of the DSCP field.

For example, you might configure the following aliases:

```
[edit class-of-service]
code-point-aliases {
  dscp {
    my1 110001;
    my2 101110;
    be 000001;
    cs7 110000;
  }
}
```

This configuration produces the following mapping:

```
user@host> show class-of-service code-point-aliases dscp
Code point type: dscp

Alias           Bit pattern
ef/my2          101110
af11             001010
af12             001100
af13             001110
```

af21	010010
af22	010100
af23	010110
af31	011010
af32	011100
af33	011110
af41	100010
af42	100100
af43	100110
be	000001
cs1	001000
cs2	010000
cs3	011000
cs4	100000
cs5	101000
nc1/cs6/cs7	110000
nc2	111000
my1	110001

The following notes explain certain results in the mapping:

- **my1 110001:**
 - 110001 was not mapped to anything before, and **my1** is a new alias.
 - Nothing in the default mapping table is changed by this statement.
- **my2 101110:**
 - 101110 is now mapped to **my2** as well as **ef**.
- **be 000001:**
 - **be** is now mapped to 000001.
 - The old value of **be**, 000000, is not associated with any alias. Packets with this DSCP value are now mapped to the default forwarding class.
- **cs7 110000:**
 - **cs7** is now mapped to 110000, as well as **nc1** and **cs6**.
 - The old value of **cs7**, 111000, is still mapped to **nc2**.

**Related
Documentation**

- [Applying DSCP IPv6 Classifiers on page 1341](#)

Configuration Statements

- [\[edit class-of-service\] Hierarchy Level on page 1377](#)

[edit class-of-service] Hierarchy Level

```

class-of-service {
  classifiers {
    type classifier-name {
      forwarding-class class-name {
        loss-priority (high | low | medium-high | medium-low) code-points [ aliases bits ];
      }
      import (classifier-name | default);
    }
  }
  code-point-aliases {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) {
      alias-name bits;
    }
  }
  drop-profiles {
    profile-name {
      fill-level percentage drop-probability percentage;
      interpolate {
        drop-probability value;
        fill-level value;
      }
    }
  }
  fabric {
    scheduler-map {
      priority (high | low) scheduler scheduler-name;
    }
  }
  forwarding-class-map {
    map-name {
      class class-name queue-num queue-number <restricted-queue queue-number>;
    }
  }
  forwarding-classes {
    class class-name policing-priority (normal | premium) queue-num queue-number
      priority (high | low);
    queue queue-number class-name policing-priority (normal | premium) priority (high |
      low);
  }
  forwarding-policy {
    class class-name {
      classification-override {
        forwarding-class class-name;
      }
    }
    next-hop-map map-name {
      forwarding-class class-name {
        discard;
        lsp-next-hop [ lsp-regular-expressions ];
        next-hop [ next-hop-names ];
        non-lsp-next-hop;
      }
    }
  }
}

```

```
}
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      drop-timeout milliseconds;
      fragment-threshold bytes;
      multilink-class number;
      no-fragmentation;
    }
  }
}
host-outbound-traffic {
  dscp-code-point value;
  forwarding-class class-name;
  ieee-802.1 {
    default value;
    rewrite-rules;
  }
  tcp {
    raise-internet-control-priority;
  }
}
interfaces {
  ... the interfaces subhierarchy appears after the main [edit class-of-service] hierarchy
  ...
}
}
restricted-queues {
  forwarding-class class-name queue-number;
}
rewrite-rules {
  (dscp | dscp-ipv6 | exp | frame-relay-de | ieee-802.1 | ieee-802.1ad | inet-precedence)
  rewrite-rule {
    forwarding-class class-name {
      loss-priority level code-point (alias | bits);
    }
    import (rewrite-rule | default);
  }
}
}
routing-instances routing-instance-name {
  classifiers {
    dscp (classifier-name | default);
    dscp-ipv6 (classifier-name | default);
    exp (classifier-name | default);
    ieee-208.1 (classifier-name | default | encapsulated | vlan-tag (inner | outer));
  }
}
scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
}
schedulers {
  scheduler-name {
    adjust-minimum value;
    adjust-percent value;
```



```

    buffer-size (exact | percent percentage | remainder);
    drop-profile-map loss-priority (any | high | low | medium-high | medium-low)
        protocol any;
    excess-priority (high | low | medium-high | medium-low);
    excess-rate (percent percentage | proportion proportion);
    priority (high | low | medium-high | medium-low | strict-high);
    shaping-rate (bps | percent percentage | burst-size size);
    transmit-rate (bps | percent percentage | remainder) <exact | rate-limit>;
}
}
traceoptions {
    file <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
traffic-control-profiles {
    profile-name {
        adjust-minimum rate;
        delay-buffer-rate (bps | cps cps | percent percentage);
        excess-rate (percent percentage | proportion value);
        guaranteed-rate (bps | percent percentage) <burst-size bytes>;
        overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
        scheduler-map map-name;
        shaping-rate (bps | percent percentage) <burst-size bytes>;
    }
}
tri-color;
}

class-of-service {
    interfaces {
        interface-name {
            excess-bandwidth-share (equal | proportional value);
            input-excess-bandwidth-share (equal | proportional value);
            input-scheduler-map map-name;
            input-shaping-rate bps;
            input-traffic-control-profile profile-name;
            output-forwarding-class-map map-name;
            output-traffic-control-profile profile-name;
            scheduler-map map-name;
            scheduler-map-chassis (map-name | derived);
            shaping-rate bps;
            unit (logical-unit-number | *) {
                classifiers {
                    dscp (classifier-name | default) {
                        family [ inet mpls ];
                    }
                    dscp-ipv6 (classifier-name | default) {
                        family [ inet mpls ];
                    }
                    exp (classifier-name | default);
                    ieee-208.1 (classifier-name | default) <vlan-tag (inner | outer)>;
                    ieee-208.1ad (classifier-name | default);
                    inet-precedence (classifier-name | default);
                }
            }
        }
    }
}

```

```

forwarding-class class-name;
input-scheduler-map map-name;
input-shaping-rate bps;
input-traffic-control-profile profile-name shared-instance instance-name;
loss-priority-maps {
    (map-name | default);
}
loss-priority-rewrites {
    (map-name | default);
}
output-forwarding-class-map map-name;
output-traffic-control-profile profile-name shared-instance instance-name;
rewrite-rules {
    dscp (rule-name | default) <protocol mpls>;
    dscp-ipv6 (rule-name | default);
    exp (rule-name | default) <protocol [ mpls-any | mpls-inet-both |
        mpls-inet-both-non-vpn ]>;
    exp-push-push-push default;
    exp-swap-push-push default;
    ieee-802.1 (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
    ieee-802.1ad (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
    inet-precedence (rewrite-name | default) <protocol mpls>;
}
scheduler-map map-name;
shaping-rate bps;
translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 |
    to-exp-from-exp | to-inet-precedence-from-inet-precedence) table-name;
}
}
interface-set interface-set-name {
    excess-bandwidth-share (equal | proportional value);
    input-excess-bandwidth-share (equal | proportional value);
    input-traffic-control-profile profile-name;
    input-traffic-control-profile-remaining profile-name;
    internal-node;
    output-traffic-control-profile profile-name;
    output-traffic-control-profile-remaining profile-name;
}
}
}

```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

code-point-aliases

Syntax	code-point-aliases { type { alias-name bits; } }
Hierarchy Level	[edit class-of-service]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define an alias for a CoS marker.
Options	<p>alias-name—Name of the code-point alias.</p> <p>bits—6-bit value of the code-point bits, in decimal form.</p> <p>type—CoS marker type.</p> <p>Values: dscp, dscp-ipv6, exp, ieee-802.1, inet-precedence</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Defining Code Point Aliases for Bit Patterns on page 1375

dscp (Rewrite Rules)

Syntax	<code>dscp (rewrite-name default);</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For IPv4 traffic, apply a Differentiated Services (DiffServ) code point (DSCP) rewrite rule.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules dscp] hierarchy level.</p> <p>default—The default mapping.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694• dscp-ipv6 (Class-of-Service) on page 1357• exp on page 1358• exp-push-push-push on page 1723• exp-swap-push-push on page 1724• ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360• ieee-802.1ad on page 1362• inet-precedence on page 1364• rewrite-rules (Definition) on page 1481

dscp-ipv6 (Class-of-Service)

Syntax	<code>dscp-ipv6 (<i>rewrite-name</i> <default>) { protocol mpls }</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For IPv6 traffic, apply a DSCP rewrite rule.
Options	<p><i>rewrite-name</i>—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules dscp-ipv6] hierarchy level.</p> <p>default— Default mapping.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Rewrite Rules on page 1694 • dscp (Rewrite Rules) on page 1382 • exp on page 1358 • exp-push-push-push on page 1723 • exp-swap-push-push on page 1724 • ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360 • ieee-802.1ad on page 1362 • inet-precedence on page 1364 • rewrite-rules (Definition) on page 1481

exp

Syntax	<code>exp (rewrite-name default) protocol protocol-types;</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS Release 12.2. for ACX series
Description	Apply an MPLS experimental (EXP) rewrite rule.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules exp] hierarchy level.</p> <p>default—The default mapping.</p> <p>By default, IP precedence rewrite rules alter the first three bits on the type-of-service (ToS) byte while leaving the last three bits unchanged. This default behavior applies to rewrite rules you configure for MPLS packets with IPv4 payloads. You configure these types of rewrite rules by including the mpls-inet-both or mpls-inet-both-non-vpn option at the [edit class-of-service interfaces <i>interface</i> <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules exp <i>rewrite-rule-name</i> protocol] hierarchy level. The IP precedence rewrite rules explanation does not apply to ACX Series Universal Access routers.</p> <p>On interfaces configured on Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) on MX Series Ethernet Services Routers and EX Series switches, we highly recommend that you configure the default option when you configure a behavior aggregate (BA) classifier that does not include a specific rewrite rule for MPLS packets. Doing so ensures that MPLS exp value is rewritten according to the BA classifier rules configured for forwarding or packet loss priority. This does not apply to ACX Series Universal Access routers.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Rewriting the EXP Bits of All Three Labels of an Outgoing Packet</i> • dscp (Rewrite Rules) on page 1382 • dscp-ipv6 (Class-of-Service) on page 1357 • exp-push-push-push on page 1723 • exp-swap-push-push on page 1724 • ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360 • ieee-802.1ad on page 1362 • inet-precedence on page 1364

- [rewrite-rules \(Definition\) on page 1481](#)

ieee-802.1 (Rewrite Rules on Logical Interface)

Syntax	ieee-802.1 (<i>rewrite-name</i> default) vlan-tag (outer outer-and-inner);
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4. vlan-tag statement introduced in Junos OS Release 8.1.
Description	Apply an IEEE-802.1 rewrite rule. For IQ PICs, you can only configure one IEEE 802.1 rewrite rule on a physical port. All logical ports (units) on that physical port should apply the same IEEE 802.1 rewrite rule.
Options	rewrite-name —Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules ieee-802.1] hierarchy level. default —The default mapping.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Rewrite Rules on page 1694 • <i>Example: Configuring CoS for a PBB Network</i> • dscp (Rewrite Rules) on page 1382 • dscp-ipv6 (Class-of-Service) on page 1357 • exp on page 1358 • exp-push-push-push on page 1723 • exp-swap-push-push on page 1724 • ieee-802.1ad on page 1362 • inet-precedence on page 1364 • rewrite-rules (Definition) on page 1481

inet-precedence

Syntax	<code>inet-precedence (rewrite-name default);</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Apply a IPv4 precedence rewrite rule.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules inet-precedence] hierarchy level.</p> <p>default—The default mapping. By default, IP precedence rewrite rules alter the first three bits on the type of service (ToS) byte while leaving the last three bits unchanged.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694• dscp (Rewrite Rules) on page 1382• dscp-ipv6 (Class-of-Service) on page 1357• exp on page 1358• exp-push-push-push on page 1723• exp-swap-push-push on page 1724• ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360• ieee-802.1ad on page 1362• rewrite-rules (Definition) on page 1481

Classifying Packets Based on Various Packet Header Fields

- [Overview on page 1386](#)
- [Configuration on page 1390](#)

Overview

- [Overview on page 1386](#)

Overview

- [Multifield Classifier Overview on page 1387](#)
- [Overview of Simple Filters on page 1388](#)

- [Two-Color Policers and Shaping Rate Changes on page 1388](#)
- [Understanding IEEE 802.1p Inheritance push and swap from a Transparent Tag on page 1389](#)

Multifield Classifier Overview

A multifield classifier is a method of classifying traffic flows. Devices that sit at the edge of a network usually classify packets according to codings that are located in multiple packet header fields. Multifield classification is normally performed at the network edge because of the general lack of DiffServ code point (DSCP) or IP precedence support in end-user applications.

In an edge router, a multifield classifier provides the filtering functionality that scans through a variety of packet fields to determine the forwarding class for a packet. Typically, a classifier performs matching operations on the selected fields against a configured value.

Unlike a behavior aggregate (BA), which classifies packets based on class-of-service (CoS) bits in the packet header, a multifield classifier can examine multiple fields in the packet header—for example, the source and destination address of the packet, and the source and destination port numbers of the packet. A multifield classifier typically matches one or more of the six packet header fields: destination address, source address, IP protocol, source port, destination port, and DSCP. Multifield classifiers are used when a simple BA classifier is insufficient to classify a packet.

In the Juniper Networks® Junos® operating system (Junos OS), you configure a multifield classifier with a firewall filter and its associated match conditions. This enables you to use any filter match criteria to locate packets that require classification. From a CoS perspective, multifield classifiers (or firewall filter rules) provide the following services:

- Classify packets to a forwarding class and loss priority. The forwarding class determines the output queue. The loss priority is used by schedulers in conjunction with the random early discard (RED) algorithm to control packet discard during periods of congestion.
- Police traffic to a specific bandwidth and burst size. Packets exceeding the policer limits can be discarded, or can be assigned to a different forwarding class, to a different loss priority, or to both.



NOTE: You *police* traffic on input to conform to established CoS parameters, setting loss handling and forwarding class assignments as needed. You *shape* traffic on output to make sure that router resources, especially bandwidth, are distributed fairly. However, input policing and output shaping are two different CoS processes, each with their own configuration statements.

Overview of Simple Filters

Simple filters are recommended for metropolitan Ethernet applications. They are supported on Gigabit Ethernet intelligent queuing 2 (IQ2) and Enhanced Queuing Dense Port Concentrator (DPC) interfaces only.

Unlike normal filters, simple filters are for IPv4 traffic only and have the following restrictions:

- The **next term** action is not supported.
- Qualifiers, such as the **except** and **protocol-except** statements, are not supported.
- Noncontiguous masks are not supported.
- Multiple source addresses and destination addresses in a single term are not supported. If you configure multiple addresses, only the last one is used.
- Ranges are only valid as source or destination ports. For example, **source-port 400-500** or **destination-port 600-700**.
- Output filters are not supported. You can apply a simple filter to ingress traffic only.
- Simple filters are not supported for interfaces in an aggregated-Ethernet bundle.
- Explicitly configurable terminating actions, such as **accept**, **reject**, and **discard**, are not supported. Simple filters always accept packets.



NOTE: On the MX Series routers with the Enhanced Queuing DPC, the forwarding class is not supported as a **from** match condition.

Two-Color Policers and Shaping Rate Changes

When you configure a change in shaping rate, it is important to consider the effect on the bandwidth limit. Whenever the shaping rate changes, the bandwidth limit is adjusted based on whether a logical interface (unit) or bandwidth percentage policer is configured.

When a logical interface bandwidth policer is configured, the order of priority for the shaping rate (if configured at that level) is:

- The shaping rate applied to the logical interface (unit).
- The shaping rate applied to the physical interface (port).
- The physical interface speed.

When a bandwidth percentage policer is configured, the order of priority for the shaping rate (if configured at that level) is:

- The shaping rate applied to the physical interface (port).
- The physical interface speed.

These guidelines must be kept in mind when calculating the logical link speed and link speed from the configured shaping rate, which determines the rate-limited bandwidth after the policer is applied.

**Related
Documentation**

- [Example: Two-Color Policers and Shaping Rate Changes on page 1401](#)

Understanding IEEE 802.1p Inheritance push and swap from a Transparent Tag

MX Series routers with Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) and M320 routers and EX Series switches support configurable IEEE 802.1p inheritance of push and swap bits from the transparent tag of each incoming packet which allows you to classify incoming packets based on the IEEE 802.1p bits from the transparent tag.

During a tagging operation, Junos OS by default inherits the IEEE 802.1p bits from incoming tags in swap and push operations from the known tags configured on the interface.

It can be useful to override the default behavior by configuring Junos OS to inherit the IEEE 802.1p bits from a transparent tag, and to classify incoming packets based on the IEEE 802.1p bits of the incoming transparent tag. The configuration statements **swap-by-poppush** and **transparent** enable Junos OS to do this.

By default, during a swap operation, the IEEE 802.1p bits of the VLAN tag remain unchanged. When the **swap-by-poppush** operation is enabled on a logical interface, the swap operation is treated as a **pop** operation followed by **push** operation. The **pop** operation removes the existing tag and the associated IEEE 802.1p bits and the push operation copies the inner VLAN IEEE 802.1p bits to the IEEE bits of the VLAN or VLANs being pushed. As a result, the IEEE 802.1p bits are inherited from the incoming transparent tag.

To classify incoming packets based on the IEEE 802.1p bits from the transparent tag, include the **transparent** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* classifiers ieee-802.1 vlan-tag]** hierarchy level.

To configure Junos OS to inherit the IEEE 802.1p bits from the transparent tag, include the **swap-by-poppush** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.



NOTE: IEEE 802.1p Inheritance push and swap is only supported on untagged and single-tagged logical interfaces, and is not supported on dual-tagged logical interfaces.

**Related
Documentation**

- *swap-by-poppush*
- [transparent on page 1406](#)
- *Understanding swap-by-poppush*
- [Configuring IEEE 802.1p Inheritance push and swap from the Transparent Tag on page 1391](#)

- [Understanding Transparent Tag Operations and IEEE 802.1p Inheritance](#)

Configuration

- [Configuration Tasks on page 1390](#)
- [Examples on page 1393](#)
- [Configuration Statements on page 1402](#)

Configuration Tasks

- [Configuring Multifield Classifiers on page 1390](#)
- [Configuring Logical Bandwidth Policers on page 1391](#)
- [Configuring IEEE 802.1p Inheritance push and swap from the Transparent Tag on page 1391](#)

Configuring Multifield Classifiers

If you configure both a behavior aggregate (BA) classifier and a multifield classifier, BA classification is performed first; then multifield classification is performed. If they conflict, any BA classification result is overridden by the multifield classifier.



NOTE: For a specified interface, you can configure both a multifield classifier and a BA classifier without conflicts. Because the classifiers are always applied in sequential order, the BA classifier followed by the multifield classifier, any BA classification result is overridden by a multifield classifier if they conflict.

To activate a multifield classifier, you must configure it on a logical interface. There is no restriction on the number of multifield classifiers you can configure.



NOTE: For MX Series routers and EX Series switches, if you configure a firewall filter with a DSCP action or traffic-class action on a DPC, the commit does not fail, but a warning displays and an entry is made in the syslog.

For an L2TP LNS on MX Series routers, you can attach firewall for static LNS sessions by configuring these at logical interfaces directly on the inline services device (si-fpc/pic/port). RADIUS-configured firewall attachments are not supported.

To configure multifield classifiers, include the following statements at the **[edit firewall]** hierarchy level:

```
[edit firewall]
family family-name {
  filter filter-name {
    term term-name {
      from {
```

```

        match-conditions;
    }
    then {
        dscp 0;
        forwarding-class class-name;
        loss-priority (high | low);
    }
}
}
simple-filter filter-name {
    term term-name {
        from {
            match-conditions;
        }
        then {
            forwarding-class class-name;
            loss-priority (high | low | medium);
        }
    }
}
}
}

```

Configuring Logical Bandwidth Policers

When you configure a policer as a percentage (using the **bandwidth-percent** statement), the bandwidth is calculated as a percentage of either the physical interface media rate or the logical interface shaping rate. To specify that the bandwidth be calculated based on the logical interface shaping rate and not the physical interface media rate, include the **logical-bandwidth-policer** statement. If a shaping rate is not configured for the logical interface, the physical interface media rate is used, even if you include the **logical-bandwidth-policer**. You can configure the shaping rate on the logical interface using class-of-service statements.

```

[edit firewall policer policer-name]
logical-bandwidth-policer;

```

Configuring IEEE 802.1p Inheritance push and swap from the Transparent Tag

To classify incoming packets based on the IEEE 802.1p bits from the transparent tag, include the **transparent** statement at the **[edit class-of-service interfaces interface-name unit logical-unit-number classifiers ieee-802.1 vlan-tag]** hierarchy level.

Tagged Interface Example

The following example configuration specifies the classification based on the transparent VLAN tag.

```

edit
class-of-service {
    interfaces {
        ge-3/0/1 {
            unit 0 {
                classifiers {
                    ieee-802.1 default vlan-tag transparent;
                }
            }
        }
    }
}

```

```
}
```

To configure Junos OS to inherit the IEEE 802.1p bits from the transparent tag, include the **swap-by-poppush** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

The following is a configuration to swap and push VLAN tags and allow inheritance of the IEEE 802.1p value from the transparent VLAN tag in incoming packets.

```
edit
ge-3/0/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  unit 0 {
    encapsulation vlan-ccc;
    vlan-id 100;
    swap-by-poppush;
    input-vlan-map {
      swap-push;
      tag-protocol-id 0x9100;
      inner-tag-protocol-id 0x9100;
      vlan-id 500;
      inner-vlan-id 400;
    }
    output-vlan-map {
      pop-swap;
      inner-vlan-id 100;
      inner-tag-protocol-id 0x88a8;
    }
  }
}
```

The **swap-by-poppush** statement causes a swap operation to be done as a pop followed by a push operation. So for the outer tag, the incoming S-Tag is popped and a new tag is pushed. As a result, the S-Tag inherits the IEEE 802.1p bits from the transparent tag. The inner tag is then pushed, which results in the inner tag inheriting the IEEE 802.1p bits from the transparent tag.

Untagged Interface Example

The following is a configuration to push two VLAN tags and allow inheritance of the IEEE 802.1p value from the transparent VLAN tag in the incoming packet.

```
[edit]
ge-3/0/1 {
  encapsulation ccc;
  unit 0 {
    input-vlan-map {
      push-push;
      tag-protocol-id 0x9100;
      inner-tag-protocol-id 0x9100;
      vlan-id 500;
      inner-vlan-id 400;
    }
    output-vlan-map {
      pop-pop;
    }
  }
}
```

```
}
```

No additional configuration is required to inherit the IEEE 802.1p value, as the **push** operation inherits the IEEE 802.1p values by default.

The following configuration specifies the classification based on the transparent VLAN tag.

```
[edit]
class-of-service {
  interfaces {
    ge-3/0/1 {
      unit 0 {
        classifiers {
          ieee-802.1 default vlan-tag transparent;
        }
      }
    }
  }
}
```

Related Documentation

- [transparent on page 1406](#)
- *swap-by-poppush*
- [Understanding IEEE 802.1p Inheritance push and swap from a Transparent Tag on page 1389](#)
- *Understanding swap-by-poppush*
- *Understanding Transparent Tag Operations and IEEE 802.1p Inheritance*

Examples

- [Example: Classifying Packets Based on Their Destination Address on page 1393](#)
- [Example: Configuring and Verifying a Complex Multifield Filter on page 1394](#)
- [Example: Writing Different DSCP and EXP Values in MPLS-Tagged IP Packets on page 1397](#)
- [Example: Configuring a Simple Filter on page 1399](#)
- [Example: Configuring a Logical Bandwidth Policer on page 1400](#)
- [Example: Two-Color Policers and Shaping Rate Changes on page 1401](#)

Example: Classifying Packets Based on Their Destination Address

Configure a multifield classifier that ensures that all IPv4 packets destined for the **10.10.10.0/24** network are placed into the **platinum** forwarding class. This assignment occurs regardless of the received CoS bit values in the packet. Apply this filter to the inbound interface **ge-1/2/2.0**.

To verify that your configuration is attached to the correct interface, issue the **show interfaces filters** command.

```
[edit]
firewall {
```

```
family inet {
  filter set-FC-to-platinum {
    term match-a-single-route {
      from {
        destination-address {
          10.10.10.0/24;
        }
      }
      then {
        forwarding-class platinum;
        accept;
      }
    }
    term accept-all {
      then accept;
    }
  }
}
interfaces {
  ge-1/2/2 {
    unit 0 {
      family inet {
        filter {
          input set-FC-to-platinum;
        }
      }
    }
  }
}
```

Example: Configuring and Verifying a Complex Multifield Filter

In this example, SIP signaling (VoIP) messages use TCP/UDP, port 5060, and RTP media channels use UDP with port assignments from 16,384 through 32,767. See the following sections:

- [Configuring a Complex Multifield Filter on page 1394](#)
- [Verifying a Complex Multifield Filter on page 1396](#)

Configuring a Complex Multifield Filter

To configure the multifield filter, perform the following actions:

- Classify SIP signaling messages (VoIP network control traffic) as NC with a firewall filter.
- Classify VoIP traffic as EF with the same firewall filter.
- Police all remaining traffic with IP precedence 0 and make it BE.
- Police BE traffic to 1 Mbps with excess data marked with PLP high.
- Apply the firewall filter with policer to the interface.

The firewall filter called **classify** matches on the transport protocol and ports identified in the incoming packets and classifies packets into the forwarding classes specified by your criteria.

The first term, **sip**, classifies SIP signaling messages as network control messages. The **port** statement matches any source port or destination port (or both) that is coded to 5060.

Classifying SIP Signaling Messages

```

firewall {
  family inet {
    filter classify {
      interface-specific;
      term sip {
        from {
          protocol [ udp tcp ];
          port 5060;
        }
        then {
          forwarding-class network-control;
          accept;
        }
      }
    }
  }
}

```

The second term, **rtp**, classifies VoIP media channels that use UDP-based transport.

Classifying VoIP Channels That Use UDP

```

term rtp {
  from {
    protocol udp;
    port 16384-32767;
  }
  then {
    forwarding-class expedited-forwarding;
    accept;
  }
}

```

The policer's burst tolerance is set to the recommended value for a low-speed interface, which is ten times the interface MTU. For a high-speed interface, the recommended burst size is the transmit rate of the interface times 3 to 5 milliseconds.

Configuring the Policer

```

policer be-policer {
  if-exceeding {
    bandwidth-limit 1m;
    burst-size-limit 15k;
  }
  then loss-priority high;
}

```

The third term, **be**, ensures that all remaining traffic is policed according to a bandwidth restriction.

Policing All Remaining Traffic

```
term be {
  then policer be-policer;
}
```

The **be** term does not include a **forwarding-class** action modifier. Furthermore, there is no explicit treatment of network control (NC) traffic provided in the **classify** filter. You can configure explicit classification of NC traffic and all remaining IP traffic, but you do not need to, because the default IP precedence classifier correctly classifies the remaining traffic.

Apply the **classify** classifier to the **fe-0/0/2** interface:

Applying the Classifier

```
interfaces {
  fe-0/0/2 {
    unit 0 {
      family inet {
        filter {
          input classify;
        }
        address 10.12.0.13/30;
      }
    }
  }
}
```

Verifying a Complex Multifield Filter

Before the configuration is committed, display the default classifiers in effect on the interface using the **show class-of-service interface interface-name** command. The display confirms that the **ipprec-compatibility** classifier is in effect by default.

Verifying Default Classification

```
user@host> show class-of-service fe-0/0/2
Physical interface: fe-0/0/2, Index: 135
Queues supported: 8, Queues in use: 4
Scheduler map: <default>, Index: 2032638653
```

```
Logical interface: fe-0/0/2.0, Index: 68
Shaping rate: 32000
Object      Name                                Type      Index
Scheduler-map <default>                               27
Rewrite     exp-default                             exp       21
Classifier   exp-default                             exp        5
Classifier   ipprec-compatibility                     ip         8
```

To view the default classifier mappings, use the **show class-of-service classifier name name** command. The highlighted output confirms that traffic with IP precedence setting of 0 is correctly classified as BE, and NC traffic, with precedence values of 6 or 7, is properly classified as NC.

Displaying Default Classifier Mappings

```

user@host> show class-of-service classifier name ipprec-compatibility
Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 12
  Code point      Forwarding class      Loss priority
  ---
  000             best-effort          low
  001             best-effort          high
  010             best-effort          low
  011             best-effort          high
  100             best-effort          low
  101             best-effort          high
  110             network-control    low
  111             network-control    high

```

After your configuration is committed, verify that your multifield classifier is working correctly. You can monitor the queue counters for the routing device's **egress** interface used when forwarding traffic received from the peer. Displaying the queue counters for the ingress interface (**fe-0/0/2**) does not allow you to check your ingress classification, because queuing generally occurs only at egress in the Junos OS. (Ingress queuing is supported on Gigabit Ethernet IQ2 PICs and Enhanced IQ2 PICs only.)

To verify the operation of the multifield filter:

1. To determine which egress interface is used for the traffic, use the **traceroute** command.
2. After you identify the egress interface, clear its associated queue counters by issuing the **clear interfaces statistics interface-name** command.
3. Confirm the default forwarding class-to-queue number assignment. This allows you to predict which queues are used by the VoIP, NC, and other traffic. To do this, issue the **show class-of-service forwarding-class** command.
4. Display the queue counts on the interface by issuing the **show interfaces queue** command.

Example: Writing Different DSCP and EXP Values in MPLS-Tagged IP Packets

On Juniper Networks M320 Multiservice Edge Routers and T Series Core Routers and on EX Series switches, you can selectively set the DSCP field of MPLS-tagged IPv4 and IPv6 packets to **000000**. In the same packets, you can set the MPLS EXP field according to a configured rewrite table, which is based on the forwarding classes that you set in incoming packets using a BA or multifield classifier.

Queue selection is based on the forwarding classes you assign in scheduler maps. This means that you can direct traffic to a single output queue, regardless of whether the DSCP field is unchanged or rewritten to **000000**. To do this, you must configure a multifield classifier that matches selected packets and modifies them with the **dscp 0** action.

Selective marking of DSCP fields to **0**, without affecting output queue assignment, can be useful. For example, suppose you need to use the MPLS EXP value to configure CoS applications for core provider routing devices. At the penultimate egress provider edge (PE) router where the MPLS labels are removed, the CoS bits need to be provided by another value, such as DSCP code points. This case illustrates why it is useful to mark both the DSCP and MPLS EXP fields in the packet. Furthermore, it is useful to be able to

mark the two fields differently, because the CoS rules of the core provider routing device might differ from the CoS rules of the egress penultimate router. At egress, as always, you can use a rewrite table to rewrite the MPLS EXP values corresponding to the forwarding classes that you need to set.



NOTE: When both customer-facing and core-facing interfaces exist, you can derive the EXP value in the following precedence order, while adding the MPLS label:

1. EXP value provided by the CoS rewrite action.
2. EXP value derived from the top label of the stack (MPLS label stacking).
3. IPv4 or IPv6 precedence (Layer 3 VPN, Layer 2 VPN, and VPLS scenarios).

For IPv4 traffic, the **dscp 0** action modifier at the **[edit firewall family inet filter *filter-name* term *term-name* then]** hierarchy level is valid. However, for IPv6 traffic, you configure this feature by including the **traffic-class 0** action modifier at the **[edit firewall family inet6 filter *filter-name* term *term-name* then]** hierarchy level.

In the following IPv4 example, term 1 of the multifield classifier matches packets with DSCP **001100** code points coming from a certain VRF, rewrites the bits to DSCP **000000**, and sets the forwarding class to **best-effort**. In term 2, the classifier matches packets with DSCP **010110** code points and sets the forwarding class to **best-effort**. Because term 2 does not include the **dscp 0** action modifier, the DSCP **010110** bits remain unchanged. Because the classifier sets the forwarding class for both code points to **best-effort**, both traffic types are directed to the same output queue.



NOTE: If you configure a bit string in a DSCP match condition in a firewall filter, then you must include the letter “b” in front of the string, or the match rule creation fails on commit.

```
[edit]
firewall {
  family inet {
    filter vrf-rewrite {
      term 1 {
        from {
          dscp b001100;
        }
        then {
          dscp 0;
          forwarding-class best-effort;
        }
      }
      term 2 {
        from {
          dscp b010110;
        }
        then {
```

```

        forwarding-class best-effort;
    }
}
}
}
}

```

Applying the Multifield Classifier

Apply the filter to an input interface corresponding to the VRF:

```

[edit]
interfaces {
  ge-0/1/0 {
    unit 0 {
      family inet {
        filter input vrf-rewrite;
      }
    }
  }
}

```



NOTE: The `dscp 0` action is supported in both input and output filters. You can use this action for non-MPLS packets as well as for IPv4 and IPv6 packets entering an MPLS network. All IPv4 and IPv6 firewall filter match conditions are supported with the `dscp 0` action.

The following limitations apply:

- You can use a multifield classifier to rewrite DSCP fields to value 0 only. Other values are not supported.
- If a packet matches a filter that has the `dscp 0` action, then the outgoing DSCP value of the packet is 0, even if the packet matches a rewrite rule, and the rewrite rule is configured to mark the packet to a non-zero value. The `dscp 0` action overrides any other rewrite rule actions configured on the routing device.
- Although you can use the `dscp 0` action on an input filter, the output filter and other classifiers do not see the packet as being marked `dscp 0`. Instead, they classify the packet based on its original incoming DSCP value. The DSCP value of the packet is set to 0 after all other classification actions have completed on the packet.

Example: Configuring a Simple Filter

This simple filter sets the loss priority to low for TCP traffic with source address 1.1.1.1, sets the loss priority to high for HTTP (web) traffic with source addresses in the 4.0.0.0/8 range, and sets the loss priority to low for all traffic with destination address 6.6.6.6. The simple filter is applied as an input filter (arriving packets are checking for destination address 6.6.6.6, not queued output packets) on interface `ge-0/0/1.0`.

```

[edit]
firewall {
  family inet {

```

```
simple-filter filter1 {
  term 1 {
    from {
      source-address {
        1.1.1.1/32;
      }
      protocol {
        tcp;
      }
    }
    then loss-priority low;
  }
  term 2 {
    from {
      source-address {
        4.0.0.0/8;
      }
      source-port {
        http;
      }
    }
    then loss-priority high;
  }
  term 3 {
    from {
      destination-address {
        6.6.6.6/32;
      }
    }
    then {
      loss-priority low;
      forwarding-class best-effort;
    }
  }
}

interfaces {
  ge-0/0/1 {
    unit 0 {
      family inet {
        simple-filter {
          input filter1;
        }
        address 10.1.2.3/30;
      }
    }
  }
}
```

Example: Configuring a Logical Bandwidth Policer

This example applies a logical bandwidth policer rate to two logical interfaces on interface **ge-0/2/7**. The policed rate on **unit 0** is 2 Mbps (50 percent of 4 Mbps) and the policed rate on **unit 1** is 1 Mbps (50 percent of 2 Mbps).

```

[edit firewall]
policer Logical_Policer {
  logical-bandwidth-policer; # This applies the policer to logical interfaces
  if-exceeding {
    bandwidth-percent 50; # This applies 50 percent to the shaping-rate
    burst-size-limit 125k;
  }
  then discard;
}

[edit class-of-service]
interfaces {
  ge-0/2/7 {
    unit 0 {
      shaping-rate 4m # This establishes the rate to be policed on unit 0
    }
    unit 1 {
      shaping-rate 2m # This establishes the rate to be policed on unit 1
    }
  }
}
[edit interfaces ge-0/2/7]
per-unit-scheduler;
vlan-tagging;
unit 0 {
  vlan-id 100;
  family inet {
    policer {
      input Logical_Policer;
      output Logical_Policer;
    }
    address 172.1.1.1/30;
  }
}
unit 1 {
  vlan-id 200;
  family inet {
    policer {
      input Logical_Policer;
      output Logical_Policer;
    }
    address 172.2.1.1/30;
  }
}

```

Example: Two-Color Policers and Shaping Rate Changes

In this example, the shaping rate has been configured for the logical interface, but a bandwidth percentage policer is also configured. Therefore policing is based on the physical interface speed of 1 Gbps.

If both a shaping rate and a bandwidth percentage policer is configured on the same logical interface, the policing is based on the physical interface speed. Here is the example configuration:

```

[edit interfaces]

```

```
ge-0/1/0 {
  per-unit-scheduler;
  vlan-tagging;
  unit 0 {
    vlan-id 1;
    family inet {
      policer {
        output policer_test;
      }
      address 10.0.7.1/24;
    }
  }
}
```

```
[edit firewall]
policer policer_test {
  if-exceeding {
    bandwidth-percent 75;
    burst-size-limit 256k;
  }
  then discard;
}
```

```
[edit]
class-of-service {
  interfaces {
    ge-0/1/0 {
      unit 0 {
        shaping-rate 15m;
      }
    }
  }
}
```

Configuration Statements

- [\[edit class-of-service\] Hierarchy Level on page 1402](#)
- [\[edit firewall\] Hierarchy Level on page 1406](#)
- [\[edit interfaces\] Hierarchy Level on page 1427](#)

[edit class-of-service] Hierarchy Level

```
class-of-service {
  classifiers {
    type classifier-name {
      forwarding-class class-name {
        loss-priority (high | low | medium-high | medium-low) code-points [ aliases bits ];
      }
      import (classifier-name | default);
    }
  }
  code-point-aliases {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) {
      alias-name bits;
    }
  }
}
```



```

}
drop-profiles {
  profile-name {
    fill-level percentage drop-probability percentage;
    interpolate {
      drop-probability value;
      fill-level value;
    }
  }
}
fabric {
  scheduler-map {
    priority (high | low) scheduler scheduler-name;
  }
}
forwarding-class-map {
  map-name {
    class class-name queue-num queue-number <restricted-queue queue-number>;
  }
}
forwarding-classes {
  class class-name policing-priority (normal | premium) queue-num queue-number
  priority (high | low);
  queue queue-number class-name policing-priority (normal | premium) priority (high |
  low);
}
forwarding-policy {
  class class-name {
    classification-override {
      forwarding-class class-name;
    }
  }
  next-hop-map map-name {
    forwarding-class class-name {
      discard;
      lsp-next-hop [ lsp-regular-expressions ];
      next-hop [ next-hop-names ];
      non-lsp-next-hop;
    }
  }
}
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      drop-timeout milliseconds;
      fragment-threshold bytes;
      multilink-class number;
      no-fragmentation;
    }
  }
}
host-outbound-traffic {
  dscp-code-point value;
  forwarding-class class-name;
  ieee-802.1 {
    default value;
  }
}

```

```

rewrite-rules;
}
tcp {
  raise-internet-control-priority;
}
}
interfaces {
  ... the interfaces subhierarchy appears after the main [edit class-of-service] hierarchy
  ...
}
}
restricted-queues {
  forwarding-class class-name queue-number;
}
rewrite-rules {
  (dscp | dscp-ipv6 | exp | frame-relay-de | ieee-802.1 | ieee-802.1ad | inet-precedence)
  rewrite-rule {
    forwarding-class class-name {
      loss-priority level code-point (alias | bits);
    }
    import (rewrite-rule | default);
  }
}
}
routing-instances routing-instance-name {
  classifiers {
    dscp (classifier-name | default);
    dscp-ipv6 (classifier-name | default);
    exp (classifier-name | default);
    ieee-208.1 (classifier-name | default | encapsulated | vlan-tag (inner | outer));
  }
}
scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
schedulers {
  scheduler-name {
    adjust-minimum value;
    adjust-percent value;
    buffer-size (exact | percent percentage | remainder);
    drop-profile-map loss-priority (any | high | low | medium-high | medium-low)
      protocol any;
    excess-priority (high | low | medium-high | medium-low);
    excess-rate (percent percentage | proportion proportion);
    priority (high | low | medium-high | medium-low | strict-high);
    shaping-rate (bps | percent percentage | burst-size size);
    transmit-rate (bps | percent percentage | remainder) <exact | rate-limit>;
  }
}
}
traceoptions {
  file <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
}

```

```

traffic-control-profiles {
  profile-name {
    adjust-minimum rate;
    delay-buffer-rate (bps | cps cps | percent percentage);
    excess-rate (percent percentage | proportion value);
    guaranteed-rate (bps | percent percentage) <burst-size bytes>;
    overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
    scheduler-map map-name;
    shaping-rate (bps | percent percentage) <burst-size bytes>;
  }
}
tri-color;
}

class-of-service {
  interfaces {
    interface-name {
      excess-bandwidth-share (equal | proportional value);
      input-excess-bandwidth-share (equal | proportional value);
      input-scheduler-map map-name;
      input-shaping-rate bps;
      input-traffic-control-profile profile-name;
      output-forwarding-class-map map-name;
      output-traffic-control-profile profile-name;
      scheduler-map map-name;
      scheduler-map-chassis (map-name | derived);
      shaping-rate bps;
      unit (logical-unit-number | *) {
        classifiers {
          dscp (classifier-name | default) {
            family [ inet mpls ];
          }
          dscp-ipv6 (classifier-name | default) {
            family [ inet mpls ];
          }
          exp (classifier-name | default);
          ieee-208.1 (classifier-name | default) <vlan-tag (inner | outer)>;
          ieee-208.1ad (classifier-name | default);
          inet-precedence (classifier-name | default);
        }
        forwarding-class class-name;
        input-scheduler-map map-name;
        input-shaping-rate bps;
        input-traffic-control-profile profile-name shared-instance instance-name;
        loss-priority-maps {
          (map-name | default);
        }
        loss-priority-rewrites {
          (map-name | default);
        }
        output-forwarding-class-map map-name;
        output-traffic-control-profile profile-name shared-instance instance-name;
        rewrite-rules {
          dscp (rule-name | default) <protocol mpls>;
          dscp-ipv6 (rule-name | default);
        }
      }
    }
  }
}

```

```

exp (rule-name | default) <protocol [ mpls-any | mpls-inet-both |
    mpls-inet-both-non-vpn ]>;
exp-push-push-push default;
exp-swap-push-push default;
ieee-802.1 (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
ieee-802.1ad (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
inet-precedence (rewrite-name | default) <protocol mpls>;
}
scheduler-map map-name;
shaping-rate bps;
translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 |
    to-exp-from-exp | to-inet-precedence-from-inet-precedence) table-name;
}
}
interface-set interface-set-name {
    excess-bandwidth-share (equal | proportional value);
    input-excess-bandwidth-share (equal | proportional value);
    input-traffic-control-profile profile-name;
    input-traffic-control-profile-remaining profile-name;
    internal-node;
    output-traffic-control-profile profile-name;
    output-traffic-control-profile-remaining profile-name;
}
}
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

transparent

Syntax	transparent;
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> classifiers ieee802.1 vlan-tag]
Release Information	Statement introduced in Junos OS Release 11.2
Description	Packet classification based on the transparent VLAN tag.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

[edit firewall] Hierarchy Level

Several statements in the **[edit firewall]** hierarchy are valid at numerous locations within the hierarchy. To make the complete hierarchy easier to read, the repeated statements are listed in the following sections, which are referenced at the appropriate locations in “Complete **[edit firewall]** Hierarchy” on page 325.

- [Common Firewall Actions on page 1407](#)
- [Common IP Firewall Actions on page 1407](#)
- [Common IPv4 Firewall Actions on page 1408](#)

- [Common IP Firewall Match Conditions on page 1408](#)
- [Common IPv4 Firewall Match Conditions on page 1409](#)
- [Common Layer 2 Firewall Match Conditions on page 1409](#)
- [Complete \[edit firewall\] Hierarchy on page 1411](#)

Common Firewall Actions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit firewall\] Hierarchy](#)” on page 325 instead of the statements being repeated.

- **[edit firewall family (any | ccc | ethernet-switching | inet | inet6 | mpls | vpls) filter *filter-name* term *term-name* then]**
- **[edit firewall filter *filter-name* term *term-name* then]**

The common firewall actions are as follows:

```
count counter-name;
forwarding-class class-name;
loss-priority (high | low | medium-high | medium-low);
next term;
policer policer-name;
three-color-policer policer-name {
    (single-rate single-rate-policer-name | two-rate two-rate-policer-name);
}
```

Common IP Firewall Actions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit firewall\] Hierarchy](#)” on page 325 instead of the statements being repeated.

- **[edit firewall family inet filter *filter-name* term *term-name* then]**
- **[edit firewall family inet6 filter *filter-name* term *term-name* then]**
- **[edit firewall filter *filter-name* term *term-name* then]**

The common IP firewall actions are as follows:

```
log;
logical-system logical-system-name <routing-instance routing-instance-name>
    <topology topology-name>;
port-mirror;
port-mirror-instance instance-name;
routing-instance routing-instance-name <topology topology-name>;
sample;
service-filter-hit;
syslog;
topology topology-name;
```

Common IPv4 Firewall Actions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- [edit firewall family inet filter *filter-name* term *term-name* then]
- [edit firewall filter *filter-name* term *term-name* then]

The common IP version 4 (IPv4) firewall actions are as follows:

```
(accept | discard <accounting collector-name> | reject <administratively-prohibited |
  bad-host-tos | bad-network-tos | fragmentation-needed | host-prohibited |
  host-unknown | host-unreachable | network-prohibited | network-unknown |
  network-unreachable | port-unreachable | precedence-cutoff | precedence-violation |
  protocol-unreachable | source-host-isolated | source-route-failed | tcp-reset>);
ipsec-sa sa-name;
load-balance sa-name;
next-hop-group group-name;
prefix-action action-name;
```

Common IP Firewall Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- [edit firewall family inet dialer-filter *filter-name* term *term-name* from] (with the exceptions noted at this level in [“Complete \[edit firewall\] Hierarchy” on page 325](#))
- [edit firewall family inet filter *filter-name* term *term-name* from]
- [edit firewall family inet6 dialer-filter *filter-name* term *term-name* from] (with the exceptions noted at this level in [“Complete \[edit firewall\] Hierarchy” on page 325](#))
- [edit firewall family inet6 filter *filter-name* term *term-name* from]
- [edit firewall filter *filter-name* term *term-name* from]

The common IP firewall match conditions are as follows:

```
address {
  ip-prefix</prefix-length> <except>;
}
destination-address {
  ip-prefix</prefix-length> <except>;
}
destination-class [ class-names ] | destination-class-except [ class-names ];
(destination-port [ port-names ] | destination-port-except [ port-names ]);
destination-prefix-list {
  list-name <except>;
}
(forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
icmp-code [ codes ] | icmp-code-except [ codes ];
icmp-type [ types ] | icmp-type-except [ types ];
interface interface-name;
```

```

(interface-group [ group-names ] | interface-group-except [ group-names ]);
interface-set set-name;
(loss-priority [ priorities ] | loss-priority-except [ priorities ]);
(packet-length [ values ] | packet-length-except [ values ]);
(port [ port-names ] | port-except [ port-names ]);
prefix-list {
    list-name <except>;
}
service-filter-hit;
source-address {
    ip-prefix/prefix-length> <except>;
}
(source-class [ class-names ] | source-class-except [ class-names ]);
(source-port [ port-names ] | source-port-except [ port-names ]);
source-prefix-list {
    list-name <except>;
}
tcp-established;
tcp-flags flag;
tcp-initial;

```

Common IPv4 Firewall Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit firewall\] Hierarchy](#)” on page 325 instead of the statements being repeated.

- [edit firewall family inet dialer-filter *filter-name* term *term-name* from] (with the exceptions noted at this level in “[Complete \[edit firewall\] Hierarchy](#)” on page 325)
- [edit firewall family inet filter *filter-name* term *term-name* from]
- [edit firewall filter *filter-name* term *term-name* from]

The common IPv4 firewall match conditions are as follows:

```

(ah-spi [ values ] | ah-spi-except [ values ]);
(dscp [ code-point-values ] | dscp-except [ code-point-values ]);
(esp-spi [ values ] | esp-spi-except [ values ]);
first-fragment;
fragment-flags flag;
(fragment-offset [ offsets ] | fragment-offset-except [ offsets ]);
(ip-options [ option-names ] | ip-options-except [ option-names ]);
is-fragment;
(precedence [ precedence-names ] | precedence-except [ precedence-names ]);
(protocol [ protocol-names ] | protocol-except [ protocol-names ]);
(ttl [ ttl-values ] | ttl-except [ ttl-values ]);

```

Common Layer 2 Firewall Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit firewall\] Hierarchy](#)” on page 325 instead of the statements being repeated.

- [edit firewall family ethernet-switching filter *filter-name* term *term-name* from]
- [edit firewall family vpls filter *filter-name* term *term-name* from]

The common Layer 2 firewall match conditions are as follows:

```
destination-mac-address {
    mac-address <except>;
}
(destination-port [ port-names ] | destination-port-except [ port-names ]);
(dscp [ code-point-values ] | dscp-except [ code-point-values ]);
(ether-type [ protocol-types ] | ether-type-except [ protocol-types ]);
(forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
icmp-code [ codes ] | icmp-code-except [ codes ];
icmp-type [ types ] | icmp-type-except [ types ];
(interface-group [ group-names ] | interface-group-except [ group-names ]);
ip-address {
    ip-prefix</prefix-length> <except>;
}
ip-destination-address {
    ip-prefix</prefix-length> <except>;
}
(ip-precedence [ precedence-names ] | ip-precedence-except [ precedence-names ]);
(ip-protocol [ protocol-names ] | ip-protocol-except [ protocol-names ]);
ip-source-address ip-prefix</prefix-length>;
(learn-vlan-lp-priority [ priorities ] | learn-vlan-lp-priority [ priorities ]);
(learn-vlan-id [ vlan-ids ] | learn-vlan-id-except [ vlan-ids ]);
(loss-priority [ priorities ] | loss-priority-except [ priorities ]);
(port [ port-names ] | port-except [ port-names ]);
source-mac-address {
    mac-address <except>;
}
(source-port [ port-names ] | source-port-except [ port-names ]);
tcp-flags flag;
(traffic-type [ broadcast known-unicast multicast unknown-unicast ] |
    traffic-type-except [ broadcast known-unicast multicast unknown-unicast ]);
(user-vlan-lp-priority [ priorities ] | user-vlan-lp-priority [ priorities ]);
(user-vlan-id [ vlan-ids ] | user-vlan-id-except [ vlan-ids ]);
(vlan-ether-type [ protocol-types ] | vlan-ether-type-except [ protocol-types ]);
```


Complete [edit firewall] Hierarchy

```

firewall {
  family (any | ccc | ethernet-switching | inet | inet6 | mpls | vpls) {
    ... the family subhierarchies appear after the main [edit firewall] hierarchy ...
  }
  filter filter-name {
    accounting-profile [ profile-names ];
    enhanced-mode;
    interface-shared-with;
    interface-specific;
    physical-interface-policer;
    term term-name {
      filter filter-name;
      from {
        ... statements in Common IP Firewall Match Conditions on page 322 AND
        statements in Common IPv4 Firewall Match Conditions on page 323 ...
      }
      then {
        ... statements in Common Firewall Actions on page 320 AND
        statements in Common IP Firewall Actions on page 321 AND
        statements in Common IPv4 Firewall Actions on page 321 ...
      }
    }
  }
}
hierarchical-policer policer-name {
  aggregate {
    if-exceeding {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    then {
      discard;
      forwarding-class class-name;
      loss-priority (high | low | medium-high | medium-low);
    }
  }
  logical-interface-policer;
  physical-interface-policer;
  premium {
    if-exceeding {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    then {
      discard;
    }
  }
}
}
shared-bandwidth-policer;
interface-set interface-set-name {
  interface-name;
}
load-balance-group group-name {
  next-hop-group [ group-names ];
}

```

```
}
policer policer-name {
  filter-specific;
  if-exceeding {
    (bandwidth-limit bps | bandwidth-percent percentage);
    burst-size-limit bytes;
  }
  logical-bandwidth-policer;
  logical-interface-policer;
  physical-interface-policer;
  then {
    discard;
    forwarding-class class-name;
    loss-priority (high | low | medium-high | medium-low);
  }
}
three-color-policer policer-name {
  action {
    loss-priority high then discard;
  }
  filter-specific;
  logical-interface-policer;
  physical-interface-policer;
  shared-bandwidth-policer;
  single-rate {
    (color-aware | color-blind);
    committed-burst-size bytes;
    committed-information-rate bps;
    excess-burst-size bytes;
  }
  two-rate {
    (color-aware | color-blind);
    committed-burst-size bytes;
    committed-information-rate bps;
    peak-burst-size bytes;
    peak-information-rate bps;
  }
}
}

firewall {
  family any {
    filter filter-name {
      interface-shared;
      term term-name {
        from {
          (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
          interface interface-name;
          interface-set set-name;
          (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
          (packet-length [ values ] | packet-length-except [ values ]);
        }
        then {
          ... statements in Common Firewall Actions on page 320 PLUS ...
          (accept | discard);
        }
      }
    }
  }
}
```

```

    }
  }
}

firewall {
  family ccc {
    filter filter-name {
      accounting-profile [ profile-names ];
      physical-interface-filter;
      interface-specific;
      term term-name {
        filter filter-name;
        from {
          (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
          (interface-group [ group-names ] | interface-group-except [ group-names ]);
          (learn-vlan-1p-priority [ priorities ] | learn-vlan-1p-priority [ priorities ]);
          (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
          (user-vlan-1p-priority [ priorities ] | user-vlan-1p-priority [ priorities ]);
        }
        then {
          ... statements in Common Firewall Actions on page 320 PLUS ...
          (accept | discard);
          port-mirror-instance instance-name;
        }
      }
    }
  }
}

firewall {
  family ethernet-switching {
    filter filter-name {
      interface-specific;
      term term-name {
        from {
          destination-address {
            ip-prefix</prefix-length>;
          }
          destination-mac-address {
            mac-address;
          }
          destination-port [ port-names ];
          destination-prefix-list {
            list-name;
          }
          dot1q-tag [ tag-values ];
          dot1q-user-priority [ priority-values ];
          dscp [ code-point-values ];
          ether-type [ protocol-names ];
          fragment-flags flag;
          icmp-code [ codes ];
          icmp-type [ types ];
          interface interface-name;
          is-fragment;

```

```

precedence [ precedence-names ];
protocol [ protocol-names ];
source-address {
    ip-prefix < / prefix-length >;
}
source-mac-address {
    mac-address;
}
source-port [ port-names ];
source-prefix-list {
    list-name;
}
tcp-established;
tcp-flags flag;
tcp-initial;
vlan [ vlan-names ];
}
then {
    (accept | discard);
    analyzer analyzer-name;
    count counter-name;
    forwarding-class class-name;
    interface interface-name;
    log;
    loss-priority (high | low);
    policer policer-name;
    syslog;
    vlan vlan-name;
}
}
}
}
}

firewall {
    family inet {
        dialer-filter filter-name {
            accounting-profile [ profile-names ];
            term term-name {
                from {
                    ... statements in Common IP Firewall Match Conditions on page 322 AND
                    statements in Common IPv4 Firewall Match Conditions on page 323 EXCEPT
                    FOR ...
                    (ah-spi [ values ] | ah-spi-except [ values ]); # NOT valid at this level
                    (destination-class [ class-names ] |
                     destination-class-except [ class-names ]); # NOT valid at this level
                    interface interface-name; # NOT valid at this level
                    (loss-priority [ priorities ] | loss-priority-except [ priorities ]); # NOT valid at
                     this level
                    service-filter-hit; # NOT valid at this level
                    (source-class [ class-names ] | source-class-except [ class-names ]); # NOT
                     valid at this level
                }
            }
            then {
                (ignore | note);
                log;
            }
        }
    }
}

```

```

        sample;
        syslog;
    }
}
filter filter-name {
    accounting-profile [ profile-names ];
    interface-specific;
    term term-name {
        filter filter-name;
        from {
            ... statements in Common IP Firewall Match Conditions on page 322 AND
               statements in Common IPv4 Firewall Match Conditions on page 323 ...
        }
        then {
            ... statements in Common Firewall Actions on page 320 AND
               statements in Common IP Firewall Actions on page 321 AND
               statements in Common IPv4 Firewall Actions on page 321 ...
        }
    }
}
prefix-action name {
    count;
    destination-prefix-length prefix-length;
    filter-specific;
    policer policer-name;
    source-prefix-length prefix-length;
    subnet-prefix-length prefix-length;
}
service-filter filter-name {
    term term-name {
        from {
            address {
                ip-prefix</prefix-length>;
            }
            (ah-spi [ values ] | ah-spi-except [ values ]);
            destination-address {
                ip-prefix</prefix-length>;
            }
            (destination-port [ port-names ] | destination-port-except [ port-names ]);
            destination-prefix-list {
                list-name;
            }
            (esp-spi [ values ] | esp-spi-except [ values ]);
            first-fragment;
            fragment-flags flag;
            (fragment-offset [ offsets ] | fragment-offset-except [ offsets ]);
            (interface-group [ group-names ] | interface-group-except [ group-names ]);
            (ip-options [ option-names ] | ip-options-except [ option-names ]);
            is-fragment;
            (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
            (port [ port-names ] | port-except [ port-names ]);
            prefix-list {
                list-name;
            }
            (protocol [ protocol-names ] | protocol-except [ protocol-names ]);
        }
    }
}

```

```

    source-address {
        ip-prefix</prefix-length>;
    }
    (source-port [ port-names ] | source-port-except [ port-names ]);
    source-prefix-list {
        list-name;
    }
    tcp-flags flag-name;
}
then {
    count counter-name;
    log;
    port-mirror;
    sample;
    (service | skip);
}
}
}
simple-filter filter-name {
    term term-name {
        from {
            destination-address ip-prefix</prefix-length>;
            destination-port port-name;
            forwarding-class [ class-names ];
            protocol protocol-name;
            source-address ip-prefix</prefix-length>;
            source-port port-name;
        }
        then {
            forwarding-class class-name;
            loss-priority (high | low | medium-high | medium-low);
            policer policer-name;
        }
    }
}
}
}
}
firewall {
    family inet6 {
        dialer-filter filter-name {
            accounting-profile [ profile-names ];
            term term-name {
                from {
                    ... statements in Common IP Firewall Match Conditions on page 322 PLUS ...
                    (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
                    ... BUT NOT ...
                    (destination-class [ class-names ] |
                     destination-class-except [ class-names ]); # NOT valid at this level
                    (forwarding-class [ class-names ] |
                     forwarding-class-except [ class-names ]); # NOT valid at this level
                    interface interface-name; # NOT valid at this level
                    (interface-group [ group-names ] | interface-group-except [ group-names ]); #
                     NOT valid at this level
                    (loss-priority [ priorities ] | loss-priority-except [ priorities ]); # NOT valid at
                     this level
                }
            }
        }
    }
}
}
```

```

        service-filter-hit; # NOT valid at this level
        (source-class [ class-names ] | source-class-except [ class-names ]); # NOT
            valid at this level
        tcp-established; # NOT valid at this level
        tcp-flags flag; # NOT valid at this level
        tcp-initial; # NOT valid at this level
    }
    then {
        (ignore | note);
        log;
        sample;
        syslog;
    }
}
}
filter filter-name {
    accounting-profile [ profile-names ];
    interface-specific;
    term term-name {
        filter filter-name;
        from {
            ... statements in Common IP Firewall Match Conditions on page 322 PLUS ...
            (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
            (traffic-class [ code-point-values ] | traffic-class-except [ code-point-values ]);
        }
        then {
            ... statements in Common Firewall Actions on page 320 AND
            statements in Common IP Firewall Actions on page 321 PLUS ...
            (accept | discard | reject <address-unreachable | administratively-prohibited |
                beyond-scope | fragmentation-needed | no-route | port-unreachable |
                tcp-reset>);
        }
    }
}
}
service-filter filter-name {
    term term-name {
        from {
            address {
                ip-prefix</prefix-length>;
            }
            (ah-spi [ values ] | ah-spi-except [ values ]);
            destination-address {
                ip-prefix</prefix-length>;
            }
            (destination-port [ port-names ] | destination-port-except [ port-names ]);
            destination-prefix-list {
                list-name;
            }
            (esp-spi [ values ] | esp-spi-except [ values ]);
            (interface-group [ group-names ] | interface-group-except [ group-names ]);
            (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
            (port [ port-names ] | port-except [ port-names ]);
            prefix-list {
                list-name;
            }
            source-address {

```

```

        ip-prefix </prefix-length>;
    }
    (source-port [ port-names ] | source-port-except [ port-names ]);
    source-prefix-list {
        list-name;
    }
    tcp-flags flag-name;
}
then {
    count counter-name;
    log;
    port-mirror;
    sample;
    (service | skip);
}
}
}
}
}

firewall {
    family mpls {
        filter filter-name {
            accounting-profile [ profile-names ];
            interface-specific;
            physical-interface-filter;
            term term-name {
                from {
                    (exp [ exp-bits ] | exp-except [ exp-bits ]);
                }
                then {
                    (ignore | note);
                    log;
                    sample;
                    syslog;
                }
            }
        }
    }
    filter filter-name {
        accounting-profile [ profile-names ];
        interface-specific;
        physical-interface-filter;
        term term-name {
            filter filter-name;
            from {
                (exp [ exp-bits ] | exp-except [ exp-bits ]);
                (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
                interface interface-name;
                interface-set set-name;
                (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
            }
            then {
                ... statements in Common Firewall Actions on page 320 PLUS ...
                (accept | discard);
                sample;
            }
        }
    }
}

```



```

    }
  }
}

firewall {
  family vpls {
    filter filter-name {
      accounting-profile [ profile-names ];
      interface-specific;
      term term-name {
        filter filter-name;
        from {
          ... statements in Common Layer 2 Firewall Match Conditions on page 323 ...
        }
        then {
          ... statements in Common Firewall Actions on page 320 PLUS ...
          (accept | discard);
          port-mirror;
          port-mirror-instance instance-name;
        }
      }
    }
  }
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

dscp (Multifield Classifier)

Syntax	<code>dscp [0 <i>value</i>];</code>
Hierarchy Level	<code>[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> then]</code>
Release Information	Statement introduced in Junos OS Release 7.4.
Description	<p>For M320 and T Series routers, set the DSCP field of incoming or outgoing packets to 000000. On the same packets, you can use a behavior aggregate (BA) classifier and a rewrite rule to rewrite the MPLS EXP field.</p> <p>For MX Series routers with MPCs and EX Series switches, the DSCP field can be set from a numeric range.</p> <p>For MX Series routers and EX Series switches, if you configure a firewall filter with a DSCP action or traffic-class action on a DPC, the commit does not fail, but the filter is not applied to the interface, a warning displays, and an entry is made in the syslog.</p>
Options	value —For MX Series routers with MPCs and EX Series switches, specify the field of incoming or outgoing packets in the range from 0 through 63 .
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Applying Tricolor Marking Policers to Firewall Filters on page 1456

family (Multifield Classifier)

```
Syntax  family family-name {
        filter filter-name {
            term term-name {
                from {
                    match-conditions;
                }
                then {
                    dscp 0;
                    forwarding-class class-name;
                    loss-priority (high | low);
                    three-color-policer {
                        (single-rate | two-rate) policer-name;
                    }
                }
            }
        }
    }
```

Hierarchy Level [edit [firewall](#)]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a firewall filter for IP version 4 (IPv4) or IP version 6 (IPv6) traffic.

Options *family-name*—Protocol family:

- **ccc**—Circuit cross-connect parameters
- **inet**—IPv4 parameters
- **inet6**—IPv6 protocol parameters
- **iso**—OSI ISO protocol parameters
- **mpls**—MPLS protocol parameters
- **tcc**—Translational cross-connect parameters
- **vpls**—Virtual private LAN service parameters.

The remaining statements are explained separately.

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

Related Documentation • [Configuring Multifield Classifiers on page 1390](#)

filter (Configuring)

Syntax	<pre>filter <i>filter-name</i> { accounting-profile <i>name</i>; enhanced-mode; interface-shared; interface-specific; physical-interface-filter; term <i>term-name</i> { filter <i>filter-name</i>; from { match-conditions; } then { actions; } } }</pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall family <i>family-name</i>], [edit firewall family <i>family-name</i>], [edit logical-systems <i>logical-system-name</i> firewall family <i>family-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Logical systems support introduced in Junos OS Release 9.3. physical-interface-filter statement introduced in Junos OS Release 9.6. Support at the [edit dynamic-profiles ... family <i>family-name</i>] hierarchy level introduced in Junos OS Release 11.4. Support for the interface-shared > statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure firewall filters.
Options	<p><i>filter-name</i>—Name that identifies the filter. This must be a non-reserved string of not more than 64 characters. To include spaces in the name, enclose it in quotation marks (" "). In Junos OS Release 9.0 and later, you can no longer use special characters within the name of a firewall filter. Firewall filter names are restricted from having the form _.* (beginning and ending with underscores) or _.* (beginning with an underscore).</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	firewall —To view this statement in the configuration. firewall-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Guidelines for Configuring Standard Firewall Filters on page 4478• Guidelines for Applying Standard Firewall Filters on page 4483• Configuring Multifield Classifiers on page 1390• Using Multifield Classifiers to Set PLP on page 1461

- [simple-filter on page 4798](#)

forwarding-class (Multifield Classifiers)

Syntax	<code>forwarding-class <i>class-name</i>;</code>
Hierarchy Level	<code>[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> then]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Set the forwarding class of incoming packets.
Options	<i>class-name</i> —Name of the forwarding class.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	• Configuring Multifield Classifiers on page 1390

from

Syntax	<code>from { applications [<i>application-name</i>]; application-sets [<i>set-name</i>]; destination-address (CoS) <i>address</i>; source-address <i>address</i>; }</code>
Hierarchy Level	<code>[edit services cos rule <i>rule-name</i> term <i>term-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 8.1.
Description	Specify input conditions for a CoS term.
Options	The remaining statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	• Configuring CoS Rule Sets

loss-priority (Normal Filter)

Syntax	loss-priority (high low);
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> then]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Set the loss priority of incoming packets.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Multifield Classifiers on page 1390

loss-priority (Simple Filter)

Syntax	loss-priority (high low medium);
Hierarchy Level	[edit firewall family <i>family-name</i> simple-filter <i>filter-name</i> term <i>term-name</i> then]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	Set the loss priority of incoming packets.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Multifield Classifiers on page 1390

simple-filter (Configuring)

Syntax	<pre> simple-filter <i>filter-name</i> { term <i>term-name</i> { from { match-conditions; } then { forwarding-class <i>class-name</i>; loss-priority (high low medium); } } } </pre>
Hierarchy Level	[edit firewall family inet filter <i>filter-name</i>]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	Define a simple filter. Simple filters are recommended for metropolitan Ethernet applications.
Options	<p>from—Match packet fields to values. If the from option is not included, all packets are considered to match and the actions and action modifiers in the then statement are taken.</p> <p>match-conditions—One or more conditions to use to make a match.</p> <p>term-name—Name that identifies the term. The name can contain letters, numbers, and hyphens (-), and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").</p> <p>then—Actions to take on matching packets. If the then option is not included and a packet matches all the conditions in the from statement, the packet is accepted.</p> <p>The remaining statements are explained separately. Only forwarding-class and loss-priority are valid in a simple filter configuration.</p>
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Multifield Classifiers on page 1390 • filter (Applying to an Interface) on page 1439 • simple-filter (Applying to an Interface) on page 1440

term (Simple Filter)

Syntax	<pre>term <i>term-name</i> { from { <i>match-conditions</i>; } then { forwarding-class <i>class-name</i>; loss-priority (high low medium); } }</pre>
Hierarchy Level	[edit firewall family inet simple-filter <i>filter-name</i>]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	Define a simple filter term.
Options	<p>from—Match packet fields to values. If the from option is not included, all packets are considered to match and the actions and action modifiers in the then statement are taken.</p> <p>match-conditions—One or more conditions to use to make a match.</p> <p>term-name—Name that identifies the term. The name can contain letters, numbers, and hyphens (-), and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").</p> <p>then—Actions to take on matching packets. If the then option is not included and a packet matches all the conditions in the from statement, the packet is accepted. For CoS, only the actions listed are allowed. These statements are explained separately.</p>
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Multifield Classification on page 4889• Simple Filter Overview on page 4529• Standard Firewall Filter Match Conditions for IPv4 Traffic on page 4707• Standard Firewall Filter Match Conditions for IPv6 Traffic on page 4716

then

Syntax	<pre> then { application-profile <i>profile-name</i>; dscp (<i>alias</i> <i>bits</i>); forwarding-class <i>class-name</i>; syslog; (reflexive reverse) { application-profile <i>profile-name</i>; dscp (<i>alias</i> <i>bits</i>); forwarding-class <i>class-name</i>; syslog; } } </pre>
Hierarchy Level	[edit services cos rule <i>rule-name</i> term <i>term-name</i>]
Release Information	Statement introduced in Junos OS Release 8.1.
Description	<p>Define the CoS term actions.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Actions in a CoS Rule</i> • <i>Configuring Actions in CoS Rules</i>

[edit interfaces] Hierarchy Level

The following statement hierarchy can also be included at the [edit logical-systems *logical-system-name*] hierarchy level.

```

interfaces {
    interface-name {
        ... the “interface-name” subhierarchy appears after the main [edit interfaces] hierarchy level ...
    }
    interface-set interface-set-name {
        interface interface-name {
            (unit unit-number | vlan-tags-outer vlan-tag);
        }
    }
    irb (Interfaces) {
        accounting-profile name;
        description text;
        disable;

        (gratuitous-arp-reply | no-gratuitous-arp-reply);
        hold-time up milliseconds down milliseconds;
        mtu bytes;
    }
}

```

```
no-gratuitous-arp-request;

traceoptions {
  flag flag;
}
(traps | no-traps);
unit logical-unit-number {
  accounting-profile name;
  bandwidth rate;
  description text;
  disable;
  encapsulation type;
  family inet {
    accounting {
      destination-class-usage;
      source-class-usage {
        input;
        output;
      }
    }
  }
  address ipv4-address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    preferred;
    primary;
    vrrp-group group-id {
      (accept-data | no-accept-data);
      advertise-interval seconds;
      advertisements-threshold number;
      authentication-key key;
      authentication-type authentication;
      fast-interval milliseconds;
      (preempt | no-preempt) {
        hold-time seconds;
      }
      priority number;
      track {
        interface interface-name {
          bandwidth-threshold bits-per-second priority-cost priority;
          priority-cost priority;
        }
        priority-hold-time seconds;
        route prefix/prefix-length routing-instance instance-name priority-cost priority;
      }
      virtual-address [ addresses ];
      vrrp-inherit-from vrrp-group;
    }
  }
}
filter {
  input filter-name;
  output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
```

```

primary;
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bandwidth priority-cost number;
                priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-inet6-address [addresses];
        virtual-link-local-address ipv6-address;
        vrrp-inherit-from {
            active-group group-number;
            active-interface interface-name;
        }
    }
}
}
(dad-disable | no-dad-disable);
filter {
    input filter-name;
    output filter-name;

```

```
    }
    mtu bytes;
    nd6-stale-time seconds;
    no-neighbor-learn;
    no-redirects;
    policer {
        input policer-name;
        output policer-name;
    }
    rpf-check {
        fail-filter filter-name;
        mode {
            loose;
        }
    }
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    filter {
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    policer {
        input policer-name;
        output policer-name;
    }
}
native-inner-vlan-id vlan-id;
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id-list [vlan-id's];
vlan-id-range [vlan-id-range];
}
}
traceoptions {
    file <filename> <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag <disable>;
    no-remote-trace;
}
}

interfaces {
    interface-name {
        disable;
        accounting-profile name;
        aggregated-ether-options {
            ethernet-switch-profile {
                tag-protocol-id [ hexadecimal-identifiers ];
            }
            (flow-control | no-flow-control);
            lacp {
```

```

(active | passive);
admin-key key;
fast-failover;
link-protection {
    disable;
    (revertive | non-revertive);
}
periodic (fast | slow);
system-id mac-address;
system-priority priority;
}
(link-protection | no-link-protection);
link-speed (100m | 1g | 8g | 10g | 40g | 50g | 80g | 100g | oc192);
logical-interface-fpc-redundancy;
(loopback | no-loopback);
mc-ae {
    chassis-id chassis-id;
    events {
        iccp-peer-down {
            force-icl-down;
            prefer-status-control-active;
        }
    }
    mc-ae-id mc-ae-id;
    mode (active-active | active-standby);
    redundancy-group group-id;
    status-control (active | standby);
}
minimum-links number;
rebalance-periodic {
    start-time time;
    interval number;
}
source-address-filter {
    mac-address;
}
(source-filtering | no-source-filtering);
}
auto-configure {
    remove-when-no-subscribers;
    stacked-vlan-ranges {
        access-profile profile-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                mac-address;
                option-82 ( circuit-id | remote-id);
                radius-realm radius-realm-string;
                user-prefix user-prefix-string;
            }
        }
    }
    dynamic-profile profile-name {

```

```

        accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
        ranges (any | low-tag-high-tag), (any | low-tag-high-tag);
    }
}
vlan-ranges {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-name;
            mac-address;
            option-82;
            radius-realm radius-realm-string;
            user-prefix user-prefix-string;
        }
    }
    dynamic-profile profile-name {
        accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
        ranges (any | low-tag)–(any | high-tag);
    }
}
override tag vlan-tag dynamic-profile profile name;
}
encapsulation (ethernet-bridge | ethernet-vpls | extended-vlan-bridge |
    extended-vlan-vpls | flexible-ethernet-services | vlan-vpls);
ether-options {
    802.3ad {
        aex;
        (backup | primary);
        lacp {
            force-up;
            port-priority
        }
    }
}
asynchronous-notification;
(auto-negotiation | no-auto-negotiation);
ethernet-switch-profile {
    ethernet-policer-profile {
        input-priority-map {
            ieee802.1p premium [ values ];
        }
        output-priority-map {
            classifier {
                premium {
                    forwarding-class class-name {
                        loss-priority (high | low);
                    }
                }
            }
        }
    }
}
policer cos-policer-name {
    aggregate {
        bandwidth-limit bps;
    }
}

```

```

        burst-size-limit bytes;
    }
    premium {
        bandwidth-limit bps;
        burst-size-limit bytes;
    }
}
tag-protocol-id;
}
(mac-learn-enable | no-mac-learn-enable);
}
(flow-control | no-flow-control);
ignore-l3-incompletes;
link-mode (automatic | full-duplex | half-duplex);
(loopback | no-loopback);
keepalives <interval seconds> <down-count number> <up-count number>;
speed (1g | 10m | 100m | 10m-100m | auto-negotiation);
source-address-filter {
    mac-address;
}
source-filtering | no-source-filtering;
}
flexible-vlan-tagging;
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time (up milliseconds | down milliseconds);
interface-transmit-statistics;
(keepalives <down-count number> <interval seconds> <up-count number> |
 no-keepalives);
layer2-policer {
    apply-groups [ group-names ];
    apply-groups-except [ group-names ];
}
link-mode (automatic | full-duplex);
mac mac-address;
mtu bytes;
multi-chassis-protection peer-ip-address {
    interface interface-name;
}
native-vlan-id number;
no-gratuitous-arp-request;
optics-options {
    alarm low-light-alarm {
        (link-down | syslog);
    }
    warning low-light-warning {
        (link-down | syslog);
    }
}
wavelength nm;
}
passive-monitor-mode;
per-unit-scheduler;
speed (10m | 100m | 1g | auto | oc3 | oc12 | oc48);
stacked-vlan-tagging;
traceoptions {
    flag flag;
}

```

```
transmit-bucket {
    overflow discard;
    rate percentage;
    threshold bytes;
}
(traps | no-traps);
unidirectional;
vlan-tagging;
}

interface-name {
    unit logical-unit-number {
        disable;
        accept-source-mac {
            mac-address mac-address {
                policer {
                    input policer-name;
                    output policer-name;
                }
            }
        }
        account-layer2-overhead (Interface Level) {
            value;
            egress bytes;
            ingress bytes;
        }
        accounting-profile name;
        advisory-options {
            downstream-rate rate;
            upstream-rate rate;
        }
        arp-resp (restricted|unrestricted);
        bandwidth rate;
        clear-dont-fragment-bit;
        copy-tos-to-outer-ip-header;
        demux-destination family;
        encapsulation (vlan-bridge | vlan-vpls);
        epd-threshold cells plp1 cells;
        filter filter-name;
        inner-vlan-id-range start start-id end end-id;
        input-vlan-map {
            (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
            inner-tag-protocol-id tpid;
            inner-vlan-id number;
            tag-protocol-id tpid;
            vlan-id number;
        }
        interface-shared-with psdnumerical-index;
        layer2-policer {
            input-hierarchical-policer policer-name;
            input-policer policer-name;
            input-three-color policer-name;
            output-policer policer-name;
            output-three-color policer-name;
        }
    }
}
```



```

multi-chassis-protection peer-ip-address {
    interface interface-name;
}
native-inner-vlan-id number;
output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
peer-interface interface-name;
peer-unit unit-number;
plp-to-clp;
proxy-arp <restricted | unrestricted>;
rpm {
    (client | server);
    twamp-server;
}
swap-by-poppush;
vlan-id number;
vlan-id-list [ vlan-id vlan-id-vlan-id ];
vlan-id-range number-number;
vlan-tags (inner <tpid.>vlan-id | inner-list [ vlan-id vlan-id-vlan-id ] |
    inner-range <tpid.>vlan-id-vlan-id) outer <tpid.>vlan-id;
}

unit logical-unit-number {
    family ethernet-switching {
        filter {
            group filter-group-number;
            (input filter-name | input-list [ filter-names ]);
            (output filter-name | output-list [ filter-names ]);
            (inner-vlan-id-list [ vlan-ids ] | vlan-id number | vlan-id-list [ number
                number-number ]);
            interface-mode (access | trunk);
            policer {
                input policer-name;
                output policer-name;
            }
            vlan-rewrite {
                translate old-vlan-id new-vlan-id;
            }
            vlan {
                members [ all vlan-identifiers ];
            }
        }
    }
    family inet {
        filter {
            group filter-group-number;
            (input filter-name | input-list [ filter-names ]);
            (output filter-name | output-list [ filter-names ]);
        }
        input-hierarchical-policer policer-name;
        mac-validate (loose | strict);
        mtu bytes;
    }
}

```

```

no-neighbor-learn;
no-redirects;
policer {
    arp policer-template-name;
    input policer-name;
    output policer-name;
}
primary;
receive-options-packets;
receive-ttl-exceeded;
rpf-check {
    fail-filter filter-name;
    mode loose;
}
sampling {
    (input | output | input output);
}
simple-filter {
    input filter-name;
}
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
}
unnumbered-address interface-name <destination address>
    <destination-profile profile-name> <preferred-source-address address>;
}

family inet6 {
    address ipv6-address {
        destination destination-address;
        eui-64;
        ndp ipv6-address <l2-interface interface-name> <(mac mac-address |
            multicast-mac multicast-mac-address) <publish>>;
        preferred;
        primary;
        vrrp-inet6-group group-number {
            (accept-data | no-accept-data);
            fast-interval milliseconds;
            inet6-advertise-interval seconds;
            (no-preempt; | ... the following preempt statement ...)
            preempt {
                hold-time seconds;
            }
            priority number;
            track {
                interface interface-name {
                    bandwidth-threshold bits-per-second priority-cost priority;
                    priority-cost priority;
                }
                priority-hold-time seconds;
                route ip-address-prefix/prefix-length routing-instance instance-name
                    priority-cost priority;
            }
        }
        virtual-inet6-address [ addresses ];
    }
}

```

```

        virtual-link-local-address ipv6-address;
        vrrp-inherit-from {
            active-group group-number;
            active-interface interface-name;
        }
    }
    (dad-disable | no-dad-disable);
    filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
    input-hierarchical-policer policer-name;
    mtu bytes;
    nd6-stale-time seconds;
    no-neighbor-learn;
    policer {
        input policer-name;
        output policer-name;
    }
    rpf-check {
        fail-filter filter-name;
        mode loose;
    }
    sampling {
        (input | output | input output);
    }
    unnumbered-address interface-name preferred-source-address address;
}

family iso {
    address iso-address;
    mtu bytes;
}

family mlfrr-end-to-end {
    bundle logical-interface-name;
}

family mpls {
    filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
    input-hierarchical-policer policer-name;
    maximum-labels maximum-labels;
    mtu bytes;
    policer {
        input policer-name;
        output policer-name;
    }
}

```

```
    }

    family vpls {
      core-facing;
      filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
      }
      policer {
        input policer-name;
        output policer-name;
      }
    }
  }
}
```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

filter (Applying to an Interface)

Syntax	<pre>filter { input <i>filter-name</i>; output <i>filter-name</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Apply a filter to an interface. You can also use filters for encrypted traffic. When you configure filters, you can configure the family inet , inet6 , mpls , or vpls only.
Options	<p>input <i>filter-name</i>—Name of one filter to evaluate when packets are received on the interface.</p> <p>output <i>filter-name</i>—Name of one filter to evaluate when packets are transmitted on the interface.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • simple-filter on page 1440 • <i>Applying Firewall Filter Tricolor Marking Policers to Interfaces</i> • <i>Example: Classifying Packets Based on Their Destination Address</i> • Example: Configuring and Verifying a Complex Multifield Filter on page 1394 • <i>Example: Writing Different DSCP and EXP Values in MPLS-Tagged IP Packets</i> • Example: Configuring a Simple Filter on page 1399 • Example: Configuring a Logical Bandwidth Policer on page 1400 • Example: Two-Color Policers and Shaping Rate Changes on page 1401

simple-filter (Applying to an Interface)

Syntax	<code>simple-filter { input <i>filter-name</i>; }</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	Apply a simple filter to an interface. You can apply simple filters to the family inet only, and only in the input direction.
Options	input <i>filter-name</i> —Name of one filter to evaluate when packets are received on the interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Multifield Classifiers on page 1390• filter on page 1439

Tricolor Marking Policers

- [Overview on page 1440](#)
- [Configuration on page 1447](#)

Overview

- [Tricolor Marking Policers on page 1440](#)

Tricolor Marking Policers

- [Traffic Policing Overview on page 1440](#)
- [Tricolor Marking Architecture on page 1444](#)
- [Tricolor Marking Limitations on page 1445](#)
- [Policer Support for Aggregated Ethernet Bundle Overview on page 1446](#)

Traffic Policing Overview

This topic covers the following information:

- [Congestion Management for IP Traffic Flows on page 1441](#)
- [Traffic Limits on page 1441](#)
- [Traffic Color Marking on page 1442](#)

- [Forwarding Classes and PLP Levels on page 1443](#)
- [Policer Application to Traffic on page 1444](#)

Congestion Management for IP Traffic Flows

Traffic policing, also known *rate limiting*, is an essential component of network access security that is designed to thwart denial-of-service (DoS) attacks. Traffic policing enables you to control the maximum rate of IP traffic sent or received on an interface and also to partition network traffic into multiple priority levels, also known as *classes of service*. A policer defines a set of traffic rate limits and sets consequences for traffic that does not conform to the configured limits. Packets in a traffic flow that does not conform to traffic limits are either discarded or marked with a different forwarding class or packet loss priority (PLP) level.

With the exception of policers configured to rate-limit aggregate traffic (all protocol families and logical interfaces configured on a physical interface), you can apply a policer to all IP packets in a Layer 2 or Layer 3 traffic flow at a logical interface.

With the exception of policers configured to rate-limit based on physical interface media rate, you can apply a policer to specific IP packets in a Layer 3 traffic flow at a logical interface by using a stateless firewall filter.

You can apply a policer to inbound or outbound interface traffic. Policers applied to inbound traffic help to conserve resources by dropping traffic that does not need to be routed through a network. Dropping inbound traffic also helps to thwart denial-of-service (DoS) attacks. Policers applied to outbound traffic control the bandwidth used.



NOTE: Traffic policers are instantiated on a per-PIC basis. Traffic policing does not work when the traffic for one local policy decision function (L-PDF) subscriber is distributed over multiple Multiservices PICs in an AMS group.

Traffic Limits

Junos[®] OS policers use the *token-bucket* algorithm to enforce a limit on average transmit or receive rate of IP traffic at an interface while allowing bursts of traffic up to a maximum value based on the overall traffic load. The token-bucket algorithm offers more flexibility than the *leaky-bucket* algorithm in that you can allow a specified amount of bursting before starting to discard packets or apply a penalty to packet output-queuing priority or packet drop priority.


In the token-bucket model, the bucket represents the policing function. Tokens are added to the bucket at a fixed rate, but only up to the specified depth of the bucket. Each token represents a “credit” for some number of bits, and tokens in the bucket are “cached in” for the ability to transmit or receive traffic at the interface. When sufficient tokens are present in the bucket, a traffic flow continues unrestricted. Otherwise, packets might be dropped or else re-marked with a lower forwarding class, a higher packet loss priority (PLP) level, or both.

- The rate at which tokens are added to the bucket represents the highest average transmit or receive rate in bits per second allowed for a given service level. You specify

this highest average traffic rate as the *bandwidth limit* of the policer. If the traffic arrival rate is so high that at some point insufficient tokens are present in the bucket, then the traffic flow is no longer conforming to the traffic limit.

- The depth of the bucket in bytes controls the amount of back-to-back bursting allowed. You specify this factor as the *burst-size limit* of the policer. This second limit affects the average transmit or receive rate by limiting the number of bytes permitted in a transmission burst for a given interval of time. Bursts exceeding the current burst-size limit are dropped until there are sufficient tokens available to permit the burst to proceed.

Figure 9: Network Traffic and Burst Rates



As shown in the figure above, a UPC bar code is a good facsimile of what traffic looks like on the line; an interface is either transmitting (bursting at full rate) or it is not. The black lines represent periods of data transmission and the white space represents periods of silence when the token bucket can replenish.

Depending on the type of policer used, packets in a policed traffic flow that surpasses the defined limits might be implicitly set to a higher PLP level, assigned to a configured forwarding class or set to a configured PLP level (or both), or simply discarded. If packets encounter downstream congestion, packets with a **low** PLP level are less likely to be discarded than those with a **medium-low**, **medium-high**, or **high** PLP level.

Traffic Color Marking

Based on the particular set of traffic limits configured, a policer identifies a traffic flow as belonging to one of either two or three categories that are similar to the colors of a traffic light used to control automobile traffic.

A *two-color-marking* policer categorizes traffic as either conforming to the traffic limits (green) or violating the traffic limits (red):

- **Green**—Two-color-marking policers implicitly set the packets in a green flow to the low PLP level, and you cannot configure any policer actions for conforming traffic.
- **Red**—Two-color-marking policers do not perform any implicit actions on packets in a red flow. Instead, those packets are handled according to the actions specified in the policer configuration. You can configure a two-color-marking policer to simply discard packets if the traffic flow is red. Alternatively, you can configure a two-color-marking policer to handle the packets in a red flow by setting the PLP level to either **low** or **high**, assigning the packets to any forwarding class already configured, or both.

On MX Series, M120, and M320 routers and M7i and M10i routers with the Enhanced CFEB (CFEB-E) and EX Series switches only, you can specify two additional PLP levels for packets in a red flow: **medium-low** or **medium-high**.

Three-color-marking policers categorize traffic as conforming to the traffic limits (green), violating the traffic limits (red), or exceeding the traffic limits but within an allowed range (yellow):

- **Green**—Like two-color-marking policers, three-color-marking policers implicitly set the packets in a green flow to the low PLP level, and you cannot configure any policer actions for conforming traffic.
- **Yellow**—Unlike two-color-marking policers, three-color-marking policers categorize a second type of nonconforming traffic: yellow.

Single-rate three-color policing categorizes as yellow traffic that exceeds the traffic limits while conforming to a second defined burst-size limit. Two-rate three-color policing categorizes as yellow traffic that exceeds the traffic limits while conforming to both a second defined burst-size limit and a second defined bandwidth limit.

Three-color-marking policers implicitly set the packets in a yellow flow to the medium-high PLP level so that the packets incur a less severe penalty than those in a red flow. You cannot configure any policer actions for yellow traffic.

- **Red**—Unlike two-color-marking policers, three-color-marking policers implicitly set the packets in a red flow to the high PLP level, which is the highest PLP value. You can also configure a three-color-marking policer to discard the packets in a red flow instead of forwarding them with a high PLP setting.

Two-color-marking policers allows bursts of traffic for short periods, whereas three-color-marking policers allow more sustained bursts of traffic.

Forwarding Classes and PLP Levels

A packet's forwarding class assignment and PLP level are used by the Junos OS class of service (CoS) features. The Junos CoS features include a set of mechanisms that you can use to provide differentiated services when best-effort traffic delivery is insufficient. For router (and switch) interfaces that carry IPv4, IPv6, and MPLS traffic, you can configure CoS features to take in a single flow of traffic entering at the edge of your network and provide different levels of service across the network—internal forwarding and scheduling (queuing) for output—based on the forwarding class assignments and PLP levels of the individual packets.



NOTE: Forwarding-class or loss-priority assignments performed by a policer or a stateless firewall filter override any such assignments performed on the ingress by the CoS default IP precedence classification at all logical interfaces or by any configured behavior aggregate (BA) classifier that is explicitly mapped to a logical interface.

Based on CoS configurations, packets of a given forwarding class are transmitted through a specific output queue, and each output queue is associated with a transmission service level defined in a *scheduler*.

Based on other CoS configurations, when packets in an output queue encounter congestion, packets with higher loss-priority values are more likely to be dropped by the random early detection (RED) algorithm. Packet loss priority values affect the scheduling of a packet without affecting the packet's relative ordering within the traffic flow.

Policer Application to Traffic

After you have defined and named a policer, it is stored as a template. You can later use the same policer name to provide the same policer configuration each time you want to use it. This eliminates the need to define the same policer values more than once.

You can apply a policer to a traffic flow in either of two ways:

- You can configure a standard stateless firewall filter that specifies the **policer *policer-name*** nonterminating action or the **three-color-policer (single-rate | two-rate) *policer-name*** nonterminating action. When you apply the standard filter to the input or output at a logical interface, the policer is applied to all packets of the filter-specific protocol family that match the conditions specified in the filter configuration.

With this method of applying a policer, you can define specific classes of traffic on an interface and apply traffic rate-limiting to each class.

- You can apply a policer directly to an interface so that traffic rate-limiting applies to all traffic on that interface, regardless of protocol family or any match conditions.

You can configure policers at the queue, logical interface, or Layer 2 (MAC) level. Only a single policer is applied to a packet at the egress queue, and the search for policers occurs in this order:

- Queue level
- Logical interface level
- Layer 2 (MAC) level

Related Documentation


- [Stateless Firewall Filter Overview on page 4463.](#)
- [Traffic Policer Types](#)
- [Order of Policer and Firewall Filter Operations on page 4810](#)
- [Packet Flow Through the CoS Process Overview on page 1273](#)

Tricolor Marking Architecture

Policers provide two functions: metering and marking.

The policer meters each packet and passes the packet and the metering result to the marker, as shown in [Figure 10 on page 1445](#).

Figure 10: Flow of Tricolor Marking Policer Operation



9017049

The meter operates in two modes. In the color-blind mode, the meter treats the packet stream as uncolored. Any preset loss priorities are ignored. In the color-aware mode, the meter inspects the packet loss priority (PLP) field, which has been set by an upstream device as PLP high, medium-high, medium-low, or low; in other words, the PLP field has already been set by a behavior aggregate (BA) or multifield classifier. The marker changes the PLP of each incoming IP packet according to the results of the meter. For more information, see [“Configuring Two-Rate Tricolor Marking” on page 1452](#).

This chapter emphasizes configuration and use of TCM policers. For more information about configuring and using two-color policers (“policers”), see the *Traffic Policers Configuration Guide*.

Single-rate TCM is so called because traffic is policed according to one rate—the CBR—and two burst sizes: the CBS and EBS. The CBS specifies the usual burst size in bytes and the EBS specifies the maximum burst size in bytes for packets that are admitted to the network. The EBS is greater than or equal to the CBS, and neither can be 0. As each packet enters the network, its bytes are counted. Packets that do not exceed the CBS are marked low PLP. Packets that exceed the CBS but are below the EBS are marked medium-high PLP. Packets that exceed the EBS are marked high PLP.

Two-rate TCM is so called because traffic is policed according to two rates: the CIR and the PIR. The PIR is greater than or equal to the CIR. The CIR specifies the average rate at which bits are admitted to the network and the PIR specifies the maximum rate at which bits are admitted to the network. As each packet enters the network, its bits are counted. Bits in packets that do not exceed the CIR have their packets marked low PLP. Bits in packets that exceed the CIR but are below the PIR have their packets marked medium-high PLP. Bits in packets that exceed the PIR have their packets marked high PLP.

For information about how to use marking policers with BA and multifield classifiers, see [“Using BA Classifiers to Set PLP” on page 1460](#) and [“Using Multifield Classifiers to Set PLP” on page 1461](#).

Tricolor Marking Limitations

Tricolor Marking (TCM) has some limitations that must be kept in mind during configuration and operation.

The following limitations apply to TCM:

- When you enable TCM on a 10-port Gigabit Ethernet PIC or a 10-Gigabit Ethernet PIC, for queues 6 and 7 only, the output of the **show interfaces queue *interface-name*** command does not display the number of queued bytes and packets, or the number of bytes and packets dropped due to RED. If you do not configure tricolor marking on the interface, these statistics are available for all queues.
- When you enable TCM, Transmission Control Protocol (TCP)-based configurations for drop profiles are rejected. In other words, you cannot include the **protocol** statement at the **[edit class-of-service schedulers *scheduler-name* drop-profile-map]** hierarchy level. The result is that drop profiles are applied to packets with the specified PLP and any protocol type.
- On Gigabit Ethernet IQ PICs, for IEEE 802.1 rewrite rules, only two loss priorities are supported. Exiting packets with medium-high loss priority are treated as high, and packets with medium-low loss priority are treated as low. In other words rewrite rules corresponding to high and low apply instead of those corresponding to medium-high and medium-low. For IQ PICs, you can only configure one IEEE 802.1 rewrite rule on a physical port. All logical ports (units) on that physical port should apply the same IEEE 802.1 rewrite rule.
- When some PICs with Frame Relay encapsulation mark a packet with high loss priority, the packet is treated as having medium-high loss priority on M320 Multiservice Edge Routers and T Series Core Routers with Enhanced II FPCs and T640 Core Routers with Enhanced Scaling FPC4.
- In a single firewall filter term, you cannot configure both the **loss-priority** action modifier and the **three-color-policer** action modifier. These statements are mutually exclusive.

Policer Support for Aggregated Ethernet Bundle Overview

Aggregated interfaces support single-rate policers, three-color marking policers, two-rate three-color marking policers, hierarchical policers, and percentage-based policers. By default, policer bandwidth and burst-size applied on aggregated bundles is not matched to the user-configured bandwidth and burst-size.

You can configure interface-specific policers applied on an aggregated Ethernet bundle to match the effective bandwidth and burst-size to user-configured values. The **shared-bandwidth-policer** statement is required to achieve this match behavior.

This capability applies to all interface-specific policers of the following types: single-rate policers, single-rate three-color marking policers, two-rate three-color marking policers, and hierarchical policers. Percentage-based policers match the bandwidth to the user-configured values by default, and do not require shared-bandwidth-policer configuration. The **shared-bandwidth-policer** statement causes a split in burst-size for percentage-based policers.



NOTE: This feature is supported on the following platforms: T Series routers, M120, M10i, M7i (CFEB-E only), M320 (SFPC only), MX240, MX480, and MX960 (DPC only), and EX Series switches.

The following usage scenarios are supported:

- Interface policers used by the following configuration:

[edit] interfaces (aeX | asX) unit *unit-num* family *family* policer [input | output | arp]

- Policers and three-color policers (both single-rate three-color marking and two-rate three-color marking) used inside interface-specific filters; that is, filters that have an interface-specific keyword and are used by the following configuration:

[edit] interfaces (aeX | asX) unit *unit-num* family *family* filter [input | output]

- Common-edge service filters, which are derived from CLI-configured filters and thus inherit interface-specific properties. All policers and three-color policers used by these filters are also affected.

The following usage scenarios are not supported:

- Policers and three-color policers used inside filters that are not interface specific; such a filter is meant to be shared across multiple interfaces.
- Any implicit policers or policers that are part of implicit filters; for example, the default ARP policer applied to an aggregate Ethernet interface. Such a policer is meant to be shared across multiple interfaces.
- Prefix-specific action policers.

To configure this feature, include the **shared-bandwidth-policer** statement at the following hierarchy levels: **[edit firewall policer *policer-name*]**, **[edit firewall three-color-policer *policer-name*]**, or **[edit firewall hierarchical-policer *policer-name*]**.

Related Documentation

- [shared-bandwidth-policer on page 1504](#)

Configuration

- [Configuration Tasks for Tricolor Marking Policers on page 1447](#)
- [Configuration Tasks for Packet Loss Priority on page 1460](#)
- [Configuration Statements for Tricolor Marking Policers on page 1463](#)

Configuration Tasks for Tricolor Marking Policers

- [Configuring Tricolor Marking on page 1447](#)
- [Configuring Single-Rate Tricolor Marking on page 1449](#)
- [Configuring Two-Rate Tricolor Marking on page 1452](#)
- [Enabling Tricolor Marking on page 1455](#)
- [Configuring Tricolor Marking Policers on page 1455](#)
- [Applying Tricolor Marking Policers to Firewall Filters on page 1456](#)
- [Applying Firewall Filter Tricolor Marking Policers to Interfaces on page 1458](#)
- [Applying Layer 2 Policers to Gigabit Ethernet Interfaces on page 1459](#)

Configuring Tricolor Marking

You configure marking policers by defining the policer and multiple levels of PLP for classifiers, rewrite rules, random early detection (RED) drop profiles, and firewall filters.

To configure marking policers, include the following statements at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
tri-color;
classifiers {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) classifier-name {
    import classifier-name | default;
    forwarding-class class-name {
      loss-priority (low | medium-low | medium-high | high) code-points [ aliases ]
        [ bit-patterns ];
    }
  }
}
rewrite-rules {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) rewrite-name {
    import (rewrite-name | default);
    forwarding-class class-name {
      loss-priority (low | medium-low | medium-high | high) code-point (aliases |
        bit-patterns;
    }
  }
}
schedulers {
  scheduler-name {
    drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol
      any drop-profile profile-name;
  }
}

[edit firewall]
policer name {
  then loss-priority (low | medium-low | medium-high | high);
}
three-color-policer policer-name {
  action {
    loss-priority high then discard; # Only for IQ2 PICs
  }
  logical-interface-policer;
  single-rate {
    (color-aware | color-blind);
    committed-information-rate bps;
    committed-burst-size bytes;
    excess-burst-size bytes;
  }
  two-rate {
    (color-aware | color-blind);
    committed-information-rate bps;
    committed-burst-size bytes;
    peak-information-rate bps;
    peak-burst-size bytes;
  }
}
filter filter-name {
  <family family> {
```

```

term rule-name {
  then {
    three-color-policer (single-rate | two-rate) policer-name;
  }
}

```

Configuring Single-Rate Tricolor Marking

With TCM, you can configure traffic policing according to two separate modes—color-blind and color-aware. In color-blind mode, the current PLP value is ignored. In color-aware mode, the current PLP values are considered by the policer and can only be increased.

- [Configuring Color-Blind Mode for Single-Rate Tricolor Marking on page 1449](#)
- [Configuring Color-Aware Mode for Single-Rate Tricolor Marking on page 1450](#)

Configuring Color-Blind Mode for Single-Rate Tricolor Marking

All packets are evaluated by the CBS. If a packet exceeds the CBS, it is evaluated by the EBS. In color-blind mode, the policer supports three loss priorities only: low, medium-high, and high.

In color-blind mode, packets that exceed the CBS but are below the EBS are marked yellow (medium-high). Packets that exceed the EBS are marked red (high), as shown in [Table 102 on page 1449](#).

Table 102: Color-Blind Mode TCM Color-to-PLP Mapping

Color	PLP	Meaning
Green	low	Packet does not exceed the CBS.
Yellow	medium-high	Packet exceeds the CBS but does not exceed the EBS.
Red	high	Packet exceeds the EBS.

If you are using color-blind mode and you wish to configure an output policer that marks packets to have medium-low loss priority, you must configure a policer at the **[edit firewall policer policer-name]** hierarchy level. For example:

```

firewall {
  policer 4PLP {
    if-exceeding {
      bandwidth-limit 40k;
      burst-size-limit 4k;
    }
    then loss-priority medium-low;
  }
}

```

Apply this policer at one or both of the following hierarchy levels:

- **[edit firewall family family filter filter-name term rule-name then policer policer-name]**

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

Configuring Color-Aware Mode for Single-Rate Tricolor Marking

In color-aware mode, the metering treatment the packet receives depends on its classification. Metering can increase a packet's preassigned PLP, but cannot decrease it, as shown in [Table 103 on page 1450](#).

Table 103: Color-Aware Mode TCM PLP Mapping

Incoming PLP	Packet Metered Against	Possible Cases	Outgoing PLP
low	CBS and EBS	Packet does not exceed the CBS.	low
		Packet exceeds the CBS but not the EBS.	medium-high
		Packet exceeds the EBS.	high
medium-low	EBS only	Packet does not exceed the CBS.	medium-low
		Packet does not exceed the EBS.	medium-low
		Packet exceeds the EBS.	high
medium-high	EBS only	Packet does not exceed the CBS.	medium-high
		Packet does not exceed the EBS.	medium-high
		Packet exceeds the EBS.	high
high	Not metered by the policer.	All cases.	high

The following sections describe single-rate color-aware PLP mapping in more detail.

Effect on Low PLP of Single-Rate Policer

Packets belonging to the green class have already been marked by a classifier with low PLP. The marking policer can leave the packet's PLP unchanged or increase the PLP to medium-high or high. Therefore, these packets are metered against both the CBS and the EBS.

For example, if a BA or multifield classifier marks a packet with low PLP according to the type-of-service (ToS) bits in the IP header, and the two-rate TCM policer is in color-aware mode, the output loss priority is as follows:

- If the rate of traffic flow is less than the CBS, packets remain marked as low PLP.
- If the rate of traffic flow is greater than the CBS but less than the EBS, some of the packets are marked as medium-high PLP, and some of the packets remain marked as low PLP.
- If the rate of traffic flow is greater than the EBS, some of the packets are marked as high PLP, and some of the packets remain marked as low PLP.

Effect on Medium-Low PLP of Single-Rate Policer

Packets belonging to the yellow class have already been marked by a classifier with medium-low or medium-high PLP. The marking policer can leave the packet's PLP unchanged or increase the PLP to high. Therefore, these packets are metered against the EBS only.

For example, if a BA or multifield classifier marks a packet with medium-low PLP according to the ToS bits in the IP header, and the two-rate TCM policer is in color-aware mode, the output loss priority is as follows:

- If the rate of traffic flow is less than the CBS, packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the CBS but less than the EBS, packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the EBS, some of the packets are marked as high PLP, and some of the packets remain marked as medium-low PLP.

Effect on Medium-High PLP of Single-Rate Policer

Packets belonging to the yellow class have already been marked by a classifier with medium-low or medium-high PLP. The marking policer can leave the packet's PLP unchanged or increase the PLP to high. Therefore, these packets are metered against the EBS only.

For example, if a BA or multifield classifier marks a packet with medium-high PLP according to the ToS bits in the IP header, and the two-rate TCM policer is in color-aware mode, the output loss priority is as follows:

- If the rate of traffic flow is less than the CBS, packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the CBS but less than the EBS, packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the EBS, some of the packets are marked as high PLP, and some of the packets remain marked as medium-high PLP.

Effect on High PLP of Single-Rate Policer

Packets belonging to the red class have already been marked by a classifier with high PLP. The marking policer can only leave the packet's PLP unchanged. Therefore, these packets are not metered against the CBS or the EBS and all the packets remain marked as high PLP.

Configuring Two-Rate Tricolor Marking

With TCM, you can configure traffic policing according to two separate modes—color-blind and color-aware. In color-blind mode, the current PLP value is ignored. In color-aware mode, the current PLP values are considered by the policer and can only be increased.

- [Configuring Color-Blind Mode for Two-Rate Tricolor Marking on page 1452](#)
- [Configuring Color-Aware Mode for Two-Rate Tricolor Marking on page 1453](#)

Configuring Color-Blind Mode for Two-Rate Tricolor Marking

All packets are evaluated by the CIR. If a packet exceeds the CIR, it is evaluated by the PIR. In color-blind mode, the policer supports three loss priorities only: low, medium-high, and high.

In color-blind mode, packets that exceed the CIR but are below the PIR are marked yellow (medium-high). Packets that exceed the PIR are marked red (high), as shown in [Table 104 on page 1452](#).

Table 104: Color-Blind Mode TCM Color-to-PLP Mapping

Color	PLP	Meaning
Green	low	Packet does not exceed the CIR.
Yellow	medium-high	Packet exceeds the CIR but does not exceed the PIR.
Red	high	Packet exceeds the PIR.

If you are using color-blind mode and you wish to configure an output policer that marks packets to have medium-low loss priority, you must configure a policer at the **[edit firewall policer *policer-name*]** hierarchy level. For example:

```
firewall {
  policer 4PLP {
    if-exceeding {
      bandwidth-limit 40k;
      burst-size-limit 4k;
    }
    then loss-priority medium-low;
  }
}
```

Apply this policer at one or both of the following hierarchy levels:

- **[edit firewall family *family* filter *filter-name* term *rule-name* then policer *policer-name*]**

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

Configuring Color-Aware Mode for Two-Rate Tricolor Marking

In color-aware mode, the metering treatment the packet receives depends on its classification. Metering can increase a packet's preassigned PLP, but cannot decrease it, as shown in [Table 105 on page 1453](#).

Table 105: Color-Aware Mode TCM Mapping

Incoming PLP	Packet Metered Against	Possible Cases	Outgoing PLP
low	CIR and PIR	Packet does not exceed the CIR.	low
		Packet exceeds the CIR but not the PIR.	medium-high
		Packet exceeds the PIR.	high
medium-low	PIR only	Packet does not exceed the CIR.	medium-low
		Packet does not exceed the PIR.	medium-low
		Packet exceeds the PIR.	high
medium-high	PIR only	Packet does not exceed the CIR.	medium-high
		Packet does not exceed the PIR.	medium-high
		Packet exceeds the PIR.	high
high	Not metered by the policer.	All cases.	high

The following sections describe color-aware two-rate PLP mapping in more detail.

Effect on Low PLP of Two-Rate Policer

Packets belonging to the green class have already been marked by a classifier with low PLP. The marking policer can leave the packet's PLP unchanged or increase the PLP to medium-high or high. Therefore, these packets are metered against both the CIR and the PIR.

For example, if a BA or multifield classifier marks a packet with low PLP according to the ToS bits in the IP header, and the two-rate TCM policer is in color-aware mode, the output loss priority is as follows:

- If the rate of traffic flow is less than the CIR, packets remain marked as low PLP.
- If the rate of traffic flow is greater than the CIR but less than the PIR, some of the packets are marked as medium-high PLP, and some of the packets remain marked as low PLP.

- If the rate of traffic flow is greater than the PIR, some of the packets are marked as high PLP, and some of the packets remain marked as low PLP.

Effect on Medium-Low PLP of Two-Rate Policer

Packets belonging to the yellow class have already been marked by a classifier with medium-low or medium-high PLP. The marking policer can leave the packet's PLP unchanged or increase the PLP to high. Therefore, these packets are metered against the PIR only.

For example, if a BA or multifield classifier marks a packet with medium-low PLP according to the ToS bits in the IP header, and the two-rate TCM policer is in color-aware mode, the output loss priority is as follows:

- If the rate of traffic flow is less than the CIR, packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the CIR/CBS but less than the PIR, packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the PIR, some of the packets are marked as high PLP, and some of the packets remain marked as medium-low PLP.

Effect on Medium-High PLP of Two-Rate Policer

Packets belonging to the yellow class have already been marked by a classifier with medium-low or medium-high PLP. The marking policer can leave the packet's PLP unchanged or increase the PLP to high. Therefore, these packets are metered against the PIR only.

For example, if a BA or multifield classifier marks a packet with medium-high PLP according to the ToS bits in the IP header, and the two-rate TCM policer is in color-aware mode, the output loss priority is as follows:

- If the rate of traffic flow is less than the CIR, packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the CIR but less than the PIR, packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the PIR, some of the packets are marked as high PLP, and some of the packets remain marked as medium-high PLP.

Effect on High PLP of Two-Rate Policer

Packets belonging to the red class have already been marked by a classifier with high PLP. The marking policer can only leave the packet's PLP unchanged. Therefore, these packets are not metered against the CIR or the PIR and all the packets remain marked as high PLP.

Enabling Tricolor Marking

By default, TCM is enabled on M120, MX Series, and T4000 routers, and EX Series switches. To enable TCM on other routers, include the **tri-color** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
tri-color;
```

This statement is necessary on the following routers:

- M320 and T Series routers with Enhanced II FPCs
- T640 routers with Enhanced Scaling FPC4s

If you do not include this statement in the configuration on platforms that require it, you cannot configure medium-low or medium-high PLP for classifiers, rewrite rules, drop profiles, or firewall filters.

Configuring Tricolor Marking Policers

A tricolor marking policer polices traffic on the basis of metering rates, including the CIR, the PIR, their associated burst sizes, and any policing actions configured for the traffic. To configure a tricolor marking policer, include the following statements at the **[edit firewall]** hierarchy level:

```
[edit firewall]
three-color-policer name {
  action {
    loss-priority high then discard; # Only for IQ2 PICs
  }
  logical-interface-policer;
  single-rate {
    (color-aware | color-blind);
    committed-information-rate bps;
    committed-burst-size bytes;
    excess-burst-size bytes;
  }
  two-rate {
    (color-aware | color-blind);
    committed-information-rate bps;
    committed-burst-size bytes;
    peak-information-rate bps;
    peak-burst-size bytes;
  }
}
```

You can configure a tricolor policer to discard high loss priority traffic on a logical interface in the ingress or egress direction. To configure a policer on a logical interface using tricolor marking policing to discard high loss priority traffic, include the **logical-interface-policer** statement and **action** statement.

In all cases, the range of allowable bits-per-second or byte values is 1500 to 100,000,000,000. You can specify the values for bps and bytes either as complete decimal numbers or as decimal numbers followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

The color-aware policer implicitly marks packets into four loss priority categories:

- Low
- Medium-low
- Medium-high
- High

The color-blind policer implicitly marks packets into three loss priority categories:

- Low
- Medium-high
- High

Table 106 on page 1456 describes all the configurable TCM statements.

Table 106: Tricolor Marking Policer Statements

Statement	Meaning	Configurable Values
single-rate	Marking is based on the CIR, CBS, and EBS.	—
two-rate	Marking is based on the CIR, PIR, and rated burst sizes.	—
color-aware	Metering depends on the packet's preclassification. Metering can increase a packet's assigned PLP, but cannot decrease it.	—
color-blind	All packets are evaluated by the CIR or CBS. If a packet exceeds the CIR or CBS, it is evaluated by the PIR or EBS.	—
committed-information-rate	Guaranteed bandwidth under normal line conditions and the average rate up to which packets are marked green.	1500 through 100,000,000,000 bps
committed-burst-size	Maximum number of bytes allowed for incoming packets to burst above the CIR, but still be marked green.	1500 through 100,000,000,000 bytes
excess-burst-size	Maximum number of bytes allowed for incoming packets to burst above the CIR, but still be marked yellow.	1500 through 100,000,000,000 bytes
peak-information-rate	Maximum achievable rate. Packets that exceed the CIR but are below the PIR are marked yellow. Packets that exceed the PIR are marked red.	1500 through 100,000,000,000 bps
peak-burst-size	Maximum number of bytes allowed for incoming packets to burst above the PIR, but still be marked yellow.	1500 through 100,000,000,000 bytes

Applying Tricolor Marking Policers to Firewall Filters

To rate-limit traffic by applying a tricolor marking policer to a firewall filter, include the **three-color-policer** statement:

```
three-color-policer {
```

```
(single-rate | two-rate) policer-name;
}
```

You can include this statement at the following hierarchy levels:

- [edit firewall family *family* filter *filter-name* term *rule-name* then]
- [edit firewall filter *filter-name* term *rule-name* then]

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

You must identify the referenced policer as a **single-rate** or **two-rate** policer, and this statement must match the configured TCM policer. Otherwise, an error message appears in the configuration listing.

For example, if you configure **srTCM** as a single-rate TCM policer and try to apply it as a two-rate policer, the following message appears:

```
[edit firewall]
user@host# show three-color-policer srTCM
single-rate {
  color-aware;
  ...
}
user@host# show filter TESTER
term A {
  then {
    three-color-policer {
      ##
      ## Warning: Referenced two-rate policer does not exist
      ##
      two-rate srTCM;
    }
  }
}
```

Example: Applying a Two-Rate Tricolor Marking Policer to a Firewall Filter

Apply the **trtcm1-cb** policer to a firewall filter:

```
firewall {
  three-color-policer trtcm1-cb { # Configure the trtcm1-cb policer.
    two-rate {
      color-blind;
      committed-information-rate 1048576;
      committed-burst-size 65536;
      peak-information-rate 10485760;
      peak-burst-size 131072;
    }
  }
  filter fil { # Configure the fil firewall filter, applying the trtcm1-cb policer.
    term default {
      then {
        three-color-policer {
          two-rate trtcm1-cb;
        }
      }
    }
  }
}
```

```
}
```

**Related
Documentation**

- *Firewall Filters Configuration Guide*

Applying Firewall Filter Tricolor Marking Policers to Interfaces

To apply a tricolor marking policer to an interface, you must reference the filter name in the interface configuration. To do this, include the **filter** statement:

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

You can include these statements at the following hierarchy levels:

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

The filter name that you reference should have an attached tricolor marking policer, as shown in [“Applying Tricolor Marking Policers to Firewall Filters” on page 1456](#).

Example: Applying a Single-Rate Tricolor Marking Policer to an Interface

Apply the **trtcm1-cb** policer to an interface:

```
firewall {  
    three-color-policer srtcm1 { # Configure the srtcm1-cb policer.  
        single-rate {  
            color-blind;  
            committed-information-rate 1048576;  
            committed-burst-size 65536;  
            excess-burst-size 131072;  
        }  
    }  
    filter fil { # Configure the fil firewall filter, applying the srtcm1-cb policer.  
        term default {  
            then {  
                three-color-policer {  
                    single-rate srtcm1-cb; # The TCM policer must be single-rate.  
                }  
            }  
        }  
    }  
    interfaces { # Configure the interface, which attaches the fil firewall filter.  
        ge-1/0/0 {  
            unit 0 {  
                family inet {  
                    filter {  
                        input fil;  
                    }  
                }  
            }  
        }  
    }  
}
```


Applying Layer 2 Policers to Gigabit Ethernet Interfaces

To rate-limit traffic by applying a policer to a Gigabit Ethernet interface (or a 10-Gigabit Ethernet interface [xe-fpc/pic/port]), include the **layer2-policer** statement with the direction, type, and name of the policer:

```
[edit interfaces ge-fpc/pic/port unit 0]
layer2-policer {
  input-policer policer-name;
  input-three-color policer-name;
  output-policer policer-name;
  output-three-color policer-name;
}
```

The direction (input or output) and type (policer or three-color) are combined into one statement and the policer named must be properly configured.

One input or output policer of either type can be configured on the interface.

Examples: Applying Layer 2 Policers to a Gigabit Ethernet Interface

Apply color-blind and color-aware two-rate TCM policers as input and output policers to a Gigabit Ethernet interface:

```
ge-1/0/0 {
  unit 0
  layer2-policer {
    input-three-color trTCM1-cb; # Apply the trTCM1-color-blind policer.
    output-three-color trTCM1-ca; # Apply the trTCM1-color-aware policer.
  }
}
```

Apply two-level and color-blind single-rate TCM policers as input and output policers to a Gigabit Ethernet interface:

```
ge-1/0/0 {
  unit 1
  layer2-policer {
    input-policer two-color-policer; # Apply a two-color policer.
    output-three-color srTCM2-cb; # Apply the srTCM1-color-blind policer.
  }
}
```

Apply a color-aware single-rate TCM policer as output policer on a Gigabit Ethernet interface:

```
ge-1/0/0 {
  unit 2
  layer2-policer {
    output-three-color srTCM3-ca { # Apply the srTCM3-color-aware policer.
  }
}
```

Configuration Tasks for Packet Loss Priority

- [Using BA Classifiers to Set PLP on page 1460](#)
- [Using Multifield Classifiers to Set PLP on page 1461](#)
- [Configuring PLP for Drop-Profile Maps on page 1462](#)
- [Configuring Rewrite Rules Based on PLP on page 1462](#)

Using BA Classifiers to Set PLP

Behavior aggregate (BA) classifiers take action on incoming packets. When TCM is enabled, Juniper Networks M320 Multiservice Edge Routers and T Series Core Routers support four classifier PLP designations: **low**, **medium-low**, **medium-high**, and **high**. To configure the PLP for a classifier, include the following statements at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
classifiers {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
      loss-priority (low | medium-low | medium-high | high) code-points [ aliases ]
        [ bit-patterns ];
    }
  }
}
```

The inputs for a classifier are the CoS values. The outputs for a classifier are the forwarding class and the loss priority (PLP). A classifier sets the forwarding class and the PLP for each packet entering the interface with a specific set of CoS values.

For example, in the following configuration, the **assured-forwarding** forwarding class and **medium-low** PLP are assigned to all packets entering the interface with the **101110** CoS values:

```
class-of-service {
  classifiers {
    dscp dscp-cl {
      forwarding-class assured-forwarding {
        loss-priority medium-low {
          code-points 101110;
        }
      }
    }
  }
}
```

To use this classifier, you must configure the settings for the **assured-forwarding** forwarding class at the **[edit class-of-service forwarding-classes queue queue-number assured-forwarding]** hierarchy level. For more information, see [“Overview of Forwarding Classes” on page 1526](#).

Using Multifield Classifiers to Set PLP

Multifield classifiers take action on incoming or outgoing packets, depending whether the firewall rule is applied as an input filter or an output filter. When TCM is enabled, Juniper Networks M320 Multiservice Edge Routers and T Series Core Routers support four multifield classifier PLP designations: **low**, **medium-low**, **medium-high**, and **high**.

To configure the PLP for a multifield classifier, include the **loss-priority** statement in a policer or firewall filter that you configure at the **[edit firewall]** hierarchy level:

```
[edit firewall]
family family-name {
  filter filter-name {
    term term-name {
      from {
        match-conditions;
      }
      then {
        loss-priority (low | medium-low | medium-high | high);
        forwarding-class class-name;
      }
    }
  }
}
```

The inputs (match conditions) for a multifield classifier are one or more of the six packet header fields: destination address, source address, IP protocol, source port, destination port, and DSCP. The outputs for a multifield classifier are the forwarding class and the loss priority (PLP). In other words, a multifield classifier sets the forwarding class and the PLP for each packet entering or exiting the interface with a specific destination address, source address, IP protocol, source port, destination port, or DSCP.

For example, in the following configuration, the forwarding class **expedited-forwarding** and PLP **medium-high** are assigned to all IPv4 packets with the **10.1.1.0/24** or **10.1.2.0/24** source address:

```
firewall {
  family inet {
    filter classify-customers {
      term isp1-customers {
        from {
          source-address 10.1.1.0/24;
          source-address 10.1.2.0/24;
        }
        then {
          loss-priority medium-high;
          forwarding-class expedited-forwarding;
        }
      }
    }
  }
}
```

To use this classifier, you must configure the settings for the **expedited-forwarding** forwarding class at the **[edit class-of-service forwarding-classes queue queue-number]**

expedited-forwarding] hierarchy level. For more information, see [“Overview of Forwarding Classes” on page 1526](#).

Configuring PLP for Drop-Profile Maps

RED drop profiles take action on outgoing packets. When TCM is enabled, M320 and T Series routers support four drop-profile map PLP designations: **low**, **medium-low**, **medium-high**, and **high**.

To configure the PLP for the drop-profile map, include the **schedulers** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
schedulers {
  scheduler-name {
    drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol
    any drop-profile profile-name;
  }
}
```

When you configure TCM, the drop-profile map's protocol type must be **any**.

The inputs for a drop-profile map are the loss priority and the protocol type. The output for a drop-profile map is the drop profile name. In other words, the map sets the drop profile for each packet with a specific PLP and protocol type exiting the interface.

For example, in the following configuration, the **dp** drop profile is assigned to all packets exiting the interface with a medium-low PLP and belonging to any protocol:

```
class-of-service {
  schedulers {
    af {
      drop-profile-map loss-priority medium-low protocol any drop-profile dp;
    }
  }
}
```

To use this drop-profile map, you must configure the settings for the **dp** drop profile at the **[edit class-of-service drop-profiles dp]** hierarchy level. For more information, see [“RED Drop Profiles Overview” on page 1674](#).

Configuring Rewrite Rules Based on PLP

Rewrite rules take action on outgoing packets. When TCM is enabled, M320 and T Series routers support four rewrite PLP designations: **low**, **medium-low**, **medium-high**, and **high**. To configure the PLP for a rewrite rule, include the following statements at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
rewrite-rules {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) rewrite-name {
    import (rewrite-name | default);
    forwarding-class class-name {
      loss-priority (low | medium-low | medium-high | high) code-point (alias | bits);
    }
  }
}
```

```
}
```

The inputs for a rewrite rule are the forwarding class and the loss priority (PLP). The output for a rewrite rule are the CoS values. In other words, a rewrite rule sets the CoS values for each packet exiting the interface with a specific forwarding class and PLP.

For example, if you configure the following, the **000000** CoS values are assigned to all packets exiting the interface with the **assured-forwarding** forwarding class and **medium-high** PLP:

```
class-of-service {
  rewrite-rules {
    dscp dscp-rw {
      forwarding-class assured-forwarding {
        loss-priority medium-high code-point 000000;
      }
    }
  }
}
```

To use this classifier, you must configure the settings for the **assured-forwarding** forwarding class at the **[edit class-of-service forwarding-classes queue *queue-number* assured-forwarding]** hierarchy level. For more information, see [“Overview of Forwarding Classes” on page 1526](#).

Configuration Statements for Tricolor Marking Policers

- [\[edit class-of-service\] Hierarchy Level on page 1463](#)
- [\[edit firewall\] Hierarchy Level on page 1483](#)
- [\[edit interfaces\] Hierarchy Level on page 1509](#)

[edit class-of-service] Hierarchy Level

```
class-of-service {
  classifiers {
    type classifier-name {
      forwarding-class class-name {
        loss-priority (high | low | medium-high | medium-low) code-points [ aliases bits ];
      }
      import (classifier-name | default);
    }
  }
  code-point-aliases {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) {
      alias-name bits;
    }
  }
  drop-profiles {
    profile-name {
      fill-level percentage drop-probability percentage;
      interpolate {
        drop-probability value;
        fill-level value;
      }
    }
  }
}
```

```

}
fabric {
  scheduler-map {
    priority (high | low) scheduler scheduler-name;
  }
}
forwarding-class-map {
  map-name {
    class class-name queue-num queue-number <restricted-queue queue-number>;
  }
}
forwarding-classes {
  class class-name policing-priority (normal | premium) queue-num queue-number
    priority (high | low);
  queue queue-number class-name policing-priority (normal | premium) priority (high |
    low);
}
forwarding-policy {
  class class-name {
    classification-override {
      forwarding-class class-name;
    }
  }
  next-hop-map map-name {
    forwarding-class class-name {
      discard;
      lsp-next-hop [ lsp-regular-expressions ];
      next-hop [ next-hop-names ];
      non-lsp-next-hop;
    }
  }
}
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      drop-timeout milliseconds;
      fragment-threshold bytes;
      multilink-class number;
      no-fragmentation;
    }
  }
}
host-outbound-traffic {
  dscp-code-point value;
  forwarding-class class-name;
  ieee-802.1 {
    default value;
    rewrite-rules;
  }
  tcp {
    raise-internet-control-priority;
  }
}
interfaces {
  ... the interfaces subhierarchy appears after the main [edit class-of-service] hierarchy
  ...

```

```

}
}
restricted-queues {
    forwarding-class class-name queue-number;
}
rewrite-rules {
    (dscp | dscp-ipv6 | exp | frame-relay-de | ieee-802.1 | ieee-802.1ad | inet-precedence)
    rewrite-rule {
        forwarding-class class-name {
            loss-priority level code-point (alias | bits);
        }
        import (rewrite-rule | default);
    }
}
routing-instances routing-instance-name {
    classifiers {
        dscp (classifier-name | default);
        dscp-ipv6 (classifier-name | default);
        exp (classifier-name | default);
        ieee-208.1 (classifier-name | default | encapsulated | vlan-tag (inner | outer));
    }
}
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
schedulers {
    scheduler-name {
        adjust-minimum value;
        adjust-percent value;
        buffer-size (exact | percent percentage | remainder);
        drop-profile-map loss-priority (any | high | low | medium-high | medium-low)
            protocol any;
        excess-priority (high | low | medium-high | medium-low);
        excess-rate (percent percentage | proportion proportion);
        priority (high | low | medium-high | medium-low | strict-high);
        shaping-rate (bps | percent percentage | burst-size size);
        transmit-rate (bps | percent percentage | remainder) <exact | rate-limit>;
    }
}
traceoptions {
    file <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
traffic-control-profiles {
    profile-name {
        adjust-minimum rate;
        delay-buffer-rate (bps | cps cps | percent percentage);
        excess-rate (percent percentage | proportion value);
        guaranteed-rate (bps | percent percentage) <burst-size bytes>;
        overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
        scheduler-map map-name;
        shaping-rate (bps | percent percentage) <burst-size bytes>;
    }
}

```

```

    }
  }
  tri-color;
}

class-of-service {
  interfaces {
    interface-name {
      excess-bandwidth-share (equal | proportional value);
      input-excess-bandwidth-share (equal | proportional value);
      input-scheduler-map map-name;
      input-shaping-rate bps;
      input-traffic-control-profile profile-name;
      output-forwarding-class-map map-name;
      output-traffic-control-profile profile-name;
      scheduler-map map-name;
      scheduler-map-chassis (map-name | derived);
      shaping-rate bps;
      unit (logical-unit-number | *) {
        classifiers {
          dscp (classifier-name | default) {
            family [ inet mpls ];
          }
          dscp-ipv6 (classifier-name | default) {
            family [ inet mpls ];
          }
          exp (classifier-name | default);
          ieee-208.1 (classifier-name | default) <vlan-tag (inner | outer)>;
          ieee-208.1ad (classifier-name | default);
          inet-precedence (classifier-name | default);
        }
        forwarding-class class-name;
        input-scheduler-map map-name;
        input-shaping-rate bps;
        input-traffic-control-profile profile-name shared-instance instance-name;
        loss-priority-maps {
          (map-name | default);
        }
        loss-priority-rewrites {
          (map-name | default);
        }
        output-forwarding-class-map map-name;
        output-traffic-control-profile profile-name shared-instance instance-name;
        rewrite-rules {
          dscp (rule-name | default) <protocol mpls>;
          dscp-ipv6 (rule-name | default);
          exp (rule-name | default) <protocol [ mpls-any | mpls-inet-both |
            mpls-inet-both-non-vpn ]>;
          exp-push-push-push default;
          exp-swap-push-push default;
          ieee-802.1 (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
          ieee-802.1ad (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
          inet-precedence (rewrite-name | default) <protocol mpls>;
        }
        scheduler-map map-name;
        shaping-rate bps;
      }
    }
  }
}

```



```


        translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 |
        to-exp-from-exp | to-inet-precedence-from-inet-precedence) table-name;
    }
}
interface-set interface-set-name {
    excess-bandwidth-share (equal | proportional value);
    input-excess-bandwidth-share (equal | proportional value);
    input-traffic-control-profile profile-name;
    input-traffic-control-profile-remaining profile-name;
    internal-node;
    output-traffic-control-profile profile-name;
    output-traffic-control-profile-remaining profile-name;
}
}
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

classifiers (Definition)

Syntax	<pre> classifiers { type classifier-name { import (classifier-name default); forwarding-class class-name { loss-priority level code-points [aliases] [bit-patterns]; } } } </pre>
Hierarchy Level	[edit class-of-service], [edit class-of-service routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. ieee-802.1ad option introduced in Junos OS Release 9.2.
Description	Define a CoS behavior aggregate (BA) classifier for classifying packets. You can associate the classifier with a forwarding class or code-point mapping, and import a default classifier or one that is previously defined.
<div>  <p>NOTE: The [edit class-of-service routing-instances <i>routing-instance-name</i>] hierarchy level and the dscp-ipv6 and ieee-802.1ad classifier types are not supported on ACX Series routers.</p> </div>	
Options	<p>classifier-name—Name of the aggregate behavior classifier.</p> <p>type—Traffic type: dscp, dscp-ipv6, exp, ieee-802.1, ieee-802.1ad, inet-precedence.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Overview of BA Classifier Types on page 1321 • Example: Configuring CoS for a PBB Network • Configuring CoS on ACX Series Universal Access Routers

code-points

Syntax	<code>code-points ([<i>aliases</i>] [<i>bit-patterns</i>]);</code>
Hierarchy Level	[edit class-of-service classifiers <i>type classifier-name</i> forwarding-class <i>class-name</i> loss-priority <i>level</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 8.5 for J Series devices. Statement introduced in Junos OS Release 9.2 for SRX Series devices.
Description	Specify one or more DSCP code-point aliases or bit sets for association with a forwarding class.
Options	<i>aliases</i> —Name of the DSCP alias. <i>bit-patterns</i> —Value of the code-point bits, in six-bit binary form.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Overview of BA Classifier Types on page 1321 • <i>Example: Configuring CoS for a PBB Network</i> • <i>Example: Configuring Behavior Aggregate Classifiers</i> • <i>Example: Configuring Forwarding Classes</i>

drop-profile (Schedulers)

Syntax	<code>drop-profile <i>profile-name</i>;</code>
Hierarchy Level	[edit class-of-service schedulers <i>scheduler-name</i> drop-profile-map loss-priority (any low medium-low medium-high high) protocol (any non-tcp tcp)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches. Statement introduced in Junos OS Release 12.2 for ACX Series Routers.
Description	Define drop profiles for RED. When a packet arrives, RED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the arriving packet.
Options	<i>profile-name</i> —Name of the drop profile.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Drop Profile Maps for Schedulers on page 1590• RED Drop Profiles Overview on page 1674

drop-profile-map (Schedulers)

Syntax	<code>drop-profile-map loss-priority (any low medium-low medium-high high) protocol(any non-tcp tcp) drop-profile (Schedulers) <i>profile-name</i>;</code>
Hierarchy Level	[edit class-of-service schedulers <i>scheduler-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches. Statement introduced in Junos OS Release 12.2 for ACX Series Routers.
Description	Define the loss-priority value for a drop profile. The statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Default Schedulers Overview on page 1575• Configuring Drop Profile Maps for Schedulers on page 1590

dscp (Multifield Classifier)

Syntax	<code>dscp [0 <i>value</i>];</code>
Hierarchy Level	<code>[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> then]</code>
Release Information	Statement introduced in Junos OS Release 7.4.
Description	<p>For M320 and T Series routers, set the DSCP field of incoming or outgoing packets to 000000. On the same packets, you can use a behavior aggregate (BA) classifier and a rewrite rule to rewrite the MPLS EXP field.</p> <p>For MX Series routers with MPCs and EX Series switches, the DSCP field can be set from a numeric range.</p> <p>For MX Series routers and EX Series switches, if you configure a firewall filter with a DSCP action or traffic-class action on a DPC, the commit does not fail, but the filter is not applied to the interface, a warning displays, and an entry is made in the syslog.</p>
Options	value —For MX Series routers with MPCs and EX Series switches, specify the field of incoming or outgoing packets in the range from 0 through 63 .
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Applying Tricolor Marking Policers to Firewall Filters on page 1456

dscp (Rewrite Rules)

Syntax	<code>dscp (rewrite-name default);</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For IPv4 traffic, apply a Differentiated Services (DiffServ) code point (DSCP) rewrite rule.
Options	rewrite-name —Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules dscp] hierarchy level. default —The default mapping.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694• dscp-ipv6 (Class-of-Service) on page 1357• exp on page 1358• exp-push-push-push on page 1723• exp-swap-push-push on page 1724• ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360• ieee-802.1ad on page 1362• inet-precedence on page 1364• rewrite-rules (Definition) on page 1481

dscp-ipv6 (Class-of-Service)

Syntax	<code>dscp-ipv6 (<i>rewrite-name</i> <default>) { protocol mpls }</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For IPv6 traffic, apply a DSCP rewrite rule.
Options	<p><i>rewrite-name</i>—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules dscp-ipv6] hierarchy level.</p> <p>default— Default mapping.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Rewrite Rules on page 1694 • dscp (Rewrite Rules) on page 1382 • exp on page 1358 • exp-push-push-push on page 1723 • exp-swap-push-push on page 1724 • ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360 • ieee-802.1ad on page 1362 • inet-precedence on page 1364 • rewrite-rules (Definition) on page 1481

exp

Syntax	<code>exp (rewrite-name default) protocol protocol-types;</code>
Hierarchy Level	<code>[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS Release 12.2. for ACX series
Description	Apply an MPLS experimental (EXP) rewrite rule.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the <code>[edit class-of-service rewrite-rules exp]</code> hierarchy level.</p> <p>default—The default mapping.</p> <p>By default, IP precedence rewrite rules alter the first three bits on the type-of-service (ToS) byte while leaving the last three bits unchanged. This default behavior applies to rewrite rules you configure for MPLS packets with IPv4 payloads. You configure these types of rewrite rules by including the mpls-inet-both or mpls-inet-both-non-vpn option at the <code>[edit class-of-service interfaces interface interface-name unit logical-unit-number rewrite-rules exp rewrite-rule-name protocol]</code> hierarchy level. The IP precedence rewrite rules explanation does not apply to ACX Series Universal Access routers.</p> <p>On interfaces configured on Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) on MX Series Ethernet Services Routers and EX Series switches, we highly recommend that you configure the default option when you configure a behavior aggregate (BA) classifier that does not include a specific rewrite rule for MPLS packets. Doing so ensures that MPLS exp value is rewritten according to the BA classifier rules configured for forwarding or packet loss priority. This does not apply to ACX Series Universal Access routers.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Rewriting the EXP Bits of All Three Labels of an Outgoing Packet</i>• dscp (Rewrite Rules) on page 1382• dscp-ipv6 (Class-of-Service) on page 1357• exp-push-push-push on page 1723• exp-swap-push-push on page 1724• ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360• ieee-802.1ad on page 1362• inet-precedence on page 1364

- [rewrite-rules \(Definition\)](#) on page 1481

forwarding-class (BA Classifiers)

Syntax	forwarding-class <i>class-name</i> { loss-priority <i>level</i> code-points [<i>aliases</i>] [<i>bit-patterns</i>]; }
Hierarchy Level	[edit class-of-service classifiers <i>type classifier-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define forwarding class name and option values.
Options	<i>class-name</i> —Name of the forwarding class. The remaining statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Defining Classifiers on page 1330 • <i>Example: Configuring CoS for a PBB Network</i>

ieee-802.1 (Rewrite Rules on Logical Interface)

Syntax	ieee-802.1 (<i>rewrite-name</i> default) vlan-tag (outer outer-and-inner);
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4. vlan-tag statement introduced in Junos OS Release 8.1.
Description	Apply an IEEE-802.1 rewrite rule. For IQ PICs, you can only configure one IEEE 802.1 rewrite rule on a physical port. All logical ports (units) on that physical port should apply the same IEEE 802.1 rewrite rule.
Options	rewrite-name —Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules ieee-802.1] hierarchy level. default —The default mapping.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694• <i>Example: Configuring CoS for a PBB Network</i>• dscp (Rewrite Rules) on page 1382• dscp-ipv6 (Class-of-Service) on page 1357• exp on page 1358• exp-push-push-push on page 1723• exp-swap-push-push on page 1724• ieee-802.1ad on page 1362• inet-precedence on page 1364• rewrite-rules (Definition) on page 1481

import (Classifiers)

Syntax	<code>import (classifier-name default);</code>
Hierarchy Level	<code>[edit class-of-service classifiers type classifier-name]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify a default or previously defined classifier.
Options	classifier-name —Name of the classifier mapping configured at the <code>[edit class-of-service classifiers]</code> hierarchy level. default —The default classifier mapping.
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Overview of BA Classifier Types on page 1321

import (Rewrite Rules)

Syntax	<code>import (rewrite-name default);</code>
Hierarchy Level	<code>[edit class-of-service rewrite-rules type rewrite-name]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify a default or previously defined rewrite-rules mapping to import.
Options	rewrite-name —Name of a rewrite-rules mapping configured at the <code>[edit class-of-service rewrite-rules]</code> hierarchy level. default —The default rewrite-rules mapping.
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694

inet-precedence

Syntax	<code>inet-precedence (rewrite-name default);</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Apply a IPv4 precedence rewrite rule.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules inet-precedence] hierarchy level.</p> <p>default—The default mapping. By default, IP precedence rewrite rules alter the first three bits on the type of service (ToS) byte while leaving the last three bits unchanged.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694• dscp (Rewrite Rules) on page 1382• dscp-ipv6 (Class-of-Service) on page 1357• exp on page 1358• exp-push-push-push on page 1723• exp-swap-push-push on page 1724• ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360• ieee-802.1ad on page 1362• rewrite-rules (Definition) on page 1481

loss-priority (Scheduler Drop Profiles)

Syntax	loss-priority (any high low medium-high medium-low);
Hierarchy Level	[edit class-of-service schedulers <i>scheduler-name</i> drop-profile-map]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches. Statement introduced in Junos OS Release 12.2 for ACX Series Routers.
Description	Specify a loss priority to which to apply a drop profile. The drop profile map sets the drop profile for a specific PLP and protocol type. The inputs for the map are the PLP designation and the protocol type. The output is the drop profile.
Options	any —The drop profile applies to packets with any PLP.



NOTE: On ACX Series Routers, only the **any** option is supported when you configure the **non-tcp** option for [protocol](#).

high—The drop profile applies to packets with high PLP.

low—The drop profile applies to packets with low PLP.

medium-high—The drop profile applies to packets with medium-high PLP.

medium-low—The drop profile applies to packets with medium-low PLP.

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

Related Documentation	<ul style="list-style-type: none"> • Default Schedulers Overview on page 1575 • Configuring Drop Profile Maps for Schedulers on page 1590 • Configuring Schedulers for Priority Scheduling on page 1594 • Configuring Tricolor Marking on page 1447 • protocol (Schedulers) on page 1480
------------------------------	---

protocol (Schedulers)

Syntax	protocol (any non-tcp tcp);
Hierarchy Level	[edit class-of-service schedulers <i>scheduler-name</i> drop-profile-map]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches. Statement introduced in Junos OS Release 12.2 for ACX Series Routers.
Description	Specify the protocol type for the specified scheduler.
Options	any —Accept any protocol type. non-tcp —(ACX Series Routers, M Series and T Series (except T4000) routers only) Accept any protocol type other than TCP/IP.



NOTE: On ACX Series Routers, when you configure the **non-tcp** option, only the **any** option is supported for [loss-priority](#).

	tcp —(ACX Series Routers, M Series and T Series (except T4000) routers only) Accept TCP/IP protocol type.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Schedulers on page 1579

rewrite-rules (Definition)

Syntax	<pre>rewrite-rules { type <i>rewrite-name</i>{ import (<i>rewrite-name</i> default); forwarding-class <i>class-name</i> { loss-priority <i>level</i> <i>code-point</i> [<i>aliases</i>] [<i>bit-patterns</i>]; } } }</pre>
Hierarchy Level	[edit class-of-service]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>ieee-802.1ad option introduced in Junos OS Release 9.2.</p>
Description	Specify a rewrite-rules mapping for the traffic that passes through all queues on the interface.
Options	<p><i>rewrite-name</i>—Name of a rewrite-rules mapping.</p> <p><i>type</i>—Traffic type.</p> <p>Values: dscp, dscp-ipv6, exp, frame-relay-de (J Series routers only), ieee-802.1, ieee-802.1ad, inet-precedence</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Rewrite Rules on page 1694 • <i>Example: Configuring CoS for a PBB Network</i> • J Series router documentation

schedulers (Class of Service)

Syntax	<pre>schedulers { scheduler-name { adjust-minimum <i>rate</i>; adjust-percent <i>percentage</i>; buffer-size (<i>seconds</i> percent <i>percentage</i> remainder temporal <i>microseconds</i>); drop-profile-map loss-priority (any low medium-low medium-high high) <i>protocol</i> (any non-tcp tcp) drop-profile <i>profile-name</i>; excess-priority [low medium-low medium-high high none]; excess-rate (percent <i>percentage</i> proportion <i>value</i>); priority <i>priority-level</i>; shaping-rate (percent <i>percentage</i> <i>rate</i>); transmit-rate (percent <i>percentage</i> <i>rate</i> remainder) <exact rate-limit>; } }</pre>
Hierarchy Level	[edit class-of-service]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series switches.
Description	Specify the scheduler name and parameter values.
Options	<i>scheduler-name</i> —Name of the scheduler to be configured. The remaining statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Schedulers Overview on page 1573• Default Schedulers Overview on page 1575• Configuring Schedulers on page 1579• Configuring a Scheduler• Example: Configuring CoS for a PBB Network

tri-color

Syntax	tri-color;
Hierarchy Level	[edit class-of-service]
Release Information	Statement introduced in Junos OS Release 7.4.
Description	For IPv4 packets on M320, MX Series, T Series routers with Enhanced II Flexible PIC Concentrators (FPCs), and EX Series switches, enable two-rate tricolor marking (TCM), as defined in RFC 2698.
Default	If you do not include this statement, tricolor marking is not enabled and the medium packet loss priority (PLP) is not configurable.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Tricolor Marking on page 1447

[edit firewall] Hierarchy Level

Several statements in the **[edit firewall]** hierarchy are valid at numerous locations within the hierarchy. To make the complete hierarchy easier to read, the repeated statements are listed in the following sections, which are referenced at the appropriate locations in “[Complete \[edit firewall\] Hierarchy](#)” on page 325.

- [Common Firewall Actions on page 1483](#)
- [Common IP Firewall Actions on page 1484](#)
- [Common IPv4 Firewall Actions on page 1484](#)
- [Common IP Firewall Match Conditions on page 1485](#)
- [Common IPv4 Firewall Match Conditions on page 1486](#)
- [Common Layer 2 Firewall Match Conditions on page 1486](#)
- [Complete \[edit firewall\] Hierarchy on page 1488](#)

Common Firewall Actions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit firewall\] Hierarchy](#)” on page 325 instead of the statements being repeated.

- **[edit firewall family (any | ccc | ethernet-switching | inet | inet6 | mpls | vpls) filter *filter-name* term *term-name* then]**
- **[edit firewall filter *filter-name* term *term-name* then]**

The common firewall actions are as follows:

```
count counter-name;  
forwarding-class class-name;  
loss-priority (high | low | medium-high | medium-low);  
next term;  
policer policer-name;  
three-color-policer policer-name {  
    (single-rate single-rate-policer-name | two-rate two-rate-policer-name);  
}
```

Common IP Firewall Actions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- [edit firewall family inet filter *filter-name* term *term-name* then]
- [edit firewall family inet6 filter *filter-name* term *term-name* then]
- [edit firewall filter *filter-name* term *term-name* then]

The common IP firewall actions are as follows:

```
log;  
logical-system logical-system-name <routing-instance routing-instance-name>  
    <topology topology-name>;  
port-mirror;  
port-mirror-instance instance-name;  
routing-instance routing-instance-name <topology topology-name>;  
sample;  
service-filter-hit;  
syslog;  
topology topology-name;
```

Common IPv4 Firewall Actions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- [edit firewall family inet filter *filter-name* term *term-name* then]
- [edit firewall filter *filter-name* term *term-name* then]

The common IP version 4 (IPv4) firewall actions are as follows:

```
(accept | discard <accounting collector-name> | reject <administratively-prohibited |  
    bad-host-tos | bad-network-tos | fragmentation-needed | host-prohibited |  
    host-unknown | host-unreachable | network-prohibited | network-unknown |  
    network-unreachable | port-unreachable | precedence-cutoff | precedence-violation |  
    protocol-unreachable | source-host-isolated | source-route-failed | tcp-reset>);  
ipsec-sa sa-name;  
load-balance sa-name;  
next-hop-group group-name;  
prefix-action action-name;
```

Common IP Firewall Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit firewall\] Hierarchy](#)” on page 325 instead of the statements being repeated.

- `[edit firewall family inet dialer-filter filter-name term term-name from]` (with the exceptions noted at this level in “[Complete \[edit firewall\] Hierarchy](#)” on page 325)
- `[edit firewall family inet filter filter-name term term-name from]`
- `[edit firewall family inet6 dialer-filter filter-name term term-name from]` (with the exceptions noted at this level in “[Complete \[edit firewall\] Hierarchy](#)” on page 325)
- `[edit firewall family inet6 filter filter-name term term-name from]`
- `[edit firewall filter filter-name term term-name from]`

The common IP firewall match conditions are as follows:

```

address {
    ip-prefix</prefix-length> <except>;
}
destination-address {
    ip-prefix</prefix-length> <except>;
}
destination-class [ class-names ] | destination-class-except [ class-names ];
(destination-port [ port-names ] | destination-port-except [ port-names ]);
destination-prefix-list {
    list-name <except>;
}
(forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
icmp-code [ codes ] | icmp-code-except [ codes ];
icmp-type [ types ] | icmp-type-except [ types ];
interface interface-name;
(interface-group [ group-names ] | interface-group-except [ group-names ]);
interface-set set-name;
(loss-priority [ priorities ] | loss-priority-except [ priorities ]);
(packet-length [ values ] | packet-length-except [ values ]);
(port [ port-names ] | port-except [ port-names ]);
prefix-list {
    list-name <except>;
}
service-filter-hit;
source-address {
    ip-prefix</prefix-length> <except>;
}
(source-class [ class-names ] | source-class-except [ class-names ]);
(source-port [ port-names ] | source-port-except [ port-names ]);
source-prefix-list {
    list-name <except>;
}
tcp-established;
tcp-flags flag;
tcp-initial;

```

Common IPv4 Firewall Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- **[edit firewall family inet dialer-filter *filter-name* term *term-name* from]** (with the exceptions noted at this level in [“Complete \[edit firewall\] Hierarchy” on page 325](#))
- **[edit firewall family inet filter *filter-name* term *term-name* from]**
- **[edit firewall filter *filter-name* term *term-name* from]**

The common IPv4 firewall match conditions are as follows:

```
(ah-spi [ values ] | ah-spi-except [ values ]);
(dscp [ code-point-values ] | dscp-except [ code-point-values ]);
(esp-spi [ values ] | esp-spi-except [ values ]);
first-fragment;
fragment-flags flag;
(fragment-offset [ offsets ] | fragment-offset-except [ offsets ]);
(ip-options [ option-names ] | ip-options-except [ option-names ]);
is-fragment;
(precedence [ precedence-names ] | precedence-except [ precedence-names ]);
(protocol [ protocol-names ] | protocol-except [ protocol-names ]);
(ttl [ tll-values ] | ttl-except [ tll-values ]);
```

Common Layer 2 Firewall Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- **[edit firewall family ethernet-switching filter *filter-name* term *term-name* from]**
- **[edit firewall family vpls filter *filter-name* term *term-name* from]**

The common Layer 2 firewall match conditions are as follows:

```
destination-mac-address {
    mac-address <except>;
}
(destination-port [ port-names ] | destination-port-except [ port-names ]);
(dscp [ code-point-values ] | dscp-except [ code-point-values ]);
(ether-type [ protocol-types ] | ether-type-except [ protocol-types ]);
(forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
(icmp-code [ codes ] | icmp-code-except [ codes ]);
(icmp-type [ types ] | icmp-type-except [ types ]);
(interface-group [ group-names ] | interface-group-except [ group-names ]);
ip-address {
    ip-prefix</prefix-length> <except>;
}
ip-destination-address {
    ip-prefix</prefix-length> <except>;
}
(ip-precedence [ precedence-names ] | ip-precedence-except [ precedence-names ]);
```

```
(ip-protocol [ protocol-names ] | ip-protocol-except [ protocol-names ] );
ip-source-address ip-prefix </prefix-length>;
(learn-vlan-1p-priority [ priorities ] | learn-vlan-1p-priority [ priorities ] );
(learn-vlan-id [ vlan-ids ] | learn-vlan-id-except [ vlan-ids ] );
(loss-priority [ priorities ] | loss-priority-except [ priorities ] );
(port [ port-names ] | port-except [ port-names ] );
source-mac-address {
    mac-address <except>;
}
(source-port [ port-names ] | source-port-except [ port-names ] );
tcp-flags flag;
(traffic-type [ broadcast known-unicast multicast unknown-unicast ] |
    traffic-type-except [ broadcast known-unicast multicast unknown-unicast ] );
(user-vlan-1p-priority [ priorities ] | user-vlan-1p-priority [ priorities ] );
(user-vlan-id [ vlan-ids ] | user-vlan-id-except [ vlan-ids ] );
(vlan-ether-type [ protocol-types ] | vlan-ether-type-except [ protocol-types ] );
```

Complete [edit firewall] Hierarchy

```
firewall {
  family (any | ccc | ethernet-switching | inet | inet6 | mpls | vpls) {
    ... the family subhierarchies appear after the main [edit firewall] hierarchy ...
  }
  filter filter-name {
    accounting-profile [ profile-names ];
    enhanced-mode;
    interface-shared-with;
    interface-specific;
    physical-interface-policer;
    term term-name {
      filter filter-name;
      from {
        ... statements in Common IP Firewall Match Conditions on page 322 AND
        statements in Common IPv4 Firewall Match Conditions on page 323 ...
      }
      then {
        ... statements in Common Firewall Actions on page 320 AND
        statements in Common IP Firewall Actions on page 321 AND
        statements in Common IPv4 Firewall Actions on page 321 ...
      }
    }
  }
  hierarchical-policer policer-name {
    aggregate {
      if-exceeding {
        bandwidth-limit bps;
        burst-size-limit bytes;
      }
      then {
        discard;
        forwarding-class class-name;
        loss-priority (high | low | medium-high | medium-low);
      }
    }
    logical-interface-policer;
    physical-interface-policer;
    premium {
      if-exceeding {
        bandwidth-limit bps;
        burst-size-limit bytes;
      }
      then {
        discard;
      }
    }
  }
  shared-bandwidth-policer;
  interface-set interface-set-name {
    interface-name;
  }
  load-balance-group group-name {
    next-hop-group [ group-names ];
  }
}
```

```

}
policer policer-name {
  filter-specific;
  if-exceeding {
    (bandwidth-limit bps | bandwidth-percent percentage);
    burst-size-limit bytes;
  }
  logical-bandwidth-policer;
  logical-interface-policer;
  physical-interface-policer;
  then {
    discard;
    forwarding-class class-name;
    loss-priority (high | low | medium-high | medium-low);
  }
}
three-color-policer policer-name {
  action {
    loss-priority high then discard;
  }
  filter-specific;
  logical-interface-policer;
  physical-interface-policer;
  shared-bandwidth-policer;
  single-rate {
    (color-aware | color-blind);
    committed-burst-size bytes;
    committed-information-rate bps;
    excess-burst-size bytes;
  }
  two-rate {
    (color-aware | color-blind);
    committed-burst-size bytes;
    committed-information-rate bps;
    peak-burst-size bytes;
    peak-information-rate bps;
  }
}
}

firewall {
  family any {
    filter filter-name {
      interface-shared;
      term term-name {
        from {
          (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
          interface interface-name;
          interface-set set-name;
          (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
          (packet-length [ values ] | packet-length-except [ values ]);
        }
        then {
          ... statements in Common Firewall Actions on page 320 PLUS ...
          (accept | discard);
        }
      }
    }
  }
}

```

```
    }  
  }  
}
```

```
firewall {  
  family ccc {  
    filter filter-name {  
      accounting-profile [ profile-names ];  
      physical-interface-filter;  
      interface-specific;  
      term term-name {  
        filter filter-name;  
        from {  
          (forwarding-class [ class-names ] | forwarding-class-except [ class-names ] );  
          (interface-group [ group-names ] | interface-group-except [ group-names ] );  
          (learn-vlan-1p-priority [ priorities ] | learn-vlan-1p-priority [ priorities ] );  
          (loss-priority [ priorities ] | loss-priority-except [ priorities ] );  
          (user-vlan-1p-priority [ priorities ] | user-vlan-1p-priority [ priorities ] );  
        }  
        then {  
          ... statements in Common Firewall Actions on page 320 PLUS ...  
          (accept | discard);  
          port-mirror-instance instance-name;  
        }  
      }  
    }  
  }  
}
```

```
firewall {  
  family ethernet-switching {  
    filter filter-name {  
      interface-specific;  
      term term-name {  
        from {  
          destination-address {  
            ip-prefix</prefix-length>;  
          }  
          destination-mac-address {  
            mac-address;  
          }  
          destination-port [ port-names ];  
          destination-prefix-list {  
            list-name;  
          }  
          dot1q-tag [ tag-values ];  
          dot1q-user-priority [ priority-values ];  
          dscp [ code-point-values ];  
          ether-type [ protocol-names ];  
          fragment-flags flag;  
          icmp-code [ codes ];  
          icmp-type [ types ];  
          interface interface-name;  
          is-fragment;  
        }  
      }  
    }  
  }  
}
```



```

precedence [ precedence-names ];
protocol [ protocol-names ];
source-address {
    ip-prefix</prefix-length>;
}
source-mac-address {
    mac-address;
}
source-port [ port-names ];
source-prefix-list {
    list-name;
}
tcp-established;
tcp-flags flag;
tcp-initial;
vlan [ vlan-names ];
}
then {
    (accept | discard);
    analyzer analyzer-name;
    count counter-name;
    forwarding-class class-name;
    interface interface-name;
    log;
    loss-priority (high | low);
    policer policer-name;
    syslog;
    vlan vlan-name;
}
}
}
}
}

firewall {
    family inet {
        dialer-filter filter-name {
            accounting-profile [ profile-names ];
            term term-name {
                from {
                    ... statements in Common IP Firewall Match Conditions on page 322 AND
                    statements in Common IPv4 Firewall Match Conditions on page 323 EXCEPT
                    FOR ...
                    (ah-spi [ values ] | ah-spi-except [ values ]); # NOT valid at this level
                    (destination-class [ class-names ] |
                     destination-class-except [ class-names ]); # NOT valid at this level
                    interface interface-name; # NOT valid at this level
                    (loss-priority [ priorities ] | loss-priority-except [ priorities ]); # NOT valid at
                     this level
                    service-filter-hit; # NOT valid at this level
                    (source-class [ class-names ] | source-class-except [ class-names ]); # NOT
                     valid at this level
                }
            }
            then {
                (ignore | note);
                log;
            }
        }
    }
}

```

```

        sample;
        syslog;
    }
}
filter filter-name {
    accounting-profile [ profile-names ];
    interface-specific;
    term term-name {
        filter filter-name;
        from {
            ... statements in Common IP Firewall Match Conditions on page 322 AND
               statements in Common IPv4 Firewall Match Conditions on page 323 ...
        }
        then {
            ... statements in Common Firewall Actions on page 320 AND
               statements in Common IP Firewall Actions on page 321 AND
               statements in Common IPv4 Firewall Actions on page 321 ...
        }
    }
}
prefix-action name {
    count;
    destination-prefix-length prefix-length;
    filter-specific;
    policer policer-name;
    source-prefix-length prefix-length;
    subnet-prefix-length prefix-length;
}
service-filter filter-name {
    term term-name {
        from {
            address {
                ip-prefix</prefix-length>;
            }
            (ah-spi [ values ] | ah-spi-except [ values ]);
            destination-address {
                ip-prefix</prefix-length>;
            }
            (destination-port [ port-names ] | destination-port-except [ port-names ]);
            destination-prefix-list {
                list-name;
            }
            (esp-spi [ values ] | esp-spi-except [ values ]);
            first-fragment;
            fragment-flags flag;
            (fragment-offset [ offsets ] | fragment-offset-except [ offsets ]);
            (interface-group [ group-names ] | interface-group-except [ group-names ]);
            (ip-options [ option-names ] | ip-options-except [ option-names ]);
            is-fragment;
            (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
            (port [ port-names ] | port-except [ port-names ]);
            prefix-list {
                list-name;
            }
            (protocol [ protocol-names ] | protocol-except [ protocol-names ]);
        }
    }
}

```

```

firewall {
  family inet6 {
    dialer-filter filter-name {
      accounting-profile [ profile-names ];
      term term-name {
        from {
          ... statements in Common IP Firewall Match Conditions on page 322 PLUS ...
            (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
          ... BUT NOT ...
            (destination-class [ class-names ] |
              destination-class-except [ class-names ]); # NOT valid at this level
            (forwarding-class [ class-names ] |
              forwarding-class-except [ class-names ]); # NOT valid at this level
          interface interface-name; # NOT valid at this level
          (interface-group [ group-names ] | interface-group-except [ group-names ]); #
            NOT valid at this level
          (loss-priority [ priorities ] | loss-priority-except [ priorities ]); # NOT valid at
            this level
        }
      }
    }
  }
}

```

```

        service-filter-hit; # NOT valid at this level
        (source-class [ class-names ] | source-class-except [ class-names ]); # NOT
            valid at this level
        tcp-established; # NOT valid at this level
        tcp-flags flag; # NOT valid at this level
        tcp-initial; # NOT valid at this level
    }
    then {
        (ignore | note);
        log;
        sample;
        syslog;
    }
}
}
filter filter-name {
    accounting-profile [ profile-names ];
    interface-specific;
    term term-name {
        filter filter-name;
        from {
            ... statements in Common IP Firewall Match Conditions on page 322 PLUS ...
            (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
            (traffic-class [ code-point-values ] | traffic-class-except [ code-point-values ]);
        }
        then {
            ... statements in Common Firewall Actions on page 320 AND
            statements in Common IP Firewall Actions on page 321 PLUS ...
            (accept | discard | reject <address-unreachable | administratively-prohibited |
                beyond-scope | fragmentation-needed | no-route | port-unreachable |
                tcp-reset>);
        }
    }
}
}
service-filter filter-name {
    term term-name {
        from {
            address {
                ip-prefix</prefix-length>;
            }
            (ah-spi [ values ] | ah-spi-except [ values ]);
            destination-address {
                ip-prefix</prefix-length>;
            }
            (destination-port [ port-names ] | destination-port-except [ port-names ]);
            destination-prefix-list {
                list-name;
            }
            (esp-spi [ values ] | esp-spi-except [ values ]);
            (interface-group [ group-names ] | interface-group-except [ group-names ]);
            (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
            (port [ port-names ] | port-except [ port-names ]);
            prefix-list {
                list-name;
            }
            source-address {

```

```

        ip-prefix </prefix-length>;
    }
    (source-port [ port-names ] | source-port-except [ port-names ]);
    source-prefix-list {
        list-name;
    }
    tcp-flags flag-name;
}
then {
    count counter-name;
    log;
    port-mirror;
    sample;
    (service | skip);
}
}
}
}
}

firewall {
    family mpls {
        filter filter-name {
            accounting-profile [ profile-names ];
            interface-specific;
            physical-interface-filter;
            term term-name {
                from {
                    (exp [ exp-bits ] | exp-except [ exp-bits ]);
                }
                then {
                    (ignore | note);
                    log;
                    sample;
                    syslog;
                }
            }
        }
    }
    filter filter-name {
        accounting-profile [ profile-names ];
        interface-specific;
        physical-interface-filter;
        term term-name {
            filter filter-name;
            from {
                (exp [ exp-bits ] | exp-except [ exp-bits ]);
                (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
                interface interface-name;
                interface-set set-name;
                (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
            }
            then {
                ... statements in Common Firewall Actions on page 320 PLUS ...
                (accept | discard);
                sample;
            }
        }
    }
}

```

```

    }
  }
}

firewall {
  family vpls {
    filter filter-name {
      accounting-profile [ profile-names ];
      interface-specific;
      term term-name {
        filter filter-name;
        from {
          ... statements in Common Layer 2 Firewall Match Conditions on page 323 ...
        }
        then {
          ... statements in Common Firewall Actions on page 320 PLUS ...
          (accept | discard);
          port-mirror;
          port-mirror-instance instance-name;
        }
      }
    }
  }
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

action

Syntax	<pre>action { loss-priority high then discard; }</pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall three-color-policer <i>name</i>], [edit firewall three-color-policer <i>name</i>], [edit logical-systems <i>logical-system-name</i> firewall three-color-policer <i>name</i>]
Release Information	Statement introduced in Junos OS Release 8.2. Logical systems support introduced in Junos OS Release 9.3. Support at the [edit dynamic-profiles ... three-color-policer] hierarchy level introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Discard traffic on a logical interface using tricolor marking policing.



NOTE: This statement is supported only on IQ2 interfaces.

The remaining statement is explained separately.

Required Privilege	firewall—To view this statement in the configuration.
Level	firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Three-Color Policer Configuration Overview on page 4817 • Basic Single-Rate Three-Color Policers on page 4915 • Basic Two-Rate Three-Color Policers on page 4921 • Two-Color and Three-Color Logical Interface Policers on page 4927 • Two-Color and Three-Color Physical Interface Policers on page 4940 • Two-Color and Three-Color Policers at Layer 2 on page 4947 • loss-priority high then discard on page 4996

family (Multifield Classifier)

Syntax `family family-name {
 filter filter-name {
 term term-name {
 from {
 match-conditions;
 }
 then {
 dscp 0;
 forwarding-class class-name;
 loss-priority (high | low);
 three-color-policer {
 (single-rate | two-rate) policer-name;
 }
 }
 }
 }
 }`

Hierarchy Level [edit [firewall](#)]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a firewall filter for IP version 4 (IPv4) or IP version 6 (IPv6) traffic.

Options *family-name*—Protocol family:

- **ccc**—Circuit cross-connect parameters
- **inet**—IPv4 parameters
- **inet6**—IPv6 protocol parameters
- **iso**—OSI ISO protocol parameters
- **mpls**—MPLS protocol parameters
- **tcc**—Translational cross-connect parameters
- **vpls**—Virtual private LAN service parameters.

The remaining statements are explained separately.

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


Related Documentation • [Configuring Multifield Classifiers on page 1390](#)

filter (Configuring)

Syntax	<pre> filter <i>filter-name</i> { accounting-profile <i>name</i>; enhanced-mode; interface-shared; interface-specific; physical-interface-filter; term <i>term-name</i> { filter <i>filter-name</i>; from { match-conditions; } then { actions; } } } </pre>
Hierarchy Level	<p>[edit dynamic-profiles <i>profile-name</i> firewall family <i>family-name</i>], [edit firewall family <i>family-name</i>], [edit logical-systems <i>logical-system-name</i> firewall family <i>family-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Logical systems support introduced in Junos OS Release 9.3. physical-interface-filter statement introduced in Junos OS Release 9.6. Support at the [edit dynamic-profiles ... family <i>family-name</i>] hierarchy level introduced in Junos OS Release 11.4. Support for the interface-shared> statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure firewall filters.</p>
Options	<p><i>filter-name</i>—Name that identifies the filter. This must be a non-reserved string of not more than 64 characters. To include spaces in the name, enclose it in quotation marks (" "). In Junos OS Release 9.0 and later, you can no longer use special characters within the name of a firewall filter. Firewall filter names are restricted from having the form _.* (beginning and ending with underscores) or _.* (beginning with an underscore).</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Guidelines for Configuring Standard Firewall Filters on page 4478 • Guidelines for Applying Standard Firewall Filters on page 4483 • Configuring Multifield Classifiers on page 1390 • Using Multifield Classifiers to Set PLP on page 1461

- [simple-filter on page 4798](#)

logical-interface-policer

Syntax	logical-interface-policer;
Hierarchy Level	<p>[edit dynamic-profiles <i>profile-name</i> firewall policer <i>policer-name</i>],</p> <p>[edit dynamic-profiles <i>profile-name</i> firewall three-color-policer <i>name</i>],</p> <p>[edit firewall atm-policer <i>atm-policer-name</i>]</p> <p>[edit firewall policer <i>policer-name</i>],</p> <p>[edit firewall policer <i>policer-template-name</i>],</p> <p>[edit firewall three-color-policer <i>policer-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> firewall policer <i>policer-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> firewall three-color-policer <i>name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support at the [edit firewall three-color-policer <i>policer-name</i>] hierarchy level introduced in Junos OS Release 8.2.</p> <p>Logical systems support introduced in Junos OS Release 9.3.</p> <p>Support at the [edit dynamic-profiles ... policer <i>policer-name</i>] and [edit dynamic-profiles ... three-color-policer <i>name</i>] hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Configure a logical interface policer.
	<div>  <p>NOTE: Starting in Junos OS Release 12.2R2, on T Series Core Routers only, you can configure an MPLS LSP policer for a specific LSP to be shared across different protocol family types. You must include the logical-interface-policer statement to do so.</p> </div>
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Two-Color and Three-Color Logical Interface Policers on page 4927 • Traffic Policer Types • Configuring Tricolor Marking Policers on page 1455 • action on page 1497 • Configuring Gigabit Ethernet Two-Color and Tricolor Policers • action

loss-priority (Normal Filter)

Syntax	loss-priority (high low);
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> then]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Set the loss priority of incoming packets.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Multifield Classifiers on page 1390

loss-priority (Simple Filter)

Syntax	loss-priority (high low medium);
Hierarchy Level	[edit firewall family <i>family-name</i> simple-filter <i>filter-name</i> term <i>term-name</i> then]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	Set the loss priority of incoming packets.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Multifield Classifiers on page 1390

policer (Configuring)

Syntax	<pre>policer <i>policer-name</i> { filter-specific; if-exceeding { bandwidth-limit <i>bps</i>; bandwidth-percent <i>number</i>; burst-size-limit <i>bytes</i>; } logical-bandwidth-policer; logical-interface-policer; physical-interface-policer; shared-bandwidth-policer; then { <i>policer-action</i>; } }</pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall], [edit firewall], [edit logical-systems <i>logical-system-name</i> firewall]
Release Information	Statement introduced before Junos OS Release 7.4. The out-of-profile policer action added in Junos OS Release 8.1. The logical-bandwidth-policer statement added in Junos OS Release 8.2. Logical systems support introduced in Junos OS Release 9.3. The physical-interface-policer statement introduced in Junos OS Release 9.6. The shared-bandwidth-policer statement added in Junos OS Release 11.2. Support at the [edit dynamic-profiles ... firewall] hierarchy level introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure policer rate limits and actions. When included at the [edit firewall] hierarchy level, the policer statement creates a template, and you do not have to configure a policer individually for every firewall filter or interface. To activate a policer, you must include the policer-action modifier in the then statement in a firewall filter term or on an interface.
Options	<i>policer-action</i> —One or more actions to take: <ul style="list-style-type: none">• discard—Discard traffic that exceeds the rate limits.• forwarding-class <i>class-name</i>—Specify the particular forwarding class.• loss-priority—Set the packet loss priority (PLP) to low, medium-low, medium-high, or high.• out-of-profile—On J Series routers with strict priority queuing, prevent starvation of other queues by rate limiting the data stream entering the strict priority queue, marking the packets that exceed the rate limit as out-of-profile, and dropping the out-of-profile packets if the physical interface is congested.

policer-name—Name that identifies the policer. The name can contain letters, numbers, and hyphens (-), and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" "). Policer names cannot begin with an underscore in the form `__.*`.


then—Actions to take on matching packets.

The remaining statements are explained separately.

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

Related Documentation	<ul style="list-style-type: none">• Bandwidth Policer Overview on page 4854• Configuring Multifield Classifiers on page 1390• Logical Interface (Aggregate) Policer Overview on page 4928• Physical Interface Policer Overview on page 4940• Statement Hierarchy for Configuring Policers on page 4811• Single-Rate Two-Color Policer Overview on page 4838• Using Multifield Classifiers to Set PLP on page 1461• filter (Configuring) on page 1422• priority (Schedulers) on page 1641
------------------------------	--

shared-bandwidth-policer

Syntax	shared-bandwidth-policer;
Hierarchy Level	[edit firewall policer <i>policer-name</i>] [edit firewall three-color-policer <i>policer-name</i>] [edit firewall hierarchical-policer <i>policer-name</i>]
Release Information	Statement introduced in Junos OS Release 11.2. Support for MX Series MPC and MIC interfaces added in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Policer instances share bandwidth. This enables configuration of interface-specific policers applied on an aggregated Ethernet bundle or an aggregated SONET bundle to match the effective bandwidth and burst-size to user-configured values. This feature is supported on the following platforms: T Series routers, M120, M10i, M7i (CFEB-E only), M320 (SFPC only), MX240, MX480, and MX960 with DPC, MPC, and MIC interfaces, and EX Series switches.
	<div> NOTE: This statement is not supported on T4000 Type 5 FPCs.</div>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Policer Support for Aggregated Ethernet Bundle Overview on page 1446

term (Normal Filter)

Syntax	<pre> term <i>term-name</i> { from { <i>match-conditions</i>; } then { forwarding-class <i>class-name</i>; loss-priority (high low); three-color-policer { (single-rate two-rate) <i>policer-name</i>; } } } </pre>
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define a firewall filter term.
Options	<p>from—Match packet fields to values. If not included, all packets are considered to match and the actions and action modifiers in the then statement are taken.</p> <p>match-conditions—One or more conditions to use to make a match.</p> <p>term-name—Name that identifies the term. The name can contain letters, numbers, and hyphens (-), and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").</p> <p>then—Actions to take on matching packets. If not included and a packet matches all the conditions in the from statement, the packet is accepted. For CoS, only the actions listed are allowed. These statements are explained separately.</p>
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Multifield Classifiers on page 1390

then

Syntax then {
 application-profile *profile-name*;
 dscp (*alias* | *bits*);
 forwarding-class *class-name*;
 syslog;
 (reflexive | reverse) {
 application-profile *profile-name*;
 dscp (*alias* | *bits*);
 forwarding-class *class-name*;
 syslog;
 }
 }

Hierarchy Level [edit services cos rule *rule-name* term *term-name*]

Release Information Statement introduced in Junos OS Release 8.1.

Description Define the CoS term actions.

The remaining statements are explained separately.

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

Related Documentation • *Configuring Actions in a CoS Rule*
 • *Configuring Actions in CoS Rules*

three-color-policer (Applying)

Syntax	three-color-policer { (single-rate two-rate) <i>policer-name</i> ; }
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> then] [edit logical-systems <i>logical-system-name</i> firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> then]
Release Information	Statement introduced in Junos OS Release 7.4. single-rate statement added in Junos OS Release 8.2. Logical systems support introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	For M320 and T Series routers with Enhanced II Flexible PIC Concentrators (FPCs) and the T640 router with Enhanced Scaling FPC4, apply a tricolor marking policer.
Options	single-rate —Named tricolor policer is a single-rate policer. two-rate —Named tricolor policer is a two-rate policer. <i>policer-name</i> —Name of a tricolor policer.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Applying Tricolor Marking Policers to Firewall Filters on page 1456 • Standard Firewall Filter Nonterminating Actions on page 4744 • Three-Color Policer Configuration Overview on page 4817

three-color-policer (Configuring)

Syntax	<pre>three-color-policer <i>policer-name</i> { action { loss-priority high then discard; } filter-specific; logical-interface-policer; physical-interface-policer; shared-bandwidth-policer; single-rate { (color-aware color-blind); committed-burst-size <i>bytes</i>; committed-information-rate <i>bps</i>; excess-burst-size <i>bytes</i>; } two-rate { (color-aware color-blind); committed-burst-size <i>bytes</i>; committed-information-rate <i>bps</i>; peak-burst-size <i>bytes</i>; peak-information-rate <i>bps</i>; } }</pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall], [edit firewall], [edit logical-systems <i>logical-system-name</i> firewall]
Release Information	Statement introduced before Junos OS Release 7.4. The action and single-rate statements added in Junos OS Release 8.2. Logical systems support introduced in Junos OS Release 9.3. Support at the [edit dynamic-profiles ... firewall] hierarchy level introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure a three-color policer.
Options	<i>policer-name</i> —Name of the three-color policer. Reference this name when you apply the policer to an interface. The remaining statements are explained separately.
Required Privilege Level	firewall —To view this statement in the configuration. firewall-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Statement Hierarchy for Configuring Policers on page 4811• Configuring Tricolor Marking Policers on page 1455• Three-Color Policer Configuration Guidelines on page 4912• Basic Single-Rate Three-Color Policers on page 4915

- [Basic Two-Rate Three-Color Policers on page 4921](#)
- [Two-Color and Three-Color Logical Interface Policers on page 4927](#)
- [Two-Color and Three-Color Physical Interface Policers on page 4940](#)
- [Two-Color and Three-Color Policers at Layer 2 on page 4947](#)

[edit interfaces] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```

interfaces {
  interface-name {
    ... the "interface-name" subhierarchy appears after the main [edit interfaces] hierarchy
    level ...
  }
  interface-set interface-set-name {
    interface interface-name {
      (unit unit-number | vlan-tags-outer vlan-tag);
    }
  }
  irb (Interfaces) {
    accounting-profile name;
    description text;
    disable;

    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;

    traceoptions {
      flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
      accounting-profile name;
      bandwidth rate;
      description text;
      disable;
      encapsulation type;
      family inet {
        accounting {
          destination-class-usage;
          source-class-usage {
            input;
            output;
          }
        }
      }
      address ipv4-address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        preferred;
      }
    }
  }
}

```

```
primary;
vrrp-group group-id {
  (accept-data | no-accept-data);
  advertise-interval seconds;
  advertisements-threshold number;
  authentication-key key;
  authentication-type authentication;
  fast-interval milliseconds;
  (preempt | no-preempt) {
    hold-time seconds;
  }
  priority number;
  track {
    interface interface-name {
      bandwidth-threshold bits-per-second priority-cost priority;
      priority-cost priority;
    }
    priority-hold-time seconds;
    route prefix/prefix-length routing-instance instance-name priority-cost priority;
  }
  virtual-address [ addresses ];
  vrrp-inherit-from vrrp-group;
}
}
filter {
  input filter-name;
  output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
  fail-filter filter-name;
  mode {
    loose;
  }
}
targeted-broadcast {
  forward-and-send-to-re;
  forward-only;
}
}
family inet6 {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
}
address address {
  eui-64;
  ndp ip-address (mac | multicast-mac) mac-address <publish>;
  preferred;
  primary;
```

```

vrp-inet6-group group-id {
  accept-data | no-accept-data;
  advertisements-threshold number;
  authentication-key key;
  authentication-type authentication;
  fast-interval milliseconds;
  inet6-advertise-interval milliseconds;
  preempt | no-preempt {
    hold-time seconds;
  }
  priority number;
  track {
    interface interface-name {
      bandwidth-threshold bandwidth priority-cost number;
      priority-cost number;
    }
    priority-hold-time seconds;
    route ip-address/mask routing-instance instance-name priority-cost cost;
  }
  virtual-inet6-address [addresses];
  virtual-link-local-address ipv6-address;
  vrrp-inherit-from {
    active-group group-number;
    active-interface interface-name;
  }
}
(dad-disable | no-dad-disable);
filter {
  input filter-name;
  output filter-name;
}
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
no-redirects;
policer {
  input policer-name;
  output policer-name;
}
rpf-check {
  fail-filter filter-name;
  mode {
    loose;
  }
}
}
family iso {
  address interface-address;
  mtu bytes;
}
family mpls {
  filter {
    input filter-name;
    output filter-name;
  }
}

```

```
    mtu bytes;
    policer {
        input policer-name;
        output policer-name;
    }
}
native-inner-vlan-id vlan-id;
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id-list [vlan-id's];
vlan-id-range [vlan-id-range];
}
}
traceoptions {
    file <filename> <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag <disable>;
    no-remote-trace;
}
}

interfaces {
    interface-name {
        disable;
        accounting-profile name;
        aggregated-ether-options {
            ethernet-switch-profile {
                tag-protocol-id [ hexadecimal-identifiers ];
            }
            (flow-control | no-flow-control);
            lacp {
                (active | passive);
                admin-key key;
                fast-failover;
                link-protection {
                    disable;
                    (revertive | non-revertive);
                }
                periodic (fast | slow);
                system-id mac-address;
                system-priority priority;
            }
            (link-protection | no-link-protection);
            link-speed (100m | 1g | 8g | 10g | 40g | 50g | 80g | 100g | oc192);
            logical-interface-fpc-redundancy;
            (loopback | no-loopback);
            mc-ae {
                chassis-id chassis-id;
                events {
                    iccp-peer-down {
                        force-icl-down;
                        prefer-status-control-active;
                    }
                }
            }
            mc-ae-id mc-ae-id;
            mode (active-active | active-standby);
        }
    }
}
```

```

    redundancy-group group-id;
    status-control (active | standby);
}
minimum-links number;
rebalance-periodic {
    start-time time;
    interval number;
}
source-address-filter {
    mac-address;
}
(source-filtering | no-source-filtering);
}
auto-configure {
    remove-when-no-subscribers;
    stacked-vlan-ranges {
        access-profile profile-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                mac-address;
                option-82 ( circuit-id | remote-id);
                radius-realm radius-realm-string;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name {
            accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
            ranges (any | low-tag-high-tag), (any | low-tag-high-tag);
        }
    }
}
vlan-ranges {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-name;
            mac-address;
            option-82;
            radius-realm radius-realm-string;
            user-prefix user-prefix-string;
        }
    }
    dynamic-profile profile-name {
        accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
        ranges (any | low-tag)—(any | high-tag);
    }
}
override tag vlan-tag dynamic-profile profile name;

```

```

}
encapsulation (ethernet-bridge | ethernet-vpls | extended-vlan-bridge |
    extended-vlan-vpls | flexible-ethernet-services | vlan-vpls);
ether-options {
    802.3ad {
        aex;
        (backup | primary);
        lacp {
            force-up;
            port-priority
        }
    }
    asynchronous-notification;
    (auto-negotiation | no-auto-negotiation);
    ethernet-switch-profile {
        ethernet-policer-profile {
            input-priority-map {
                ieee802.1p premium [ values ];
            }
            output-priority-map {
                classifier {
                    premium {
                        forwarding-class class-name {
                            loss-priority (high | low);
                        }
                    }
                }
            }
        }
        policer cos-policer-name {
            aggregate {
                bandwidth-limit bps;
                burst-size-limit bytes;
            }
            premium {
                bandwidth-limit bps;
                burst-size-limit bytes;
            }
        }
        tag-protocol-id;
    }
    (mac-learn-enable | no-mac-learn-enable);
}
(flow-control | no-flow-control);
ignore-l3-incompletes;
link-mode (automatic | full-duplex | half-duplex);
(loopback | no-loopback);
keepalives <interval seconds> <down-count number> <up-count number>;
speed (1g | 10m | 100m | 10m-100m | auto-negotiation);
source-address-filter {
    mac-address;
}
source-filtering | no-source-filtering;
}
flexible-vlan-tagging;
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time (up milliseconds | down milliseconds);

```



```

interface-transmit-statistics;
(keepalives <down-count number> <interval seconds> <up-count number> |
  no-keepalives);
layer2-policer {
  apply-groups [ group-names ];
  apply-groups-except [ group-names ];
}
link-mode (automatic | full-duplex);
mac mac-address;
mtu bytes;
multi-chassis-protection peer-ip-address {
  interface interface-name;
}
native-vlan-id number;
no-gratuitous-arp-request;
optics-options {
  alarm low-light-alarm {
    (link-down | syslog);
  }
  warning low-light-warning {
    (link-down | syslog);
  }
  wavelength nm;
}
passive-monitor-mode;
per-unit-scheduler;
speed (10m | 100m | 1g | auto | oc3 | oc12 | oc48);
stacked-vlan-tagging;
traceoptions {
  flag flag;
}
transmit-bucket {
  overflow discard;
  rate percentage;
  threshold bytes;
}
(traps | no-traps);
unidirectional;
vlan-tagging;
}

interface-name {
  unit logical-unit-number {
    disable;
    accept-source-mac {
      mac-address mac-address {
        policer {
          input policer-name;
          output policer-name;
        }
      }
    }
  }
  account-layer2-overhead (Interface Level) {
    value;
    egress bytes;
  }
}

```

```

    ingress bytes;
}
accounting-profile name;
advisory-options {
    downstream-rate rate;
    upstream-rate rate;
}
arp-resp (restricted|unrestricted);
bandwidth rate;
clear-dont-fragment-bit;
copy-tos-to-outer-ip-header;
demux-destination family;
encapsulation (vlan-bridge | vlan-vpls);
epd-threshold cells plp1 cells;
filter filter-name;
inner-vlan-id-range start start-id end end-id;
input-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
interface-shared-with psd numerical-index;
layer2-policer {
    input-hierarchical-policer policer-name;
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
}
multi-chassis-protection peer-ip-address {
    interface interface-name;
}
native-inner-vlan-id number;
output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
peer-interface interface-name;
peer-unit unit-number;
plp-to-clp;
proxy-arp <restricted | unrestricted>;
rpm {
    (client | server);
    twamp-server;
}
swap-by-poppush;
vlan-id number;
vlan-id-list [ vlan-id vlan-id-vlan-id ];
vlan-id-range number-number;
vlan-tags (inner <tpid.>vlan-id | inner-list [ vlan-id vlan-id-vlan-id ] |
    inner-range <tpid.>vlan-id-vlan-id) outer <tpid.>vlan-id;

```

```

}

unit logical-unit-number {
  family ethernet-switching {
    filter {
      group filter-group-number;
      (input filter-name | input-list [ filter-names ]);
      (output filter-name | output-list [ filter-names ]);
      (inner-vlan-id-list [ vlan-ids ] | vlan-id number | vlan-id-list [ number
        number-number ]);
      interface-mode (access | trunk);
      policer {
        input policer-name;
        output policer-name;
      }
      vlan-rewrite {
        translate old-vlan-id new-vlan-id;
      }
      vlan {
        members [ all vlan-identifiers ];
      }
    }
  }
  family inet {
    filter {
      group filter-group-number;
      (input filter-name | input-list [ filter-names ]);
      (output filter-name | output-list [ filter-names ]);
    }
    input-hierarchical-policer policer-name;
    mac-validate (loose | strict);
    mtu bytes;
    no-neighbor-learn;
    no-redirects;
    policer {
      arp policer-template-name;
      input policer-name;
      output policer-name;
    }
    primary;
    receive-options-packets;
    receive-ttl-exceeded;
    rpf-check {
      fail-filter filter-name;
      mode loose;
    }
    sampling {
      (input | output | input output);
    }
    simple-filter {
      input filter-name;
    }
    targeted-broadcast {
      forward-and-send-to-re;
      forward-only;
    }
  }
}

```

```

unnumbered-address interface-name <destination address>
    <destination-profile profile-name> <preferred-source-address address>;

}

family inet6 {
    address ipv6-address {
        destination destination-address;
        eui-64;
        ndp ipv6-address <l2-interface interface-name> <(mac mac-address |
            multicast-mac multicast-mac-address) <publish>>;
        preferred;
        primary;
        vrrp-inet6-group group-number {
            (accept-data | no-accept-data);
            fast-interval milliseconds;
            inet6-advertise-interval seconds;
            (no-preempt; | ... the following preempt statement ...)
            preempt {
                hold-time seconds;
            }
            priority number;
            track {
                interface interface-name {
                    bandwidth-threshold bits-per-second priority-cost priority;
                    priority-cost priority;
                }
                priority-hold-time seconds;
                route ip-address-prefix/prefix-length routing-instance instance-name
                    priority-cost priority;
            }
            virtual-inet6-address [ addresses ];
            virtual-link-local-address ipv6-address;
            vrrp-inherit-from {
                active-group group-number;
                active-interface interface-name;
            }
        }
    }
    (dad-disable | no-dad-disable);
    filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
    input-hierarchical-policer policer-name;
    mtu bytes;
    nd6-stale-time seconds;
    no-neighbor-learn;
    policer {
        input policer-name;
        output policer-name;
    }
    rpf-check {
        fail-filter filter-name;
        mode loose;
    }
}

```

```

    }
    sampling {
        (input | output | input output);
    }
    unnumbered-address interface-name preferred-source-address address;
}

family iso {
    address iso-address;
    mtu bytes;
}

family mlfrr-end-to-end {
    bundle logical-interface-name;
}

family mpls {
    filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
    input-hierarchical-policer policer-name;
    maximum-labels maximum-labels;
    mtu bytes;
    policer {
        input policer-name;
        output policer-name;
    }
}

family vpls {
    core-facing;
    filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
    policer {
        input policer-name;
        output policer-name;
    }
}
}
}
}

```

Related Documentation • *Notational Conventions Used in Junos OS Configuration Hierarchies*

filter (Applying to an Interface)

Syntax	<pre>filter { input <i>filter-name</i>; output <i>filter-name</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Apply a filter to an interface. You can also use filters for encrypted traffic. When you configure filters, you can configure the family inet , inet6 , mpls , or vpls only.
Options	<p>input <i>filter-name</i>—Name of one filter to evaluate when packets are received on the interface.</p> <p>output <i>filter-name</i>—Name of one filter to evaluate when packets are transmitted on the interface.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• simple-filter on page 1440• <i>Applying Firewall Filter Tricolor Marking Policers to Interfaces</i>• <i>Example: Classifying Packets Based on Their Destination Address</i>• Example: Configuring and Verifying a Complex Multifield Filter on page 1394• <i>Example: Writing Different DSCP and EXP Values in MPLS-Tagged IP Packets</i>• Example: Configuring a Simple Filter on page 1399• Example: Configuring a Logical Bandwidth Policer on page 1400• Example: Two-Color Policers and Shaping Rate Changes on page 1401

input-policer

Syntax	<code>input-policer <i>policer-name</i>;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer]</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer]</code>
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Apply a single-rate two-color policer to the Layer 2 input traffic at the logical interface. The input-policer and input-three-color statements are mutually exclusive.
Options	<i>policer-name</i> —Name of the single-rate two-color policer that you define at the [edit firewall] hierarchy level.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Two-Color and Three-Color Policers at Layer 2 on page 4947 • Applying Layer 2 Policers to Gigabit Ethernet Interfaces on page 1459 • Configuring a Gigabit Ethernet Policer • input-three-color on page 1522 • layer2-policer on page 1523 • logical-interface-policer on page 1500 • output-policer on page 1524 • output-three-color on page 1525

input-three-color

Syntax	<code>input-three-color <i>policer-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer] [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Apply a single-rate or two-rate three-color policer to the Layer 2 input traffic at the logical interface. The input-three-color and input-policer statements are mutually exclusive.
Options	<i>policer-name</i> —Name of the single-rate or two-rate three-color policer.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Two-Color and Three-Color Policers at Layer 2 on page 4947• Applying Layer 2 Policers to Gigabit Ethernet Interfaces on page 1459• Configuring a Gigabit Ethernet Policer• input-policer on page 1521• layer2-policer on page 1523• logical-interface-policer on page 1500• output-policer on page 1524• output-three-color on page 1525

layer2-policer

Syntax	<pre> layer2-policer { input-policer <i>policer-name</i>; input-three-color <i>policer-name</i>; output-policer <i>policer-name</i>; output-three-color <i>policer-name</i>; } </pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</p>
Release Information	<p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>For 1-Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series, MX Series, and T Series routers, and for aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces on EX Series switches, apply Layer 2 logical interface policers. The following policers are supported:</p> <ul style="list-style-type: none"> • Two-color • Single-rate tricolor marking (srTCM) • Two-rate tricolor marking (trTCM) <p>Two-color and tricolor policers are configured at the [edit firewall] hierarchy level.</p>
Options	<p>input-policer <i>policer-name</i>—Two-color input policer to associate with the interface. This statement is mutually exclusive with the input-three-color statement.</p> <p>input-three-color <i>policer-name</i>—Tricolor input policer to associate with the interface. This statement is mutually exclusive with the input-policer statement.</p> <p>output-policer <i>policer-name</i>—Two-color output policer to associate with the interface. This statement is mutually exclusive with the output-three-color statement.</p> <p>output-three-color <i>policer-name</i>—Tricolor output policer to associate with the interface. This statement is mutually exclusive with the output-policer statement.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Applying Layer 2 Policers to Gigabit Ethernet Interfaces on page 1459 • <i>Configuring Gigabit Ethernet Two-Color and Tricolor Policers</i>

output-policer

Syntax	<code>output-policer <i>policer-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Apply a single-rate two-color policer to the Layer 2 output traffic at the logical interface. The output-policer and output-three-color statements are mutually exclusive.
Options	<i>policer-name</i> —Name of the single-rate two-color policer that you define at the [edit firewall] hierarchy level.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Two-Color and Three-Color Policers at Layer 2 on page 4947• Applying Layer 2 Policers to Gigabit Ethernet Interfaces on page 1459• Configuring a Gigabit Ethernet Policer• input-policer on page 1521• input-three-color on page 1522• layer2-policer on page 1523• logical-interface-policer on page 1500• output-three-color on page 1525

output-three-color

Syntax	<code>output-three-color <i>policer-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer] [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Apply a single-rate or two-rate three-color policer to the Layer 2 output traffic at the logical interface. The output-three-color and output-policer statements are mutually exclusive.
Options	<i>policer-name</i> —Name of the single-rate or two-rate three-color policer.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Two-Color and Three-Color Policers at Layer 2 on page 4947 • Applying Layer 2 Policers to Gigabit Ethernet Interfaces on page 1459 • Configuring a Gigabit Ethernet Policer • input-three-color on page 1522 • input-policer on page 1521 • layer2-policer on page 1523 • logical-interface-policer on page 1500 • output-policer on page 1524

Forwarding Classes

- [Overview on page 1525](#)
- [Configuration on page 1530](#)

Overview

- [Forwarding Classes on page 1525](#)

Forwarding Classes

- [Overview of Forwarding Classes on page 1526](#)
- [Default Forwarding Classes on page 1528](#)

Overview of Forwarding Classes

This topic covers the following information:

- [Output Queue Assignments Based on Forwarding Class on page 1526](#)
- [Devices That Support Up to Four Forwarding Classes on page 1526](#)
- [Devices That Support Up to 16 Forwarding Classes on page 1527](#)
- [Default and Configurable Packet Loss Priority Values on page 1527](#)
- [Configuration Statements Used to Configure and Apply Forwarding Classes on page 1527](#)

Output Queue Assignments Based on Forwarding Class

It is helpful to think of forwarding classes as output queues. In effect, the end result of classification is the identification of an output queue for a particular packet.

CoS packet classification assigns an incoming packet to an output queue based on the packet's forwarding class. Each packet is associated with one of the following default forwarding classes:

- Expedited forwarding (EF)—Provides a low-loss, low-latency, low-jitter, assured bandwidth, end-to-end service.
- Assured forwarding (AF)—Provides a group of values you can define and includes four subclasses: AF1, AF2, AF3, and AF4, each with three drop probabilities: low, medium, and high.
- Best effort (BE)—Provides no service profile. For the best effort forwarding class, loss priority is typically not carried in a class-of-service (CoS) value and random early detection (RED) drop profiles are more aggressive.
- Network control (NC)—This class is typically high priority because it supports protocol control.

Devices That Support Up to Four Forwarding Classes

Some of the Juniper Networks routing platforms support up to four forwarding classes for classifying customer traffic. On these platforms, you can configure one of each type of default forwarding class. The following Juniper Networks platforms support up to four forwarding classes:

- M7i Multiservice Edge Routers with Compact Forwarding Engine Boards (CFEBs)
- M10i Multiservice Edge Routers with CFEBs



NOTE: This list does not reference any Juniper Networks device that has reached its End of Life (EOL) period and its End of Support (EOS) milestone date.

Devices That Support Up to 16 Forwarding Classes

Other Juniper Networks routing platforms support up to 16 forwarding classes, which enables you to classify packets more granularly. For example, you can configure multiple classes of EF traffic: EF, EF1, and EF2. On these platforms, the Junos OS software supports up to eight output queues; therefore, if you configure more than eight forwarding classes, you must map multiple forwarding classes to single output queues. The following Juniper Networks routing and switching platforms support up to 16 forwarding classes and up to 8 output queues:

- EX Series switches
- M7i Multiservices Edge Routers with Enhanced Compact Forwarding Engine Boards (CFEB-Es)
- M10i Multiservices Edge Routers with CFEB-Es
- M120 Multiservices Edge Routers
- M320 Multiservices Edge Routers
- MX Series 3D Universal Edge Routers
- T Series Core Routers
- PTX Packet Transport Switches

Default and Configurable Packet Loss Priority Values

By default, the loss priority is low. On most devices, you can configure high or low loss priority. On the following devices, you can configure high, low, medium-high, or medium-low loss priority:

- J Series Services Router interfaces
- M320 routers and T Series routers with Enhanced II Flexible PIC Concentrators (FPCs)
- T640 routers with Enhanced Scaling FPC4s
- PTX Series Packet Transport Switches

Configuration Statements Used to Configure and Apply Forwarding Classes

To configure CoS forwarding classes, include the **forwarding-classes** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
forwarding-classes {
  class class-name queue-num queue-number priority (high | low);
  queue queue-number class-name priority (high | low);
}
forwarding-classes-interface-specific forwarding-class-map-name {
  class class-name queue-num queue-number [ restricted-queue queue-number ];
}
interfaces {
  interface-name {
    unit logical-unit-number {
      forwarding-class class-name;
```

```

        forwarding-classes-interface-specific forwarding-class-map-name;
    }
}
restricted-queues {
    forwarding-class class-name queue queue-number;
}

```

Related Documentation

- *Default Forwarding Classes*
- [Configuring Forwarding Classes on page 1530](#)
- [Applying Forwarding Classes to Interfaces on page 1531](#)
- *Configuring Up to 16 Forwarding Classes*
- *Policer Overview*

Default Forwarding Classes

By default, four queues are assigned to four forwarding classes, each with a queue number, name, and abbreviation.

These default mappings apply to all routers. The four forwarding classes defined by default are shown in [Table 107 on page 1528](#).

If desired, you can rename the forwarding classes associated with the queues supported on your hardware. Assigning a new class name to an output queue does not alter the default classification or scheduling that is applicable to that queue. CoS configurations can be quite complicated, so unless it is required by your scenario, we recommend that you not alter the default class names or queue number associations.

Some routers support eight queues. Queues 4 through 7 have no default mappings to forwarding classes. To use queues 4 through 7, you must create custom forwarding class names and map them to the queues. For more information, see the Juniper Networks J Series Services Router documentation.

Table 107: Default Forwarding Classes

Queue	Forwarding Class Name	Comments
Queue 0	best-effort (be)	The software does not apply any special CoS handling to packets with 000000 in the DiffServ field, a backward compatibility feature. These packets are usually dropped under congested network conditions.
Queue 1	expedited-forwarding (ef)	<p>The software delivers assured bandwidth, low loss, low delay, and low delay variation (jitter) end-to-end for packets in this service class.</p> <p>Routers accept excess traffic in this class, but in contrast to assured forwarding, out-of-profile expedited-forwarding packets can be forwarded out of sequence or dropped.</p>

Table 107: Default Forwarding Classes (*continued*)

Queue	Forwarding Class Name	Comments
Queue 2	assured-forwarding (af)	<p>The software offers a high level of assurance that the packets are delivered as long as the packet flow from the customer stays within a certain service profile that you define.</p> <p>The software accepts excess traffic, but applies a RED drop profile to determine if the excess packets are dropped and not forwarded.</p> <p>Depending on router type, up to four drop probabilities (low, medium-low, medium-high, and high) are defined for this service class.</p>
Queue 3	network-control (nc)	<p>The software delivers packets in this service class with a low priority. (These packets are not delay sensitive.)</p> <p>Typically, these packets represent routing protocol hello or keepalive messages. Because loss of these packets jeopardizes proper network operation, delay is preferable to discard.</p>

The following rules govern queue assignment:

- If classifiers fail to classify a packet, the packet always receives the default classification to the class associated with queue 0.
- The number of queues is dependent on the hardware plugged into the chassis. CoS configurations are inherently contingent on the number of queues on the system. Only two classes, **best-effort** and **network-control**, are referenced in the default configuration. The default configuration works on all routers.
- CoS configurations that specify more queues than the router can support are not accepted. The commit fails with a detailed message that states the total number of queues available.
- All default CoS configuration is based on queue number. The name of the forwarding class that shows up when the default configuration is displayed is the forwarding class currently associated with that queue.

This is the default configuration for the **forwarding-classes** statement:

```
[edit class-of-service]
forwarding-classes {
  queue 0 best-effort;
  queue 1 expedited-forwarding;
  queue 2 assured-forwarding;
  queue 3 network-control;
}
```

If you reassign the forwarding-class names, the **best-effort** forwarding-class name appears in the locations in the configuration previously occupied by **network-control** as follows:

```
[edit class-of-service]
forwarding-classes {
  queue 0 network-control;
```

```
queue 1 assured-forwarding;  
queue 2 expedited-forwarding;  
queue 3 best-effort;  
}
```

All the default rules of classification and scheduling that applied to Queue 3 still apply. Queue 3 is simply now renamed **best-effort**.

On Juniper Networks M320 Multiservice Edge Routers and T Series Core Routers, you can assign multiple forwarding classes to a single queue. If you do so, the first forwarding class that you assign to queue 0 acquires the default BE classification and scheduling. The first forwarding class that you assign to queue 1 acquires the default EF classification and scheduling. The first forwarding class that you assign to queue 2 acquires the default AF classification and scheduling. The first forwarding class that you assign to queue 3 acquires the default NC classification and scheduling. For more information, see *Configuring Up to 16 Forwarding Classes*.

- In the current default configuration:
 - Only IP precedence classifiers are associated with interfaces.
 - The only classes designated are **best-effort** and **network-control**.
 - Schedulers are not defined for the **expedited-forwarding** or **assured-forwarding** forwarding classes.
- You must explicitly classify packets to the **expedited-forwarding** or **assured-forwarding** forwarding class and define schedulers for these classes.

For more information, see *Hardware Capabilities and Limitations*.

Configuration

- [Configuration Tasks on page 1530](#)
- [Configuration Statements on page 1542](#)

Configuration Tasks

- [Configuring Forwarding Classes on page 1530](#)
- [Applying Forwarding Classes to Interfaces on page 1531](#)
- [Classifying Packets by Egress Interface on page 1532](#)
- [Assigning Forwarding Class and DSCP Value for Routing Engine-Generated Traffic on page 1534](#)
- [Overriding Fabric Priority Queuing on page 1535](#)
- [Configuring Up to 16 Forwarding Classes on page 1535](#)

Configuring Forwarding Classes

You assign each forwarding class to an internal queue number by including the **forwarding-classes** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]  
forwarding-classes {  
  class queue-num queue-number priority (high | low);
```



```

queue queue-number class-name priority (high | low) [ policing-priority (premium |
normal) ];
}

```

You cannot commit a configuration that assigns the same forwarding class to two different queues.



CAUTION: We do not recommend classifying packets into a forwarding class that has no associated scheduler on the egress interface. Such a configuration can cause unnecessary packet drops because an unconfigured scheduling class might lack adequate buffer space. For example, if you configure a custom scheduler map that does not define queue 0, and the default classifier assigns incoming packets to the best-effort class (queue 0), the unconfigured egress queue for the best-effort forwarding class might not have enough space to accommodate even short packet bursts.

A default congestion and transmission control mechanism is used when an output interface is not configured for a certain forwarding class, but receives packets destined for that unconfigured forwarding class. This default mechanism uses the delay buffer and weighted round robin (WRR) credit allocated to the designated forwarding class, with a default drop profile. Because the buffer and WRR credit allocation is minimal, packets might be lost if a larger number of packets are forwarded without configuring the forwarding class for the interface.

Applying Forwarding Classes to Interfaces

You can configure *fixed classification* on a logical interface by specifying a forwarding class to be applied to all packets received by the logical interface, regardless of the packet contents.



NOTE: On the T4000 router, BA classification and fixed classification are mutually exclusive. That is, only one of the classifications can be configured.

To apply a forwarding class configuration to the input logical interface, include the **forwarding-class** statement at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number*] hierarchy level:

```

[edit class-of-service interfaces interface-name unit logical-unit-number]
  forwarding-class class-name;

```

You can include interface wildcards for *interface-name* and *logical-unit-number*.

In the following example, all packets coming into the router from the **ge-3/0/0.0** interface are assigned to the **assured-forwarding** forwarding class:

```

[edit class-of-service]
interfaces {
  ge-3/0/0 {
    unit 0 {

```

```

        forwarding-class assured-forwarding;
    }
}

```

Related Documentation

- [forwarding-class on page 1547](#)

Classifying Packets by Egress Interface

For Juniper Networks EX Series switches, M320 Multiservice Edge Routers, and T Series Core Routers with the Intelligent Queuing (IQ), IQ2, Enhanced IQ (IQE), or Multiservices link services intelligent queuing (LSQ) interfaces,, you can classify unicast and multicast packets based on the egress interface. For unicast traffic, you can also use a multifield filter, but only egress interface classification applies to multicast traffic as well as unicast traffic. If you configure egress classification of an interface, you cannot perform Differentiated Services code point (DSCP) rewrites on the interface. By default, the system will not perform any classification based on the egress interface.

To enable packet classification by the egress interface, you first configure a forwarding class map and one or more queue numbers for the egress interface at the **[edit class-of-service forwarding-classes-interface-specific forwarding-class-map-name]** hierarchy level:

```

[edit class-of-service]
  forwarding-classes-interface-specific forwarding-class-map-name {
    class class-name queue-num queue-number [ restricted-queue queue-number ];
  }

```

For T Series routers that are restricted to only four queues, you can control the queue assignment with the **restricted-queue** option, or you can allow the system to automatically determine the queue in a modular fashion. For example, a map assigning packets to queue 6 would map to queue 2 on a four-queue system.



NOTE: If you configure an output forwarding class map associating a forwarding class with a queue number, this map is not supported on multiservices link services intelligent queuing (lsq-) interfaces.

Once the forwarding class map has been configured, you apply the map to the logical interface by using the **output-forwarding-class-map** statement at the **[edit class-of-service interfaces interface-name unit logical-unit-number]** hierarchy level:

```

[edit class-of-service interfaces interface-name unit logical-unit-number]
  output-forwarding-class-map forwarding-class-map-name;

```

All parameters relating to the queues and forwarding class must be configured as well. For more information about configuring forwarding classes and queues, see [“Configuring Forwarding Classes” on page 1530](#).

This example shows how to configure an interface-specific forwarding-class map named **FCMAP1** that restricts queues 5 and 6 to different queues on four-queue systems and then applies **FCMAP1** to **unit 0** of interface **ge-6/0/0**:

```

[edit class-of-service]

```

```

forwarding-classes-interface-specific FCMAP1 {
  class FC1 queue-num 6 restricted-queue 3;
  class FC2 queue-num 5 restricted-queue 2;
  class FC3 queue-num 3;
  class FC4 queue-num 0;
  class FC3 queue-num 0;
  class FC4 queue-num 1;
}

```

```

[edit class-of-service]
interfaces {
  ge-6/0/0 unit 0 {
    output-forwarding-class-map FCMAP1;
  }
}

```

Note that without the **restricted-queue** option in **FCMAP1**, the example would assign **FC1** and **FC2** to queues 2 and 1, respectively, on a system restricted to four queues.

Use the **show class-of-service forwarding-class *forwarding-class-map-name*** command to display the forwarding-class map queue configuration:

```
user@host> show class-of-service forwarding-class FCMAP2
```

Forwarding class	ID	Queue	Restricted queue
FC1	0	6	3
FC2	1	5	2
FC3	2	3	3
FC4	3	0	0
FC5	4	0	0
FC6	5	1	1
FC7	6	6	2
FC8	7	7	3

Use the **show class-of-service interface *interface-name*** command to display the forwarding-class maps (and other information) assigned to a logical interface:

```
user@host> show class-of-service interface ge-6/0/0
```

```

Physical interface: ge-6/0/0, Index: 128
Queues supported: 8, Queues in use: 8
Scheduler map: <default>, Index: 2
Input scheduler map: <default>, Index: 3
Chassis scheduler map: <default-chassis>, Index: 4

```

```

Logical interface: ge-6/0/0.0, Index: 67
Object      Name      Type      Index
Scheduler-map  sch-map1  Output    6998
Scheduler-map  sch-map1  Input     6998
Classifier     dot1p     ieee8021p 4906
forwarding-class-map  FCMAP1    Output    1221

```

```

Logical interface: ge-6/0/0.1, Index 68
Object      Name      Type      Index
Scheduler-map  <default>  Output    2
Scheduler-map  <default>  Input     3

```

```
Logical interface: ge-6/0/0.32767, Index 69
```

Object	Name	Type	Index
Scheduler-map	<default>	Output	2
Scheduler-map	<default>	Input	3

Assigning Forwarding Class and DSCP Value for Routing Engine–Generated Traffic

You can set the forwarding class and differentiated service code point (DSCP) value for traffic originating in the Routing Engine. To configure forwarding class and DSCP values that apply to Routing Engine–generated traffic only, apply an output filter to the loopback (lo.0) interface and set the appropriate forwarding class and DSCP bit configuration for various protocols. For example, you can set the DSCP value on OSPF packets that originate in the Routing Engine to **10** and assign them to the AF (assured forwarding) forwarding class while the DSCP value on ping packets are set to **0** and use forwarding class BE (best effort).

This particular classification ability applies to packets generated by the Routing Engine only.

The following example assigns Routing Engine sourced ping packets (using ICMP) a DSCP value of **38** and a forwarding class of **af17**, OSPF packets a DSCP value of **12** and a forwarding class of **af11**, and BGP packets (using TCP) a DSCP value of **10** and a forwarding class of **af16**.

```
[edit class-of-service]
forwarding-classes {
  class af11 queue-num 7;
  class af12 queue-num 1;
  class af13 queue-num 2;
  class af14 queue-num 4;
  class af15 queue-num 5;
  class af16 queue-num 4;
  class af17 queue-num 6;
  class af18 queue-num 7;
}

[edit firewall filter family inet]
filter loopback-filter {
  term t1 {
    from {
      protocol icmp; # For pings
    }
    then {
      forwarding-class af17;
      dscp 38;
    }
  }
  term t2 {
    from {
      protocol ospf; # For OSPF
    }
    then {
      forwarding-class af11;
      dscp 12;
    }
  }
}
```

```

term t3 {
  from {
    protocol tcp; # For BGP
  }
  then {
    forwarding-class af16;
    dscp 10;
  }
}
term t4 {
  then accept; # Do not forget!
}
}

[edit interfaces]
lo0 {
  unit 0 {
    family inet {
      filter {
        output loopback-filter;
      }
    }
  }
}

```



NOTE: This is not a complete router configuration. You still have to assign resources to the queues, configure the routing protocols, addresses, and so on.

Overriding Fabric Priority Queuing

On EX Series switches, and on M320 and T Series routers, the default behavior is for fabric priority queuing on egress interfaces to match the scheduling priority you assign. High-priority egress traffic is automatically assigned to high-priority fabric queues. Likewise, low-priority egress traffic is automatically assigned to low-priority fabric queues.

You can override the default fabric priority queuing of egress traffic by including the **priority** statement at the **[edit class-of-service forwarding-classes queue queue-number class-name]** hierarchy level:

```

[edit class-of-service forwarding-classes queue queue-number class-name]
priority (high | low);

```

For information about associating a scheduler with a fabric priority, see [“Associating Schedulers with Fabric Priorities” on page 1612](#).

Configuring Up to 16 Forwarding Classes

By default on all routers and switches, four output queues are mapped to four forwarding classes, as shown in the topic *Default Forwarding Classes*. On Juniper Networks J Series Services Routers, M120 and M320 Multiservice Edge Routers, and T Series Core Routers, you can configure more than four forwarding classes and queues. For information about configuring J Series routers, see the J Series router documentation.



NOTE: You cannot use CoS-based forwarding features if you configure more than eight forwarding classes on the device.

On M120, M320, MX Series, T Series routers, EX Series switches and PTX Series Packet Transport Switches, you can configure up to 16 forwarding classes and eight queues, with multiple forwarding classes assigned to single queues. The concept of assigning multiple forwarding classes to a queue is sometimes referred to as creating *forwarding-class aliases*. This section explains how to configure M320 and T Series routers.

Mapping multiple forwarding classes to single queues is useful. Suppose, for example, that forwarding classes are set based on multifield packet classification, and the multifield classifiers are different for core-facing interfaces and customer-facing interfaces. Suppose you need four queues for a core-facing interface and five queues for a customer-facing interface, where **fc0** through **fc4** correspond to the classifiers for the customer-facing interface, and **fc5** through **fc8** correspond to classifiers for the core-facing interface, as shown in [Figure 11 on page 1536](#).

Figure 11: Customer-Facing and Core-Facing Forwarding Classes



9016702

In this example, there are nine classifiers and, therefore, nine forwarding classes. The forwarding class-to-queue mapping is shown in [Table 108 on page 1536](#).

Table 108: Sample Forwarding Class-to-Queue Mapping

Forwarding Class Names	Queue Number
fc0	0
fc5	
fc1	1
fc6	
fc2	2
fc7	
fc3	3
fc8	
fc4	4

To configure up to 16 forwarding classes, include the **class** and **queue-num** statements at the **[edit class-of-service forwarding-classes]** hierarchy level:

```
[edit class-of-service forwarding-classes]
class class-name queue-num queue-number;
```

You can configure up to 16 different forwarding-class names. The corresponding output queue number can be from 0 through 7. Therefore, you can map multiple forwarding classes to a single queue. If you map multiple forwarding classes to a queue, the multiple forwarding classes must refer to the same scheduler (at the **[edit class-of-service scheduler-maps *map-name* forwarding-class *class-name* scheduler *scheduler-name*]** hierarchy level).

When you configure up to 16 forwarding classes, you can use them as you can any other forwarding class—in classifiers, schedulers, firewall filters (multifield classifiers), policers, and rewrite rules.

When you configure up to 16 forwarding classes, the following limitations apply:

- The **class** and **queue** statements at the **[edit class-of-service forwarding-classes]** hierarchy level are mutually exclusive. In other words, you can include one or the other of the following configurations, but not both:

```
[edit class-of-service forwarding-classes]
queue queue-number class-name;
```

```
[edit class-of-service forwarding-classes]
class class-name queue-num queue-number;
```

- On T Series routers only, when you configure IEEE 802.1p rewrite marking on Gigabit Ethernet IQ, Gigabit Ethernet IQ2, Gigabit Ethernet Enhanced IQ (IQE), and Gigabit Ethernet Enhanced IQ2 (IQ2E) PICs, you cannot configure more than eight forwarding classes. This limitation does not apply to M Series routers. On M Series routers, you can configure up to 16 forwarding classes when you configure IEEE 802.1p rewrite marking on any of these PICs.
- For GRE and IP-IP tunnels, IP precedence and DSCP rewrite marking of the inner header do not work with more than eight forwarding classes.
- When you use CoS-based forwarding features, you cannot configure more than eight forwarding classes with a forwarding policy. However, if you try to configure CoS-based forwarding with more than eight forwarding classes configured, commit fails with a message. Therefore, you can configure CBF on a router with eight or less than eight forwarding classes only. Under this condition, the forwarding class to queue mapping can be either one-to-one or one-to-many.
- A scheduler map that maps eight different forwarding classes to eight different schedulers can only be applied to interfaces that support eight queues. If you apply this type of scheduler map to an interface that only supports four queues, then the commit will fail.
- We recommend that you configure the statements changing PICs to support eight queues and then applying an eight queue scheduler map in two separate steps. Otherwise, the commit might succeed but the PIC might not have eight queues when the scheduler map is applied, generating an error.

You can determine the ID number assigned to a forwarding class by issuing the **show class-of-service forwarding-class** command. You can determine whether the classification is fixed by issuing the **show class-of-service forwarding-table classifier mapping** command. In the command output, if the **Table Type** field appears as **Fixed**, the classification is fixed. For more information about fixed classification, see [“Applying Forwarding Classes to Interfaces” on page 1531](#).

This section discusses the following topics:

- [Enabling Eight Queues on Interfaces on page 1538](#)
- [Multiple Forwarding Classes and Default Forwarding Classes on page 1539](#)
- [PICs Restricted to Four Queues on page 1540](#)
- [Examples: Configuring Up to 16 Forwarding Classes on page 1541](#)

Enabling Eight Queues on Interfaces

By default, Intelligent Queuing (IQ), Intelligent Queuing 2 (IQ2), Intelligent Queuing Enhanced (IQE), and Intelligent Queuing 2 Enhanced (IQ2E) PICs on M320 and T Series routers are restricted to a maximum of four egress queues per interface. To configure a maximum of eight egress queues on these interfaces, include the **max-queues-per-interface** statement at the **[edit chassis fpc slot-number pic pic-number]** hierarchy level:

```
[edit chassis fpc slot-number pic pic-number]
max-queues-per-interface (4 | 8);
```

On a TX Matrix or TX Matrix Plus router, include the **max-queues-per-interface** statement at the **[edit chassis lcc number fpc slot-number pic pic-number]** hierarchy level:

```
[edit chassis lcc number fpc slot-number pic pic-number]
max-queues-per-interface (4 | 8);
```

The numerical value can be 4 or 8.

For Juniper Networks J Series routers, this statement is not supported. J Series routers always have eight queues available.



NOTE: In addition to configuring eight queues at the **[edit chassis]** hierarchy level, the configuration at the **[edit class-of-service]** hierarchy level must support eight queues per interface.

The maximum number of queues per IQ PIC can be 4 or 8. If you include the **max-queues-per-interface** statement, all ports on the IQ PIC use configured mode and all interfaces on the IQ PIC have the same maximum number of queues.

To determine how many queues an interface supports, you can check the **CoS queues** output field of the **show interfaces interface-name extensive** command:

```
user@host> show interfaces ge-1/0/0 extensive
CoS queues: 8 supported
```


If you include the **max-queues-per-interface 4** statement, you can configure all four ports and configure up to four queues per port.

For 4-port OC3c/STM1 Type I and Type II PICs on M320 and T Series routers, when you include the **max-queues-per-interface 8** statement, you can configure up to eight queues on ports 0 and 2. After you commit the configuration, the PIC goes offline and comes back online with only ports 0 and 2 operational. No interfaces can be configured on ports 1 and 3.

For Quad T3 and Quad E3 PICs, when you include the **max-queues-per-interface 8** statement, you can configure up to eight queues on ports 0 and 2. After you commit the configuration, the PIC goes offline and comes back online with only ports 0 and 2 operational. No interfaces can be configured on ports 1 and 3.

When you include the **max-queues-per-interface** statement and commit the configuration, all physical interfaces on the IQ PIC are deleted and readded. Also, the PIC is taken offline and then brought back online immediately. You do not need to take the PIC offline and online manually. You should change modes between four queues and eight queues only when there is no active traffic going to the IQ PIC.

Multiple Forwarding Classes and Default Forwarding Classes

For queues 0 through 3, if you assign multiple forwarding classes to a single queue, default forwarding class assignment works as follows:

- The first forwarding class that you assign to queue 0 acquires the default BE classification and scheduling.
- The first forwarding class that you assign to queue 1 acquires the default EF classification and scheduling.
- The first forwarding class that you assign to queue 2 acquires the default AF classification and scheduling.
- The first forwarding class that you assign to queue 3 acquires the default NC classification and scheduling.

Of course you can override the default classification and scheduling by configuring custom classifiers and schedulers.

If you do not explicitly map forwarding classes to queues 0 through 3, then the respective default classes are automatically assigned to those queues. When you are counting the 16 forwarding classes, you must include in the total any default forwarding classes automatically assigned to queues 0 through 3. As a result, you can map up to 13 forwarding classes to a single queue when the single queue is queue 0, 1, 2, or 3. You can map up to 12 forwarding classes to a single queue when the single queue is queue 4, 5, 6, or 7. In summary, there must be at least one forwarding class each (default or otherwise) assigned to queue 0 through 3, and you can assign the remaining 12 forwarding classes (16–4) to any queue.

For example, suppose you assign two forwarding classes to queue 0 and you assign no forwarding classes to queues 1 through 3. The software automatically assigns one default

forwarding class each to queues 1 through 3. This means 11 forwarding classes (16–5) are available for you to assign to queues 4 through 7.

For more information about forwarding class defaults, see *Default Forwarding Classes*.

PICs Restricted to Four Queues

Some Juniper Networks T Series Core Router PICs support up to 16 forwarding classes and are restricted to 4 queues. Contact Juniper Networks customer support for a current list of T Series router PICs that are restricted to four queues. To determine how many queues an interface supports, you can check the **CoS queues** output field of the **show interfaces interface-name extensive** command:

```
user@host> show interfaces ge-1/0/0 extensive
CoS queues: 8 supported
```

By default, for T Series router PICs that are restricted to four queues, the router overrides the global configuration based on the following formula:

$$Q_r = Q_d \bmod R_{max}$$

Q_r is the queue number assigned if the PIC is restricted to four queues.

Q_d is the queue number that would have been mapped if this PIC were not restricted.

R_{max} is the maximum number of restricted queues available. Currently, this is four.

For example, assume you map the forwarding class **ef** to queue 6. For a PIC restricted to four queues, the queue number for forwarding class **ef** is **Q_r = 6 mod 4 = 2**.

To determine which queue is assigned to a forwarding class, use the **show class-of-service forwarding-class** command from the top level of the CLI. The output shows queue assignments for both global queue mappings and restricted queue mappings:

```
user@host> show class-of-service forwarding-class
```

Forwarding class	Queue	Restricted Queue	Fabric priority	
be		0	2	low
ef		1	2	low
assured-forwarding		2	2	low
network-control		3	3	low

For T Series router PICs restricted to four queues, you can override the formula-derived queue assignment by including the **restricted-queues** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
restricted-queues {
  forwarding-class class-name queue queue-number;
}
```

You can configure up to 16 forwarding classes. The output queue number can be from 0 through 3. Therefore, for PICs restricted to four queues, you can map multiple forwarding classes to single queues. If you map multiple forwarding classes to a queue, the multiple forwarding classes must refer to the same scheduler. This requirement applies to all PICs. The class name you configure at the **[edit class-of-service restricted-queues]** hierarchy level must be either a default forwarding class name or a forwarding class you configure at the **[edit class-of-service forwarding-classes]** hierarchy level.

Examples: Configuring Up to 16 Forwarding Classes

Configure 16 forwarding classes:

**Configuring 16
Forwarding Classes**

```
[edit class-of-service]
forwarding-classes {
  class fc0 queue-num 0;
  class fc1 queue-num 0;
  class fc2 queue-num 1;
  class fc3 queue-num 1;
  class fc4 queue-num 2;
  class fc5 queue-num 2;
  class fc6 queue-num 3;
  class fc7 queue-num 3;
  class fc8 queue-num 4;
  class fc9 queue-num 4;
  class fc10 queue-num 5;
  class fc11 queue-num 5;
  class fc12 queue-num 6;
  class fc13 queue-num 6;
  class fc14 queue-num 7;
  class fc15 queue-num 7;
}
```

For PICs restricted to four queues, map four forwarding classes to each queue:

**Restricted Queues:
Mapping Two
Forwarding Classes to
Each Queue**

```
[edit class-of-service]
restricted-queues {
  forwarding-class fc0 queue 0;
  forwarding-class fc1 queue 0;
  forwarding-class fc2 queue 0;
  forwarding-class fc3 queue 0;
  forwarding-class fc4 queue 1;
  forwarding-class fc5 queue 1;
  forwarding-class fc6 queue 1;
  forwarding-class fc7 queue 1;
  forwarding-class fc8 queue 2;
  forwarding-class fc9 queue 2;
  forwarding-class fc10 queue 2;
  forwarding-class fc11 queue 2;
  forwarding-class fc12 queue 3;
  forwarding-class fc13 queue 3;
  forwarding-class fc14 queue 3;
  forwarding-class fc15 queue 3;
}
```

If you map multiple forwarding classes to a queue, the multiple forwarding classes must refer to the same scheduler:

**Configuring a
Scheduler Map
Applicable to an
Interface Restricted to
Four Queues**

```
[edit class-of-service]
scheduler-maps {
  interface-restricted {
    forwarding-class be scheduler Q0;
    forwarding-class ef scheduler Q1;
    forwarding-class ef1 scheduler Q1;
    forwarding-class ef2 scheduler Q1;
    forwarding-class af1 scheduler Q2;
  }
}
```

```
    forwarding-class af scheduler Q2;
    forwarding-class nc scheduler Q3;
    forwarding-class nc1 scheduler Q3;
  }
}
[edit class-of-service]
restricted-queues {
  forwarding-class be queue 0;
  forwarding-class ef queue 1;
  forwarding-class ef1 queue 1;
  forwarding-class ef2 queue 1;
  forwarding-class af queue 2;
  forwarding-class af1 queue 2;
  forwarding-class nc queue 3;
  forwarding-class nc1 queue 3;
}
```

Configuration Statements

- [\[edit class-of-service\] Hierarchy Level on page 1542](#)

[edit class-of-service] Hierarchy Level

```
class-of-service {
  classifiers {
    type classifier-name {
      forwarding-class class-name {
        loss-priority (high | low | medium-high | medium-low) code-points [ aliases bits ];
      }
      import (classifier-name | default);
    }
  }
  code-point-aliases {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) {
      alias-name bits;
    }
  }
  drop-profiles {
    profile-name {
      fill-level percentage drop-probability percentage;
      interpolate {
        drop-probability value;
        fill-level value;
      }
    }
  }
  fabric {
    scheduler-map {
      priority (high | low) scheduler scheduler-name;
    }
  }
  forwarding-class-map {
    map-name {
      class class-name queue-num queue-number <restricted-queue queue-number>;
    }
  }
}
```

```

forwarding-classes {
  class class-name policing-priority (normal | premium) queue-num queue-number
    priority (high | low);
  queue queue-number class-name policing-priority (normal | premium) priority (high |
    low);
}
forwarding-policy {
  class class-name {
    classification-override {
      forwarding-class class-name;
    }
  }
  next-hop-map map-name {
    forwarding-class class-name {
      discard;
      lsp-next-hop [ lsp-regular-expressions ];
      next-hop [ next-hop-names ];
      non-lsp-next-hop;
    }
  }
}
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      drop-timeout milliseconds;
      fragment-threshold bytes;
      multilink-class number;
      no-fragmentation;
    }
  }
}
host-outbound-traffic {
  dscp-code-point value;
  forwarding-class class-name;
  ieee-802.1 {
    default value;
    rewrite-rules;
  }
  tcp {
    raise-internet-control-priority;
  }
}
interfaces {
  ... the interfaces subhierarchy appears after the main [edit class-of-service] hierarchy
  ...
}
restricted-queues {
  forwarding-class class-name queue-number;
}
rewrite-rules {
  (dscp | dscp-ipv6 | exp | frame-relay-de | ieee-802.1 | ieee-802.1ad | inet-precedence)
  rewrite-rule {
    forwarding-class class-name {
      loss-priority level code-point (alias | bits);
    }
  }
}

```

```

        import (rewrite-rule | default);
    }
}
routing-instances routing-instance-name {
    classifiers {
        dscp (classifier-name | default);
        dscp-ipv6 (classifier-name | default);
        exp (classifier-name | default);
        ieee-208.1 (classifier-name | default | encapsulated | vlan-tag (inner | outer));
    }
}
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
schedulers {
    scheduler-name {
        adjust-minimum value;
        adjust-percent value;
        buffer-size (exact | percent percentage | remainder);
        drop-profile-map loss-priority (any | high | low | medium-high | medium-low)
            protocol any;
        excess-priority (high | low | medium-high | medium-low);
        excess-rate (percent percentage | proportion proportion);
        priority (high | low | medium-high | medium-low | strict-high);
        shaping-rate (bps | percent percentage | burst-size size);
        transmit-rate (bps | percent percentage | remainder) <exact | rate-limit>;
    }
}
traceoptions {
    file <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
traffic-control-profiles {
    profile-name {
        adjust-minimum rate;
        delay-buffer-rate (bps | cps cps | percent percentage);
        excess-rate (percent percentage | proportion value);
        guaranteed-rate (bps | percent percentage) <burst-size bytes>;
        overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
        scheduler-map map-name;
        shaping-rate (bps | percent percentage) <burst-size bytes>;
    }
}
tri-color;
}

class-of-service {
    interfaces {
        interface-name {
            excess-bandwidth-share (equal | proportional value);
            input-excess-bandwidth-share (equal | proportional value);
            input-scheduler-map map-name;

```

```

input-shaping-rate bps;
input-traffic-control-profile profile-name;
output-forwarding-class-map map-name;
output-traffic-control-profile profile-name;
scheduler-map map-name;
scheduler-map-chassis (map-name | derived);
shaping-rate bps;
unit (logical-unit-number | *){
  classifiers {
    dscp (classifier-name | default) {
      family [ inet mpls ];
    }
    dscp-ipv6 (classifier-name | default) {
      family [ inet mpls ];
    }
    exp (classifier-name | default);
    ieee-208.1 (classifier-name | default) <vlan-tag (inner | outer)>;
    ieee-208.1ad (classifier-name | default);
    inet-precedence (classifier-name | default);
  }
  forwarding-class class-name;
  input-scheduler-map map-name;
  input-shaping-rate bps;
  input-traffic-control-profile profile-name shared-instance instance-name;
  loss-priority-maps {
    (map-name | default);
  }
  loss-priority-rewrites {
    (map-name | default);
  }
  output-forwarding-class-map map-name;
  output-traffic-control-profile profile-name shared-instance instance-name;
  rewrite-rules {
    dscp (rule-name | default) <protocol mpls>;
    dscp-ipv6 (rule-name | default);
    exp (rule-name | default) <protocol [ mpls-any | mpls-inet-both |
      mpls-inet-both-non-vpn ]>;
    exp-push-push-push default;
    exp-swap-push-push default;
    ieee-802.1 (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
    ieee-802.1ad (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
    inet-precedence (rewrite-name | default) <protocol mpls>;
  }
  scheduler-map map-name;
  shaping-rate bps;
  translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 |
    to-exp-from-exp | to-inet-precedence-from-inet-precedence) table-name;
}
}
interface-set interface-set-name {
  excess-bandwidth-share (equal | proportional value);
  input-excess-bandwidth-share (equal | proportional value);
  input-traffic-control-profile profile-name;
  input-traffic-control-profile-remaining profile-name;
  internal-node;
  output-traffic-control-profile profile-name;

```

```
        output-traffic-control-profile-remaining profile-name;  
    }  
}  
}
```

Related Documentation • *Notational Conventions Used in Junos OS Configuration Hierarchies*

class (Forwarding Classes)

Syntax	<code>class <i>class-name</i> queue-num <i>queue-number</i> priority (high low);</code>
Hierarchy Level	[edit class-of-service forwarding-classes]
Release Information	Statement introduced in Junos OS Release 8.1.
Description	<p>On M120 , M320, MX Series routers, T Series routers and EX Series switches only, specify the output transmission queue to which to map all input from an associated forwarding class.</p> <p>This statement enables you to configure up to 16 forwarding classes with multiple forwarding classes mapped to single queues. If you want to configure up to eight forwarding classes with one-to-one mapping to output queues, use the queue statement instead of the class statement at the [edit class-of-service forwarding-classes] hierarchy level.</p>
Options	<p><i>class-name</i>—Name of forwarding class.</p> <p><i>queue-number</i>—Output queue number.</p> <p>Range: 0 through 15. Some T Series router PICs are restricted to 0 through 3.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Forwarding Classes on page 1530• queue (Global Queues) on page 1554


forwarding-class (Interfaces)

Syntax	<code>forwarding-class class-name;</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.2 for ACX Series routers.
Description	Associate a forwarding class configuration or default mapping with a specific interface.
Options	<i>class-name</i> —Name of the forwarding class.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Applying Forwarding Classes to Interfaces on page 1531 • Configuring Fixed Classification on an ATM IMA Pseudowire • Example: Configuring Fixed Classification on an ATM IMA Pseudowire

forwarding-class (Restricted Queues)

Syntax	<code>forwarding-class class-name queue queue-number;</code>
Hierarchy Level	[edit class-of-service restricted-queues]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For M320 and T Series routers only, map forwarding classes to restricted queues. You can map up to eight forwarding classes to restricted queues.
Options	<i>class-name</i> —Name of the forwarding class. The remaining statement is explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

forwarding-classes (Class-of-Service)

Syntax	<pre>forwarding-classes { class queue-num queue-number priority (high low); queue queue-number class-name priority (high low) [policing-priority (premium normal)]; }</pre>
Hierarchy Level	[edit class-of-service]
Release Information	Statement introduced before Junos OS Release 7.4. policing-priority option introduced in Junos OS Release 9.5. Statement introduced on PTX Series Packet Transport Switches in Junos OS Release 12.1.
Description	Associate the forwarding class with a queue name and number. For M320, MX Series, T Series routers and EX Series switches only, you can configure fabric priority queuing by including the priority statement. For Enhanced IQ PICs, you can include the policing-priority option.
<div> NOTE: The priority and policing-priority options are not supported on PTX Series Packet Transport Switches.</div>	
<p>The statements are explained separately.</p> <p>See “Configuring Forwarding Classes” on page 1530, “Overriding Fabric Priority Queuing” on page 1535, and <i>Example: Configuring CoS for a PBB Network</i>. For the policing-priority option, see <i>Configuring Layer 2 Policers on IQE PICs</i>. For classification by egress interface, see <i>Classifying Packets by Egress Interface</i>.</p>	
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

forwarding-classes-interface-specific

Syntax	forwarding-classes-interface-specific <i>forwarding-class-map-name</i> { class <i>class-name</i> queue-num <i>queue-number</i> [restricted-queue <i>queue-number</i>]; }
Hierarchy Level	[edit class-of-service]
Release Information	Statement introduced in Junos OS Release 9.6.
Description	For the IQ, IQE, LSQ and ATM2 PICs in the T Series routers only and for EX Series switches, configure a forwarding class map for unicast and multicast traffic and a user-configured queue number for an egress interface.
Options	<p><i>class-name</i>—Name of the forwarding class.</p> <p><i>forwarding-class-map-name</i>—Name of the forwarding class map for traffic.</p> <p><i>queue-number</i>—Number of the egress queue.</p> <p>Range: 0 through 3 or 7, depending on chassis and configuration</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Forwarding Classes on page 1530 • Classifying Packets by Egress Interface • output-forwarding-class-map on page 1552

interfaces

```
Syntax interfaces {
    interface-name {
        classifiers{
            dscp(classifier-name | default) {
            }
            ieee-802.1 (classifier-name | default) vlan-tag (inner | outer | classifier-name);
            inet-precedence (rewrite-name | default);
        }
        input-scheduler-map map-name;
        input-shaping-rate rate;
        irb {
            unit logical-unit-number {
                classifiers {
                    type (classifier-name | default);
                }
                rewrite-rules {
                    dscp (rewrite-name | default);
                    dscp-ipv6 (rewrite-name | default);
                    exp (rewrite-name | default) protocol protocol-types;
                    ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                    inet-precedence (rewrite-name | default);
                }
            }
        }
        member-link-scheduler (replicate | scale);
        rewrite-rules {
            dscp (rewrite-name | default);
            ieee-802.1 (rewrite-name | default) vlan-tag (outer);
            inet-precedence (rewrite-name | default);
        }
        scheduler-map map-name;
        scheduler-map-chassis map-name;
        shaping-rate rate;
        unit logical-unit-number {
            classifiers {
                type (classifier-name | default) family (mpls | inet);
            }
            forwarding-class class-name;
            fragmentation-map map-name;
            input-shaping-rate (percent percentage | rate);
            input-traffic-control-profile profile-name shared-instance instance-name;
            output-traffic-control-profile profile-name shared-instance instance-name;
            per-session-scheduler;
            rewrite-rules {
                dscp (rewrite-name | default);
                dscp-ipv6 (rewrite-name | default);
                exp (rewrite-name | default) protocol protocol-types;
                exp-push-push-push default;
                exp-swap-push-push default;
                ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                inet-precedence (rewrite-name | default);
            }
        }
    }
}
```

```

    }
    scheduler-map map-name;
    shaping-rate rate;
    translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 | to-exp-from-exp
    | to-inet-precedence-from-inet-precedence) table-name;
  }
}
interface-set interface-set-name {
  excess-bandwidth-share;
  internal-node;
  output-traffic-control-profile profile-name;
  output-traffic-control-profile-remaining profile-name;
}
}

```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 7.4.
Interface-set level added in Junos OS Release 8.5.

Description Configure interface-specific CoS properties for incoming packets.



NOTE: The dscp-ipv6 and ieee-802.1ad classifier types are not supported on ACX Series routers. For further information about support on ACX Series routers, see *Understanding CoS CLI Configuration Statements on ACX Series Universal Access Routers*.

Options The remaining statements are explained separately.

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

Related Documentation

- [Overview of BA Classifier Types on page 1321](#)
- [Configuring Rewrite Rules on page 1694](#)
- *Understanding CoS CLI Configuration Statements on ACX Series Universal Access Routers*

output-forwarding-class-map

Syntax	<code>output-forwarding-class-map <i>forwarding-class-map-name</i>;</code>
Hierarchy Level	[edit class-of-service forwarding-classes-interface-specific]
Release Information	Statement introduced in Junos OS Release 9.6.
Description	Apply a configured forwarding class map to a logical interface.
Options	<i>forwarding-class-map-name</i> —Name of a forwarding class mapping configured at the [edit class-of-service forwarding-classes-interface-specific] hierarchy level.
Required Privilege Level	interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Classifying Packets by Egress Interface</i>• forwarding-classes-interface-specific on page 1549

priority (Fabric Priority)

Syntax	priority (high low);
Hierarchy Level	[edit class-of-service forwarding-classes class class-name queue-num queue-number], [edit class-of-service forwarding-classes queue queue-number class-name]
Release Information	Statement introduced before Junos OS Release 7.4. [edit class-of-service forwarding-classes class class-name queue-num queue-number] hierarchy level added in Junos OS Release 8.1.
Description	<p>For M320 routers, MX Series routers, T Series routers and EX Series switches only, specify a fabric priority value.</p> <p>The two hierarchy levels are mutually exclusive. To configure up to eight forwarding classes with one-to-one mapping between forwarding classes and output queues, include this statement at the [edit class-of-service forwarding-classes queue queue-number class-name] hierarchy level. To configure up to 16 forwarding classes with multiple forwarding classes mapped to single queues, include this statement at the [edit class-of-service forwarding-classes class class-name queue-num queue-number] hierarchy level.</p>
Options	<p>low—Forwarding class's fabric queuing has low priority.</p> <p>high—Forwarding class's fabric queuing has high priority.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Overriding Fabric Priority Queuing on page 1535 • Configuring Up to 16 Forwarding Classes

queue (Global Queues)

Syntax	<code>queue <i>queue-number</i> <i>class-name</i>;</code>
Hierarchy Level	[edit class-of-service forwarding-classes]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Specify the output transmission queue to which to map all input from an associated forwarding class.</p> <p>On M120, M320, MX Series, T Series routers and on EX Series switches, this statement enables you to configure up to eight forwarding classes with one-to-one mapping to output queues. If you want to configure up to 16 forwarding classes with multiple forwarding classes mapped to single output queues, include the class statement instead of the queue statement at the [edit class-of-service forwarding-classes] hierarchy level.</p>
Options	<p><i>class-name</i>—Name of forwarding class.</p> <p><i>queue-number</i>—Output queue number.</p> <p>Range: For M Series routers, 0 through 3. For M120, M320, MX Series, T Series routers and EX Series switches, 0 through 7. Some T Series router PICs are restricted to 0 through 3.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Forwarding Classes on page 1530• class (Forwarding Classes) on page 1546

queue (Restricted Queues)

Syntax	<code>queue <i>queue-number</i>;</code>
Hierarchy Level	[edit class-of-service restricted-queues forwarding-class <i>class-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For M320, MX Series, T Series routers and EX Series switches only, map forwarding classes to restricted queues.
Options	<p><i>queue-number</i>—Output queue number.</p> <p>Range: 0 through 3.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

restricted-queues

Syntax	<pre>restricted-queues { forwarding-class <i>class-name</i> queue <i>queue-number</i>; }</pre>
Hierarchy Level	[edit class-of-service]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>For M320, MX Series, T Series routers and EX Series switches only, map forwarding classes to restricted queues.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

unit

Syntax	<pre> unit <i>logical-unit-number</i> { classifiers { type (<i>classifier-name</i> default) family (mpls all); } forwarding-class <i>class-name</i>; fragmentation-map <i>map-name</i>; input-traffic-control-profile <i>profile-name</i> shared-instance <i>instance-name</i>; output-traffic-control-profile <i>profile-name</i> shared-instance <i>instance-name</i>; per-session-scheduler; rewrite-rules { dscp (<i>rewrite-name</i> default); dscp-ipv6 (<i>rewrite-name</i> default); exp (<i>rewrite-name</i> default) <i>protocol</i> <i>protocol-types</i>; exp-push-push-push default; exp-swap-push-push default; ieee-802.1 (<i>rewrite-name</i> default) <i>vlan-tag</i> (outer outer-and-inner); inet-precedence (<i>rewrite-name</i> default); } scheduler-map <i>map-name</i>; shaping-rate <i>rate</i>; } </pre>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.
Options	<p><i>logical-unit-number</i>—Number of the logical unit.</p> <p>Range: 0 through 16,384</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Overview of BA Classifier Types on page 1321 • Configuring Rewrite Rules on page 1694

Forwarding Policy Options

- [Overview on page 1556](#)
- [Configuration on page 1557](#)

Overview

- [Forwarding Policy on page 1557](#)

Forwarding Policy

- [Forwarding Policy Options Overview on page 1557](#)

Forwarding Policy Options Overview

Class-of-service (CoS)-based forwarding (CBF) enables you to control next-hop selection based on a packet's class of service and, in particular, the value of the IP packet's precedence bits.

For example, you might want to specify a particular interface or next hop to carry high-priority traffic while all best-effort traffic takes some other path. When a routing protocol discovers equal-cost paths, it can pick a path at random or load-balance across the paths through either hash selection or round robin. CBF allows path selection based on class.

To configure CBF properties, include the following statements at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
forwarding-policy {
  next-hop-map map-name {
    forwarding-class class-name {
      next-hop [ next-hop-name ];
      lsp-next-hop [ lsp-regular-expression ];
      non-lsp-next-hop;
      discard;
    }
  }
  class class-name {
    classification-override {
      forwarding-class class-name;
    }
  }
}
```

Configuration

- [Configuration Tasks on page 1557](#)
- [Examples on page 1561](#)
- [Configuration Statements on page 1565](#)

Configuration Tasks

- [Configuring CoS-Based Forwarding on page 1557](#)
- [Overriding the Input Classification on page 1560](#)

Configuring CoS-Based Forwarding

You can apply CoS-based forwarding (CBF) only to a defined set of routes. Therefore you must configure a policy statement as in the following example:

```
[edit policy-options]
policy-statement my-cos-forwarding {
  from {
```

```

    route-filter destination-prefix match-type;
  }
  then {
    cos-next-hop-map map-name;
  }
}

```

This configuration specifies that routes matching the route filter are subject to the CoS next-hop mapping specified by *map-name*. For more information about configuring policy statements, see the *Routing Policy Configuration Guide*.



NOTE: On M Series routers (except the M120 and M320 routers), forwarding-class-based matching and CBF do not work as expected if the forwarding class has been set with a multifield filter on an input interface.

You can configure CBF on a routing device with eight or less than eight forwarding classes only. Under this condition, the forwarding class to queue mapping can be either one-to-one or one-to-many. However, you cannot configure CBF when the number of forwarding classes configured exceeds eight. Similarly, with CBF configured, you cannot configure more than eight forwarding classes.

To specify a CoS next-hop map, include the **forwarding-policy** statement at the **[edit class-of-service]** hierarchy level:

```

[edit class-of-service]
forwarding-policy {
  next-hop-map map-name {
    forwarding-class class-name {
      next-hop [ next-hop-name ];
      lsp-next-hop [ lsp-regular-expression ];
      discard;
    }
  }
}

```

When you configure CBF with OSPF as the interior gateway protocol (IGP), you must specify the next hop as an interface name or next-hop alias, not as an IP address. This is true because OSPF adds routes with the interface as the next hop for point-to-point interfaces; the next hop does not contain the IP address. For an example configuration, see *Example: Configuring CoS-Based Forwarding*.

For Layer 3 VPNs, when you use class-based forwarding for the routes received from the far-end provider-edge (PE) router within a VRF instance, the software can match the routes based on the attributes that come with the received route only. In other words, the matching can be based on the route within RIB-in. In this case, the **route-filter** statement you include at the **[edit policy-options policy-statement my-cos-forwarding from]** hierarchy level has no effect because the policy checks the **bgp.l3vpn.0** table, not the **vrf.inet.0** table.

The Junos OS applies the CoS next-hop map to the set of next hops previously defined; the next hops themselves can be located across any outgoing interfaces on the routing

device. For example, the following configuration associates a set of forwarding classes and next-hop identifiers:

```
[edit class-of-service forwarding-policy]
next-hop-map map1 {
  forwarding-class expedited-forwarding {
    next-hop next-hop1;
    next-hop next-hop2;
  }
  forwarding-class best-effort {
    next-hop next-hop3;
    lsp-next-hop lsp-next-hop4;
  }
}
```

In this example, **next-hop N** is either an IP address or an egress interface for some next hop, and **lsp-next-hop4** is a regular expression corresponding to any next hop with that label. Q1 through QN are a set of forwarding classes that map to the specific next hop. That is, when a packet is switched with Q1 through QN, it is forwarded out the interface associated with the associated next hop.

This configuration has the following implications:

- A single forwarding class can map to multiple standard next hops or LSP next hops. This implies that load sharing is done across standard next hops or LSP next hops servicing the same class value. To make this work properly, the Junos OS creates a list of the equal-cost next hops and forwards packets according to standard load-sharing rules for that forwarding class.
- If a forwarding class configuration includes LSP next hops and standard next hops, the LSP next hops are preferred over the standard next hops. In the preceding example, if both **next-hop3** and **lsp-next-hop4** are valid next hops for a route to which **map1** is applied, the forwarding table includes entry **lsp-next-hop4** only.
- If **next-hop-map** does not specify all possible forwarding classes, the default forwarding class is selected as the default. If the default forwarding class is not specified in the next-hop map, a default is designated randomly. The default forwarding class is the class associated with queue 0.
- For LSP next hops, the Junos OS uses UNIX **regex(3)**-style regular expressions. For example, if the following labels exist: **lsp**, **lsp1**, **lsp2**, **lsp3**, the statement **lsp-next-hop lsp** matches **lsp**, **lsp1**, **lsp2**, and **lsp3**. If you do not desire this behavior, you must use the anchor characters **lsp-next-hop " ^lsp\$"**, which match **lsp** only.
- The route filter does not work because the policy checks against the **bgp.l3vpn.0** table instead of the **vrf.inet.0** table.

The final step is to apply the route filter to routes exported to the forwarding engine. This is shown in the following example:

```
routing-options {
  forwarding-table {
    export my-cos-forwarding;
  }
}
```

This configuration instructs the routing process to insert routes to the forwarding engine matching **my-cos-forwarding** with the associated next-hop CBF rules.

The following algorithm is used when you apply a configuration to a route:

- If the route is a single next-hop route, all traffic goes to that route; that is, no CBF takes effect.
- For each next hop, associate the proper forwarding class. If a next hop appears in the route but not in the **cos-next-hop** map, it does not appear in the forwarding table entry.
- The default forwarding class is used if all forwarding classes are not specified in the next-hop map. If the default is not specified, one is chosen randomly.

Overriding the Input Classification

For IPv4 or IPv6 packets, you can override the incoming classification, assigning them to the same forwarding class based on their input interface, input precedence bits, or destination address. You do so by defining a policy class when configuring CoS properties and referencing this class when configuring a routing policy.

When you override the classification of incoming packets, any mappings you configured for associated precedence bits or incoming interfaces to output transmission queues are ignored. Also, if the packet loss priority (PLP) bit was set in the packet by the incoming interface, the PLP bit is cleared.

To override the input packet classification, do the following:

1. Define the policy class by including the **class** statement at the **[edit class-of-service policy]** hierarchy level:

```
[edit class-of-service]
forwarding-policy {
  class class-name {
    classification-override {
      forwarding-class class-name;
    }
  }
}
```

class-name is a name that identifies the class.

2. Associate the policy class with a routing policy by including it in a **policy-statement** statement at the **[edit policy-options]** hierarchy level. Specify the destination prefixes in the **route-filter** statement and the CoS policy class name in the **then** statement.

```
[edit policy-options]
policy-statement policy-name {
  term term-name {
    from {
      route-filter destination-prefix match-type <class class-name>
    }
    then class class-name;
  }
}
```

3. Apply the policy by including the **export** statement at the **[edit routing-options]** hierarchy level:

```
[edit routing-options]
forwarding-table {
  export policy-name;
}
```

Examples


- [Example: Configuring CoS-Based Forwarding on page 1561](#)
- [Example: Configuring CoS-Based Forwarding for Different Traffic Types on page 1563](#)
- [Example: Configuring CoS-Based Forwarding for IPv6 on page 1564](#)

Example: Configuring CoS-Based Forwarding

Router A has two routes to destination **10.255.71.208** on Router D. One route goes through Router B, and the other goes through Router C, as shown in [Figure 12 on page 1561](#).

Configure Router A with CBF to select Router B for queue 0 and queue 2, and Router C for queue 1 and queue 3.

Figure 12: Sample CoS-Based Forwarding



When you configure CBF with OSPF as the IGP, you must specify the next hop as an interface name, not as an IP address. The next hops in this example are specified as **ge-2/0/0.0** and **ge-0/3/0.0**.

```
[edit class-of-service]
forwarding-policy {
  next-hop-map my_cbf {
    forwarding-class be {
      next-hop ge-2/0/0.0;
    }
    forwarding-class ef {
      next-hop ge-0/3/0.0;
    }
    forwarding-class af {
```

```
        next-hop ge-2/0/0.0;
      }
      forwarding-class nc {
        next-hop ge-0/3/0.0;
      }
    }
  }
  classifiers {
    inet-precedence inet {
      forwarding-class be {
        loss-priority low code-points [ 000 100 ];
      }
      forwarding-class ef {
        loss-priority low code-points [ 001 101 ];
      }
      forwarding-class af {
        loss-priority low code-points [ 010 110 ];
      }
      forwarding-class nc {
        loss-priority low code-points [ 011 111 ];
      }
    }
  }
  forwarding-classes {
    queue 0 be;
    queue 1 ef;
    queue 2 af;
    queue 3 nc;
  }
  interfaces {
    at-4/2/0 {
      unit 0 {
        classifiers {
          inet-precedence inet;
        }
      }
    }
  }
}

[edit policy-options]
policy-statement cbf {
  from {
    route-filter 10.255.71.208/32 exact;
  }
  then cos-next-hop-map my_cbf;
}

[edit routing-options]
graceful-restart;
forwarding-table {
  export cbf;
}

[edit interfaces]
traceoptions {
  file trace-intf size 5m world-readable;
```



```

    flag all;
  }
  ge-0/3/0 {
    unit 0 {
      family inet {
        address 10.40.13.1/30;
      }
      family iso;
      family mpls;
    }
  }
  ge-2/0/0 {
    unit 0 {
      family inet {
        address 10.40.12.1/30;
      }
      family iso;
      family mpls;
    }
  }
  at-4/2/0 {
    atm-options {
      vpi 1 {
        maximum-vcs 1200;
      }
    }
    unit 0 {
      vci 1.100;
      family inet {
        address 10.40.11.2/30;
      }
      family iso;
      family mpls;
    }
  }
}

```

Example: Configuring CoS-Based Forwarding for Different Traffic Types

One common use for CoS-based forwarding and next-hop maps is to enforce different handling for different traffic types, such as voice and video. For example, an LSP-based next hop can be used for voice and video, and a non-LSP next-hop can be used for best effort traffic.

Only the forwarding policy is shown in this example:

```

[edit class-of-service]
forwarding-policy {
  next-hop-map ldp-map {
    forwarding-class expedited-forwarding {
      lsp-next-hop voice;
      non-lsp-next-hop;
    }
    forwarding-class assured-forwarding {
      lsp-next-hop video;
      non-lsp-next-hop;
    }
  }
}

```

```
        forwarding-class best-effort {
            non-lsp-next-hop;
            discard;
        }
    }
}
```

Example: Configuring CoS-Based Forwarding for IPv6

This example configures CoS-based forwarding (CBF) next-hop maps and CBF LSP next-hop maps for IPv6 addresses.

You can configure a next-hop map with both IPv4 and IPv6 addresses, or you can configure separate next-hop maps for IPv4 and IPv6 addresses and include the **from family (inet | inet6)** statements at the **[edit policy-options policy-options policy-statement *policy-name* term *term-name*]** hierarchy level to ensure that only next-hop maps of a specified protocol are applied to a specified route.

If you do not configure separate next-hop maps and include the **from family (inet | inet6)** statements in the configuration, when a route uses two next hops (whether IPv4, IPv6, interface, or LSP next hop) in at least two of the specified forwarding classes, CBF is used for the route; otherwise, the CBF policy is ignored.

1. Define the CBF next-hop map:

```
[edit class-of-service]
forwarding-policy {
    next-hop-map cbf-map {
        forwarding-class best-effort {
            next-hop [ ::192.168.139.38 192.168.139.38 ];
        }
        forwarding-class expedited-forwarding {
            next-hop [ ::192.168.140.5 192.168.140.5 ];
        }
        forwarding-class assured-forwarding {
            next-hop [ ::192.168.145.5 192.168.145.5 ];
        }
        forwarding-class network-control {
            next-hop [ ::192.168.141.2 192.168.141.2 ];
        }
    }
}
```

2. Define the CBF forwarding policy:

```
[edit policy-options]
policy-statement ls {
    then cos-next-hop-map cbf-map;
}
```

3. Export the CBF forwarding policy:

```
[edit routing-options]
forwarding-table {
    export ls;
}
```

Configuration Statements

- [\[edit class-of-service\] Hierarchy Level on page 1565](#)

[edit class-of-service] Hierarchy Level

```

class-of-service {
  classifiers {
    type classifier-name {
      forwarding-class class-name {
        loss-priority (high | low | medium-high | medium-low) code-points [ aliases bits ];
      }
      import (classifier-name | default);
    }
  }
  code-point-aliases {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) {
      alias-name bits;
    }
  }
  drop-profiles {
    profile-name {
      fill-level percentage drop-probability percentage;
      interpolate {
        drop-probability value;
        fill-level value;
      }
    }
  }
  fabric {
    scheduler-map {
      priority (high | low) scheduler scheduler-name;
    }
  }
  forwarding-class-map {
    map-name {
      class class-name queue-num queue-number <restricted-queue queue-number>;
    }
  }
  forwarding-classes {
    class class-name policing-priority (normal | premium) queue-num queue-number
      priority (high | low);
    queue queue-number class-name policing-priority (normal | premium) priority (high |
      low);
  }
  forwarding-policy {
    class class-name {
      classification-override {
        forwarding-class class-name;
      }
    }
  }
  next-hop-map map-name {
    forwarding-class class-name {
      discard;
      lsp-next-hop [ lsp-regular-expressions ];
      next-hop [ next-hop-names ];
    }
  }
}

```

```

        non-lsp-next-hop;
    }
}
fragmentation-maps {
    map-name {
        forwarding-class class-name {
            drop-timeout milliseconds;
            fragment-threshold bytes;
            multilink-class number;
            no-fragmentation;
        }
    }
}
host-outbound-traffic {
    dscp-code-point value;
    forwarding-class class-name;
    ieee-802.1 {
        default value;
        rewrite-rules;
    }
    tcp {
        raise-internet-control-priority;
    }
}
interfaces {
    ... the interfaces subhierarchy appears after the main [edit class-of-service] hierarchy
    ...
}
restricted-queues {
    forwarding-class class-name queue-number;
}
rewrite-rules {
    (dscp | dscp-ipv6 | exp | frame-relay-de | ieee-802.1 | ieee-802.1ad | inet-precedence)
    rewrite-rule {
        forwarding-class class-name {
            loss-priority level code-point (alias | bits);
        }
        import (rewrite-rule | default);
    }
}
routing-instances routing-instance-name {
    classifiers {
        dscp (classifier-name | default);
        dscp-ipv6 (classifier-name | default);
        exp (classifier-name | default);
        ieee-208.1 (classifier-name | default | encapsulated | vlan-tag (inner | outer));
    }
}
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
schedulers {

```

```

scheduler-name {
    adjust-minimum value;
    adjust-percent value;
    buffer-size (exact | percent percentage | remainder);
    drop-profile-map loss-priority (any | high | low | medium-high | medium-low)
        protocol any;
    excess-priority (high | low | medium-high | medium-low);
    excess-rate (percent percentage | proportion proportion);
    priority (high | low | medium-high | medium-low | strict-high);
    shaping-rate (bps | percent percentage | burst-size size);
    transmit-rate (bps | percent percentage | remainder) <exact | rate-limit>;
}
}
traceoptions {
    file <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
traffic-control-profiles {
    profile-name {
        adjust-minimum rate;
        delay-buffer-rate (bps | cps cps | percent percentage);
        excess-rate (percent percentage | proportion value);
        guaranteed-rate (bps | percent percentage) <burst-size bytes>;
        overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
        scheduler-map map-name;
        shaping-rate (bps | percent percentage) <burst-size bytes>;
    }
}
tri-color;
}

class-of-service {
    interfaces {
        interface-name {
            excess-bandwidth-share (equal | proportional value);
            input-excess-bandwidth-share (equal | proportional value);
            input-scheduler-map map-name;
            input-shaping-rate bps;
            input-traffic-control-profile profile-name;
            output-forwarding-class-map map-name;
            output-traffic-control-profile profile-name;
            scheduler-map map-name;
            scheduler-map-chassis (map-name | derived);
            shaping-rate bps;
            unit (logical-unit-number | *) {
                classifiers {
                    dscp (classifier-name | default) {
                        family [ inet mpls ];
                    }
                    dscp-ipv6 (classifier-name | default) {
                        family [ inet mpls ];
                    }
                    exp (classifier-name | default);
                    ieee-208.1 (classifier-name | default) <vlan-tag (inner | outer)>;
                }
            }
        }
    }
}

```

```

        ieee-208.1ad (classifier-name | default);
        inet-precedence (classifier-name | default);
    }
    forwarding-class class-name;
    input-scheduler-map map-name;
    input-shaping-rate bps;
    input-traffic-control-profile profile-name shared-instance instance-name;
    loss-priority-maps {
        (map-name | default);
    }
    loss-priority-rewrites {
        (map-name | default);
    }
    output-forwarding-class-map map-name;
    output-traffic-control-profile profile-name shared-instance instance-name;
    rewrite-rules {
        dscp (rule-name | default) <protocol mpls>;
        dscp-ipv6 (rule-name | default);
        exp (rule-name | default) <protocol [ mpls-any | mpls-inet-both |
            mpls-inet-both-non-vpn ]>;
        exp-push-push-push default;
        exp-swap-push-push default;
        ieee-802.1 (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
        ieee-802.1ad (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
        inet-precedence (rewrite-name | default) <protocol mpls>;
    }
    scheduler-map map-name;
    shaping-rate bps;
    translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 |
        to-exp-from-exp | to-inet-precedence-from-inet-precedence) table-name;
    }
}
interface-set interface-set-name {
    excess-bandwidth-share (equal | proportional value);
    input-excess-bandwidth-share (equal | proportional value);
    input-traffic-control-profile profile-name;
    input-traffic-control-profile-remaining profile-name;
    internal-node;
    output-traffic-control-profile profile-name;
    output-traffic-control-profile-remaining profile-name;
}
}
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

class (CoS-Based Forwarding)

Syntax	<pre>class <i>class-name</i> { classification-override { forwarding-class <i>class-name</i>; } }</pre>
Hierarchy Level	[edit class-of-service forwarding-policy]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure CoS-based forwarding class.
Options	<p><i>class-name</i>—Name of the routing policy class.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Overriding the Input Classification on page 1560

classification-override

Syntax	<pre>classification-override { forwarding-class <i>class-name</i>; }</pre>
Hierarchy Level	[edit class-of-service forwarding-policy class <i>class-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For IPv4 packets, override the incoming packet classification, assigning all packets sent to a destination prefix to the same output transmission queue.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Overriding the Input Classification on page 1560 • policy-statement in the <i>Junos OS Routing Protocols Configuration Guide</i>

discard (Forwarding Class)

Syntax	discard;
Hierarchy Level	[edit class-of-service forwarding-policy next-hop-map map-name forwarding-class class-name]
Release Information	Statement introduced in Junos OS Release 9.1.
Description	Discard traffic sent to this forwarding class for the next-hop map referenced by this forwarding policy.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring CoS-Based Forwarding on page 1557• non-lsp-next-hop on page 1573

forwarding-class (Forwarding Policy)

Syntax	<pre>forwarding-class class-name { next-hop [next-hop-name]; lsp-next-hop [lsp-regular-expression]; non-lsp-next-hop; discard; }</pre>
Hierarchy Level	[edit class-of-service forwarding-policy next-hop-map map-name]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define forwarding class name and associated next hops.
Options	<p>class-name—Name of the forwarding class.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Overriding the Input Classification on page 1560

forwarding-policy

```
Syntax forwarding-policy {
    next-hop-map map-name {
        forwarding-class class-name {
            next-hop [ next-hop-name ];
            lsp-next-hop [ lsp-regular-expression ];
            non-lsp-next-hop;
            discard;
        }
    }
    class class-name {
        classification-override {
            forwarding-class class-name;
        }
    }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define CoS-based forwarding policy options.

The statements are explained separately.

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

Related Documentation • [Configuring CoS-Based Forwarding on page 1557](#)

lsp-next-hop (CoS-Based Forwarding)

```
Syntax lsp-next-hop [ lsp-regular-expression ];
```

Hierarchy Level [edit class-of-service forwarding-policy next-hop-map map-name forwarding-class class-name]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the LSP regular expression to which to map forwarded traffic.

Options *lsp-regular-expression*—Next-hop LSP label.

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

Related Documentation • [Configuring CoS-Based Forwarding on page 1557](#)

next-hop (Class-Of-Service)

Syntax	<code>next-hop [<i>next-hop-name</i>];</code>
Hierarchy Level	[edit class-of-service forwarding-policy next-hop-map <i>map-name</i> forwarding-class <i>class-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the next-hop name or address to which to map forwarded traffic.
Options	<i>next-hop-name</i> —Next-hop alias or IP address.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring CoS-Based Forwarding on page 1557

next-hop-map

Syntax	<pre>next-hop-map <i>map-name</i> { forwarding-class <i>class-name</i> { next-hop <i>next-hop-name</i>; lsp-next-hop [<i>lsp-regular-expression</i>]; non-lsp-next-hop; discard; } }</pre>
Hierarchy Level	[edit class-of-service forwarding-policy]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the map for CoS forwarding routes.
Options	<i>map-name</i> —Map that defines next-hop routes.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring CoS-Based Forwarding on page 1557

non-lsp-next-hop

Syntax	<code>non-lsp-next-hop;</code>
Hierarchy Level	<code>[edit class-of-service forwarding-policy next-hop-map map-name forwarding-class class-name]</code>
Release Information	Statement introduced before Junos OS Release 9.0.
Description	Use a non-LSP next hop for traffic sent to this forwarding class next-hop map of this forwarding policy.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring CoS-Based Forwarding on page 1557

Schedulers

- [Overview on page 1573](#)
- [Configuration on page 1578](#)

Overview

- [Schedulers on page 1573](#)

Schedulers

- [Schedulers Overview on page 1573](#)
- [Default Schedulers Overview on page 1575](#)
- [Priority Scheduling Overview on page 1576](#)
- [Applying Scheduler Maps Overview on page 1577](#)
- [Applying a Shaping Rate to Physical Interfaces Overview on page 1578](#)
- [Default Fabric Priority Queuing on page 1578](#)

Schedulers Overview

You use *schedulers* to define the properties of output queues. These properties include the amount of interface bandwidth assigned to the queue, the size of the memory buffer allocated for storing packets, the priority of the queue, and the random early detection (RED) drop profiles associated with the queue.

You associate the schedulers with forwarding classes by means of *scheduler maps*. You can then associate each scheduler map with an interface, thereby configuring the hardware queues, packet schedulers, and RED processes that operate according to this mapping.

To configure class-of-service (CoS) schedulers, include the following statements at the `[edit class-of-service]` hierarchy level:



NOTE: For PTX Series Packet Transport Switches:

- The fabric and traffic-control-profiles statements at the [edit class-of-service] hierarchy level are not supported.

```
[edit class-of-service]
interfaces {
  interface-name {
    scheduler-map map-name;
    scheduler-map-chassis map-name;
    shaping-rate rate;
    unit {
      output-traffic-control-profile profile-name;
      scheduler-map map-name;
      shaping-rate rate;
    }
  }
}
fabric {
  scheduler-map {
    priority (high | low) scheduler scheduler-name;
  }
}
scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
schedulers {
  scheduler-name {
    buffer-size (percent percentage | remainder | temporal microseconds );
    drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol
      (any | non-tcp | tcp) drop-profile profile-name;
    excess-priority (low | high);
    excess-rate (percent percentage | proportion value);
    priority priority-level;
    transmit-rate (rate | percent percentage remainder) <exact | rate-limit>;
  }
}
traffic-control-profiles profile-name {
  delay-buffer-rate (percent percentage | rate);
  excess-rate percent percentage;
  guaranteed-rate (percent percentage | rate);
  scheduler-map map-name;
  shaping-rate (percent percentage | rate);
}
```

You cannot configure both the **shaping-rate** statement at the [edit class-of-service interfaces *interface-name*] hierarchy level and the **transmit-rate rate-limit** statement and option at the [edit class-of-service schedulers *scheduler-name*] hierarchy level. These statements are mutually exclusive. If you do configure both, you will not be able to commit the configuration:

```
[edit class-of-service]
'shaping-rate'
only one option (shaping-rate or transmit-rate rate-limit) can be configured at a time
error: commit failed (statements constraint check failed)
```

Default Schedulers Overview

Each forwarding class has an associated scheduler priority. Only two forwarding classes, best effort and network control (queue 0 and queue 3), are used in the Junos default scheduler configuration.

By default, the best effort forwarding class (queue 0) receives 95 percent of the bandwidth and buffer space for the output link, and the network control forwarding class (queue 3) receives 5 percent. The default drop profile causes the buffer to fill and then discard all packets until it has space.

The expedited-forwarding and assured-forwarding classes have no schedulers because, by default, no resources are assigned to queue 1 and queue 2. However, you can manually configure resources for the expedited-forwarding and assured-forwarding classes.

Also by default, each queue can exceed the assigned bandwidth if additional bandwidth is available from other queues. When a forwarding class does not fully use the allocated transmission bandwidth, the remaining bandwidth can be used by other forwarding classes if they receive a larger amount of the offered load than the bandwidth allocated. For more information, see [“Allocation of Leftover Bandwidth” on page 1593](#).

The following default scheduler is provided when you install the Junos OS. These settings are not visible in the output of the **show class-of-service** command; rather, they are implicit.

```
[edit class-of-service]
schedulers {
  network-control {
    transmit-rate percent 5;
    buffer-size percent 5;
    priority low;
    drop-profile-map loss-priority any protocol any drop-profile terminal;
  }
  best-effort {
    transmit-rate percent 95;
    buffer-size percent 95;
    priority low;
    drop-profile-map loss-priority any protocol any drop-profile terminal;
  }
}
drop-profiles {
  terminal {
    fill-level 100 drop-probability 100;
  }
}
```

Priority Scheduling Overview

The Junos OS supports multiple levels of transmission priority, which in order of increasing priority are **low**, **medium-low**, **medium-high**, and **high**, and **strict-high**. This allows the software to service higher-priority queues before lower-priority queues.

Priority scheduling determines the order in which an output interface transmits traffic from the queues, thus ensuring that queues containing important traffic are provided better access to the outgoing interface. This is accomplished through a procedure in which the software examines the priority of the queue. In addition, the software determines if the individual queue is within its defined bandwidth profile. The bandwidth profile is discussed in [“Configuring Scheduler Transmission Rate” on page 1591](#). This binary decision, which is reevaluated on a regular time cycle, compares the amount of data transmitted by the queue against the amount of bandwidth allocated to it by the scheduler. When the transmitted amount is less than the allocated amount, the queue is considered to be in profile. A queue is out of profile when its transmitted amount is larger than its allocated amount.

The queues for a given output physical interface (or output logical interface if per-unit scheduling is enabled on that interface) are divided into sets based on their priority. Any such set contains queues of the same priority.

The software traverses the sets in descending order of priority. If at least one of the queues in the set has a packet to transmit, the software selects that set. A queue from the set is selected based on the weighted round robin (WRR) algorithm, which operates within the set.

The Junos OS performs priority queuing using the following steps:

1. The software locates all high-priority queues that are currently in profile. These queues are serviced first in a weighted round-robin fashion.
2. The software locates all medium-high priority queues that are currently in profile. These queues are serviced second in a weighted round-robin fashion.
3. The software locates all medium-low priority queues that are currently in profile. These queues are serviced third in a weighted round-robin fashion.
4. The software locates all low-priority queues that are currently in profile. These queues are serviced fourth in a weighted round-robin fashion.
5. The software locates all high-priority queues that are currently out of profile and are not rate limited. The weighted round-robin algorithm is applied to these queues for servicing.
6. The software locates all medium-high priority queues that are currently out of profile and are not rate limited. The weighted round-robin algorithm is applied to these queues for servicing.

7. The software locates all medium-low priority queues that are currently out of profile and are not rate limited. The weighted round-robin algorithm is applied to these queues for servicing.
8. The software locates all low-priority queues that are currently out of profile and are also not rate limited. These queues are serviced last in a weighted round-robin manner.

Applying Scheduler Maps Overview

Physical interfaces (for example, **t3-0/0/0**, **t3-0/0/0:0**, and **ge-0/0/0**) support scheduling with any encapsulation type pertinent to that physical interface. For a single port, you cannot apply scheduling to the physical interface if you have applied scheduling to one or more of the associated logical interfaces.

Logical interfaces (for example, **t3-0/0/0 unit 0** and **ge-0/0/0 unit 0**) support scheduling on data link connection identifiers (DLCIs) or VLANs only.

In the Junos OS implementation, the term *logical interfaces* generally refers to interfaces you configure by including the **unit** statement at the **[edit interfaces interface-name]** hierarchy level. Logical interfaces have the **.logical** descriptor at the end of the interface name, as in **ge-0/0/0.1** or **t1-0/0/0:0.1**, where the logical unit number is 1.

Although channelized interfaces are generally thought of as logical or virtual, the Junos OS sees T3, T1, and NxDSO interfaces within a channelized IQ PIC as physical interfaces. For example, both **t3-0/0/0** and **t3-0/0/0:1** are treated as physical interfaces by the Junos OS. In contrast, **t3-0/0/0.2** and **t3-0/0/0:1.2** are considered logical interfaces because they have the **.2** at the end of the interface names.

Within the **[edit class-of-service]** hierarchy level, you cannot use the **.logical** descriptor when you assign properties to logical interfaces. Instead, you must include the **unit** statement in the configuration. For example:

```
[edit class-of-service]
user@host# set interfaces t3-0/0/0 unit 0 scheduler-map map1
```

Related Documentation

To apply a scheduler map to network traffic, you associate the map with an interface. See the following topics:

- [Applying Scheduler Maps to Physical Interfaces on page 1613](#)
- [Applying Scheduler Maps and Shaping Rate to Physical Interfaces on IQ PICs](#)
- [Applying Scheduler Maps and Shaping Rate to DLCIs and VLANs](#)
- [Oversubscribing Interface Bandwidth on page 1600](#)
- [Providing a Guaranteed Minimum Rate on page 1608](#)
- [Applying Scheduler Maps to Packet Forwarding Component Queues](#)
- [Default Fabric Priority Queuing on page 1578](#)
- [Associating Schedulers with Fabric Priorities on page 1612](#)

Applying a Shaping Rate to Physical Interfaces Overview

On T4000 routers with Type 5 FPCs and on EX Series switches, you can configure physical interfaces to shape traffic based on the rate-limited bandwidth of the total interface bandwidth. This allows you to shape the output of the physical interface, so that the interface transmits less traffic than it is physically capable of carrying.

If you do not configure a shaping rate on the physical interface, the default physical interface bandwidth is based on the channel bandwidth and the time slot allocation.

In general, the physical interface speed is the basis for calculating the various queue parameters for a physical interface such as delay buffer size, weighted round-robin (WRR) weight, drop profile, and so forth. However, when you apply a shaping rate by including the **shaping-rate** statement, the shaping rate on that physical interface becomes the basis for calculating all the queue parameters for that physical interface.

On T4000 routers with Type 5 FPCs, the shaping rate value for the physical interface must be a minimum of 292 Kbps. The maximum value of shaping rate is limited by the maximum transmission rate of the interface.

Related Documentation

- [Configuring the Shaping Rate for Physical Interfaces on page 1599](#)
- [show class-of-service interface](#)
- [show interfaces extensive on page 2450](#)

Default Fabric Priority Queuing

On Juniper Networks EX Series switches, M320 Multiservice Edge Routers, and Juniper Networks T Series Core Routers, the default behavior is for fabric priority queuing on egress interfaces to match the scheduling priority you assign. High-priority egress traffic is automatically assigned to high-priority fabric queues. Likewise, low-priority egress traffic is automatically assigned to low-priority fabric queues.

For information about overriding automatic fabric priority queuing, see [“Overriding Fabric Priority Queuing” on page 1535](#) and [“Associating Schedulers with Fabric Priorities” on page 1612](#).

Configuration

- [Configuration Tasks for Schedulers on page 1578](#)
- [Configuration Tasks for Scheduler Maps on page 1613](#)
- [Configuration Statements for Schedulers on page 1624](#)

Configuration Tasks for Schedulers

- [Configuring Schedulers on page 1579](#)
- [Configuring the Scheduler Buffer Size on page 1579](#)
- [Configuring Drop Profile Maps for Schedulers on page 1590](#)
- [Configuring Scheduler Transmission Rate on page 1591](#)
- [Configuring Schedulers for Priority Scheduling on page 1594](#)

- [Configuring Per-Unit Schedulers for Channelized Interfaces on page 1596](#)
- [Configuring the Shaping Rate for Physical Interfaces on page 1599](#)
- [Oversubscribing Interface Bandwidth on page 1600](#)
- [Providing a Guaranteed Minimum Rate on page 1608](#)
- [Associating Schedulers with Fabric Priorities on page 1612](#)

Configuring Schedulers

You configure a scheduler by including the **scheduler** statement at the **[edit class-of-service]** hierarchy level:

```
schedulers {
  scheduler-name {
    buffer-size (percent percentage | remainder | temporal microseconds);
    drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol
      (any | non-tcp | tcp) drop-profile profile-name;
    priority priority-level;
    transmit-rate (rate | percent percentage remainder) <exact | rate-limit>;
  }
}
```

For detailed information about scheduler configuration statements, see the indicated topics:

- [Configuring the Scheduler Buffer Size on page 1579](#)
- [Configuring Drop Profile Maps for Schedulers on page 1590](#)
- [Configuring Scheduler Transmission Rate on page 1591](#)
- [Configuring Schedulers for Priority Scheduling on page 1594](#)

Configuring the Scheduler Buffer Size

To control congestion at the output stage, you can configure the delay-buffer bandwidth. The delay-buffer bandwidth provides packet buffer space to absorb burst traffic up to the specified duration of delay. Once the specified delay buffer becomes full, packets with 100 percent drop probability are dropped from the head of the buffer.

The default scheduler transmission rate for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent of the total available bandwidth.

The default buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent of the total available buffer. The total available buffer per queue differs by PIC type, as shown in [Table 109 on page 1580](#).

To configure the buffer size, include the **buffer-size** statement at the **[edit class-of-service schedulers**

scheduler-name] hierarchy level:

```
[edit class-of-service schedulers scheduler-name]
  buffer-size (percent percentage | remainder | temporal microseconds);
```

For each scheduler, you can configure the buffer size as one of the following:

- A percentage of the total buffer. The total buffer per queue is based on microseconds and differs by routing device type, as shown in [Table 109 on page 1580](#).
- The remaining buffer available. The remainder is the buffer percentage that is not assigned to other queues. For example, if you assign 40 percent of the delay buffer to queue 0, allow queue 3 to keep the default allotment of 5 percent, and assign the remainder to queue 7, then queue 7 uses approximately 55 percent of the delay buffer.
- A temporal value, in microseconds. For the temporal setting, the queuing algorithm starts dropping packets when it queues more than a computed number of bytes. This maximum is computed by multiplying the transmission rate of the queue by the configured temporal value. The buffer size temporal value per queue differs by routing device type, as shown in [Table 109 on page 1580](#). The maximums apply to the logical interface, not each queue.

For information about configuring large buffer sizes on IQ PICs, see “[Configuring Large Delay Buffers for Slower Interfaces](#)” on page 1581.

Table 109: Buffer Size Temporal Value Ranges by Routing Device Type

Routing Devices	Temporal Value Ranges
M320 and T Series router FPCs, Type 1 and Type 2	1 through 80,000 microseconds
M320 and T Series router FPCs, Type 3. All ES cards (Type 1, 2, 3, and 4).	1 through 50,000 microseconds For PICs with greater than 40 Gbps of total bandwidth, the maximum temporal buffer size that can be configured for a scheduler is 40,000 microseconds instead of 50,000 microseconds.
M120 router FEBs and MX Series router nonenhanced Queuing DPCs, and EX Series switches	1 through 100,000 microseconds
M5, M7i, M10, and M10i router FPCs	1 through 100,000 microseconds
Other M Series router FPCs	1 through 200,000 microseconds
PTX Series Packet Transport Switches	1 through 100,000 microseconds
IQ PICs on all routers	1 through 100,000 microseconds
With Large Buffer Sizes Enabled	
IQ PICs on all routers	1 through 500,000 microseconds
Gigabit Ethernet IQ VLANs	
With shaping rate up to 10 Mbps	1 through 400,000 microseconds
With shaping rate up to 20 Mbps	1 through 300,000 microseconds

Table 109: Buffer Size Temporal Value Ranges by Routing Device Type (*continued*)

Routing Devices	Temporal Value Ranges
With shaping rate up to 30 Mbps	1 through 200,000 microseconds
With shaping rate up to 40 Mbps	1 through 150,000 microseconds
With shaping rate above 40 Mbps	1 through 100,000 microseconds

For more information about configuring delay buffers, see the following subtopics:

- [Configuring Large Delay Buffers for Slower Interfaces on page 1581](#)
- [Enabling and Disabling the Memory Allocation Dynamic per Queue on page 1589](#)

Configuring Large Delay Buffers for Slower Interfaces

By default, T1, E1, and NxDS0 interfaces and DLCIs configured on channelized IQ PICs are limited to 100,000 microseconds of delay buffer. (The default average packet size on the IQ PIC is 40 bytes.) For these interfaces, it might be necessary to configure a larger buffer size to prevent congestion and packet dropping. You can do so on the following PICs:

- Channelized IQ
- 4-port E3 IQ
- Gigabit Ethernet IQ and IQ2

Congestion and packet dropping occur when large bursts of traffic are received by slower interfaces. This happens when faster interfaces pass traffic to slower interfaces, which is often the case when edge devices receive traffic from the core of the network. For example, a 100,000-microsecond T1 delay buffer can absorb only 20 percent of a 5000-microsecond burst of traffic from an upstream OC3 interface. In this case, 80 percent of the burst traffic is dropped.

[Table 110 on page 1581](#) shows some recommended buffer sizes needed to absorb typical burst sizes from various upstream interface types.

Table 110: Recommended Delay Buffer Sizes

Length of Burst	Upstream Interface	Downstream Interface	Recommended Buffer on Downstream Interface
5000 microseconds	OC3	E1 or T1	500,000 microseconds
5000 microseconds	E1 or T1	E1 or T1	100,000 microseconds
1000 microseconds	T3	E1 or T1	100,000 microseconds

To ensure that traffic is queued and transmitted properly on E1, T1, and NxDS0 interfaces and DLCIs, you can configure a buffer size larger than the default maximum. To enable

larger buffer sizes to be configured, include the **q-pic-large-buffer (large-scale | small-scale)** statement at the **[edit chassis fpc slot-number pic pic-number]** hierarchy level:

```
[edit chassis fpc slot-number pic pic-number]
q-pic-large-buffer large-scale;
```

If you specify the **large-scale** option, the feature supports a larger number of interfaces. If you specify **small-scale**, the default, then the feature supports a smaller number of interfaces.

When you include the **q-pic-large-buffer** statement in the configuration, the larger buffer is transparently available for allocation to scheduler queues. The larger buffer maximum varies by interface type, as shown in [Table 111 on page 1582](#).

Table 111: Maximum Delay Buffer with q-pic-large-buffer Enabled by Interface

Platform, PIC, or Interface Type	Maximum Buffer Size
With Large Buffer Sizes Not Enabled	
M320 and T Series router FPCs, Type 1 and Type 2	80,000 microseconds
M320 and T Series router FPCs, Type 3	50,000 microseconds
Other M Series router FPCs	200,000 microseconds
IQ PICs on all routers	100,000 microseconds
With Large Buffer Sizes Enabled	
Channelized T3 and channelized OC3 DLCIs—Maximum sizes vary by shaping rate:	
With shaping rate from 64,000 through 255,999 bps	4,000,000 microseconds
With shaping rate from 256,000 through 511,999 bps	2,000,000 microseconds
With shaping rate from 512,000 through 1,023,999 bps	1,000,000 microseconds
With shaping rate from 1,024,000 through 2,048,000 bps	500,000 microseconds
With shaping rate from 2,048,001 bps through 10 Mbps	400,000 microseconds
With shaping rate from 10,000,001 bps through 20 Mbps	300,000 microseconds
With shaping rate from 20,000,001 bps through 30 Mbps	200,000 microseconds
With shaping rate from 30,000,001 bps through 40 Mbps	150,000 microseconds
With shaping rate up to 40,000,001 bps and above	100,000 microseconds
NxDSO IQ Interfaces—Maximum sizes vary by channel size:	

Table 111: Maximum Delay Buffer with q-pic-large-buffer Enabled by Interface (*continued*)

Platform, PIC, or Interface Type	Maximum Buffer Size
1xDSO through 3xDSO	4,000,000 microseconds
4xDSO through 7xDSO	2,000,000 microseconds
8xDSO through 15xDSO	1,000,000 microseconds
16xDSO through 32xDSO	500,000 microseconds
Other IQ interfaces	500,000 microseconds

If you configure a delay buffer larger than the new maximum, the candidate configuration can be committed successfully. However, the setting is rejected by the packet forwarding component and a system log warning message is generated.

For interfaces that support DLCI queuing, the large buffer is supported for DLCIs on which the configured shaping rate is less than or equal to the physical interface bandwidth. For instance, when you configure a Frame Relay DLCI on a Channelized T3 IQ PIC, and you configure the shaping rate to be 1.5 Mbps, the amount of delay buffer that can be allocated to the DLCI is 500,000 microseconds, which is equivalent to a T1 delay buffer. For more information about DLCI queuing, see *Applying Scheduler Maps and Shaping Rate to DLCIs and VLANs*.

For NxDSO interfaces, the larger buffer sizes can be up to 4,000,000 microseconds, depending on the number of DSO channels in the NxDSO interface. For slower NxDSO interfaces with fewer channels, the delay buffer can be relatively larger than for faster NxDSO interfaces with more channels. This is shown in [Table 113 on page 1585](#). To calculate specific buffer sizes for various NxDSO interfaces, see “[Maximum Delay Buffer for NxDSO Interfaces](#)” on page 1584.

You can allocate the delay buffer as either a percentage or a temporal value. The resulting delay buffer is calculated differently depending how you configure the delay buffer, as shown in [Table 112 on page 1584](#).

Table 112: Delay-Buffer Calculations

Delay Buffer Configuration	Formula	Example
Percentage	$\text{available interface bandwidth} * \text{configured percentage buffer-size} = \text{queue buffer}$	<p>If you configure a queue on a T1 interface to use 30 percent of the available delay buffer, the queue receives 28,125 bytes of delay buffer:</p> <pre> sched-expedited { transmit-rate percent 30; buffer-size percent 30; } </pre> <p>$1.5 \text{ Mbps} * 0.3 * 500,000 \text{ microseconds} = 225,000 \text{ bits}$ $= 28,125 \text{ bytes}$</p>
Temporal	$\text{available interface bandwidth} * \text{configured percentage transmit-rate} * \text{configured temporal buffer-size} = \text{queue buffer}$	<p>If you configure a queue on a T1 interface to use 500,000 microseconds of delay buffer and you configure the transmission rate to be 20 percent, the queue receives 18,750 bytes of delay buffer:</p> <pre> sched-best { transmit-rate percent 20; buffer-size temporal 500000; } </pre> <p>$1.5 \text{ Mbps} * 0.2 * 500,000 \text{ microseconds} = 150,000 \text{ bits}$ $= 18,750 \text{ bytes}$</p>
Percentage, with buffer size larger than transmit rate		<p>In this example, the delay buffer is allocated twice the transmit rate. Maximum delay buffer latency can be up to twice the 500,000-microsecond delay buffer if the queue's transmit rate cannot exceed the allocated transmit rate.</p> <pre> sched-extra-buffer { transmit-rate percent 10; buffer-size percent 20; } </pre>
FRF.16 LSQ bundles	<p>For total bundle bandwidth < T1 bandwidth, the delay-buffer rate is 1 second.</p> <p>For total bundle bandwidth >= T1 bandwidth, the delay-buffer rate is 200 milliseconds (ms).</p>	

For more information, see the following sections:

- [Maximum Delay Buffer for NxDSO Interfaces on page 1584](#)
- [Example: Configuring Large Delay Buffers for Slower Interfaces on page 1586](#)

Maximum Delay Buffer for NxDSO Interfaces

Because NxDSO interfaces carry less bandwidth than a T1 or E1 interface, the buffer size on an NxDSO interface can be relatively larger, depending on the number of DSO channels combined. The maximum delay buffer size is calculated with the following formula:

$$\text{Interface Speed} * \text{Maximum Delay Buffer Time} = \text{Delay Buffer Size}$$

For example, a 1xDS0 interface has a speed of 64 kilobits per second (Kbps). At this rate, the maximum delay buffer time is 4,000,000 microseconds. Therefore, the delay buffer size is 32 kilobytes (KB):

$$64 \text{ Kbps} * 4,000,000 \text{ microseconds} = 32 \text{ KB}$$

Table 113 on page 1585 shows the delay-buffer calculations for 1xDS0 through 32xDS0 interfaces.

Table 113: NxDS0 Transmission Rates and Delay Buffers

Interface Speed	Delay Buffer Size
1xDS0 Through 4xDS0: Maximum Delay Buffer Time Is 4,000,000 Microseconds	
1xDS0: 64 Kbps	32 KB
2xDS0: 128 Kbps	64 KB
3xDS0: 192 Kbps	96 KB
4xDS0 Through 7xDS0: Maximum Delay Buffer Time Is 2,000,000 Microseconds	
4xDS0: 256 Kbps	64 KB
5xDS0: 320 Kbps	80 KB
6xDS0: 384 Kbps	96 KB
7xDS0: 448 Kbps	112 KB
8xDS0 Through 15xDS0: Maximum Delay Buffer Time Is 1,000,000 Microseconds	
8xDS0: 512 Kbps	64 KB
9xDS0: 576 Kbps	72 KB
10xDS0: 640 Kbps	80 KB
11xDS0: 704 Kbps	88 KB
12xDS0: 768 Kbps	96 KB
13xDS0: 832 Kbps	104 KB
14xDS0: 896 Kbps	112 KB
15xDS0: 960 Kbps	120 KB
16xDS0 Through 32xDS0: Maximum Delay Buffer Time Is 500,000 Microseconds	
16xDS0: 1024 Kbps	64 KB
17xDS0: 1088 Kbps	68 KB

Table 113: NxDSO Transmission Rates and Delay Buffers (*continued*)

Interface Speed	Delay Buffer Size
18xDSO: 1152 Kbps	72 KB
19xDSO: 1216 Kbps	76 KB
20xDSO: 1280 Kbps	80 KB
21xDSO: 1344 Kbps	84 KB
22xDSO: 1408 Kbps	88 KB
23xDSO: 1472 Kbps	92 KB
24xDSO: 1536 Kbps	96 KB
25xDSO: 1600 Kbps	100 KB
26xDSO: 1664 Kbps	104 KB
27xDSO: 1728 Kbps	108 KB
28xDSO: 1792 Kbps	112 KB
29xDSO: 1856 Kbps	116 KB
30xDSO: 1920 Kbps	120 KB
31xDSO: 1984 Kbps	124 KB
32xDSO: 2048 Kbps	128 KB

Example: Configuring Large Delay Buffers for Slower Interfaces

Set large delay buffers on interfaces configured on a Channelized OC12 IQ PIC. The CoS configuration binds a scheduler map to the interface specified in the chassis configuration. For information about the delay-buffer calculations in this example, see [Table 112 on page 1584](#).

```
chassis {
  fpc 0 {
    pic 0 {
      q-pic-large-buffer; # Enabling large delay buffer
      max-queues-per-interface 8; # Eight queues (M320, T Series, and TX Matrix routers)
    }
  }
}
```


Configuring the Delay Buffer Value for a Scheduler

You can assign to a physical or logical interface a scheduler map that is composed of different schedulers (or queues). The physical interface's large delay buffer can be distributed to the different schedulers (or queues) using the **transmit-rate** and **buffer-size** statements at the `[edit class-of-service schedulers scheduler-name]` hierarchy level.

The example shows two schedulers, **sched-best** and **sched-exped**, with the delay buffer size configured as a percentage (20 percent) and temporal value (300,000 microseconds), respectively. The **sched-best** scheduler has a transmit rate of 10 percent. The **sched-exped** scheduler has a transmit rate of 20 percent.

The **sched-best** scheduler's delay buffer is twice that of the specified transmit rate of 10 percent. Assuming that the **sched-best** scheduler is assigned to a T1 interface, this scheduler receives 20 percent of the total 500,000 microseconds of the T1 interface's delay buffer. Therefore, the scheduler receives 18,750 bytes of delay buffer:

$$\text{available interface bandwidth} * \text{configured percentage buffer-size} * \text{maximum buffer} = \text{queue buffer}$$

$$1.5 \text{ Mbps} * 0.2 * 500,000 \text{ microseconds} = 150,000 \text{ bits} = 18,750 \text{ bytes}$$

Assuming that the **sched-exped** scheduler is assigned to a T1 interface, this scheduler receives 300,000 microseconds of the T1 interface's 500,000-microsecond delay buffer with the traffic rate at 20 percent. Therefore, the scheduler receives 11,250 bytes of delay buffer:

$$\text{available interface bandwidth} * \text{configured percentage transmit-rate} * \text{configured temporal buffer-size} = \text{queue buffer}$$

$$1.5 \text{ Mbps} * 0.2 * 300,000 \text{ microseconds} = 90,000 \text{ bits} = 11,250 \text{ bytes}$$

```
[edit]
class-of-service {
  schedulers {
    sched-best {
      transmit-rate percent 10;
      buffer-size percent 20;
    }
    sched-exped {
      transmit-rate percent 20;
      buffer-size temporal 300000;
    }
  }
}
```

Configuring the Physical Interface Shaping Rate

In general, the physical interface speed is the basis for calculating the delay buffer size. However, when you include the **shaping-rate** statement, the shaping rate becomes the basis for calculating the delay buffer size. This example configures the shaping rate on a T1 interface to 200 Kbps, which means that the T1 interface bandwidth is set to 200 Kbps instead of 1.5 Mbps. Because 200 Kbps is less than 4xDS0, this interface receives 4 seconds of delay buffer, or 800 Kbps of traffic, which is 800 KB for a full second. For more information, see [Table 113 on page 1585](#).

```
class-of-service {
  interfaces {
    t1-0/0/0:1 {
      shaping-rate 200k;
    }
  }
}
```

```
    }  
  }
```

Complete Configuration

This example shows a Channelized OC12 IQ PIC in FPC slot 0, PIC slot 0 and a channelized T1 interface with Frame Relay encapsulation. It also shows a scheduler map configuration on the physical interface.

```
chassis {  
  fpc 0 {  
    pic 0 {  
      q-pic-large-buffer;  
      max-queues-per-interface 8;  
    }  
  }  
}  
interfaces {  
  coc12-0/0/0 {  
    partition 1 oc-slice 1 interface-type coc1;  
  }  
  coc1-0/0/0:1 {  
    partition 1 interface-type t1;  
  }  
  t1-0/0/0:1:1 {  
    encapsulation frame-relay;  
    unit 0 {  
      family inet {  
        address 1.1.1.1/24;  
      }  
      dlci 100;  
    }  
  }  
}  
class-of-service {  
  interfaces {  
    t1-0/0/0:1:1 {  
      scheduler-map smap-1;  
    }  
  }  
}  
scheduler-maps {  
  smap-1 {  
    forwarding-class best-effort scheduler sched-best;  
    forwarding-class expedited-forwarding scheduler sched-exped;  
    forwarding-class assured-forwarding scheduler sched-assure;  
    forwarding-class network-control scheduler sched-network;  
  }  
}  
schedulers {  
  sched-best {  
    transmit-rate percent 40;  
    buffer-size percent 40;  
  }  
  sched-exped {  
    transmit-rate percent 30;  
    buffer-size percent 30;  
  }  
  sched-assure {  
    transmit-rate percent 30;  
    buffer-size percent 30;  
  }  
}
```

```

        transmit-rate percent 20;
        buffer-size percent 20;
    }
    sched-network {
        transmit-rate percent 10;
        buffer-size percent 10;
    }
}

```

Enabling and Disabling the Memory Allocation Dynamic per Queue

In the Junos OS, the memory allocation dynamic (MAD) is a mechanism that dynamically provisions extra delay buffer when a queue is using more bandwidth than it is allocated in the transmit rate setting. With this extra buffer, queues absorb traffic bursts more easily, thus avoiding packet drops. The MAD mechanism can provision extra delay buffer only when extra transmission bandwidth is being used by a queue. This means that the queue might have packet drops if there is no surplus transmission bandwidth available.

For Juniper Networks M320 Multiservice Edge Routers, MX Services Ethernet Services Routers, and T Series Core Routers and EX Series switches only, the MAD mechanism is enabled unless the delay buffer is configured with a temporal setting for a given queue. The MAD mechanism is particularly useful for forwarding classes carrying latency-immune traffic for which the primary requirement is maximum bandwidth utilization. In contrast, for latency-sensitive traffic, you might wish to disable the MAD mechanism because large delay buffers are not optimum.

MAD support is dependent on the FPC and Packet Forwarding Engine, not the PIC. All M320, MX Series, and T Series router and EX Series switches' FPCs and Packet Forwarding Engines support MAD. No Modular Port Concentrators (MPCs) and IQ, IQ2, IQ2E or IQE PICs support MAD.

To enable the MAD mechanism on supported hardware, include the **buffer-size percent** statement at the **[edit class-of-service schedulers *scheduler-name*]** hierarchy level:

```

[edit class-of-service schedulers scheduler-name]
  buffer-size percent percentage;

```

If desired, you can configure a buffer size that is greater than the configured transmission rate. The buffer can accommodate packet bursts that exceed the configured transmission rate, if sufficient excess bandwidth is available:

```

class-of-service {
  schedulers {
    sched-best {
      transmit-rate percent 20;
      buffer-size percent 30;
    }
  }
}

```

As stated previously, you can use a temporal delay buffer configuration to disable the MAD mechanism on a queue, thus limiting the size of the delay buffer. However, the effective buffer latency for a temporal queue is bounded not only by the buffer size value

but also by the associated drop profile. If a drop profile specifies a drop probability of 100 percent at a fill-level less than 100 percent, the effective maximum buffer latency is smaller than the buffer size setting. This is because the drop profile specifies that the queue drop packets before the queue's delay buffer is 100 percent full.

Such a configuration might look like the following example:

```
class-of-service {
  drop-profiles {
    plp-high {
      fill-level 70 drop-probability 100;
    }
    plp-low {
      fill-level 80 drop-probability 100;
    }
  }
  schedulers {
    sched {
      buffer-size temporal 500000;
      drop-profile-map loss-priority low protocol any drop-profile plp-low;
      drop-profile-map loss-priority high protocol any drop-profile plp-high;
      transmit-rate percent 20;
    }
  }
}
```

Configuring Drop Profile Maps for Schedulers

Drop-profile maps associate drop profiles with a scheduler. The map examines the current loss priority setting of the packet (high, low, or any) and assigns a drop profile according to these values. For example, you can specify that all TCP packets with low loss priority are assigned a drop profile that you name **low-drop**. You can associate multiple drop-profile maps with a single queue.

The scheduler drop profile defines the drop probabilities across the range of delay-buffer occupancy, thereby supporting the RED process. Depending on the drop probabilities, RED might drop packets aggressively long before the buffer becomes full, or it might drop only a few packets even if the buffer is almost full. For information on how to configure drop profiles, see [“RED Drop Profiles Overview” on page 1674](#).

By default, the drop profile is mapped to packets with low PLP and any protocol type. To configure how packet types are mapped to a specified drop profile, include the **drop-profile-map** statement at the **[edit class-of-service schedulers scheduler-name]** hierarchy level:

```
[edit class-of-service schedulers scheduler-name ]
  drop-profile-map loss-priority (any | low | medium-low | medium-high | high) protocol
    (any | non-tcp | tcp) drop-profile profile-name;
```

The map sets the drop profile for a specific PLP and protocol type. The inputs for the map are the PLP and the protocol type. The output is the drop profile. For more information about how CoS maps work, see [“CoS Inputs and Outputs Overview” on page 1272](#).



NOTE: On Juniper Network MX Series 3D Universal Edge Routers, T4000 Core Routers, EX Series switches, and PTX Series Packet Transport Switches, you can configure only the `any` option for the protocol statement.

For each scheduler, you can configure separate drop profile maps for each loss priority.

You can configure a maximum of 32 different drop profiles.

Related Documentation

- [Configuring RED Drop Profiles on page 1677](#)

Configuring Scheduler Transmission Rate

The transmission rate control determines the actual traffic bandwidth from each forwarding class you configure. The rate is specified in bits per second (bps). Each queue is allocated some portion of the bandwidth of the outgoing interface.

This bandwidth amount can be a fixed value, such as 1 megabit per second (Mbps), a percentage of the total available bandwidth, or the rest of the available bandwidth. You can limit the transmission bandwidth to the exact value you configure, or allow it to exceed the configured rate if additional bandwidth is available from other queues. This property allows you to ensure that each queue receives the amount of bandwidth appropriate to its level of service.

On M Series routers other than the M120 and M320 routers, you should not configure a **buffer-size** larger than the **transmit-rate** for a rate-limited queue in a scheduler. If you do, the Packet Forwarding Engine will reject the CoS configuration. However, you can achieve the same effect by removing the **exact** option from the transmit rate or specifying the buffer size using the **temporal** option.



NOTE: For 8-port, 12-port, and 48-port Fast Ethernet PICs, transmission scheduling is not supported.

On Juniper Networks J Series Services Routers, you can include the **transmit-rate** statement described in this section to assign the WRR weights within a given priority level and not between priorities. For more information, see [“Configuring Schedulers for Priority Scheduling” on page 1594](#).

To configure transmission scheduling, include the **transmit-rate** statement at the **[edit class-of-service schedulers scheduler-name]** hierarchy level:

```
[edit class-of-service schedulers scheduler-name]
  transmit-rate (rate | percent percentage | remainder) <exact | rate-limit>;
```

You can specify the transmit rate as follows:

- **rate**—Transmission rate, in bits per second. For all MX Series router and EX Series switch interfaces, the rate can be from 65,535 through 160,000,000,000 bps. On all other platforms, the rate can be from 3200 through 160,000,000,000 bps.
- **percent *percentage***—Percentage of transmission capacity.
- **remainder**—Use remaining rate available. In the configuration, you cannot combine the **remainder** and **exact** options.
- **exact**—(Optional) Enforce the exact transmission rate or percentage you configure with the **transmit-rate *rate*** or **transmit-rate percent** statement. Under sustained congestion, a rate-controlled queue that goes into negative credit fills up and eventually drops packets. You specify the **exact** option as follows:

```
[edit class-of-service schedulers scheduler-name]
transmit-rate rate exact;
```

```
[edit class-of-service schedulers scheduler-name]
transmit-rate percent percentage exact;
```

In the configuration, you cannot combine the **remainder** and **exact** options.



NOTE:

- Including the **exact** option is not supported on Enhanced Queuing Dense Port Concentrators (DPCs) on Juniper Network MX Series 3D Universal Edge Routers.
- The configuration of the **transmit-rate percent 0 exact** statement at the `[edit class-of-service schedulers scheduler-name]` hierarchy is ineffective on T4000 routers with Type 5 FPCs.

- **rate-limit**—(Optional) Limit the transmission rate to the specified amount. You can configure this option for all 8 queues of a logical interface (unit) and apply it to shaped or unshaped logical interfaces. If you configure a zero rate-limited transmit rate, all packets belonging to that queue are dropped. On IQE PICs, the **rate-limit** option for the schedulers' transmit rate is implemented as a static policer. Therefore, these schedulers are not aware of congestion and the maximum rate possible on these schedulers is limited by the value specified in the **transmit-rate** statement. Even if there is no congestion, the queue cannot send traffic above the transmit rate due to the static policer.



NOTE: You can apply a transmit rate limit to logical interfaces on Multiservices 100, 400, or 500 PICs. Typically, rate limits are used to prevent a strict-high queue (such as voice) from starving lower priority queues. You can only rate-limit one queue per logical interface. To apply a rate-limit to a Multiservices PIC interface, configure the rate limit in a scheduler and apply the scheduler map to the Multiservices (lsq-) interface at the `[edit class-of-service interfaces]` hierarchy level. For information about configuring other scheduler components, see [“Configuring Schedulers” on page 1579](#).

For more information about scheduler transmission rate, see the following sections:

- [Example: Configuring Scheduler Transmission Rate on page 1593](#)
- [Allocation of Leftover Bandwidth on page 1593](#)

Example: Configuring Scheduler Transmission Rate

Configure the **best-effort** scheduler to use the remainder of the bandwidth on any interface to which it is assigned:

```
class-of-service {
  schedulers {
    best-effort {
      transmit-rate remainder;
    }
  }
}
```

Allocation of Leftover Bandwidth

The allocation of leftover bandwidth is a complex topic. It is difficult to predict and to test, because the behavior of the software varies depending on the traffic mix.

If a queue receives offered loads in excess of the queue's bandwidth allocation, the queue has negative bandwidth credit, and receives a share of any available leftover bandwidth. Negative bandwidth credit means the queue has used up its allocated bandwidth. If a queue's bandwidth credit is positive, meaning it is not receiving offered loads in excess of its bandwidth configuration, then the queue does not receive a share of leftover bandwidth. If the credit is positive, then the queue does not need to use leftover bandwidth, because it can use its own allocation.

This use of leftover bandwidth is the default. If you do not want a queue to use any leftover bandwidth, you must configure it for strict allocation by including the **transmit-rate** statement with the **exact** option at the **[edit class-of-service schedulers scheduler-name]** hierarchy level. With rate control in place, the specified bandwidth is strictly observed. (On Juniper Networks J Series routers, the **exact** option is useful within a given priority, but not between the priorities. For more information, see [“Configuring Schedulers for Priority Scheduling” on page 1594.](#))

On J Series routers, leftover bandwidth is allocated to queues with negative credit in proportion to the configured transmit rate of the queues within a given priority level.

Juniper Networks M Series Multiservice Edge Routers and T Series Core Routers do not distribute leftover bandwidth in proportion to the configured transmit rate of the queues. Instead, the scheduler distributes the leftover bandwidth equally in round-robin fashion to queues that have negative bandwidth credit. All negative-credit queues can take the leftover bandwidth in equal share. This description suggests a simple round-robin distribution process among the queues with negative credits. In actual operation, a queue might change its bandwidth credit status from positive to negative and from negative to positive instantly while the leftover bandwidth is being distributed. Lower-rate queues tend to be allocated a larger share of leftover bandwidth, because their bandwidth credit is more likely to be negative at any given time, if they are overdriven persistently. Also, if there is a large packet size difference, (for example, queue 0 receives 64-byte packets,

whereas queue 1 receives 1500-byte packets), then the actual leftover bandwidth distribution ratio can be skewed substantially, because each round-robin turn allows exactly one packet to be transmitted by a negative-credit queue, regardless of the packet size.

By default, on MX Series routers, the M320 Enhanced Type 4 FPCs, and T4000 routers with Type 5 FPCs and EX Series switches, excess bandwidth is shared in the ratio of the transmit rates. You can adjust this distribution by configuring the **excess-rate** statement at the **[edit class-of-service schedulers scheduler-name]** hierarchy level. You can specify the excess rate sharing by percentage or by proportion.

In summary, J Series routers distribute leftover bandwidth in proportion to the configured rates of the negative-credit queues within a given priority level. M Series and T Series routers distribute leftover bandwidth in equal shares for the queues with the same priority and same negative-credit status. MX Series routers and M320 Enhanced Type 4 FPCs, and EX Series switches, share excess bandwidth in the ratio of the transmit rates, but you can adjust this distribution.

Related Documentation

- [Configuring Schedulers for Priority Scheduling on page 1594](#)
- [Schedulers Overview on page 1573](#)
- [Configuring a Scheduler](#)
- [excess-rate on page 1633](#)
- [schedulers on page 1482](#)

Configuring Schedulers for Priority Scheduling

To configure priority scheduling, include the **priority** statement at the **[edit class-of-service schedulers scheduler-name]** hierarchy level:

```
[edit class-of-service schedulers scheduler-name]  
priority priority-level;
```

The priority level can be **low**, **medium-low**, **medium-high**, **high**, or **strict-high**. The priorities map to numeric priorities in the underlying hardware. In some cases, different priorities behave similarly, because two software priorities behave differently only if they map to two distinct hardware priorities. For more information, see *Platform Support for Priority Scheduling*.

Higher-priority queues transmit packets ahead of lower priority queues as long as the higher-priority forwarding classes retain enough bandwidth credit. When you configure a higher-priority queue with a significant fraction of the transmission bandwidth, the queue might lock out (or *starve*) lower priority traffic.

Strict-high priority queuing works differently on different platforms. For information about strict-high priority queuing on J Series Services Routers, see the J Series router documentation.

The following sections discuss priority scheduling:

- [Example: Configuring Priority Scheduling on page 1595](#)
- [Configuring Strict-High Priority on M Series and T Series Routers on page 1595](#)

Example: Configuring Priority Scheduling

Configure priority scheduling, as shown in the following example:

1. Configure a scheduler, **be-sched**, with **medium-low** priority.

```
[edit class-of-service]
schedulers {
  be-sched {
    priority medium-low;
  }
}
```

2. Configure a scheduler map, **be-map**, that associates **be-sched** with the **best-effort** forwarding class.

```
[edit class-of-service]
scheduler-maps {
  be-map {
    forwarding-class best-effort scheduler be-sched;
  }
}
```

3. Assign **be-map** to a Gigabit Ethernet interface, **ge-0/0/0**.

```
[edit class-of-service]
interfaces {
  ge-0/0/0 {
    scheduler-map be-map;
  }
}
```

Configuring Strict-High Priority on M Series and T Series Routers

On M Series Multiservice Edge Routers and T Series Core Routers, you can configure one queue per interface to have **strict-high** priority, which works the same as **high** priority, but provides unlimited transmission bandwidth. As long as the queue with **strict-high** priority has traffic to send, it receives precedence over all other queues, except queues with **high** priority. Queues with **strict-high** and **high** priority take turns transmitting packets until the **strict-high** queue is empty, the **high** priority queues are empty, or the **high** priority queues run out of bandwidth credit. Only when these conditions are met can lower priority queues send traffic.

When you configure a queue to have **strict-high** priority, you do not need to include the **transmit-rate** statement in the queue configuration at the **[edit class-of-service schedulers scheduler-name]** hierarchy level because the transmission rate of a **strict-high** priority queue is not limited by the WRR configuration. If you do configure a transmission rate on a **strict-high** priority queue, it does not affect the WRR operation. The transmission rate only serves as a placeholder in the output of commands such as the **show interface queue** command.

strict-high priority queues might starve **low** priority queues. The **high** priority allows you to protect traffic classes from being starved by traffic in a **strict-high** queue. For example, a network-control queue might require a small bandwidth allocation (say, 5 percent). You can assign **high** priority to this queue to prevent it from being underserved.

A queue with **strict-high** priority supersedes bandwidth guarantees for queues with lower priority; therefore, we recommend that you use the **strict-high** priority to ensure proper ordering of special traffic, such as voice traffic. You can preserve bandwidth guarantees for queues with lower priority by allocating to the queue with **strict-high** priority only the amount of bandwidth that it generally requires. For example, consider the following allocation of transmission bandwidth:

- Q0 BE—20 percent, low priority
- Q1 EF—30 percent, strict-high priority
- Q2 AF—40 percent, low priority
- Q3 NC—10 percent, low priority

This bandwidth allocation assumes that, in general, the EF forwarding class requires only 30 percent of an interface's transmission bandwidth. However, if short bursts of traffic are received on the EF forwarding class, 100 percent of the bandwidth is given to the EF forwarding class because of the **strict-high** setting.

**Related
Documentation**

- [Schedulers Overview on page 1573](#)
- *Platform Support for Priority Scheduling*

Configuring Per-Unit Schedulers for Channelized Interfaces

You can configure per-unit scheduling on T1 and DS0 physical interfaces configured on channelized DS3 and STM1 IQ PICs. To enable per-unit scheduling, configure the **per-unit-scheduler** statements at the **[edit interfaces *interface-name*]** hierarchy level.

When per-unit scheduling is enabled on the channelized PICs, you can associate a scheduler map with the physical interface. For more information about configuring scheduler maps, see [“Configuring Scheduler Maps” on page 1613](#).



NOTE: If you configure the **per-unit-scheduler** statement on the physical interface of a 4-port channelized OC-12 IQ PIC and configure 975 logical interfaces or data link connection identifiers (DLCIs), some of the logical interfaces or DLCIs will drop all packets intermittently.

The following example configures per-unit scheduling on a channelized DS3 PIC and an STM1 IQ PIC.

```
[edit interfaces]
ct3-5/3/1 {
  partition 1 interface-type t1;
}
t1-5/3/1:1 {
```

```

per-unit-scheduler; # This enables per-unit scheduling
encapsulation frame-relay;
unit 0 {
    dlci 1;
    family inet {
        address 10.0.0.2/32;
    }
}
}
ct3-5/3/0 {
    partition 1 interface-type ct1;
}
ct1-5/3/0:1 {
    partition 1 timeslots 1 interface-type ds;
}
ds-5/3/0:1:1 {
    per-unit-scheduler; # This enables per-unit scheduling
    encapsulation frame-relay;
    unit 0 {
        dlci 1;
        family inet {
            address 10.0.0.1/32;
        }
    }
}
cau4-3/0/0 {
    partition 1 interface-type cel;
}
cstm1-3/0/0 {
    no-partition 1 interface-type cau4;
}
cel-3/0/0:1 {
    partition 1 timeslots 1 interface-type ds;
}
ds-3/0/0:1:1 {
    per-unit-scheduler; # This enables per-unit scheduling
    encapsulation frame-relay;
    unit 0 {
        dlci 1;
        family inet {
            address 10.1.1.1/32;
        }
    }
}
}

```

```

[edit class-of-service]
classifiers {
    dscp all-traffic-dscp {
        forwarding-class assured-forwarding {
            loss-priority low code-points 001010;
        }
        forwarding-class expedited-forwarding {
            loss-priority low code-points 101110;
        }
        forwarding-class best-effort {

```

```
        loss-priority low code-points 101010;
    }
    forwarding-class network-control {
        loss-priority low code-points 000110;
    }
}
forwarding-classes {
    queue 0 best-effort;
    queue 1 assured-forwarding;
    queue 2 expedited-forwarding;
    queue 3 network-control;
}
interfaces {
    ds-3/0/0:1 {
        unit 0 {
            scheduler-map schedule-mlppp;
        }
    }
    ds-5/3/0:1 {
        unit 0 {
            scheduler-map schedule-mlppp;
        }
    }
    t1-5/3/1:1 {
        unit 0 {
            scheduler-map schedule-mlppp;
        }
    }
}
scheduler-maps {
    schedule-mlppp {
        forwarding-class expedited-forwarding scheduler expedited-forwarding;
        forwarding-class assured-forwarding scheduler assured-forwarding;
        forwarding-class best-effort scheduler best-effort;
        forwarding-class network-control scheduler network-control;
    }
}
schedulers {
    best-effort {
        transmit-rate percent 2;
        buffer-size percent 5;
        priority low;
    }
    assured-forwarding {
        transmit-rate percent 7;
        buffer-size percent 30;
        priority low;
    }
    expedited-forwarding {
        transmit-rate percent 90 exact;
        buffer-size percent 60;
        priority high;
    }
    network-control {
        transmit-rate percent 1;
    }
}
```

```

    buffer-size percent 5;
    priority strict-high;
  }
}

```

Configuring the Shaping Rate for Physical Interfaces

To configure the shaping rate on the physical interface, either include the **shaping-rate** statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level or include the **output-traffic-control-profile** statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level.

You can specify a peak bandwidth rate in bps, either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000). For physical interfaces, the range is from 1000 through 160,000,000,000 bps.

For physical interfaces on T4000 routers with Type 5 FPCs, the shaping rate value for the physical interface must be a minimum of 292 Kbps. The maximum value of **shaping-rate** is limited by the maximum transmission rate of the interface.

The following are two example configurations for applying a shaping rate of 5 Gbps on a T4000 12x10 Gbps physical interface (xe-4/0/0):

Applying a shaping rate at the **[edit class-of-service interfaces *interface-name*] hierarchy:**

```

[edit class-of-service]
interfaces {
  xe-4/0/0 {
    shaping-rate 5g;
  }
}

```

Applying a shaping rate using traffic-control-profiles:

```

[edit class-of-service]
traffic-control-profiles {
  output {
    shaping-rate 5g;
  }
}
interfaces {
  xe-4/0/0 {
    output-traffic-control-profile output;
  }
}

```

To view the results of your configuration, issue the following **show** commands:

- **show class-of-service interface *interface-name***
- **show interfaces *interface-name* extensive**

Related Documentation

- [Applying a Shaping Rate to Physical Interfaces Overview on page 1578](#)

Oversubscribing Interface Bandwidth

The term *oversubscribing interface bandwidth* means configuring shaping rates (peak information rates [PIRs]) so that their sum exceeds the interface bandwidth.

On Channelized IQ PICs, Gigabit Ethernet IQ PICs, and FRF.15 and FRF.16 link services IQ (LSQ) interfaces on AS PICs, Multiservices PICs, and Multiservices DPCs, you can oversubscribe interface bandwidth. This means that the logical interfaces (and DLCIs within an FRF.15 or FRF.16 bundle) can be oversubscribed when there is leftover bandwidth. In the case of FRF.16 bundle interfaces, the physical interface can be oversubscribed. The oversubscription is capped to the configured PIR. Any unused bandwidth is distributed equally among oversubscribed logical interfaces or DLCIs, or physical interfaces.

For networks that are not likely to experience congestion, oversubscribing interface bandwidth improves network utilization, thereby allowing more customers to be provisioned on a single interface. If the actual data traffic does not exceed the interface bandwidth, oversubscription allows you to sell more bandwidth than the interface can support.

We recommend avoiding oversubscription in networks that are likely to experience congestion. Be cautious not to oversubscribe a service by too much, because this can cause degradation in the performance of the routing platform during congestion. When you configure oversubscription, starvation of some output queues can occur if the actual data traffic exceeds the physical interface bandwidth. You can prevent degradation by using statistical multiplexing to ensure that the actual data traffic does not exceed the interface bandwidth.



NOTE: You cannot oversubscribe interface bandwidth when you configure traffic shaping using the method described in *Applying Scheduler Maps and Shaping Rate to DLCIs and VLANs*.

When configuring oversubscription for FRF.16 bundle interfaces, you can assign traffic control profiles that apply on a physical interface basis. When you apply traffic control profiles to FRF.16 bundles at the *logical* interface level, member link interface bandwidth is underutilized when there is a small proportion of traffic or no traffic at all on an individual DLCI. Support for traffic control features on the FRF.16 bundle physical interface level addresses this limitation.

To configure oversubscription of the interface, perform the following steps:

1. Include the **shaping-rate** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

```
[edit class-of-service traffic-control-profiles profile-name]  
  shaping-rate (percent percentage | rate);
```



NOTE: When configuring oversubscription for FRF.16 bundle interfaces on a physical interface basis, you *must* specify **shaping-rate** as a percentage.

On LSQ interfaces, you can configure the shaping rate as a percentage from 1 through 100.

On IQ and IQ2 interfaces, you can configure the shaping rate as an absolute rate from 1000 through 160,000,000,000 bps.

For all MX Series router and EX Series switch interfaces, the shaping rate can be from 65,535 through 160,000,000,000 bps.

Alternatively, you can configure a shaping rate for a logical interface and oversubscribe the physical interface by including the **shaping-rate** statement at the **[edit class-of-service interfaces interface-name unit logical-unit-number]** hierarchy level. However, with this configuration approach, you cannot independently control the delay-buffer rate, as described in Step 2.



NOTE: For channelized and Gigabit Ethernet IQ interfaces, the **shaping-rate** and **guaranteed-rate** statements are mutually exclusive. You cannot configure some logical interfaces to use a shaping rate and others to use a guaranteed rate. This means there are no service guarantees when you configure a PIR. For these interfaces, you can configure either a PIR or a committed information rate (CIR), but not both.

This restriction does not apply to Gigabit Ethernet IQ2 PICs or LSQ interfaces on AS PICs. For LSQ and Gigabit Ethernet IQ2 interfaces, you can configure both a PIR and a CIR on an interface. For more information about CIRs, see [“Providing a Guaranteed Minimum Rate” on page 1608](#).

For more information about Gigabit Ethernet IQ2 PICs, see *CoS on Enhanced IQ2 PICs Overview*.

2. Optionally, you can base the delay-buffer calculation on a delay-buffer rate. To do this, include the **delay-buffer-rate** statement at the **[edit class-of-service traffic-control-profiles profile-name]** hierarchy level:



NOTE: When configuring oversubscription for FRF.16 bundle interfaces on a physical interface basis, you *must* specify **delay-buffer-rate** as a percentage.

```
[edit class-of-service traffic-control-profiles profile-name]
  delay-buffer-rate (percent percentage | rate);
```

The delay-buffer rate overrides the shaping rate as the basis for the delay-buffer calculation. In other words, the shaping rate or scaled shaping rate is used for delay-buffer calculations only when the delay-buffer rate is not configured.

For LSQ interfaces, if you do not configure a delay-buffer rate, the guaranteed rate (CIR) is used to assign buffers. If you do not configure a guaranteed rate, the shaping rate (PIR) is used in the undersubscribed case, and the scaled shaping rate is used in the oversubscribed case.

On LSQ interfaces, you can configure the delay-buffer rate as a percentage from 1 through 100.

On IQ and IQ2 interfaces, you can configure the delay-buffer rate as an absolute rate from 1000 through 160,000,000,000 bps.

The actual delay buffer is based on the calculations described in [“Configuring Large Delay Buffers for Slower Interfaces” on page 1581](#) and [“Maximum Delay Buffer for NxDS0 Interfaces” on page 1584](#). For an example showing how the delay-buffer rates are applied, see [“Examples: Oversubscribing Interface Bandwidth” on page 1605](#).

Configuring large buffers on relatively slow-speed links can cause packet aging. To help prevent this problem, the software requires that the sum of the delay-buffer rates be less than or equal to the port speed.

This restriction does not eliminate the possibility of packet aging, so you should be cautious when using the **delay-buffer-rate** statement. Though some amount of extra buffering might be desirable for burst absorption, delay-buffer rates should not far exceed the service rate of the logical interface.

If you configure delay-buffer rates so that the sum exceeds the port speed, the configured delay-buffer rate is not implemented for the last logical interface that you configure. Instead, that logical interface receives a delay-buffer rate of zero, and a warning message is displayed in the CLI. If bandwidth becomes available (because another logical interface is deleted or deactivated, or the port speed is increased), the configured delay-buffer-rate is reevaluated and implemented if possible.

If you do not configure a delay-buffer rate or a guaranteed rate, the logical interface receives a delay-buffer rate in proportion to the shaping rate and the remaining delay-buffer rate available. In other words, the delay-buffer rate for each logical interface with no configured delay-buffer rate is equal to:

$$(\text{remaining delay-buffer rate} * \text{shaping rate}) / (\text{sum of shaping rates})$$

where the remaining delay-buffer rate is equal to:

$$(\text{interface speed}) - (\text{sum of configured delay-buffer rates})$$

3. To assign a scheduler map to the logical interface, include the **scheduler-map** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

```
[edit class-of-service traffic-control-profiles profile-name]  
  scheduler-map map-name;
```

For information about configuring schedulers and scheduler maps, see [“Configuring Schedulers” on page 1579](#) and [“Configuring Scheduler Maps” on page 1613](#).

4. Optionally, you can enable large buffer sizes to be configured. To do this, include the **q-pic-large-buffer** statement at the **[edit chassis fpc *slot-number* pic *pic-number*]** hierarchy level:


```
[edit chassis fpc slot-number pic pic-number]
q-pic-large-buffer;
```

If you do not include this statement, the delay-buffer size is more restricted. We recommend restricted buffers for delay-sensitive traffic, such as voice traffic. For more information, see [“Configuring Large Delay Buffers for Slower Interfaces” on page 1581](#).

- To enable scheduling on logical interfaces, include the **per-unit-scheduler** statement at the **[edit interfaces interface-name]** hierarchy level:

```
[edit interfaces interface-name]
per-unit-scheduler;
```

When you include this statement, the maximum number of VLANs supported is 768 on a single-port Gigabit Ethernet IQ PIC. On a dual-port Gigabit Ethernet IQ PIC, the maximum number is 384.

- To enable scheduling for FRF.16 bundles physical interfaces, include the **no-per-unit-scheduler** statement at the **[edit interfaces interface-name]** hierarchy level:

```
[edit interfaces interface-name]
no-per-unit-scheduler;
```

- To apply the traffic-scheduling profile, include the **output-traffic-control-profile** statement at the **[edit class-of-service interfaces interface-name unit logical-unit-number]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
output-traffic-control-profile profile-name;
```

You cannot include the **output-traffic-control-profile** statement in the configuration if any of the following statements are included in the logical interface configuration: **scheduler-map**, **shaping-rate**, **adaptive-shaper**, or **virtual-channel-group** (the last two are valid on Juniper Networks J Series Services Routers only).

[Table 114 on page 1603](#) shows how the bandwidth and delay buffer are allocated in various configurations.

Table 114: Bandwidth and Delay Buffer Allocations by Configuration Scenario

Configuration Scenario	Delay Buffer Allocation
You do not oversubscribe the interface. You do not configure a guaranteed rate. You do not configure a shaping rate. You do not configure a delay-buffer rate.	Logical interface receives the remaining bandwidth and receives a delay buffer in proportion to the remaining bandwidth.
You do not oversubscribe the interface. You do not configure a guaranteed rate. You configure a shaping rate at the [edit class-of-service interfaces interface-name unit logical-unit-number] hierarchy level.	<p>For backward compatibility, the shaped logical interface receives a delay buffer based on the shaping rate. The multiplicative factor depends on whether you include the q-pic-large-buffer statement. For more information, see “Configuring Large Delay Buffers for Slower Interfaces” on page 1581.</p> <p>Unshaped logical interfaces receive the remaining bandwidth and a delay buffer in proportion to the remaining bandwidth.</p>

Table 114: Bandwidth and Delay Buffer Allocations by Configuration Scenario (*continued*)

Configuration Scenario	Delay Buffer Allocation
You oversubscribe the interface. You do not configure a guaranteed rate. You do not configure a shaping rate. You do not configure a delay-buffer rate.	Logical interface receives minimal bandwidth with no guarantees and receives a minimal delay buffer equal to four MTU-sized packets.
You oversubscribe the interface. You configure a shaping rate. You do not configure a guaranteed rate. You do not configure a delay-buffer rate.	<p>Logical interface receives a delay buffer based on the scaled shaping rate:</p> $\text{scaled shaping rate} = (\text{shaping-rate} * [\text{physical interface bandwidth}]) / \text{SUM}(\text{shaping-rates of all logical interfaces on the physical interface})$ <p>The logical interface receives variable bandwidth, depending on how much oversubscription and statistical multiplexing is present. If the amount of oversubscription is low enough that statistical multiplexing does not make all logical interfaces active at the same time and the physical interface bandwidth is not exceeded, the logical interface receives bandwidth equal to the shaping rate. Otherwise, the logical interface receives a smaller amount of bandwidth. In either case, the logical interface bandwidth does not exceed the shaping rate.</p>
You oversubscribe the interface. You configure a shaping rate. You configure a delay-buffer rate.	<p>Logical interface receives a delay buffer based on the delay-buffer rate. For example, on IQ and IQ2 interfaces:</p> <p>delay-buffer-rate <= 10 Mbps: 400-millisecond (ms) delay buffer delay-buffer-rate <= 20 Mbps: 300-ms delay buffer delay-buffer-rate <= 30 Mbps: 200-ms delay buffer delay-buffer-rate <= 40 Mbps: 150-ms delay buffer delay-buffer-rate > 40 Mbps: 100-ms delay buffer</p> <p>On LSQ DLCIs, if total bundle bandwidth < T1 bandwidth:</p> <p>delay-buffer-rate = 1 second</p> <p>On LSQ DLCIs, if total bundle bandwidth >= T1 bandwidth:</p> <p>delay-buffer-rate = 200 ms</p> <p>The multiplicative factor depends on whether you include the q-pic-large-buffer statement. For more information, see "Configuring Large Delay Buffers for Slower Interfaces" on page 1581.</p> <p>The logical interface receives variable bandwidth, depending on how much oversubscription and statistical multiplexing is present. If the amount of oversubscription is low enough that statistical multiplexing does not make all logical interfaces active at the same time and the physical interface bandwidth is not exceeded, the logical interface receives bandwidth equal to the shaping rate. Otherwise, the logical interface receives a smaller amount of bandwidth. In either case, the logical interface bandwidth does not exceed the shaping rate.</p>
You oversubscribe the interface. You do not configure a shaping rate. You configure a guaranteed rate. You configure a delay-buffer rate.	Logical interface receives a delay buffer based on the delay-buffer rate.
You oversubscribe the interface. You do not configure a shaping rate. You do not configure a guaranteed rate. You configure a delay-buffer rate.	This scenario is not allowed. If you configure a delay-buffer rate, the traffic-control profile must also include either a shaping rate or a guaranteed rate.

Table 114: Bandwidth and Delay Buffer Allocations by Configuration Scenario (*continued*)

Configuration Scenario	Delay Buffer Allocation
You oversubscribe the interface. You configure a shaping rate. You configure a guaranteed rate. You do not configure a delay-buffer rate.	<p>Logical interface receives a delay buffer based on the guaranteed rate.</p> <p>This configuration is valid on LSQ interfaces and Gigabit Ethernet IQ2 interfaces only. On channelized interfaces, you cannot configure both a shaping rate (PIR) and a guaranteed rate (CIR).</p>

Verifying Configuration of Bandwidth Oversubscription

To verify your configuration, you can issue this following operational mode commands:

- **show class-of-service interfaces**
- **show class-of-service traffic-control-profile *profile-name***

Examples: Oversubscribing Interface Bandwidth

This section provides two examples: oversubscription of a channelized interface and oversubscription of an LSQ interface.

Oversubscribing a Channelized Interface

Two logical interface units, 0 and 1, are shaped to rates 2 Mbps and 3 Mbps, respectively. The delay-buffer rates are 750 Kbps and 500 Kbps, respectively. The actual delay buffers allocated to each logical interface are 1 second of 750 Kbps and 2 seconds of 500 Kbps, respectively. The 1-second and 2-second values are based on the following calculations:

delay-buffer-rate < [16 x 64 Kbps]): 1 second of delay-buffer-rate
 delay-buffer-rate < [8 x 64 Kbps]): 2 seconds of delay-buffer-rate

For more information about these calculations, see [“Maximum Delay Buffer for NxDSO Interfaces” on page 1584](#).

```
chassis {
  fpc 3 {
    pic 0 {
      q-pic-large-buffer;
    }
  }
}
interfaces {
  t1-3/0/0 {
    per-unit-scheduler;
  }
}
class-of-service {
  traffic-control-profiles {
    tc-profile1 {
      shaping-rate 2m;
      delay-buffer-rate 750k; # 750 Kbps is less than 16 x 64 Kbps
      scheduler-map sched-map1;
    }
    tc-profile2 {
      shaping-rate 3m;
      delay-buffer-rate 500k; # 500 Kbps is less than 8 x 64 Kbps
    }
  }
}
```

```
        scheduler-map sched-map2;
    }
}
interfaces {
    t1-3/0/0 {
        unit 0 {
            output-traffic-control-profile tc-profile1;
        }
        unit 1 {
            output-traffic-control-profile tc-profile2;
        }
    }
}
}
```

**Oversubscribing an
LSQ Interface with
Scheduling Based on
the Logical Interface**

Apply a traffic-control profile to a logical interface representing a DLCI on an FRF.16 bundle:

```
interfaces {
    lsq-1/3/0:0 {
        per-unit-scheduler;
        unit 0 {
            dlci 100;
        }
        unit 1 {
            dlci 200;
        }
    }
}

class-of-service {
    traffic-control-profiles {
        tc_0 {
            shaping-rate percent 100;
            guaranteed-rate percent 60;
            delay-buffer-rate percent 80;
        }
        tc_1 {
            shaping-rate percent 80;
            guaranteed-rate percent 40;
        }
    }
}
interfaces {
    lsq-1/3/0 {
        unit 0 {
            output-traffic-control-profile tc_0;
        }
        unit 1 {
            output-traffic-control-profile tc_1;
        }
    }
}
}
```

Oversubscribing an LSQ Interface with Scheduling Based on the Physical Interface

Apply a traffic-control profile to the physical interface representing an FRF.16 bundle:

```

interfaces {
  lsq-0/2/0:0 {
    no-per-unit-scheduler;
    encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
      dlci 100;
      family inet {
        address 18.18.18.2/24;
      }
    }
  }
}
class-of-service {
  traffic-control-profiles {
    rlsq_tc {
      scheduler-map rlsq;
      shaping-rate percent 60;
      delay-buffer-rate percent 10;
    }
  }
  interfaces {
    lsq-0/2/0:0 {
      output-traffic-control-profile rlsq_tc;
    }
  }
}
scheduler-maps {
  rlsq {
    forwarding-class best-effort scheduler rlsq_scheduler;
    forwarding-class expedited-forwarding scheduler rlsq_scheduler1;
  }
}
schedulers {
  rlsq_scheduler {
    transmit-rate percent 20;
    priority low;
  }
  rlsq_scheduler1 {
    transmit-rate percent 40;
    priority high;
  }
}

```

On an FRF.15 bundle, apply the following configuration:

```
class-of-service {
  traffic-control-profiles {
    rlsq {
      scheduler-map sched_0;
      shaping-rate percent 40;
      delay-buffer-rate percent 50;
    }
  }
  interfaces lsq-2/0/0 {
    unit 0 {
      output-traffic-control-profile rlsq;
    }
  }
}
interfaces lsq-2/0/0 {
  per-unit-scheduler;
  unit 0 {
    encapsulation multilink-frame-relay-end-to-end;
    family inet {
      address 10.1.1.2/32;
    }
  }
}
```

Providing a Guaranteed Minimum Rate

On Gigabit Ethernet IQ PICs, EQ DPCs, Trio MPC/MIC modules, Channelized IQ PICs, and FRF.16 LSQ interfaces on AS PICs, you can configure guaranteed bandwidth, also known as a committed information rate (CIR). This allows you to specify a guaranteed rate for each logical interface. The guaranteed rate is a minimum. If excess physical interface bandwidth is available for use, the logical interface receives more than the guaranteed rate provisioned for the interface.

You cannot provision the sum of the guaranteed rates to be more than the physical interface bandwidth, or the bundle bandwidth for LSQ interfaces. If the sum of the guaranteed rates exceeds the interface or bundle bandwidth, the commit operation does not fail, but the software automatically decreases the rates so that the sum of the guaranteed rates is equal to the available bundle bandwidth.

To configure a guaranteed minimum rate, perform the following steps:

1. Include the **guaranteed-rate** statement at the **[edit class-of-service traffic-control-profile *profile-name*]** hierarchy level:

```
[edit class-of-service traffic-control-profiles profile-name]
  guaranteed-rate (percent percentage | rate) <burst-size bytes>;
```

On LSQ interfaces, you can configure the guaranteed rate as a percentage from 1 through 100.

On IQ and IQ2 interfaces, you can configure the guaranteed rate as an absolute rate from 1000 through 160,000,000,000 bps.



NOTE: For channelized and Gigabit Ethernet IQ interfaces, the **shaping-rate** and **guaranteed-rate** statements are mutually exclusive. You cannot configure some logical interfaces to use a shaping rate and others to use a guaranteed rate. This means there are no service guarantees when you configure a PIR. For these interfaces, you can configure either a PIR or a CIR, but not both.

This restriction does not apply to Gigabit Ethernet IQ2 PICs or LSQ interfaces on AS PICs. For LSQ and Gigabit Ethernet IQ2 interfaces, you can configure both a PIR and a CIR on an interface.

For more information about Gigabit Ethernet IQ2 PICs, see *CoS on Enhanced IQ2 PICs Overview*.

2. Optionally, you can base the delay-buffer calculation on a delay-buffer rate. To do this, include the **delay-buffer-rate** statement [edit class-of-service traffic-control-profiles *profile-name*] hierarchy level:

```
[edit class-of-service traffic-control-profiles profile-name]
  delay-buffer-rate (percent percentage | rate);
```

On LSQ interfaces, you can configure the delay-buffer rate as a percentage from 1 through 100.

On IQ and IQ2 interfaces, you can configure the delay-buffer rate as an absolute rate from 1000 through 160,000,000,000 bps.

The actual delay buffer is based on the calculations described in “[Configuring Large Delay Buffers for Slower Interfaces](#)” on page 1581 and “[Maximum Delay Buffer for NxDSO Interfaces](#)” on page 1584. For an example showing how the delay-buffer rates are applied, see “[Example: Providing a Guaranteed Minimum Rate](#)” on page 1611.

If you do not include the **delay-buffer-rate** statement, the delay-buffer calculation is based on the guaranteed rate, the shaping rate if no guaranteed rate is configured, or the scaled shaping rate if the interface is oversubscribed.

If you do not specify a shaping rate or a guaranteed rate, the logical interface receives a minimal delay-buffer rate and minimal bandwidth equal to four MTU-sized packets.

You can configure a rate for the delay buffer that is higher than the guaranteed rate. This can be useful when the traffic flow might not require much bandwidth in general, but in some cases traffic can be bursty and therefore needs a large buffer.

Configuring large buffers on relatively slow-speed links can cause packet aging. To help prevent this problem, the software requires that the sum of the delay-buffer rates be less than or equal to the port speed. This restriction does not eliminate the possibility of packet aging, so you should be cautious when using the **delay-buffer-rate** statement. Though some amount of extra buffering might be desirable for burst absorption, delay-buffer rates should not far exceed the service rate of the logical interface.

If you configure delay-buffer rates so that the sum exceeds the port speed, the configured delay-buffer rate is not implemented for the last logical interface that you

configure. Instead, that logical interface receives a delay-buffer rate of 0, and a warning message is displayed in the CLI. If bandwidth becomes available (because another logical interface is deleted or deactivated, or the port speed is increased), the configured delay-buffer-rate is reevaluated and implemented if possible.

If the guaranteed rate of a logical interface cannot be implemented, that logical interface receives a delay-buffer rate of 0, even if the configured delay-buffer rate is within the interface speed. If at a later time the guaranteed rate of the logical interface can be met, the configured delay-buffer rate is reevaluated and if the delay-buffer rate is within the remaining bandwidth, it is implemented.

If any logical interface has a configured guaranteed rate, all other logical interfaces on that port that do not have a guaranteed rate configured receive a delay-buffer rate of 0. This is because the absence of a guaranteed rate configuration corresponds to a guaranteed rate of 0 and, consequently, a delay-buffer rate of 0.

3. To assign a scheduler map to the logical interface, include the **scheduler-map** statement at the **[edit class-of-service traffic-control-profiles *profile-name*]** hierarchy level:

```
[edit class-of-service traffic-control-profiles profile-name]  
  scheduler-map map-name;
```

For information about configuring schedulers and scheduler maps, see [“Configuring Schedulers” on page 1579](#) and [“Configuring Scheduler Maps” on page 1613](#).

4. To enable large buffer sizes to be configured, include the **q-pic-large-buffer** statement at the **[edit chassis fpc *slot-number* pic *pic-number*]** hierarchy level:

```
[edit chassis fpc slot-number pic pic-number]  
  q-pic-large-buffer;
```

If you do not include this statement, the delay-buffer size is more restricted. For more information, see [“Configuring Large Delay Buffers for Slower Interfaces” on page 1581](#).

5. To enable scheduling on logical interfaces, include the **per-unit-scheduler** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]  
  per-unit-scheduler;
```

When you include this statement, the maximum number of VLANs supported is 768 on a single-port Gigabit Ethernet IQ PIC. On a dual-port Gigabit Ethernet IQ PIC, the maximum number is 384.

6. To apply the traffic-scheduling profile to the logical interface, include the **output-traffic-control-profile** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number]  
  output-traffic-control-profile profile-name;
```

[Table 115 on page 1611](#) shows how the bandwidth and delay buffer are allocated in various configurations.

Table 115: Bandwidth and Delay Buffer Allocations by Configuration Scenario

Configuration Scenario	Delay Buffer Allocation
You do not configure a guaranteed rate. You do not configure a delay-buffer rate.	Logical interface receives minimal bandwidth with no guarantees and receives a minimal delay buffer equal to 4 MTU-sized packets.
You configure a guaranteed rate. You do not configure a delay-buffer rate.	Logical interface receives bandwidth equal to the guaranteed rate and a delay buffer based on the guaranteed rate. The multiplicative factor depends on whether you include the q-pic-large-buffer statement. For more information, see “Configuring Large Delay Buffers for Slower Interfaces” on page 1581 .
You configure a guaranteed rate. You configure a delay-buffer rate.	Logical interface receives bandwidth equal to the guaranteed rate and a delay buffer based on the delay-buffer rate. The multiplicative factor depends on whether you include the q-pic-large-buffer statement. For more information, see “Configuring Large Delay Buffers for Slower Interfaces” on page 1581 .

Verifying Configuration of Guaranteed Minimum Rate

To verify your configuration, you can issue this following operational mode commands:

- **show class-of-service interfaces**
- **show class-of-service traffic-control-profile *profile-name***

Example: Providing a Guaranteed Minimum Rate

Two logical interface units, **0** and **1**, are provisioned with a guaranteed minimum of 750 Kbps and 500 Kbps, respectively. For logical unit **1**, the delay buffer is based on the guaranteed rate setting. For logical unit **0**, a delay-buffer rate of 500 Kbps is specified. The actual delay buffers allocated to each logical interface are 2 seconds of 500 Kbps. The 2-second value is based on the following calculation:

delay-buffer-rate < [8 x 64 Kbps]): 2 seconds of delay-buffer-rate

For more information about this calculation, see [“Maximum Delay Buffer for NxDS0 Interfaces” on page 1584](#).

```
chassis {
  fpc 3 {
    pic 0 {
      q-pic-large-buffer;
    }
  }
}
interfaces {
  t1-3/0/1 {
    per-unit-scheduler;
  }
}
class-of-service {
  traffic-control-profiles {
    tc-profile3 {
      guaranteed-rate 750k;
      scheduler-map sched-map3;
      delay-buffer-rate 500k; # 500 Kbps is less than 8 x 64 Kbps
    }
  }
}
```

```

    }
    tc-profile4 {
        guaranteed-rate 500k; # 500 Kbps is less than 8 x 64 Kbps
        scheduler-map sched-map4;
    }
}
interfaces {
    t1-3/0/1 {
        unit 0 {
            output-traffic-control-profile tc-profile3;
        }
        unit 1 {
            output-traffic-control-profile tc-profile4;
        }
    }
}
}

```

Associating Schedulers with Fabric Priorities

On Juniper Networks M320 Multiservice Edge Routers and T Series Core Routers only, you can associate a scheduler with a class of traffic that has a specific priority while transiting the fabric. Traffic transiting the fabric can have two priority values: **low** or **high**. To associate a scheduler with a fabric priority, include the **priority** and **scheduler** statements at the **[edit class-of-service fabric scheduler-map]** hierarchy level:

```

[edit class-of-service fabric scheduler-map]
priority (high | low) scheduler scheduler-name;

```



NOTE: For a scheduler that you associate with a fabric priority, include only the **drop-profile-map** statement at the **[edit class-of-service schedulers scheduler-name]** hierarchy level. You cannot include the **buffer-size**, **transmit-rate**, and **priority** statements at that hierarchy level.

For information about associating a forwarding class with a fabric priority, see [“Overriding Fabric Priority Queuing” on page 1535](#).

Example: Associating a Scheduler with a Fabric Priority

Associate a scheduler with a class of traffic that has a specific priority while transiting the fabric:

```

[edit class-of-service]
schedulers {
    fab-be-scheduler {
        drop-profile-map loss-priority low protocol any drop-profile fab-profile-1;
        drop-profile-map loss-priority high protocol any drop-profile fab-profile-2;
    }
    fab-ef-scheduler {
        drop-profile-map loss-priority low protocol any drop-profile fab-profile-3;
        drop-profile-map loss-priority high protocol any drop-profile fab-profile-4;
    }
}
drop-profiles {
    fab-profile-1 {

```

```

        fill-level 100 drop-probability 100;
        fill-level 85 drop-probability 50;
    }
    fab-profile-2 {
        fill-level 100 drop-probability 100;
        fill-level 95 drop-probability 50;
    }
    fab-profile-3 {
        fill-level 75 drop-probability 100;
        fill-level 95 drop-probability 50;
    }
    fab-profile-4 {
        fill-level 100 drop-probability 100;
        fill-level 80 drop-probability 50;
    }
}
fabric {
    scheduler-map {
        priority low scheduler fab-be-scheduler;
        priority high scheduler fab-ef-scheduler;
    }
}

```

Configuration Tasks for Scheduler Maps

- [Configuring Scheduler Maps on page 1613](#)
- [Applying Scheduler Maps to Physical Interfaces on page 1613](#)
- [Applying Scheduler Maps and Shaping Rate to DLCIs and VLANs on page 1614](#)
- [Applying Scheduler Maps to Packet Forwarding Component Queues on page 1620](#)

Configuring Scheduler Maps

After defining a scheduler, you can associate it with a specified forwarding class by including it in a *scheduler map*. To do this, include the **scheduler-maps** statement at the **[edit class-of-service]** hierarchy level:

```

[edit class-of-service]
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}

```

Applying Scheduler Maps to Physical Interfaces

After you have defined a scheduler map, as described in “[Configuring Scheduler Maps on page 1613](#)”, you can apply it to an output interface. Include the **scheduler-map** statement at the **[edit class-of-service interfaces interface-name]** hierarchy level:

```

[edit class-of-service interfaces interface-name]
scheduler-map map-name;

```

Interface wildcards are supported. However, scheduler maps using wildcard interfaces are not checked against routing device interfaces at commit time and can result in a configuration that is incompatible with installed hardware. Fully specified interfaces, on

the other hand, check the configuration against the hardware and report errors or warning if the hardware does not support the configuration.

Generally, you can associate schedulers with physical interfaces only. For some IQ interfaces, you can also associate schedulers with the logical interface. For more information, see *Applying Scheduler Maps and Shaping Rate to DLCIs and VLANs*.



NOTE: For original Channelized OC12 PICs, limited CoS functionality is supported. For more information, contact Juniper Networks customer support.

When you apply a scheduler map to a physical interface, or when you modify the configuration of a scheduler map that is already applied to a physical interface, packets already in the output queues of the interface might get dropped. The amount of packet loss is not deterministic and depends on the offered traffic load at the time you apply or modify the scheduler map.

Applying Scheduler Maps and Shaping Rate to DLCIs and VLANs

By default, output scheduling is not enabled on logical interfaces. Logical interfaces without shaping configured share a default scheduler. This scheduler has a committed information rate (CIR) that equals 0. (The CIR is the guaranteed rate.) The default scheduler has a peak information rate (PIR) that equals the physical interface shaping rate.



NOTE: If you apply a shaping rate, you must keep in mind that the transit statistics for physical interfaces are obtained from the packet forwarding engine, but the traffic statistics are supplied by the PIC. Therefore, if shaping is applied to the PIC, the count of packets in the transit statistics fields do not always agree with the counts in the traffic statistics. For example, the IPv6 transit statistics will not necessarily match the traffic statistics on the interface. However, at the logical interface (DLCI) level, both transit and traffic statistics are obtained from the Packet Forwarding Engine and will not show any difference.

Logical interface scheduling (also called *per-unit scheduling*) allows you to enable multiple output queues on a logical interface and associate an output scheduler and shaping rate with the queues. You can configure logical interface scheduling on the following PICs:

- Adaptive Services PIC, on link services IQ (**lsq-**) interfaces
- Channelized E1 IQ PIC
- Channelized OC3 IQ PIC
- Channelized OC12 IQ PIC (Per-unit scheduling is not supported on T1 interfaces configured on this PIC.)
- Channelized STM1 IQ PIC
- Channelized T3 IQ PIC

- E3 IQ PIC
- Gigabit Ethernet IQ PIC
- Gigabit Ethernet IQ2 PIC
- IQE PICs
- Link services PIM (**ls-** interfaces) on J Series routers

For Juniper Networks J Series Services Routers only, you can configure per-unit scheduling for virtual channels. For more information, see the J Series router documentation.

For Channelized and Gigabit Ethernet IQ PICs only, you can configure a shaping rate for a VLAN or DLCI and oversubscribe the physical interface by including the **shaping-rate** statement at the **[edit class-of-service traffic-control-profiles]** hierarchy level. With this configuration approach, you can independently control the delay-buffer rate, as described in [“Oversubscribing Interface Bandwidth” on page 1600](#).

Physical interfaces (for example, **t3-0/0/0**, **t3-0/0/0:0**, and **ge-0/0/0**) support scheduling with any encapsulation type pertinent to that physical interface. For a single port, you cannot apply scheduling to the physical interface if you apply scheduling to one or more of the associated logical interfaces.

For Gigabit Ethernet IQ2 PICs only, you can configure hierarchical traffic shaping, meaning the shaping is performed on both the physical interface and the logical interface. You can also configure input traffic scheduling and shared scheduling. For more information, see *CoS on Enhanced IQ2 PICs Overview*.

Logical interfaces (for example, **t3-0/0/0.0**, **ge-0/0/0.0**, and **t1-0/0/0:0.1**) support scheduling on DLCIs or VLANs only. Furthermore, logical interface scheduling is not supported on PICs that do not have IQ.



NOTE: In the Junos OS implementation, the term *logical interfaces* generally refers to interfaces you configure by including the unit statement at the **[edit interfaces interface-name]** hierarchy level. As such, logical interfaces have the *logical* descriptor at the end of the interface name, as in **ge-0/0/0.1** or **t1-0/0/0:0.1**, where the logical unit number is 1.

Although channelized interfaces are generally thought of as logical or virtual, the Junos OS sees T3, T1, and NxDS0 interfaces within a channelized IQ PIC as physical interfaces. For example, both **t3-0/0/0** and **t3-0/0/0:1** are treated as physical interfaces by the Junos OS. In contrast, **t3-0/0/0.2** and **t3-0/0/0:1.2** are considered logical interfaces because they have the .2 at the end of the interface names.

Within the **[edit class-of-service]** hierarchy level, you cannot use the *.logical* descriptor when you assign properties to logical interfaces. Instead, you must include the unit statement in the configuration. For example:

```
[edit class-of-service]
user@host# set interfaces t3-0/0/0 unit 0 scheduler-map map1
```

Table 116 on page 1616 shows the interfaces that support transmission scheduling.

Table 116: Transmission Scheduling Support by Interfaces Type

Interface Type	PIC Type	Supported	Examples
IQ PICs			
Channelized interfaces configured on IQ PICs	Channelized DS3 IQ	Yes	Example of supported configuration: [edit class-of-service interfaces t1-0/0/0:1] scheduler-map map-1;
Logical interfaces (DLCIs and VLANs only) configured on IQ PICs	Gigabit Ethernet IQ with VLAN tagging enabled	Yes	Example of supported configuration: [edit class-of-service interfaces ge-0/0/0 unit 1] scheduler-map map-1;
	E3 IQ with Frame Relay encapsulation	Yes	Example of supported configuration: [edit class-of-service interfaces e3-0/0/0 unit 1] scheduler-map map-1;
	Channelized OC3 IQ with Frame Relay encapsulation	Yes	Example of supported configuration: [edit class-of-service interfaces t1-1/0/0:1 unit 0] scheduler-map map-1;
	Channelized STM1 IQ with Frame Relay encapsulation	Yes	Example of supported configuration: [edit class-of-service interfaces e1-0/0/0:1 unit 1] scheduler-map map-1;
	Channelized T3 IQ with Frame Relay encapsulation	Yes	Example of supported configuration: [edit class-of-service interfaces t1-0/0/0 unit 1] scheduler-map map-1;
Logical interfaces configured on IQ PICs (interfaces that are not DLCIs or VLANs)	E3 IQ PIC with Cisco HDLC encapsulation	No	Example of unsupported configuration: [edit class-of-service interfaces e3-0/0/0 unit 1] scheduler-map map-1;
Non-IQ PICs			
Physical interfaces	T3	Yes	Example of supported configuration: [edit class-of-service interfaces t3-0/0/0] scheduler-map map-1;
Channelized OC12 PIC	Channelized OC12	Yes	Example of supported configuration: [edit class-of-service interfaces t3-0/0/0:1] scheduler-map map-1;

Table 116: Transmission Scheduling Support by Interfaces Type (*continued*)

Interface Type	PIC Type	Supported	Examples
Channelized interfaces (except the Channelized OC12 PIC)	Channelized STM1	No	Example of unsupported configuration: [edit class-of-service interfaces e1-0/0/0:1] scheduler-map map-1;
Logical interfaces	Fast Ethernet	No	Example of unsupported configuration: [edit class-of-service interfaces fe-0/0/0 unit 1] scheduler-map map-1;
	Gigabit Ethernet	No	Example of unsupported configuration: [edit class-of-service interfaces ge-0/0/0 unit 0] scheduler-map map-1;
	Channelized OC12	No	Example of unsupported configuration: [edit class-of-service interfaces t3-0/0/0:0 unit 2] scheduler-map map-1;

To configure transmission scheduling on logical interfaces, perform the following steps:

1. Enable scheduling on the interface by including the **per-unit-scheduler** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]  
per-unit-scheduler;
```

When you include this statement, the maximum number of VLANs supported is 768 on a single-port Gigabit Ethernet IQ PIC. On a dual-port Gigabit Ethernet IQ PIC, the maximum number is 384.

2. Associate a scheduler with the interface by including the **scheduler-map** statement at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number*] hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number]  
scheduler-map map-name;
```

3. Configure shaping on the interface by including the **shaping-rate** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number]  
  shaping-rate rate;
```

By default, the logical interface bandwidth is the average of unused bandwidth for the number of logical interfaces that require default bandwidth treatment. You can specify a peak bandwidth rate in bps, either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000). The range is from 1000 through 32,000,000,000 bps. For the IQ2 Gigabit Ethernet PIC, the minimum is 80,000 bps, and for the IQ2 10 Gigabit Ethernet PIC, the minimum is 160,000 bps.

For FRF.16 bundles on link services interfaces, only shaping rates based on percentage are supported.



NOTE: If you apply a shaping rate, you must keep in mind that the transit statistics for physical interfaces are obtained from the packet forwarding engine, but the traffic statistics are supplied by the PIC. Therefore, if shaping is applied to the PIC, the count of packets in the transit statistics fields do not always agree with the counts in the traffic statistics. For example, the IPv6 transit statistics will not necessarily match the traffic statistics on the interface. However, at the logical interface (DLCI) level, both transit and traffic statistics are obtained from the Packet Forwarding Engine and will not show any difference.

Example: Applying Scheduler Maps and Shaping Rate to DLCIs and VLANs

Associate the scheduler **sched-map-logical-0** with logical interface **unit 0** on physical interface **t3-1/0/0**, and allocate 10 Mbps of transmission bandwidth to the logical interface.

Associate the scheduler **sched-map-logical-1** with logical interface **unit 1** on physical interface **t3-1/0/0**, and allocate 20 Mbps of transmission bandwidth to the logical interface.

The allocated bandwidth is shared among the individual forwarding classes in the scheduler map. Although these schedulers are configured on a single physical interface, they are independent from each other. Traffic on one logical interface unit does not affect the transmission priority, bandwidth allocation, or drop behavior on the other logical interface unit.

For another example, see the *Junos OS Feature Guides*.

```
[edit interfaces]  
t3-1/0/0:1 {  
  encapsulation frame-relay;  
  per-unit-scheduler;  
}
```



```

[edit class-of-service]
interfaces {
  t3-1/0/0:1 {
    unit 0 {
      scheduler-map sched-map-logical-0;
      shaping-rate 10m;
    }
    unit 1 {
      scheduler-map sched-map-logical-1;
      shaping-rate 20m;
    }
  }
}
scheduler-maps {
  sched-map-logical-0 {
    forwarding-class best-effort scheduler sched-best-effort-0;
    forwarding-class assured-forwarding scheduler sched-bronze-0;
    forwarding-class expedited-forwarding scheduler sched-silver-0;
    forwarding-class network-control scheduler sched-gold-0;
  }
  sched-map-logical-1 {
    forwarding-class best-effort scheduler sched-best-effort-1;
    forwarding-class assured-forwarding scheduler sched-bronze-1;
    forwarding-class expedited-forwarding scheduler sched-silver-1;
    forwarding-class network-control scheduler sched-gold-1;
  }
}
schedulers {
  sched-best-effort-0 {
    transmit-rate 4m;
  }
  sched-bronze-0 {
    transmit-rate 3m;
  }
  sched-silver-0 {
    transmit-rate 2m;
  }
  sched-gold-0 {
    transmit-rate 1m;
  }
  sched-best-effort-1 {
    transmit-rate 8m;
  }
  sched-bronze-1 {
    transmit-rate 6m;
  }
  sched-silver-1 {
    transmit-rate 4m;
  }
  sched-gold-1 {
    transmit-rate 2m;
  }
}

```

Applying Scheduler Maps to Packet Forwarding Component Queues

On Intelligent Queuing (IQ) and Intelligent Queuing 2 (IQ2) interfaces, the traffic that is fed from the packet forwarding components into the PIC uses low packet loss priority (PLP) by default and is distributed evenly across the four chassis queues (not PIC queues), regardless of the scheduling configuration for each logical interface. This default behavior can cause traffic congestion.

The default chassis scheduler allocates resources for queue 0 through queue 3, with 25 percent of the bandwidth allocated to each queue. When you configure the chassis to use more than four queues, you must configure and apply a custom chassis scheduler to override the default. To apply a custom chassis scheduler, include the **scheduler-map-chassis** statement at the **[edit class-of-service interfaces at-*fpc/pic*/*]** hierarchy level.

To control the aggregated traffic transmitted from the chassis queues into the PIC, you can configure the chassis queues to derive their scheduling configuration from the associated logical interface's. Include the **scheduler-map-chassis derived** statement at the **[edit class-of-service interfaces *type-fpc/pic*/*]** hierarchy level:

```
[edit class-of-service interfaces type-fpc/pic/*]  
scheduler-map-chassis derived;
```



CAUTION: If you include the **scheduler-map-chassis derived** statement in the configuration, packet loss might occur when you subsequently add or remove logical interfaces at the **[edit interfaces *interface-name*]** hierarchy level.

When fragmentation occurs on the egress interface, the first set of packet counters displayed in the output of the **show interfaces queue** command show the post-fragmentation values. The second set of packet counters (under the **Packet Forwarding Engine Chassis Queues** field) show the pre-fragmentation values. For more information about the **show interfaces queue** command, see the *Junos OS Operational Mode Commands*.

You can include both the **scheduler-map** and the **scheduler-map-chassis derived** statements in the same interface configuration. The **scheduler-map** statement controls the scheduler inside the PIC, while the **scheduler-map-chassis derived** statement controls the aggregated traffic transmitted into the entire PIC. For the Gigabit Ethernet IQ PIC, include both statements.

For more information about the **scheduler-map** statement, see [“Applying Scheduler Maps to Physical Interfaces” on page 1613](#). For information about logical interface scheduling configuration, see *Applying Scheduler Maps and Shaping Rate to DLCIs and VLANs*.

Generally, when you include the **scheduler-map-chassis** statement in the configuration, you must use an interface wildcard for the interface name, as in ***type-fpc/pic*/***. The wildcard must use this format—for example, **ge-1/2/***, which means all interfaces on FPC slot 1, PIC slot 2. There is one exception—you can apply the chassis scheduler map to a specific interface on the Gigabit Ethernet IQ PIC only.

According to Junos OS wildcard rules, specific interface configurations override wildcard configurations. For chassis scheduler map configuration, this rule does not apply; instead, specific interface CoS configurations are added to the chassis scheduler map configuration. For more information about how wildcards work with chassis scheduler maps, see [“Examples: Scheduling Packet Forwarding Component Queues” on page 1621](#). For general information about wildcards, see the *Junos OS System Basics Configuration Guide*.



NOTE: The interface applies wildcard configuration only if you do not add any specific configuration. If you add the specific interface configuration, then the interface deletes the applied wildcard configuration and applies the specified configuration.

For more information, see the following sections:

- [Applying Custom Schedulers to Packet Forwarding Component Queues on page 1621](#)
- [Examples: Scheduling Packet Forwarding Component Queues on page 1621](#)

Applying Custom Schedulers to Packet Forwarding Component Queues

Optionally, you can apply a custom scheduler to the chassis queues instead of configuring the chassis queues to automatically derive their scheduling configuration from the logical interfaces on the PIC.

To assign a custom scheduler to the packet forwarding component queues, include the **scheduler-map-chassis** statement at the **[edit class-of-service interfaces type-fpc/pic]** hierarchy level:

```
[edit class-of-service interfaces type-fpc/pic/*]
  scheduler-map-chassis map-name;
```

When you apply a custom scheduler map to packet forwarding component queues, or when you modify the configuration of a custom scheduler map that is already applied to packet forwarding component queues, packets already in the chassis queues might be dropped. The amount of packet loss is not deterministic and depends on the offered traffic load at the time you apply or modify the custom scheduler map.

For information about defining the scheduler map referenced by **map-name**, see [“Configuring Scheduler Maps” on page 1613](#).

Examples: Scheduling Packet Forwarding Component Queues

Applying a Chassis Scheduler Map to a 2-Port IQ PIC

Apply a chassis scheduler map to interfaces **ge-0/1/0** and **ge-0/1/1**.

According to customary wildcard rules, the **ge-0/1/0** configuration overrides the **ge-0/1/*** configuration, implying that the chassis scheduler map **MAP1** is not applied to **ge-0/1/0**. However, the wildcard rule is not obeyed in this case; the chassis scheduler map applies to both interfaces **ge-0/1/0** and **ge-0/1/1**.

```
[edit]
class-of-service {
  interfaces {
    ge-0/1/0 {
```

```
    unit 0 {
      classifiers {
        inet-precedence default;
      }
    }
  }
  ge-0/1/* {
    scheduler-map-chassis derived;
  }
}
```

Not Recommended:
Using a Wildcard for
Gigabit Ethernet IQ
Interfaces When
Applying a Chassis
Scheduler Map

On a Gigabit Ethernet IQ PIC, you can apply the chassis scheduler map at both the specific interface level and the wildcard level. We do not recommend this because the wildcard chassis scheduler map takes precedence, which might not be the desired effect. For example, if you want to apply the chassis scheduler map MAP1 to port 0 and MAP2 to port 1, we do not recommend the following:

```
[edit class-of-service]
interfaces {
  ge-0/1/0 {
    scheduler-map-chassis MAP1;
  }
  ge-0/1/* {
    scheduler-map-chassis MAP2;
  }
}
```

Recommended:
Identifying Gigabit
Ethernet IQ Interfaces
Individually When
Applying a Chassis
Scheduler Map

Instead, we recommend this configuration:

```
[edit class-of-service]
interfaces {
  ge-0/1/0 {
    scheduler-map-chassis MAP1;
  }
  ge-0/1/1 {
    scheduler-map-chassis MAP2;
  }
}
```

Configuring Two T3
Interfaces on a
Channelized DS3 IQ
PIC

```
[edit interfaces]
ct3-3/0/0 {
  no-partition interface-type t3; # use entire port 0 as T3
}
ct3-3/0/1 {
  no-partition interface-type t3; # use entire port 1 as T3
}
t3-3/0/0 {
  unit 0 {
    family inet {
      address 10.0.100.1/30;
    }
  }
}
t3-3/0/1 {
  unit 0 {
```

```

        family inet {
            address 10.0.101.1/30;
        }
    }
}

```

Applying Normal Schedulers to Two T3 Interfaces

Configure a scheduler for the aggregated traffic transmitted into both T3 interfaces.

```

[edit class-of-service]
interfaces {
    t3-3/0/0 {
        scheduler-map sched-qct3-0;
    }
    t3-3/0/1 {
        scheduler-map sched-qct3-1;
    }
}
scheduler-maps {
    sched-qct3-0 {
        forwarding-class best-effort scheduler be-qct3-0;
        forwarding-class expedited-forwarding scheduler ef-qct3-0;
        forwarding-class assured-forwarding scheduler as-qct3-0;
        forwarding-class network-control scheduler nc-qct3-0;
    }
    sched-qct3-1 {
        forwarding-class best-effort scheduler be-qct3-1;
        forwarding-class expedited-forwarding scheduler ef-qct3-1;
        forwarding-class assured-forwarding scheduler as-qct3-1;
        forwarding-class network-control scheduler nc-qct3-1;
    }
    sched-chassis-to-q {
        forwarding-class best-effort scheduler be-chassis;
        forwarding-class expedited-forwarding scheduler ef-chassis;
        forwarding-class assured-forwarding scheduler as-chassis;
        forwarding-class network-control scheduler nc-chassis;
    }
}
schedulers {
    be-qct3-0 {
        transmit-rate percent 40;
    }
    ef-qct3-0 {
        transmit-rate percent 30;
    }
    as-qct3-0 {
        transmit-rate percent 20;
    }
    nc-qct3-0 {
        transmit-rate percent 10;
    }
    ...
}

```

Applying a Chassis Scheduler to Two T3 Interfaces

Bind a scheduler to the aggregated traffic transmitted into the entire PIC. The chassis scheduler controls the traffic from the packet forwarding components feeding the interface **t3-3/0/***.

```
[edit class-of-service]
interfaces {
  t3-3/0/* {
    scheduler-map-chassis derived;
  }
}
```

Not Recommended: Using a Wildcard for Logical Interfaces When Applying a Scheduler

Do not apply a scheduler to a logical interface using a wildcard. For example, if you configure a logical interface (unit) with one parameter, and apply a scheduler map to the interface using a wildcard, the logical interface will not apply the scheduler. The following configuration will commit correctly but will not apply the scheduler map to interface **ge-3/0/0.0**:

```
[edit class-of-service]
interfaces {
  ge-3/0/* {
    unit 0 {
      scheduler-map MY_SCHED_MAP;
    }
  }
  ge-3/0/0 {
    unit 0 {
      shaping-rate 100m;
    }
  }
}
```

Recommended: Identifying Logical Interfaces Individually When Applying a Scheduler

Always apply the scheduler to a logical interface without the wildcard:

```
[edit class-of-service]
interfaces {
  ge-3/0/0 {
    unit 0 {
      scheduler-map MY_SCHED_MAP;
      shaping-rate 100m;
    }
  }
}
```



NOTE: This same wildcard behavior applies to classifiers and rewrites as well as schedulers.

Configuration Statements for Schedulers

- [\[edit class-of-service\] Hierarchy Level on page 1624](#)
- [\[edit interfaces\] Hierarchy Level on page 1654](#)

[edit class-of-service] Hierarchy Level

```
class-of-service {
```

```

classifiers {
  type classifier-name {
    forwarding-class class-name {
      loss-priority (high | low | medium-high | medium-low) code-points [ aliases bits ];
    }
    import (classifier-name | default);
  }
}
code-point-aliases {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) {
    alias-name bits;
  }
}
drop-profiles {
  profile-name {
    fill-level percentage drop-probability percentage;
    interpolate {
      drop-probability value;
      fill-level value;
    }
  }
}
fabric {
  scheduler-map {
    priority (high | low) scheduler scheduler-name;
  }
}
forwarding-class-map {
  map-name {
    class class-name queue-num queue-number <restricted-queue queue-number>;
  }
}
forwarding-classes {
  class class-name policing-priority (normal | premium) queue-num queue-number
  priority (high | low);
  queue queue-number class-name policing-priority (normal | premium) priority (high |
  low);
}
forwarding-policy {
  class class-name {
    classification-override {
      forwarding-class class-name;
    }
  }
  next-hop-map map-name {
    forwarding-class class-name {
      discard;
      lsp-next-hop [ lsp-regular-expressions ];
      next-hop [ next-hop-names ];
      non-lsp-next-hop;
    }
  }
}
fragmentation-maps {
  map-name {
    forwarding-class class-name {

```

```

        drop-timeout milliseconds;
        fragment-threshold bytes;
        multilink-class number;
        no-fragmentation;
    }
}
host-outbound-traffic {
    dscp-code-point value;
    forwarding-class class-name;
    ieee-802.1 {
        default value;
        rewrite-rules;
    }
    tcp {
        raise-internet-control-priority;
    }
}
interfaces {
    ... the interfaces subhierarchy appears after the main [edit class-of-service] hierarchy
    ...
}
restricted-queues {
    forwarding-class class-name queue-number;
}
rewrite-rules {
    (dscp | dscp-ipv6 | exp | frame-relay-de | ieee-802.1 | ieee-802.1ad | inet-precedence)
    rewrite-rule {
        forwarding-class class-name {
            loss-priority level code-point (alias | bits);
        }
        import (rewrite-rule | default);
    }
}
routing-instances routing-instance-name {
    classifiers {
        dscp (classifier-name | default);
        dscp-ipv6 (classifier-name | default);
        exp (classifier-name | default);
        ieee-208.1 (classifier-name | default | encapsulated | vlan-tag (inner | outer));
    }
}
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
schedulers {
    scheduler-name {
        adjust-minimum value;
        adjust-percent value;
        buffer-size (exact | percent percentage | remainder);
        drop-profile-map loss-priority (any | high | low | medium-high | medium-low)
            protocol any;
        excess-priority (high | low | medium-high | medium-low);
    }
}

```



```

    excess-rate (percent percentage | proportion proportion);
    priority (high | low | medium-high | medium-low | strict-high);
    shaping-rate (bps | percent percentage | burst-size size);
    transmit-rate (bps | percent percentage | remainder) <exact | rate-limit>;
  }
}
traceoptions {
  file <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
traffic-control-profiles {
  profile-name {
    adjust-minimum rate;
    delay-buffer-rate (bps | cps cps | percent percentage);
    excess-rate (percent percentage | proportion value);
    guaranteed-rate (bps | percent percentage) <burst-size bytes>;
    overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
    scheduler-map map-name;
    shaping-rate (bps | percent percentage) <burst-size bytes>;
  }
}
tri-color;
}

class-of-service {
  interfaces {
    interface-name {
      excess-bandwidth-share (equal | proportional value);
      input-excess-bandwidth-share (equal | proportional value);
      input-scheduler-map map-name;
      input-shaping-rate bps;
      input-traffic-control-profile profile-name;
      output-forwarding-class-map map-name;
      output-traffic-control-profile profile-name;
      scheduler-map map-name;
      scheduler-map-chassis (map-name | derived);
      shaping-rate bps;
      unit (logical-unit-number | *) {
        classifiers {
          dscp (classifier-name | default) {
            family [ inet mpls ];
          }
          dscp-ipv6 (classifier-name | default) {
            family [ inet mpls ];
          }
          exp (classifier-name | default);
          ieee-208.1 (classifier-name | default) <vlan-tag (inner | outer)>;
          ieee-208.1ad (classifier-name | default);
          inet-precedence (classifier-name | default);
        }
        forwarding-class class-name;
        input-scheduler-map map-name;
        input-shaping-rate bps;
        input-traffic-control-profile profile-name shared-instance instance-name;
      }
    }
  }
}


```

```
    loss-priority-maps {
      (map-name | default);
    }
    loss-priority-rewrites {
      (map-name | default);
    }
    output-forwarding-class-map map-name;
    output-traffic-control-profile profile-name shared-instance instance-name;
    rewrite-rules {
      dscp (rule-name | default) <protocol mpls>;
      dscp-ipv6 (rule-name | default);
      exp (rule-name | default) <protocol [ mpls-any | mpls-inet-both |
        mpls-inet-both-non-vpn ]>;
      exp-push-push-push default;
      exp-swap-push-push default;
      ieee-802.1 (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
      ieee-802.1ad (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
      inet-precedence (rewrite-name | default) <protocol mpls>;
    }
    scheduler-map map-name;
    shaping-rate bps;
    translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 |
      to-exp-from-exp | to-inet-precedence-from-inet-precedence) table-name;
  }
}
interface-set interface-set-name {
  excess-bandwidth-share (equal | proportional value);
  input-excess-bandwidth-share (equal | proportional value);
  input-traffic-control-profile profile-name;
  input-traffic-control-profile-remaining profile-name;
  internal-node;
  output-traffic-control-profile profile-name;
  output-traffic-control-profile-remaining profile-name;
}
}
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

buffer-size (Schedulers)

Syntax	buffer-size (percent <i>percentage</i> remainder temporal <i>microseconds</i>);
Hierarchy Level	[edit class-of-service schedulers <i>scheduler-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches. Statement introduced in Junos OS Release 12.2 for ACX Series Routers.
Description	Specify buffer size.
	<div>  <p>NOTE: On PTX Series Packet Transport Switches, buffer-size cannot be configured on rate-limited queues.</p> </div>
Default	If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent, respectively.
Options	<p>percent <i>percentage</i>—Buffer size as a percentage of the total buffer. Range: 0 through 100</p> <p>remainder—Remaining buffer available.</p> <p>temporal <i>microseconds</i>—Buffer size as a temporal value. The queuing algorithm starts dropping packets when it queues more than a computed number of bytes. This maximum is computed by multiplying the logical interface speed by the configured temporal value. Range: The ranges vary by platform. See Table 109 on page 1580.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Scheduler Buffer Size on page 1579 • <i>Example: Configuring CoS for a PBB Network</i>

delay-buffer-rate

Syntax	delay-buffer-rate (percent <i>percentage</i> <i>rate</i>);
Hierarchy Level	[edit class-of-service traffic-control-profiles <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	For Gigabit Ethernet IQ, Channelized IQ PICs, and FRF.15 and FRF.16 LSQ interfaces only, base the delay-buffer calculation on a delay-buffer rate.
Default	If you do not include this statement, the delay-buffer calculation is based on the guaranteed rate if one is configured, or the shaping rate if no guaranteed rate is configured. For more information, see Table 114 on page 1603 .
Options	<p>percent<i>percentage</i>—For LSQ interfaces, delay-buffer rate as a percentage of the available interface bandwidth.</p> <p>Range: 1 through 100 percent</p> <p>rate—For IQ and IQ2 interfaces, delay-buffer rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).</p> <p>Range: 1000 through 160,000,000,000 bps</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Oversubscribing Interface Bandwidth on page 1600• Providing a Guaranteed Minimum Rate on page 1608• Configuring Traffic Control Profiles for Shared Scheduling and Shaping• output-traffic-control-profile on page 1639

drop-profile-map (Schedulers)

Syntax	drop-profile-map loss-priority (any low medium-low medium-high high) protocol (any non-tcp tcp) drop-profile (Schedulers) <i>profile-name</i> ;
Hierarchy Level	[edit class-of-service schedulers <i>scheduler-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches. Statement introduced in Junos OS Release 12.2 for ACX Series Routers.
Description	Define the loss-priority value for a drop profile. The statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Default Schedulers Overview on page 1575 • Configuring Drop Profile Maps for Schedulers on page 1590

excess-priority


Syntax	<code>excess-priority [low medium-low medium-high high none];</code>
Hierarchy Level	[edit class-of-service schedulers <i>scheduler-name</i>]
Release Information	Statement introduced in Junos OS Release 9.3. Option none introduced in Junos OS Release 11.4.
Description	Determine the priority of excess bandwidth traffic on a scheduler.



NOTE: For Link Services IQ (LSQ) PICs or Multiservices PIC (MS-PICs), the **excess-priority** statement is allowed for consistency, but ignored. If an explicit priority is not configured for these interfaces, a default low priority is used. This default priority is also used in the excess region.

Options	<p>low—Excess traffic for this scheduler has low priority.</p> <p>medium-low—Excess traffic for this scheduler has medium-low priority.</p> <p>medium-high—Excess traffic for this scheduler has medium-high priority.</p> <p>high—Excess traffic for this scheduler has high priority.</p> <p>none—System does not demote the priority of guaranteed traffic when the bandwidth exceeds the shaping rate or the guaranteed rate.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Excess Bandwidth Sharing on IQE PICs</i>• Bandwidth Sharing on Nonqueueing Packet Forwarding Engines Overview on page 1666• <i>Managing Excess Bandwidth Distribution on Static Interfaces on MICs and MPCs</i>

excess-rate

Syntax	<code>excess-rate (percent <i>percentage</i> proportion <i>value</i>);</code>
Hierarchy Level	[edit class-of-service schedulers <i>scheduler-name</i>], [edit class-of-service traffic-control-profiles <i>traffic-control-profile-name</i>]
Release Information	Statement introduced in Junos OS Release 9.3. Application to the Multiservices PIC added in Junos OS Release 9.5. Application to the Trio MPC/MIC interfaces added in Junos OS Release 10.1. Statement introduced in Junos OS Release 12.1X48R2 for PTX Series Packet Transport Switches.
Description	For an Enhanced IQ PIC, Multiservices PIC, or MX Series routers with MPC/MIC interface, and T4000 routers with Type 5 FPCs and EX Series switches, determine the percentage or proportion of excess bandwidth traffic to share.
<div>  <p>NOTE: The proportion option provides a greater range of values over the percent option and hence influences the priorities assigned to the queues.</p> </div>	
Options	<p>percentage—Percentage of the excess bandwidth to share. Range: 0 through 100 percent Default: Excess bandwidth is shared in proportion to the configured transmit rate of each queue.</p> <p>value—(M Series, MX Series, T Series routers and EX Series switches only) Proportion of the excess bandwidth to share. Option available at the [edit class-of-service traffic-class-profiles <i>traffic-control-profile-name</i>] hierarchy level only. Range: 0 through 1000</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Scheduler Transmission Rate on page 1591 • Configuring Excess Bandwidth Sharing on IQE PICs • Allocating Excess Bandwidth Among Frame Relay DLCIs on Multiservices PICs • Managing Excess Bandwidth Distribution on Static Interfaces on MICs and MPCs

fabric (Class-of-Service)

Syntax	<pre>fabric { scheduler-map { priority (high low) scheduler scheduler-name; } }</pre>
Hierarchy Level	[edit class-of-service]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS 11.4 for EX Series switches.
Description	<p>Define CoS parameters of the switch fabric. For M320 and T Series routers only, associate a scheduler with a fabric priority.</p> <p>On EX Series switches, this statement is supported only on EX8200 standalone switches and EX8200 Virtual Chassis.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• See Associating Schedulers with Fabric Priorities on page 1612.

forwarding-class (Interfaces)

Syntax	<pre>forwarding-class class-name;</pre>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.2 for ACX Series routers.
Description	Associate a forwarding class configuration or default mapping with a specific interface.
Options	<i>class-name</i> —Name of the forwarding class.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Applying Forwarding Classes to Interfaces on page 1531• <i>Configuring Fixed Classification on an ATM IMA Pseudowire</i>• <i>Example: Configuring Fixed Classification on an ATM IMA Pseudowire</i>

guaranteed-rate

Syntax	<code>guaranteed-rate (percent <i>percentage</i> <i>rate</i>) <burst-size <i>bytes</i>>;</code>
Hierarchy Level	[edit class-of-service traffic-control-profiles <i>profile-name</i>]
Release Information	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Option burst-size introduced for Enhanced Queuing (EQ) DPCs in Junos OS Release 9.4.</p> <p>Option burst-size introduced for MPC/MIC modules in Junos OS Release 11.4.</p> <p>Option burst-size introduced for IQ2 and IQ2E interfaces in Junos OS Release 12.3</p>
Description	For Gigabit Ethernet IQ, Channelized IQ PICs, AS PIC FRF.16 LSQ interfaces, and EQ DPCs only, configure a guaranteed minimum rate. You can also configure an optional burst size for a logical interface on EQ DPCs and on IQ2 and IQ2E PICs. This can help to ensure that higher priority services do not starve lower priority services.
Default	If you do not include this statement and you do not include the delay-buffer-rate statement, the logical interface receives a minimal delay-buffer rate and minimal bandwidth equal to 2 MTU-sized packets.
Options	<p>percent <i>percentage</i>—For LSQ interfaces, guaranteed rate as a percentage of the available interface bandwidth.</p> <p>Range: 1 through 100 percent</p> <p><i>rate</i>—For IQ and IQ2 interfaces, guaranteed rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).</p> <p>Range: 1000 through 160,000,000,000 bps</p> <p>burst-size <i>bytes</i>—(Optional) Maximum burst size, in bytes.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Providing a Guaranteed Minimum Rate on page 1608 • Configuring Traffic Control Profiles for Shared Scheduling and Shaping • output-traffic-control-profile on page 1639

interfaces

```
Syntax  interfaces {
        interface-name {
            classifiers{
                dscp(classifier-name | default) {
                }
                ieee-802.1 (classifier-name | default) vlan-tag (inner | outer | classifier-name);
                inet-precedence (rewrite-name | default);
            }
            input-scheduler-map map-name;
            input-shaping-rate rate;
            irb {
                unit logical-unit-number {
                    classifiers {
                        type (classifier-name | default);
                    }
                    rewrite-rules {
                        dscp (rewrite-name | default);
                        dscp-ipv6 (rewrite-name | default);
                        exp (rewrite-name | default) protocol protocol-types;
                        ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                        inet-precedence (rewrite-name | default);
                    }
                }
            }
            member-link-scheduler (replicate | scale);
            rewrite-rules {
                dscp (rewrite-name | default);
                ieee-802.1 (rewrite-name | default) vlan-tag (outer);
                inet-precedence (rewrite-name | default);
            }
            scheduler-map map-name;
            scheduler-map-chassis map-name;
            shaping-rate rate;
            unit logical-unit-number {
                classifiers {
                    type (classifier-name | default) family (mpls | inet);
                }
                forwarding-class class-name;
                fragmentation-map map-name;
                input-shaping-rate (percent percentage | rate);
                input-traffic-control-profile profile-name shared-instance instance-name;
                output-traffic-control-profile profile-name shared-instance instance-name;
                per-session-scheduler;
                rewrite-rules {
                    dscp (rewrite-name | default);
                    dscp-ipv6 (rewrite-name | default);
                    exp (rewrite-name | default) protocol protocol-types;
                    exp-push-push-push default;
                    exp-swap-push-push default;
                    ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                    inet-precedence (rewrite-name | default);
                }
            }
        }
    }
```

```

    }
    scheduler-map map-name;
    shaping-rate rate;
    translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 | to-exp-from-exp
    | to-inet-precedence-from-inet-precedence) table-name;
  }
}
interface-set interface-set-name {
  excess-bandwidth-share;
  internal-node;
  output-traffic-control-profile profile-name;
  output-traffic-control-profile-remaining profile-name;
}
}

```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 7.4.
Interface-set level added in Junos OS Release 8.5.

Description Configure interface-specific CoS properties for incoming packets.



NOTE: The dscp-ipv6 and ieee-802.1ad classifier types are not supported on ACX Series routers. For further information about support on ACX Series routers, see *Understanding CoS CLI Configuration Statements on ACX Series Universal Access Routers*.

Options The remaining statements are explained separately.

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

Related Documentation

- [Overview of BA Classifier Types on page 1321](#)
- [Configuring Rewrite Rules on page 1694](#)
- *Understanding CoS CLI Configuration Statements on ACX Series Universal Access Routers*

loss-priority (Scheduler Drop Profiles)

Syntax	loss-priority (any high low medium-high medium-low);
Hierarchy Level	[edit class-of-service schedulers <i>scheduler-name</i> drop-profile-map]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches. Statement introduced in Junos OS Release 12.2 for ACX Series Routers.
Description	Specify a loss priority to which to apply a drop profile. The drop profile map sets the drop profile for a specific PLP and protocol type. The inputs for the map are the PLP designation and the protocol type. The output is the drop profile.
Options	any —The drop profile applies to packets with any PLP.



NOTE: On ACX Series Routers, only the **any** option is supported when you configure the **non-tcp** option for **protocol**.

high—The drop profile applies to packets with high PLP.

low—The drop profile applies to packets with low PLP.

medium-high—The drop profile applies to packets with medium-high PLP.

medium-low—The drop profile applies to packets with medium-low PLP.

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

Related Documentation	<ul style="list-style-type: none">• Default Schedulers Overview on page 1575• Configuring Drop Profile Maps for Schedulers on page 1590• Configuring Schedulers for Priority Scheduling on page 1594• Configuring Tricolor Marking on page 1447• protocol (Schedulers) on page 1480
------------------------------	---


output-traffic-control-profile

Syntax	<code>output-traffic-control-profile <i>profile-name</i> shared-instance <i>instance-name</i>;</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit class-of-service interfaces <i>interface-name</i> interface-set <i>interface-set-name</i>]
Release Information	Statement introduced in Junos OS Release 7.6. interface-set option added for Enhanced Queuing DPCs on MX Series routers in Junos OS Release 8.5. interface-set option added for MPC/MIC interfaces on MX Series routers in Junos OS Release 10.2.
Description	For Channelized IQ PICs, Gigabit Ethernet IQ, Gigabit Ethernet IQ2, and IQ2E PICs, link services IQ (LSQ) interfaces on AS PICs, and Enhanced Queuing DPCs and MPC/MIC interfaces on MX Series routers and on EX Series switches, apply an output traffic scheduling and shaping profile to the logical interface. The shared-instance statement is supported on Gigabit Ethernet IQ2 PICs only.
Options	<i>profile-name</i> —Name of the traffic-control profile to be applied to this interface
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Oversubscribing Interface Bandwidth on page 1600 • Configuring Traffic Control Profiles for Shared Scheduling and Shaping • Example: Configuring CoS for a PBB Network • Configuring Hierarchical Schedulers for CoS (Enhanced Queuing DPCs and MPC/MIC interfaces on MX Series routers) • Configuring Interface Sets (Enhanced Queuing DPCs and MPC/MIC interfaces on MX Series routers) • output-traffic-control-profile-remaining • traffic-control-profiles on page 1651

priority (Fabric Queues, Schedulers)

Syntax	<code>priority (high low)scheduler scheduler-name;</code>
Hierarchy Level	[edit class-of-service fabric scheduler-map]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS 11.4 for EX Series switches.
Description	<p>Define Fabric traffic priority. For M320, MX Series, T Series routers and EX Series switches only, specify the fabric priority with which a scheduler is associated.</p> <p>For a scheduler that you associate with a fabric priority, you cannot include the buffer-size, transmit-rate, or priority statements at the [edit class-of-service schedulers scheduler-name] hierarchy level.</p> <p>On EX Series switches, this statement is supported only on EX8200 standalone switches and EX8200 Virtual Chassis.</p>
Options	<p>high—Scheduler has high priority.</p> <p>low—Scheduler has low priority.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• See Associating Schedulers with Fabric Priorities on page 1612.• <i>Understanding Junos OS CoS Components for EX Series Switches</i>

priority (Schedulers)

Syntax	<code>priority <i>priority-level</i>;</code>
Hierarchy Level	[edit class-of-service schedulers <i>scheduler-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches. Statement introduced in Junos OS Release 12.2 for ACX Series Routers.
Description	Specify the packet-scheduling priority value.
Options	<p><i>priority-level</i> can be one of the following:</p> <ul style="list-style-type: none"> • low—Scheduler has low priority. • medium-low—Scheduler has medium-low priority. • medium-high—Scheduler has medium-high priority. • high—Scheduler has high priority. Assigning high priority to a queue prevents the queue from being underserved. • strict-high—Scheduler has strictly high priority. Configure a high priority queue with unlimited transmission bandwidth available to it. As long as it has traffic to send, the strict-high priority queue receives precedence over low, medium-low, and medium-high priority queues, but not high priority queues. You can configure strict-high priority on only one queue per interface.
<div>  <p>NOTE: The strict-high priority level is the only priority level supported on ACX Series Routers. However, multiple strict-high priority queues can be configured per interface on ACX Series Routers.</p> </div>	
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Schedulers for Priority Scheduling on page 1594

protocol (Schedulers)

Syntax	protocol (any non-tcp tcp);
Hierarchy Level	[edit class-of-service schedulers <i>scheduler-name</i> drop-profile-map]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches. Statement introduced in Junos OS Release 12.2 for ACX Series Routers.
Description	Specify the protocol type for the specified scheduler.
Options	any —Accept any protocol type. non-tcp —(ACX Series Routers, M Series and T Series (except T4000) routers only) Accept any protocol type other than TCP/IP.



NOTE: On ACX Series Routers, when you configure the **non-tcp** option, only the **any** option is supported for [loss-priority](#).

	tcp —(ACX Series Routers, M Series and T Series (except T4000) routers only) Accept TCP/IP protocol type.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Schedulers on page 1579

scheduler (Fabric Queues)

Syntax	<code>scheduler <i>scheduler-name</i>;</code>
Hierarchy Level	[edit class-of-service fabric scheduler-map <i>priority</i> (high low)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS 11.4 for EX Series switches.
Description	Define scheduler name. For M320, MX Series, T Series routers and for EX Series switches only, specify a scheduler to associate with a fabric queue. For fabric CoS configuration, schedulers are restricted to transmit rates and drop profiles.
Options	<i>scheduler-name</i> —Name of the scheduler configuration block.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • See Associating Schedulers with Fabric Priorities on page 1612. • <i>Understanding Junos OS CoS Components for EX Series Switches</i>

scheduler (Scheduler Map)

Syntax	<code>scheduler <i>scheduler-name</i>;</code>
Hierarchy Level	[edit class-of-service scheduler-maps <i>map-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.2 for ACX Series Routers.
Description	Associate a scheduler with a scheduler map.
Options	<i>scheduler-name</i> —Name of the scheduler configuration block.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Schedulers on page 1579 • <i>Example: Configuring CoS for a PBB Network</i>

scheduler-map (Fabric Queues)

Syntax	<code>scheduler-map priority (high low) scheduler <i>scheduler-name</i>;</code>
Hierarchy Level	<code>[edit class-of-service fabric]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS 11.4 for EX Series switches.
Description	<p>Mapping of fabric traffic to packet schedulers. For M320, MX Series, T Series routers, and for EX Series switches only, associate a scheduler with a fabric priority.</p> <p>On EX Series switches, this statement is supported only on EX8200 standalone switches and EX8200 Virtual Chassis.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• See Associating Schedulers with Fabric Priorities on page 1612.• <i>Understanding Junos OS CoS Components for EX Series Switches</i>

scheduler-map (Interfaces and Traffic-Control Profiles)

Syntax	<code>scheduler-map <i>map-name</i>;</code>
Hierarchy Level	<code>[edit class-of-service interfaces <i>interface-name</i>],</code> <code>[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</code> <code>[edit class-of-service traffic-control-profiles]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>For Gigabit Ethernet IQ, Channelized IQ PICs, and FRF.15 and FRF.16 LSQ interfaces only, associate a scheduler map name with an interface or with a traffic-control profile.</p> <p>For channelized OC12 intelligent queuing (IQ), channelized T3 IQ, channelized E1 IQ, and Gigabit Ethernet IQ interfaces only, you can associate a scheduler map name with a logical interface.</p>
Options	<i>map-name</i> —Name of the scheduler map.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Schedulers on page 1579• Oversubscribing Interface Bandwidth on page 1600• output-traffic-control-profile on page 1639

scheduler-map-chassis

Syntax	<code>scheduler-map-chassis (derived <i>map-name</i>);</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-type-fpc/pic/*</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For IQ and IQ2 interfaces, assign a custom scheduler to the packet forwarding component queues that control the aggregated traffic transmitted into the entire PIC.
Default	On Intelligent Queuing (IQ) and Intelligent Queuing 2 (IQ2) interfaces, the traffic that is fed from the packet forwarding components into the PIC uses low packet loss priority (PLP) by default and is distributed evenly across the four chassis queues (not PIC queues), regardless of the scheduling configuration for each logical interface. This default behavior can cause traffic congestion.
Options	<p>derived—Sets the chassis queues to derive their scheduling configuration from the associated logical interface scheduling configuration.</p> <p><i>map-name</i>—Name of the scheduler map configured at the [edit class-of-service scheduler-maps] hierarchy level.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Applying Scheduler Maps to Packet Forwarding Component Queues</i> • scheduler-map (Fabric Queues) on page 1644

scheduler-maps (For Most Interface Types)

Syntax	<pre>scheduler-maps { map-name { forwarding-class class-name scheduler scheduler-name; } }</pre>
Hierarchy Level	[edit class-of-service]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify a scheduler map name and associate it with the scheduler configuration and forwarding class.
Options	<p>map-name—Name of the scheduler map.</p> <p>The remaining statements are explained separately.</p> <p>See “Configuring Schedulers” on page 1579 and <i>Example: Configuring CoS for a PBB Network</i>.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

schedulers (Class of Service)

Syntax	<pre> schedulers { scheduler-name { adjust-minimum <i>rate</i>; adjust-percent <i>percentage</i>; buffer-size (<i>seconds</i> percent <i>percentage</i> remainder temporal <i>microseconds</i>); drop-profile-map loss-priority (any low medium-low medium-high high) <i>protocol</i> (any non-tcp tcp) drop-profile <i>profile-name</i>; excess-priority [low medium-low medium-high high none]; excess-rate (percent <i>percentage</i> proportion <i>value</i>); priority <i>priority-level</i>; shaping-rate (percent <i>percentage</i> <i>rate</i>); transmit-rate (percent <i>percentage</i> <i>rate</i> remainder) <exact rate-limit>; } } </pre>
Hierarchy Level	[edit class-of-service]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1X48 for PTX Series switches.</p>
Description	Specify the scheduler name and parameter values.
Options	<p><i>scheduler-name</i>—Name of the scheduler to be configured.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Schedulers Overview on page 1573 • Default Schedulers Overview on page 1575 • Configuring Schedulers on page 1579 • Configuring a Scheduler • Example: Configuring CoS for a PBB Network

shaping-rate (Applying to an Interface)

Syntax	<code>shaping-rate rate;</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i>], [edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. [edit class-of-service interfaces <i>interface-name</i>] hierarchy level added in Junos OS Release 7.5.
Description	For logical interfaces on which you configure packet scheduling, configure traffic shaping by specifying the amount of bandwidth to be allocated to the logical interface. For physical interfaces on IQ PICs and T4000 routers with Type 5 FPCs, configure traffic shaping based on the rate-limited bandwidth of the total interface bandwidth.



NOTE: The `shaping-rate` statement cannot be applied to a physical interface on J Series routers.

Logical and physical interface traffic shaping rates are mutually exclusive. This means you can include the `shaping-rate` statement at the [edit class-of-service interfaces *interface-name*] hierarchy level or the [edit class-of-service interfaces *interface-name* **unit** *logical-unit-number*] hierarchy level, but not both.



NOTE: For MX Series routers and for EX Series switches, the shaping rate value for the physical interface at the [edit class-of-service interfaces *interface-name*] hierarchy level must be a minimum of 160 Kbps. If the value is less than the sum of the logical interface guaranteed rates, the user is not allowed to apply the shaping rate to a physical interface.

For T4000 routers with Type 5 FPCs, the shaping rate value for the physical interface must be a minimum of 292 Kbps. The maximum value of `shaping-rate` is limited by the maximum transmission rate of the interface.

Alternatively, you can configure a shaping rate for a logical interface and oversubscribe the physical interface by including the `shaping-rate` statement at the [edit class-of-service **traffic-control-profiles**] hierarchy level. With this configuration approach, you can independently control the delay-buffer rate, as described in [“Oversubscribing Interface Bandwidth” on page 1600](#).

For FRF.15 and FRF.16 bundles on link services interfaces, only shaping rates based on percentage are supported.

Default If you do not include this statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level, the default logical interface bandwidth is the average of unused bandwidth for the number of logical interfaces that require default bandwidth treatment. If you do not include this statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level, the default physical interface bandwidth is the average of unused bandwidth for the number of physical interfaces that require default bandwidth treatment.

Options **rate**—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: For logical interfaces, 1000 through 32,000,000,000 bps. For physical interfaces, 1000 through 160,000,000,000 bps.



NOTE: For all MX Series and EX series interfaces, the rate can be from 65,535 through 160,000,000,000 bps.



NOTE: For T4000 physical interfaces, the rate can be from 1000 through 160,000,000,000 bps.

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

Related Documentation

- [Applying Scheduler Maps Overview on page 1577](#)

shaping-rate (Oversubscribing an Interface)

Syntax	<code>shaping-rate (percent <i>percentage</i> <i>rate</i>) <burst-size <i>bytes</i>>;</code>
Hierarchy Level	[edit class-of-service traffic-control-profiles <i>profile-name</i>]
Release Information	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Option burst-size introduced for Enhanced Queuing (EQ) DPCs in Junos OS Release 9.4.</p> <p>Option burst-size option introduced for MPC/MIC modules in Junos OS Release 11.4.</p> <p>Option burst-size introduced for IQ2 and IQ2E interfaces in Junos OS Release 12.3.</p>
Description	<p>For Gigabit Ethernet IQ, Channelized IQ PICs, FRF.15 and FRF.16 LSQ interfaces, EQ DPCs, and MPC/MIC modules only, configure a shaping rate for a logical interface. You can also configure an optional burst size for a logical interface on EQ DPCs and on IQ2 and IQ2E PICs. This can help to ensure that higher-priority services do not starve lower-priority services.</p> <p>For physical interfaces on T4000 routers with Type 5 FPCs, configure traffic shaping rate.</p> <p>The sum of the shaping rates for all logical interfaces on the physical interface can exceed the physical interface bandwidth. This practice is known as oversubscription of the peak information rate (PIR).</p>
Default	The default behavior depends on various factors. For more information, see Table 114 on page 1603 .
Options	<p>percent <i>percentage</i>—For LSQ interfaces, shaping rate as a percentage of the available interface bandwidth.</p> <p>Range: 1 through 100 percent</p> <p><i>rate</i>—For IQ and IQ2 interfaces, and T4000 routers with Type 5 FPCs, peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).</p> <p>Range: IQ and IQ2 interfaces—1000 through 160,000,000,000 bps</p> <p>Range: T4000 routers with Type 5 FPCs—the shaping rate value for the physical interface must be a minimum of 292 Kbps. The maximum value of shaping-rate is limited by the maximum transmission rate of the interface.</p> <p>burst-size <i>bytes</i>—(Optional) Maximum burst size, in bytes.</p> <p>Range: 0 through 1,000,000,000</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Traffic Control Profiles for Shared Scheduling and Shaping• Oversubscribing Interface Bandwidth on page 1600

- [output-traffic-control-profile on page 1639](#)

traffic-control-profiles

Syntax	<pre> traffic-control-profiles <i>profile-name</i> { adjust-minimum <i>rate</i>; atm-service (cbr rtvbr nrtvbr); delay-buffer-rate (percent <i>percentage</i> <i>rate</i>); peak-rate <i>rate</i>; sustained-rate <i>rate</i>; max-burst-size <i>cells</i>; excess-rate (percent <i>percentage</i> proportion <i>value</i>); excess-rate-high (percent <i>percentage</i> proportion <i>value</i>); excess-rate-low (percent <i>percentage</i> proportion <i>value</i>); guaranteed-rate (percent <i>percentage</i> <i>rate</i>) <burst-size <i>bytes</i>>; overhead-accounting (frame-mode cell-mode) <bytes <i>byte-value</i>>; scheduler-map <i>map-name</i>; shaping-rate (percent <i>percentage</i> <i>rate</i>) <burst-size <i>bytes</i>>; shaping-rate-excess-high <i>rate</i> [burst-size <i>bytes</i>]; shaping-rate-excess-low <i>rate</i> [burst-size <i>bytes</i>]; shaping-rate-priority-high <i>rate</i> [burst-size <i>bytes</i>]; shaping-rate-priority-low <i>rate</i> [burst-size <i>bytes</i>]; shaping-rate-priority-medium <i>rate</i> [burst-size <i>bytes</i>]; } </pre>
Hierarchy Level	[edit class-of-service]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	For Gigabit Ethernet IQ, Channelized IQ PICs, FRF.15 and FRF.16 LSQ interfaces, and Enhanced Queuing (EQ) DPCs only, configure traffic shaping and scheduling profiles. For Enhanced EQ PICs and EQ DPCs only, you can include the excess-rate statement.
Options	<p><i>profile-name</i>—Name of the traffic-control profile.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Oversubscribing Interface Bandwidth on page 1600 • Example: Configuring CoS for a PBB Network • output-traffic-control-profile on page 1639

transmit-rate (Schedulers)

Syntax	<code>transmit-rate (rate percent <i>percentage</i> remainder) <exact rate-limit>;</code>
Hierarchy Level	[edit class-of-service schedulers <i>scheduler-name</i>]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>rate-limit option introduced in Junos OS Release 8.3. Applied to the Multiservices PICs in Junos OS Release 9.4.</p> <p>Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Switches.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Routers.</p>
Description	Specify the transmit rate or percentage for a scheduler.
Default	If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 5, 0, 0, 0, and 0 percent, respectively.
Options	<p>exact—(Optional) Enforce the exact transmission rate. Under sustained congestion, a rate-controlled queue that goes into negative credit fills up and eventually drops packets. This value should never exceed the rate-controlled amount. For PTX Series Packet Transport Switches, this option is allowed only on the non-strict-high (high, medium-high, medium-low, or low) queues.</p> <p>percent <i>percentage</i>—Percentage of transmission capacity. A percentage of zero drops all packets in the queue.</p> <p>Range: 0 through 100 percent for M, MX, T Series routers and EX Series switches; 1 through 100 percent for PTX Series Packet Transport Switches; 0 through 200 percent for the SONET/SDH OC48/STM16 IQE PIC</p>



NOTE:

- On M Series Multiservice Edge Routers, for interfaces configured on 4-port E1 and 4-port T1 PICs only, you can configure a *percentage* value only from 11 through 100. These two PICs do not support transmission rates less than 11 percent.
 - The configuration of the `transmit-rate percent 0 exact` statement at the [edit class-of-service `schedulers` *scheduler-name*] hierarchy is ineffective on T4000 routers with Type 5 FPC.
 - On Trio MPC/MIC interfaces, when the transmit rate is configured as a percentage and **exact** or **rate-limit** is enabled on a queue, the shaping rate of the parent node is used to compute the transmit rate. If **exact** or **rate-limit** is not configured, the guaranteed rate of the parent node is used to compute the transmit rate.
-

rate—Transmission rate, in bps. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 3200 through 160,000,000,000 bps



NOTE: For all MX Series interfaces, the rate can be from 65,535 through 160,000,000,000 bps.

rate-limit—(Optional) Limit the transmission rate to the rate-controlled amount. In contrast to the **exact** option, the scheduler with the **rate-limit** option shares unused bandwidth above the rate-controlled amount.



NOTE: For PTX Series Packet Transport Switches, this option is allowed only on the strict-high queue. We recommend that you configure rate limit on strict-high queues because the other queues may not meet their guaranteed bandwidths.



NOTE: The configuration of the **rate-limit** statement is supported on T4000 routers only with a Type 5 FPC.

remainder—Use the remaining rate available.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Schedulers on page 1579 • Configuring Scheduler Transmission Rate on page 1591 • <i>Example: Configuring CoS for a PBB Network</i>

unit

Syntax	<pre> unit <i>logical-unit-number</i> { classifiers { type (<i>classifier-name</i> default) family (mpls all); } forwarding-class <i>class-name</i>; fragmentation-map <i>map-name</i>; input-traffic-control-profile <i>profile-name</i> shared-instance <i>instance-name</i>; output-traffic-control-profile <i>profile-name</i> shared-instance <i>instance-name</i>; per-session-scheduler; rewrite-rules { dscp (<i>rewrite-name</i> default); dscp-ipv6 (<i>rewrite-name</i> default); exp (<i>rewrite-name</i> default) <i>protocol</i> <i>protocol-types</i>; exp-push-push-push default; exp-swap-push-push default; ieee-802.1 (<i>rewrite-name</i> default) <i>vlan-tag</i> (outer outer-and-inner); inet-precedence (<i>rewrite-name</i> default); } scheduler-map <i>map-name</i>; shaping-rate <i>rate</i>; } </pre>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.
Options	<p><i>logical-unit-number</i>—Number of the logical unit.</p> <p>Range: 0 through 16,384</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Overview of BA Classifier Types on page 1321 • Configuring Rewrite Rules on page 1694

[edit interfaces] Hierarchy Level

The following statement hierarchy can also be included at the [edit logical-systems *logical-system-name*] hierarchy level.

```

interfaces {
  interface-name {
    ... the "interface-name" subhierarchy appears after the main [edit interfaces] hierarchy level ...
  }
}

```

```

interface-set interface-set-name {
  interface interface-name {
    (unit unit-number | vlan-tags-outer vlan-tag);
  }
}
irb (Interfaces) {
  accounting-profile name;
  description text;
  disable;

  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  hold-time up milliseconds down milliseconds;
  mtu bytes;
  no-gratuitous-arp-request;

  traceoptions {
    flag flag;
  }
  (traps | no-traps);
  unit logical-unit-number {
    accounting-profile name;
    bandwidth rate;
    description text;
    disable;
    encapsulation type;
    family inet {
      accounting {
        destination-class-usage;
        source-class-usage {
          input;
          output;
        }
      }
    }
    address ipv4-address {
      arp ip-address (mac | multicast-mac) mac-address <publish>;
      broadcast address;
      preferred;
      primary;
      vrrp-group group-id {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        (preempt | no-preempt) {
          hold-time seconds;
        }
        priority number;
        track {
          interface interface-name {
            bandwidth-threshold bits-per-second priority-cost priority;
            priority-cost priority;
          }
          priority-hold-time seconds;
        }
      }
    }
  }
}

```

```

    route prefix/prefix-length routing-instance instance-name priority-cost priority;
  }
  virtual-address [ addresses ];
  vrrp-inherit-from vrrp-group;
}
filter {
  input filter-name;
  output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
  fail-filter filter-name;
  mode {
    loose;
  }
}
targeted-broadcast {
  forward-and-send-to-re;
  forward-only;
}
}
family inet6 {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
}
address address {
  eui-64;
  ndp ip-address (mac | multicast-mac) mac-address <publish>;
  preferred;
  primary;
  vrrp-inet6-group group-id {
    accept-data | no-accept-data;
    advertisements-threshold number;
    authentication-key key;
    authentication-type authentication;
    fast-interval milliseconds;
    inet6-advertise-interval milliseconds;
    preempt | no-preempt {
      hold-time seconds;
    }
  }
  priority number;
  track {
    interface interface-name {
      bandwidth-threshold bandwidth priority-cost number;
      priority-cost number;
    }
    priority-hold-time seconds;
    route ip-address/mask routing-instance instance-name priority-cost cost;
  }
}

```

```

    }
    virtual-inet6-address [addresses];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
    }
}
}
(dad-disable | no-dad-disable);
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
no-redirects;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    filter {
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    policer {
        input policer-name;
        output policer-name;
    }
}
native-inner-vlan-id vlan-id;
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id-list [vlan-id's];
vlan-id-range [vlan-id-range];
}
}
traceoptions {
    file <filename> <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag <disable>;
    no-remote-trace;
}

```

```
    }  
  }  
  
  interfaces {  
    interface-name {  
      disable;  
      accounting-profile name;  
      aggregated-ether-options {  
        ethernet-switch-profile {  
          tag-protocol-id [ hexadecimal-identifiers ];  
        }  
      }  
      (flow-control | no-flow-control);  
      lacp {  
        (active | passive);  
        admin-key key;  
        fast-failover;  
        link-protection {  
          disable;  
          (revertive | non-revertive);  
        }  
        periodic (fast | slow);  
        system-id mac-address;  
        system-priority priority;  
      }  
      (link-protection | no-link-protection);  
      link-speed (100m | 1g | 8g | 10g | 40g | 50g | 80g | 100g | oc192);  
      logical-interface-fpc-redundancy;  
      (loopback | no-loopback);  
      mc-ae {  
        chassis-id chassis-id;  
        events {  
          iccp-peer-down {  
            force-icl-down;  
            prefer-status-control-active;  
          }  
        }  
        mc-ae-id mc-ae-id;  
        mode (active-active | active-standby);  
        redundancy-group group-id;  
        status-control (active | standby);  
      }  
      minimum-links number;  
      rebalance-periodic {  
        start-time time;  
        interval number;  
      }  
      source-address-filter {  
        mac-address;  
      }  
      (source-filtering | no-source-filtering);  
    }  
    auto-configure {  
      remove-when-no-subscribers;  
      stacked-vlan-ranges {  
        access-profile profile-name;  
        authentication {
```



```

password password-string;
username-include {
    circuit-type;
    delimiter delimiter-character;
    domain-name domain-name-string;
    interface-name;
    mac-address;
    option-82 ( circuit-id | remote-id);
    radius-realm radius-realm-string;
    user-prefix user-prefix-string;
}
}
dynamic-profile profile-name {
    accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
    ranges (any | low-tag-high-tag), (any | low-tag-high-tag);
}
}
vlan-ranges {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-name;
            mac-address;
            option-82;
            radius-realm radius-realm-string;
            user-prefix user-prefix-string;
        }
    }
    dynamic-profile profile-name {
        accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
        ranges (any | low-tag)—(any | high-tag);
    }
}
override tag vlan-tag dynamic-profile profile name;
}
encapsulation (ethernet-bridge | ethernet-vpls | extended-vlan-bridge |
    extended-vlan-vpls | flexible-ethernet-services | vlan-vpls);
ether-options {
    802.3ad {
        aex;
        (backup | primary);
        lacp {
            force-up;
            port-priority
        }
    }
}
asynchronous-notification;
(auto-negotiation | no-auto-negotiation);
ethernet-switch-profile {
    ethernet-policer-profile {
        input-priority-map {
            ieee802.1p premium [ values ];

```

```

    }
    output-priority-map {
        classifier {
            premium {
                forwarding-class class-name {
                    loss-priority (high | low);
                }
            }
        }
    }
    policer cos-policer-name {
        aggregate {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
        premium {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
    }
    tag-protocol-id;
}
(mac-learn-enable | no-mac-learn-enable);
}
(flow-control | no-flow-control);
ignore-l3-incompletes;
link-mode (automatic | full-duplex | half-duplex);
(lloopback | no-loopback);
keepalives <interval seconds> <down-count number> <up-count number>;
speed (1g | 10m | 100m | 10m-100m | auto-negotiation);
source-address-filter {
    mac-address;
}
source-filtering | no-source-filtering;
}
flexible-vlan-tagging;
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time (up milliseconds | down milliseconds);
interface-transmit-statistics;
(keepalives <down-count number> <interval seconds> <up-count number> |
no-keepalives);
layer2-policer {
    apply-groups [ group-names ];
    apply-groups-except [ group-names ];
}
link-mode (automatic | full-duplex);
mac mac-address;
mtu bytes;
multi-chassis-protection peer-ip-address {
    interface interface-name;
}
native-vlan-id number;
no-gratuitous-arp-request;
optics-options {
    alarm low-light-alarm {
        (link-down | syslog);
    }
}

```

```

    }
    warning low-light-warning {
        (link-down | syslog);
    }
    wavelength nm;
}
passive-monitor-mode;
per-unit-scheduler;
speed (10m | 100m | 1g | auto | oc3 | oc12 | oc48);
stacked-vlan-tagging;
traceoptions {
    flag flag;
}
transmit-bucket {
    overflow discard;
    rate percentage;
    threshold bytes;
}
(traps | no-traps);
unidirectional;
vlan-tagging;
}

interface-name {
    unit logical-unit-number {
        disable;
        accept-source-mac {
            mac-address mac-address {
                policer {
                    input policer-name;
                    output policer-name;
                }
            }
        }
    }
    account-layer2-overhead (Interface Level) {
        value;
        egress bytes;
        ingress bytes;
    }
    accounting-profile name;
    advisory-options {
        downstream-rate rate;
        upstream-rate rate;
    }
    arp-resp (restricted|unrestricted);
    bandwidth rate;
    clear-dont-fragment-bit;
    copy-tos-to-outer-ip-header;
    demux-destination family;
    encapsulation (vlan-bridge | vlan-vpls);
    epd-threshold cells plp1 cells;
    filter filter-name;
    inner-vlan-id-range start start-id end end-id;
    input-vlan-map {
        (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    }
}

```

```

    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
interface-shared-with psd numerical-index;
layer2-policer {
    input-hierarchical-policer policer-name;
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
}
multi-chassis-protection peer-ip-address {
    interface interface-name;
}
native-inner-vlan-id number;
output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
peer-interface interface-name;
peer-unit unit-number;
plp-to-clp;
proxy-arp <restricted | unrestricted>;
rpm {
    (client | server);
    twamp-server;
}
swap-by-poppush;
vlan-id number;
vlan-id-list [ vlan-id vlan-id-vlan-id ];
vlan-id-range number-number;
vlan-tags (inner <tpid.>vlan-id | inner-list [ vlan-id vlan-id-vlan-id ] |
    inner-range <tpid.>vlan-id-vlan-id) outer <tpid.>vlan-id;
}

unit logical-unit-number {
    family ethernet-switching {
        filter {
            group filter-group-number;
            (input filter-name | input-list [ filter-names ]);
            (output filter-name | output-list [ filter-names ]);
            (inner-vlan-id-list [ vlan-ids ] | vlan-id number | vlan-id-list [ number
                number-number ]);
            interface-mode (access | trunk);
        }
        policer {
            input policer-name;
            output policer-name;
        }
        vlan-rewrite {
            translate old-vlan-id new-vlan-id;
        }
    }
}

```

```

    vlan {
        members [ all vlan-identifiers ];
    }
}
family inet {
    filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
    input-hierarchical-policer policer-name;
    mac-validate (loose | strict);
    mtu bytes;
    no-neighbor-learn;
    no-redirects;
    policer {
        arp policer-template-name;
        input policer-name;
        output policer-name;
    }
    primary;
    receive-options-packets;
    receive-ttl-exceeded;
    rpf-check {
        fail-filter filter-name;
        mode loose;
    }
    sampling {
        (input | output | input output);
    }
    simple-filter {
        input filter-name;
    }
    targeted-broadcast {
        forward-and-send-to-re;
        forward-only;
    }
    unnumbered-address interface-name <destination address>
        <destination-profile profile-name> <preferred-source-address address>;
}

family inet6 {
    address ipv6-address {
        destination destination-address;
        eui-64;
        ndp ipv6-address <l2-interface interface-name> <(mac mac-address |
            multicast-mac multicast-mac-address) <publish>>;
        preferred;
        primary;
        vrrp-inet6-group group-number {
            (accept-data | no-accept-data);
            fast-interval milliseconds;
            inet6-advertise-interval seconds;
            (no-preempt; | ... the following preempt statement ...)
            preempt {

```

```

        hold-time seconds;
    }
    priority number;
    track {
        interface interface-name {
            bandwidth-threshold bits-per-second priority-cost priority;
            priority-cost priority;
        }
        priority-hold-time seconds;
        route ip-address-prefix/prefix-length routing-instance instance-name
            priority-cost priority;
    }
    virtual-inet6-address [ addresses ];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
    }
}
(dad-disable | no-dad-disable);
filter {
    group filter-group-number;
    (input filter-name | input-list [ filter-names ]);
    (output filter-name | output-list [ filter-names ]);
}
input-hierarchical-policer policer-name;
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode loose;
}
sampling {
    (input | output | input output);
}
unnumbered-address interface-name preferred-source-address address;
}

family iso {
    address iso-address;
    mtu bytes;
}

family mlfrr-end-to-end {
    bundle logical-interface-name;
}

```

```

family mpls {
  filter {
    group filter-group-number;
    (input filter-name | input-list [ filter-names ]);
    (output filter-name | output-list [ filter-names ]);
  }
  input-hierarchical-policer policer-name;
  maximum-labels maximum-labels;
  mtu bytes;
  policer {
    input policer-name;
    output policer-name;
  }
}

family vpls {
  core-facing;
  filter {
    group filter-group-number;
    (input filter-name | input-list [ filter-names ]);
    (output filter-name | output-list [ filter-names ]);
  }
  policer {
    input policer-name;
    output policer-name;
  }
}
}
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

schedulers (Interfaces)

Syntax	<code>schedulers <i>number</i>;</code>
Hierarchy Level	[edit interfaces]
Release Information	Statement introduced in Junos OS Release 8.2.
Description	Specify number of schedulers for Ethernet IQ2 PIC port interfaces.
Default	If you omit this statement, the 1024 schedulers are distributed equally over all ports in multiples of 4.
Options	<i>number</i> —Number of schedulers to configure on the port. Range: 1 through 1024
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Number of Schedulers for Ethernet IQ2 PICs</i>

Queue-Level Bandwidth Sharing

- [Overview on page 1666](#)
- [Configuration on page 1667](#)

Overview

- [Bandwidth Sharing on page 1666](#)

Bandwidth Sharing

- [Bandwidth Sharing on Nonqueuing Packet Forwarding Engines Overview on page 1666](#)

Bandwidth Sharing on Nonqueuing Packet Forwarding Engines Overview

You can configure bandwidth sharing rate limits, excess rate, and excess priority at the queue level on the following Juniper Networks routers and switches:

- EX Series switches
- M120 Multiservice Edge Router (rate limit and excess priority only; excess rate is not configured by the user)
- M320 router with Enhanced FPCs (rate limit, excess rate, and excess priority)
- MX Series 3D Universal Edge Router with nonqueuing DPCs (rate limit, excess rate, and excess priority)

You configure rate limits when you have a concern that low-latency packets (such as high or strict-high priority packets for voice) might starve low-priority and medium-priority

packets. In Junos OS, the low latency queue is implemented by rate-limiting packets to the transmit bandwidth. The rate-limiting is performed immediately before queuing the packet for transmission. All packets that exceed the rate limit are not queued, but dropped.

By default, if the excess priority is not configured for a queue, the excess priority will be the same as the normal queue priority. If none of the queues have an excess rate configured, then the excess rate will be the same as the transmit rate percentage. If at least one of the queues has an excess rate configured, then the excess rate for the queues that do not have an excess rate configured will be set to zero.

When the physical interface is on queuing hardware such as the IQ, IQ2, or IQE PICs, or MX Series routers queuing DPCs or EX Series switches, these features are dependent on the PIC (or queuing DPC in the case of the MX Series router) configuration.

You cannot configure both rate limits and buffer sizes on these Packet Forwarding Engines.

Four levels of excess priorities are supported: low, medium-low, medium-high, and high.



NOTE: Rate limiting is implemented differently on Enhanced Queuing DPCs and non-queuing Packet Forwarding Engines. On Enhanced Queuing DPCs, rate-limiting is implemented using a single rate two color policer. On non-queuing Packet Forwarding Engines, rate-limiting is achieved by shaping the queue to the transmit rate and keeping the queue delay buffers small to prevent too many packets from being queued once the shaping rate is reached.

Configuration

- [Configuration Task on page 1667](#)
- [Example on page 1668](#)

Configuration Task

- [Configuring Rate Limits on Nonqueuing Packet Forwarding Engines on page 1667](#)

Configuring Rate Limits on Nonqueuing Packet Forwarding Engines

To configure rate limits for nonqueuing Packet Forwarding Engines, include the **transmit-rate** statement at the `[edit class-of-service schedulers scheduler-name]` hierarchy level.



NOTE: Rate limiting is implemented differently on Enhanced Queuing DPCs and non-queuing Packet Forwarding Engines. On Enhanced Queuing DPCs, rate-limiting is implemented using a single rate two color policer. On non-queuing Packet Forwarding Engines, rate-limiting is achieved by shaping the queue to the transmit rate and keeping the queue delay buffers small to prevent too many packets from being queued once the shaping rate is reached.

Configuring the Schedulers	<p>The following example configures schedulers, forwarding classes, and a scheduler map for a rate-limited interface.</p> <pre>[edit class-of-service schedulers] scheduler-1 { transmit-rate percent 20 rate-limit; priority high; } scheduler-2 { transmit-rate percent 10 rate-limit; priority strict-high; } scheduler-3 { transmit-rate percent 40; priority medium-high; } scheduler-4 { transmit-rate percent 30; priority medium-high; }</pre>
Configuring the Forwarding Classes	<pre>[edit class-of-service] forwarding-classes { class cp_000 queue-num 0; class cp_001 queue-num 1; class cp_010 queue-num 2; class cp_011 queue-num 3; class cp_100 queue-num 4; class cp_101 queue-num 5; class cp_110 queue-num 6; class cp_111 queue-num 7; }</pre>
Configuring the Scheduler Map	<pre>[edit class-of-service scheduler-maps] scheduler-map-1 { forwarding-class cp_000 scheduler scheduler-1; forwarding-class cp_001 scheduler scheduler-2; forwarding-class cp_010 scheduler scheduler-3; forwarding-class cp_011 scheduler scheduler-4; }</pre>
Applying the Scheduler Map to the Interface	<pre>[edit interfaces] ge-1/0/0 { scheduler-map scheduler-map-1; unit 0 { family inet { address 192.168.1.1/32; } } }</pre>

Example

- [Excess Rate and Excess Priority Configuration Examples on page 1669](#)

Excess Rate and Excess Priority Configuration Examples

To configure the excess rate for nonqueueing Packet Forwarding Engines, include the `excess-rate` statement at the `[edit class-of-service schedulers scheduler-name]` hierarchy level.

To configure the excess priority for nonqueueing Packet Forwarding Engines, include the `excess-priority` statement at the `[edit class-of-service schedulers scheduler-name]` hierarchy level.

The relationship between the configured guaranteed rate, excess rate, guaranteed priority, excess priority, and offered load is not always obvious. The following tables show the expected throughput of a Gigabit Ethernet port with various bandwidth-sharing parameters configured on the queues.

The default behavior of a nonqueueing Gigabit Ethernet interface with multiple priority levels is shown in [Table 117 on page 1669](#). All queues in the table get their guaranteed rate. The excess bandwidth is first offered to the excess high-priority queues. Because these use all available bandwidth, there is no remaining excess bandwidth for the low-priority queues.

Table 117: Current Behavior with Multiple Priority Levels

Queue	Guaranteed (Transmit) Rate	Guaranteed Priority	Excess Priority	Offered Load	Expected Throughput
Q0	20%	high	high	600 Mbps	$200 + 366.67 = 566.67$ Mbps
Q1	10%	high	high	500 Mbps	$100 + 183.33 = 283.33$ Mbps
Q2	10%	low	low	500 Mbps	$100 + 0 = 100$ Mbps
Q3	5%	low	low	500 Mbps	$50 + 0 = 50$ Mbps

The default behavior of a nonqueueing Gigabit Ethernet interface with the same priority levels is shown in [Table 118 on page 1669](#). All queues in the table get their guaranteed rate. Because all queues have the same excess priority, they share the excess bandwidth and each queue gets excess bandwidth in proportion to the transmit rate.

Table 118: Current Behavior with Same Priority Levels

Queue	Guaranteed (Transmit) Rate	Guaranteed Priority	Excess Priority	Offered Load	Expected Throughput
Q0	20%	high	high	500 Mbps	$200 + 244.44 = 444.44$ Mbps
Q1	10%	high	high	500 Mbps	$100 + 122.22 = 222.22$ Mbps
Q2	10%	high	high	500 Mbps	$100 + 122.22 = 222.22$ Mbps
Q3	5%	high	high	500 Mbps	$50 + 61.11 = 111.11$ Mbps

The default behavior of a nonqueuing Gigabit Ethernet interface with the at least one strict-high priority level is shown in [Table 119 on page 1670](#). First the high priority and strict-high are serviced in a weighted round-robin fashion. The high priority queue gets its guaranteed bandwidth and the strict-high queue gets what remains. The high excess priority queue gets all the excess bandwidth.

Table 119: Current Behavior with Strict-High Priority

Queue	Guaranteed (Transmit) Rate	Guaranteed Priority	Excess Priority	Offered Load	Expected Throughput
Q0	20%	strict-high	X	500 Mbps	500 Mbps
Q1	10%	high	high	500 Mbps	100 + 250 = 350 Mbps
Q2	10%	low	low	500 Mbps	100 + 0 = 100 Mbps
Q3	5%	low	low	500 Mbps	50 + 0 = 50 Mbps

The default behavior of a nonqueuing Gigabit Ethernet interface with the at least one strict-high priority level and a higher offered load on Q0 is shown in [Table 120 on page 1670](#). First the high priority and strict-high are serviced in a weighted round-robin fashion. The high priority queue gets its guaranteed bandwidth and the strict-high queue gets what remains. There is no excess bandwidth.

Table 120: Strict-High Priority with Higher Load

Queue	Guaranteed (Transmit) Rate	Guaranteed Priority	Excess Priority	Offered Load	Expected Throughput
Q0	20%	strict-high	X	1 Gbps	900 Mbps
Q1	10%	high	high	500 Mbps	100 + 0 = 100 Mbps
Q2	10%	low	low	500 Mbps	0 + 0 = 0 Mbps
Q3	5%	low	low	500 Mbps	0 + 0 = 0 Mbps

Now consider the behavior of the queues with configured excess rates and excess priorities.

The behavior with multiple priority levels is shown in [Table 121 on page 1671](#). All queues get the guaranteed rate. The excess bandwidth is first offered to the excess high priority queues and these consume all the bandwidth. There is no remaining excess bandwidth for low priority queues.

Table 121: Sharing with Multiple Priority Levels

Queue	Guaranteed (Transmit) Rate	Excess Rate	Guaranteed Priority	Excess Priority	Offered Load	Expected Throughput
Q0	20%	10%	high	high	500 Mbps	$200 + 275 = 475$ Mbps
Q1	10%	20%	high	low	500 Mbps	$100 + 0 = 100$ Mbps
Q2	10%	10%	low	high	500 Mbps	$100 + 275 = 275$ Mbps
Q3	5%	20%	low	low	500 Mbps	$50 + 0 = 50$ Mbps

The behavior with the same (high) priority levels is shown in [Table 122 on page 1671](#). All queues get the guaranteed rate. Because all queues have the same excess priority, they share the excess bandwidth in proportion to their transmit rate.

Table 122: Sharing with the Same Priority Levels

Queue	Guaranteed (Transmit) Rate	Excess Rate	Guaranteed Priority	Excess Priority	Offered Load	Expected Throughput
Q0	20%	10%	high	high	500 Mbps	$200 + 91.67 = 291.67$ Mbps
Q1	10%	20%	high	high	500 Mbps	$100 + 183.33 = 283.33$ Mbps
Q2	10%	10%	high	high	500 Mbps	$100 + 91.67 = 191.67$ Mbps
Q3	5%	20%	high	high	500 Mbps	$50 + 183.33 = 233.33$ Mbps

The behavior with at least one strict-high priority level is shown in [Table 123 on page 1671](#). The high priority and strict-high queues are serviced in a weighted round-robin fashion. The high priority queue gets its guaranteed rate and the strict-high queue gets the rest. The excess high-priority queue get all the excess bandwidth.

Table 123: Sharing with at Least One Strict-High Priority

Queue	Guaranteed (Transmit) Rate	Excess Rate	Guaranteed Priority	Excess Priority	Offered Load	Expected Throughput
Q0	20%	X	strict-high	X	500 Mbps	500 Mbps
Q1	10%	20%	high	low	500 Mbps	$100 + 0 = 100$ Mbps
Q2	10%	10%	low	high	500 Mbps	$100 + 250 = 350$ Mbps
Q3	5%	20%	low	low	500 Mbps	$50 + 0 = 50$ Mbps

The behavior with at least one strict-high priority level and a higher offered load is shown in [Table 124 on page 1672](#). The high priority and strict-high queues are serviced in a weighted round-robin fashion. The high priority queue gets its guaranteed rate and the strict-high queue gets the rest. There is no excess bandwidth.

Table 124: Sharing with at Least One Strict-High Priority and Higher Load

Queue	Guaranteed (Transmit) Rate	Excess Rate	Guaranteed Priority	Excess Priority	Offered Load	Expected Throughput
Q0	20%	X	strict-high	X	900 Mbps	900 Mbps
Q1	10%	20%	high	low	500 Mbps	100 + 0 = 100 Mbps
Q2	10%	10%	low	high	500 Mbps	0 + 0 = 0 Mbps
Q3	5%	20%	low	low	500 Mbps	0 + 0 = 0 Mbps

The behavior with at least one strict-high priority level and a rate limit is shown in [Table 125 on page 1672](#). Queue 0 and Queue 2 are rate limited, so the maximum bandwidth they are offered is the transmit bandwidth and they will not be offered any excess bandwidth. All other queues are offered the guaranteed bandwidth and the excess is shared by the non-rate-limited queues.

Table 125: Sharing with at Least One Strict-High Priority and Rate Limit

Queue	Guaranteed (Transmit) Rate	Rate Limit	Excess Rate	Guaranteed Priority	Excess Priority	Offered Load	Expected Throughput
Q0	20%	Yes	X	strict-high	X	500 Mbps	200 + 0 = 200 Mbps
Q1	10%	No	20%	high	low	500 Mbps	100 + 275 = 375 Mbps
Q2	10%	Yes	10%	low	high	500 Mbps	100 + 0 = 100 Mbps
Q3	5%	No	20%	low	low	500 Mbps	50 + 275 = 325 Mbps

Configuring the Schedulers

The following example configures schedulers, forwarding classes, and a scheduler map for an interface with excess rates and excess priorities.

```
[edit class-of-service schedulers]
scheduler-1 {
  transmit-rate percent 20;
  priority high;
  excess-rate percent 10;
  excess-priority low;
}
```

	<pre> scheduler-2 { transmit-rate percent 10; priority strict-high; } scheduler-3 { transmit-rate percent 10; priority medium-high; excess-rate percent 20; excess-priority high; } scheduler-4 { transmit-rate percent 5; priority medium-high; excess-rate percent 30; excess-priority low; } </pre>
Configuring the Forwarding Classes	<pre> [edit class-of-service] forwarding-classes { class cp_000 queue-num 0; class cp_001 queue-num 1; class cp_010 queue-num 2; class cp_011 queue-num 3; class cp_100 queue-num 4; class cp_101 queue-num 5; class cp_110 queue-num 6; class cp_111 queue-num 7; } </pre>
Configuring the Scheduler Map	<pre> [edit class-of-service scheduler-maps] scheduler-map-1 { forwarding-class cp_000 scheduler scheduler-1; forwarding-class cp_001 scheduler scheduler-2; forwarding-class cp_010 scheduler scheduler-3; forwarding-class cp_011 scheduler scheduler-4; } </pre>
Applying the Scheduler Map to the Interface	<pre> [edit interfaces] ge-1/1/0 { scheduler-map scheduler-map-1; unit 0 { family inet { address 192.168.1.2/32; } } } </pre>

RED Drop Profiles

- [Overview on page 1674](#)
- [Configuration on page 1677](#)

Overview

- [RED Drop Profiles on page 1674](#)

RED Drop Profiles

- [RED Drop Profiles Overview on page 1674](#)
- [Default Drop Profile on page 1676](#)
- [Packet Loss Priority Configuration Overview on page 1676](#)

RED Drop Profiles Overview

You can configure two parameters to control congestion at the output stage. The first parameter defines the delay-buffer bandwidth, which provides packet buffer space to absorb burst traffic up to the specified duration of delay. Once the specified delay buffer becomes full, packets with 100 percent drop probability are dropped from the head of the buffer. For more information, see [“Configuring the Scheduler Buffer Size” on page 1579](#).

The second parameter defines the drop probabilities across the range of delay-buffer occupancy, supporting the random early detection (RED) process. When the number of packets queued is greater than the ability of the router or switch to empty a queue, the queue requires a method for determining which packets to drop from the network. To address this, the Junos OS provides the option of enabling RED on individual queues.

Depending on the drop probabilities, RED might drop many packets long before the buffer becomes full, or it might drop only a few packets even if the buffer is almost full.

A *drop profile* is a mechanism of RED that defines parameters that allow packets to be dropped from the network. Drop profiles define the meanings of the loss priorities.

When you configure drop profiles, there are two important values: the queue fullness and the drop probability. The *queue fullness* represents a percentage of the memory used to store packets in relation to the total amount that has been allocated for that specific queue. Similarly, the *drop probability* is a percentage value that correlates to the likelihood that an individual packet is dropped from the network. These two variables are combined in a graph-like format, as shown in [Figure 13 on page 1675](#).




NOTE: You can only specify two fill levels for interpolated drop profiles on the Enhanced Queuing DPC for Juniper Network MX Series Ethernet Services Routers and EX Series switches. For more information about interpolated drop profiles on the Enhanced Queuing DPC for MX Series routers and EX Series switches, see *Configuring WRED on Enhanced Queuing DPCs*.

[Figure 13 on page 1675](#) shows both a segmented and an interpolated graph. Although the formation of these graph lines is different, the application of the profile is the same. When a packet reaches the head of the queue, a random number between 0 and 100 is calculated by the router or switch. This random number is plotted against the drop profile using the current queue fullness of that particular queue. When the random number falls

above the graph line, the packet is transmitted onto the physical media. When the number falls below graph the line, the packet is dropped from the network.

Figure 13: Segmented and Interpolated Drop Profiles



By defining multiple fill levels and drop probabilities, you create a segmented drop profile. The line segments are defined in terms of the following graphical model: in the first quadrant, the x axis represents the fill level, and the y axis represents the drop probability. The initial line segment spans from the origin (0,0) to the point ($\langle l1 \rangle$, $\langle p1 \rangle$); a second line runs from ($\langle l1 \rangle$, $\langle p1 \rangle$) to ($\langle l2 \rangle$, $\langle p2 \rangle$) and so forth, until a final line segment connects (100, 100). The software automatically constructs a drop profile containing 64 fill levels at drop probabilities that approximate the calculated line segments.



NOTE: If you configure the `interpolate` statement, you can specify more than 64 pairs, but the system generates only 64 discrete entries.

You specify drop probabilities in the drop profile section of the class-of-service (CoS) configuration hierarchy and reference them in each scheduler configuration. For each scheduler, you can configure multiple separate drop profiles, one for each combination of loss priority (low, medium-low, medium-high, or high) and protocol.

You can configure a maximum of 32 different drop profiles.

To configure RED drop profiles, include the following statements at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
drop-profiles {
  profile-name {
    fill-level percentage drop-probability percentage;
    interpolate {
      drop-probability [ values ];
      fill-level [ values ];
    }
  }
}
```

Default Drop Profile

By default, if you configure no drop profiles, RED is still in effect and functions as the primary mechanism for managing congestion. In the default RED drop profile, when the fill-level is 0 percent, the drop probability is 0 percent. When the fill-level is 100 percent, the drop probability is 100 percent.

As a backup method for managing congestion, tail dropping takes effect when congestion of small packets occurs. On Juniper Networks M320 Multiservice Edge Routers and T Series Core Routers, the software supports *tail-RED*, which means that when tail dropping occurs, the software uses RED to execute intelligent tail drops. On other routers, the software executes tail drops unconditionally.

Packet Loss Priority Configuration Overview

Loss priority settings help determine which packets are dropped from the network during periods of congestion. The software supports multiple packet loss priority (PLP) designations: **low** and **high**. (In addition, **medium-low** and **medium-high** PLPs are supported when you configure tricolor marking, as discussed in [“Configuring Tricolor Marking” on page 1447](#).) You can set PLP by configuring a behavior aggregate or multifield classifier, as discussed in [“Setting Packet Loss Priority” on page 1337](#) and [“Configuring Multifield Classifiers” on page 1390](#).



NOTE: On T Series routers with different Packet Forwarding Engines (non-Enhanced Scaling and Enhanced Scaling FPCs), you can configure PLP bit copying for ingress and egress unicast and multicast traffic. To configure, include the **copy-plp-all** statement at the [edit class-of-service] hierarchy level.

A drop-profile map examines the loss priority setting of an outgoing packet: **high**, **medium-high**, **medium-low**, **low**, or any.

Obviously, *low*, *medium-low*, *medium-high*, and *high* are relative terms, which by themselves have no meaning. Drop profiles define the meanings of the loss priorities. In the following example, the **low-drop** drop profile defines the meaning of **low** PLP as a 10 percent drop probability when the fill level is 75 percent and a 40 percent drop probability when the fill level is 95 percent. The **high-drop** drop profile defines the meaning of **high** PLP as a 50 percent drop probability when the fill level is 25 percent and a 90 percent drop probability when the fill level is 50 percent.

In this example, the scheduler includes two drop-profile maps, which specify that packets are evaluated by the **low-drop** drop profile if they have a **low** loss priority and are from any protocol. Packets are evaluated by the **high-drop** drop profile if they have a **high** loss priority and are from any protocol.

```
[edit class-of-service]
drop-profiles {
  low-drop {
    interpolate {
      drop-probability [ 10 40];
```

```

        fill-level [ 75 95];
    }
}
high-drop {
    interpolate {
        drop-probability [ 50 90];
        fill-level [ 25 50];
    }
}
}
schedulers {
    best-effort {
        drop-profile-map loss-priority low protocol any drop-profile low-drop;
        drop-profile-map loss-priority high protocol any drop-profile high-drop;
    }
}
}

```

Related Documentation

- [Configuring Schedulers on page 1579](#)
- [Setting Packet Loss Priority on page 1337](#)
- [copy-plp-all on page 1355](#)
- [Configuring Multifield Classifiers on page 1390](#)

Configuration

- [Configuration Tasks on page 1677](#)
- [Examples on page 1679](#)
- [Configuration Statements on page 1681](#)

Configuration Tasks

- [Configuring RED Drop Profiles on page 1677](#)
- [Configuring Weighted RED Buffer Occupancy on page 1678](#)

Configuring RED Drop Profiles

You enable RED by applying a drop profile to a scheduler. When RED is operational on an interface, the queue no longer drops packets from the tail of the queue. Rather, packets are dropped after they reach the head of the queue.

To configure a drop profile, include the **drop-profiles** statement at the **[edit class-of-service]** hierarchy level:

```

[edit class-of-service]
drop-profiles {
    profile-name {
        fill-level percentage drop-probability percentage;
        interpolate {
            drop-probability [ values ];
            fill-level [ values ];
        }
    }
}
}

```

In this configuration, include either the **interpolate** statement and its options, or the fill-level and drop-probability **percentage** values. These two alternatives enable you to configure either each drop probability at up to 64 fill-level/drop-probability paired values, or a profile represented as a series of line segments, as discussed in [“RED Drop Profiles Overview” on page 1674](#).

After you configure a drop profile, you must assign the drop profile to a drop-profile map, and assign the drop-profile map to a scheduler, as discussed in [“Configuring Drop Profile Maps for Schedulers” on page 1590](#).

Configuring Weighted RED Buffer Occupancy

By default, RED is performed based on instantaneous buffer occupancy information. However, IQ-PICs can be configured to use *weighted average* buffer occupancy information. This option is configured on a per-PIC basis and applies to the following IQ-PICs:

- Channelized T1/T3
- Channelized E1/E3
- Channelized OC3/STM1
- Channelized OC12

If you configure this feature on an unsupported PIC, you see an error message.

If you configure this feature on a channelized OC12 intelligent queuing (IQ) PIC, the PIC reboots.

When weighted average buffer occupancy is configured, you configure a weight value for averaged buffer occupancy calculations. This weight value is expressed as a negative exponential value of 2 in a fractional expression. For example, a configured weight value of 2 would be expressed as $1/(2^2) = 1/4$. If a configured weight value was configured as 1 (the default), the value would be expressed as $1/(2^1) = 1/2$.

This calculated weight value is applied to the instantaneous buffer occupancy value to determine the new value of the weighted average buffer occupancy. The formula to derive the new weighted average buffer occupancy is:

new average buffer occupancy = weight value * instantaneous buffer occupancy + (1 – weight value) * current average buffer occupancy

For example, if the weight exponent value is configured as 3 (giving a weight value of $1/2^3 = 1/8$), the formula used to determine the new average buffer occupancy based on the instant buffer usage is:

new average buffer occupancy = 1/8 * instantaneous buffer occupancy + (7/8) * current average buffer occupancy

The valid operational range for the weight value on IQ-PICs is 0 through 31. A value of 0 results in the average buffer occupancy being the same as the instantaneous buffer occupancy calculations. Values higher than 31 can be configured, but in these cases the current maximum *operational* value of 31 is used for buffer occupancy calculations.



NOTE: The `show interfaces` command with the `extensive` option displays the *configured* value for the RED buffer occupancy weight exponent. However, in all such cases, the current *operational* maximum value of 31 is used internally.

To configure a Q-PIC for RED weighted average buffer occupancy calculations, include the `red-buffer-occupancy` statement with the `weighted-averaged` option at the `[edit chassis fpc slot-number pic pic-number]` hierarchy level:

```
[edit chassis]
fpc slot-number {
  pic pic-number {
    red-buffer-occupancy {
      weighted-averaged [ instant-usage-weight-exponent exponent-number ];
    }
  }
}
```

Related Documentation

- [Example: Configuring Weighted RED Buffer Occupancy on page 1680](#)
- `red-buffer-occupancy`


Examples

- [Example: Configuring RED Drop Profiles on page 1679](#)
- [Example: Configuring Weighted RED Buffer Occupancy on page 1680](#)

Example: Configuring RED Drop Profiles

Create a segmented configuration and an interpolated configuration that correspond to the graphs in [Figure 14 on page 1679](#). The values defined in the configuration are matched to represent the data points in the graph line. In this example, the drop probability is 25 percent when the queue is 50 percent full. The drop probability increases to 50 percent when the queue is 75 percent full.

Figure 14: Segmented and Interpolated Drop Profiles



Creating a Segmented Configuration

```
class-of-service {
  drop-profiles {
    segmented-style-profile {
```

```

        fill-level 25 drop-probability 25;
        fill-level 50 drop-probability 50;
        fill-level 75 drop-probability 75;
        fill-level 95 drop-probability 100;
    }
}

```

To create the profile's graph line, the software begins at the bottom-left corner, representing a 0 percent fill level and a 0 percent drop probability. This configuration draws a line directly to the right until it reaches the first defined fill level, 25 percent for this configuration. The software then continues the line vertically until the first drop probability is reached. This process is repeated for all of the defined levels and probabilities until the top-right corner of the graph is reached.

Create a smoother graph line by configuring the profile with the **interpolate** statement. This allows the software to automatically generate 64 data points on the graph beginning at (0, 0) and ending at (100, 100). Along the way, the graph line intersects specific data points, which you define as follows:

Creating an Interpolated Configuration

```

class-of-service {
  drop-profiles {
    interpolated-style-profile {
      interpolate {
        fill-level [ 50 75 ];
        drop-probability [ 25 50 ];
      }
    }
  }
}

```

Example: Configuring Weighted RED Buffer Occupancy

Configure the Q-PIC to use a weight value of 1/2 in average buffer occupancy calculations.

```

[edit chassis]
fpc 0 {
  pic 1 {
    red-buffer-occupancy {
      weighted-averaged instant-usage-weight-exponent 1;
    }
  }
}

```

or

```

[edit chassis]
fpc 0 {
  pic 1 {
    red-buffer-occupancy {
      weighted-averaged; # the default value is 1 if not specified
    }
  }
}

```

Configure the Q-PIC to use a weight value of 1/4 in average buffer occupancy calculations.

```
[edit chassis]
fpc 0 {
  pic 1 {
    red-buffer-occupancy {
      weighted-averaged instant-usage-weight-exponent 2;
    }
  }
}
```

**Related
Documentation**

- [Configuring Weighted RED Buffer Occupancy on page 1678](#)
- *red-buffer-occupancy*

Configuration Statements

- [\[edit class-of-service\] Hierarchy Level on page 1681](#)

[edit class-of-service] Hierarchy Level

```
class-of-service {
  classifiers {
    type classifier-name {
      forwarding-class class-name {
        loss-priority (high | low | medium-high | medium-low) code-points [ aliases bits ];
      }
      import (classifier-name | default);
    }
  }
  code-point-aliases {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) {
      alias-name bits;
    }
  }
  drop-profiles {
    profile-name {
      fill-level percentage drop-probability percentage;
      interpolate {
        drop-probability value;
        fill-level value;
      }
    }
  }
  fabric {
    scheduler-map {
      priority (high | low) scheduler scheduler-name;
    }
  }
  forwarding-class-map {
    map-name {
      class class-name queue-num queue-number <restricted-queue queue-number>;
    }
  }
  forwarding-classes {
```

```

class class-name policing-priority (normal | premium) queue-num queue-number
  priority (high | low);
queue queue-number class-name policing-priority (normal | premium) priority (high |
  low);
}
forwarding-policy {
  class class-name {
    classification-override {
      forwarding-class class-name;
    }
  }
  next-hop-map map-name {
    forwarding-class class-name {
      discard;
      lsp-next-hop [ lsp-regular-expressions ];
      next-hop [ next-hop-names ];
      non-lsp-next-hop;
    }
  }
}
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      drop-timeout milliseconds;
      fragment-threshold bytes;
      multilink-class number;
      no-fragmentation;
    }
  }
}
host-outbound-traffic {
  dscp-code-point value;
  forwarding-class class-name;
  ieee-802.1 {
    default value;
    rewrite-rules;
  }
  tcp {
    raise-internet-control-priority;
  }
}
interfaces {
  ... the interfaces subhierarchy appears after the main [edit class-of-service] hierarchy
  ...
}
}
restricted-queues {
  forwarding-class class-name queue-number;
}
rewrite-rules {
  (dscp | dscp-ipv6 | exp | frame-relay-de | ieee-802.1 | ieee-802.1ad | inet-precedence)
  rewrite-rule {
    forwarding-class class-name {
      loss-priority level code-point (alias | bits);
    }
    import (rewrite-rule | default);
  }
}

```



```

    }
  }
  routing-instances routing-instance-name {
    classifiers {
      dscp (classifier-name | default);
      dscp-ipv6 (classifier-name | default);
      exp (classifier-name | default);
      ieee-208.1 (classifier-name | default | encapsulated | vlan-tag (inner | outer));
    }
  }
  scheduler-maps {
    map-name {
      forwarding-class class-name scheduler scheduler-name;
    }
  }
  schedulers {
    scheduler-name {
      adjust-minimum value;
      adjust-percent value;
      buffer-size (exact | percent percentage | remainder);
      drop-profile-map loss-priority (any | high | low | medium-high | medium-low)
        protocol any;
      excess-priority (high | low | medium-high | medium-low);
      excess-rate (percent percentage | proportion proportion);
      priority (high | low | medium-high | medium-low | strict-high);
      shaping-rate (bps | percent percentage | burst-size size);
      transmit-rate (bps | percent percentage | remainder) <exact | rate-limit>;
    }
  }
  traceoptions {
    file <files number> <match regular-expression> <size maximum-file-size>
      <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
  traffic-control-profiles {
    profile-name {
      adjust-minimum rate;
      delay-buffer-rate (bps | cps cps | percent percentage);
      excess-rate (percent percentage | proportion value);
      guaranteed-rate (bps | percent percentage) <burst-size bytes>;
      overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
      scheduler-map map-name;
      shaping-rate (bps | percent percentage) <burst-size bytes>;
    }
  }
  tri-color;
}

class-of-service {
  interfaces {
    interface-name {
      excess-bandwidth-share (equal | proportional value);
      input-excess-bandwidth-share (equal | proportional value);
      input-scheduler-map map-name;
      input-shaping-rate bps;
    }
  }
}

```

```

input-traffic-control-profile profile-name;
output-forwarding-class-map map-name;
output-traffic-control-profile profile-name;
scheduler-map map-name;
scheduler-map-chassis (map-name | derived);
shaping-rate bps;
unit (logical-unit-number | *) {
  classifiers {
    dscp (classifier-name | default) {
      family [ inet mpls ];
    }
    dscp-ipv6 (classifier-name | default) {
      family [ inet mpls ];
    }
    exp (classifier-name | default);
    ieee-208.1 (classifier-name | default) <vlan-tag (inner | outer)>;
    ieee-208.1ad (classifier-name | default);
    inet-precedence (classifier-name | default);
  }
  forwarding-class class-name;
  input-scheduler-map map-name;
  input-shaping-rate bps;
  input-traffic-control-profile profile-name shared-instance instance-name;
  loss-priority-maps {
    (map-name | default);
  }
  loss-priority-rewrites {
    (map-name | default);
  }
  output-forwarding-class-map map-name;
  output-traffic-control-profile profile-name shared-instance instance-name;
  rewrite-rules {
    dscp (rule-name | default) <protocol mpls>;
    dscp-ipv6 (rule-name | default);
    exp (rule-name | default) <protocol [ mpls-any | mpls-inet-both |
      mpls-inet-both-non-vpn ]>;
    exp-push-push-push default;
    exp-swap-push-push default;
    ieee-802.1 (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
    ieee-802.1ad (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
    inet-precedence (rewrite-name | default) <protocol mpls>;
  }
  scheduler-map map-name;
  shaping-rate bps;
  translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 |
    to-exp-from-exp | to-inet-precedence-from-inet-precedence) table-name;
}
}
interface-set interface-set-name {
  excess-bandwidth-share (equal | proportional value);
  input-excess-bandwidth-share (equal | proportional value);
  input-traffic-control-profile profile-name;
  input-traffic-control-profile-remaining profile-name;
  internal-node;
  output-traffic-control-profile profile-name;
  output-traffic-control-profile-remaining profile-name;

```

```

    }
  }
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

drop-probability (Interpolated Value)

Syntax	drop-probability [<i>values</i>];
Hierarchy Level	[edit class-of-service drop-profiles <i>profile-name</i> interpolate]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS 11.4 for EX Series switches.
Description	Define up to 64 values for interpolating drop probabilities. On EX Series switches, this statement is supported only on the EX9200 switch, EX8200 standalone switches, and EX8200 Virtual Chassis.
Options	percentage —The probability (expressed in percentage) for a packet to be dropped from the queue. Range: 0 through 100
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Default Drop Profile on page 1676 • <i>Configuring CoS Tail Drop Profiles (CLI Procedure)</i>

drop-profiles (Class-of-Service)

Syntax	<pre>drop-profiles { profile-name { fill-level percentage drop-probability percentage; interpolate { drop-probability [values]; fill-level [values] } } }</pre>
Hierarchy Level	[edit class-of-service]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS 11.4 for EX Series switches.
Description	<p>Define drop profiles for RED.</p> <p>For a packet to be dropped, it must match the drop profile. When a packet arrives, RED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the packet.</p>
Options	<p><i>profile-name</i>—Name of the drop profile.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RED Drop Profiles on page 1677• <i>Understanding CoS Tail Drop Profiles</i>• <i>Example: Configuring CoS on EX Series Switches</i>• <i>Configuring CoS Tail Drop Profiles (CLI Procedure)</i>

fill-level (Interpolated Value)

Syntax	fill-level [<i>values</i>];
Hierarchy Level	[edit class-of-service drop-profiles <i>profile-name</i> interpolate]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS 11.4 for EX Series switches.
Description	Define up to 64 values for interpolating queue fill level. On EX Series switches, this statement is supported only on EX8200 standalone switches and EX8200 Virtual Chassis.
Options	values —Data points for mapping queue fill percentage. Range: 0 through 100 Default: In the default tail drop profile, when the fill level is 0 percent, the drop probability is 0 percent. When the fill level is 100 percent, the drop probability is 100 percent.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • RED Drop Profiles Overview on page 1674 • Configuring RED Drop Profiles on page 1677. • <i>Understanding CoS Tail Drop Profiles</i> • <i>Configuring CoS Tail Drop Profiles (CLI Procedure)</i>

fill-level (Drop Profiles)

Syntax	fill-level <i>percentage</i> ;
Hierarchy Level	[edit class-of-service drop-profiles <i>profile-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS 11.4 for EX Series switches.
Description	When configuring RED, map the fullness of a queue to a drop probability.
Options	<p>percentage—How full the queue is, expressed as a percentage. You configure the fill-level and drop-probability statements in pairs. To specify multiple fill levels, include multiple fill-level and drop-probability statements. The values you assign to each statement pair must increase relative to the previous pair's values. This is shown in the segmented graph in "RED Drop Profiles Overview" on page 1674.</p> <p>Range: 0 through 100 percent</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• RED Drop Profiles Overview on page 1674• Configuring RED Drop Profiles on page 1677• <i>Understanding CoS Tail Drop Profiles</i>• <i>Configuring CoS Tail Drop Profiles (CLI Procedure)</i>

interpolate

Syntax	<pre> interpolate { drop-probability [values]; fill-level [values]; } </pre>
Hierarchy Level	[edit class-of-service drop-profiles <i>profile-name</i>]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced before Junos OS 11.4 for EX Series switches.</p>
Description	<p>Specify values for interpolating relationship between queue fill level and drop probability.</p> <p>On EX Series switches, this statement is supported only on EX8200 standalone switches and EX8200 Virtual Chassis.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> See Configuring RED Drop Profiles on page 1677. <i>Understanding Junos OS CoS Components for EX Series Switches</i>

Rewriting Packet Header Information

- [Overview on page 1689](#)
- [Configuration on page 1692](#)

Overview

- [Rewriting Packet Header Information on page 1689](#)

Rewriting Packet Header Information

- [Rewriting Packet Header Information Overview on page 1690](#)
- [Header Bits Preserved, Cleared, and Rewritten on page 1691](#)
- [Setting IPv6 DSCP and MPLS EXP Values Independently on page 1691](#)
- [Classifiers and Rewrite Rules at the Global and Physical Interface Levels Overview on page 1692](#)

Rewriting Packet Header Information Overview

As packets enter or exit a network, edge routers might be required to alter the class-of-service (CoS) settings of the packets. Rewrite rules set the value of the CoS bits within the packet's header. Each rewrite rule reads the current forwarding class and loss priority information associated with the packet, locates the chosen CoS value from a table, and writes this CoS value into the packet header.


In effect, the rewrite rule performs the opposite function of the behavior aggregate (BA) classifier used when the packet enters the routing device. As the packet leaves the routing platform, the final CoS action is generally the application of a rewrite rule.

You configure rewrite rules to alter CoS values in outgoing packets on the outbound interfaces of an edge router to meet the policies of a targeted peer. This allows the downstream routing device in a neighboring network to classify each packet into the appropriate service group.

In addition, you often need to rewrite a given marker (IP precedence, Differentiated Services code point [DSCP], IEEE 802.1p, or MPLS EXP settings) at the inbound interfaces of an edge router to accommodate BA classification by core devices.

Figure 15 on page 1690 shows a flow of packets through four routing devices. Router A rewrites the CoS bits in incoming packet to accommodate the BA classification performed by Routers B and C. Router D alters the CoS bits of the packets before transmitting them to the neighboring network.

Figure 15: Packet Flow Across the Network



To configure CoS rewrite rules, you define the rewrite rule and apply it to an interface. Include the following statements at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
interfaces {
  interface-name {
    unit logical-unit-number {
      rewrite-rules {
        dscp (rewrite-name | default) protocol protocol-types;
        dscp-ipv6 (rewrite-name | default);
        exp (rewrite-name | default) protocol protocol-types;
        exp-push-push-push default;
        exp-swap-push-push default;
        ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
        ieee-802.1ad (rewrite-name | default) vlan-tag (outer | outer-and-inner);
        inet-precedence (rewrite-name | default) protocol protocol-types;
      }
    }
  }
}
```



```

    }
  }
}
rewrite-rules {
  (dscp | dscp-ipv6 | exp | frame-relay-de | ieee-802.1 | inet-precedence) rewrite-name {
    import (rewrite-name | default);
    forwarding-class class-name {
      loss-priority level code-point (alias | bits);
    }
  }
}
}

```

Header Bits Preserved, Cleared, and Rewritten

For every incoming packet, the ingress classifier decodes the ingress CoS bits into a forwarding class and packet loss priority (PLP) combination.

The egress CoS information depends on which type of rewrite marker is active, as follows:

- For Multiprotocol Label Switching (MPLS) EXP and IEEE 802.1 rewrite markers, values are derived from the forwarding class and PLP values in rewrite rules. MPLS EXP and IEEE 802.1 markers are not preserved because they are part of the Layer 2 encapsulation.
- For IP precedence and DiffServ code point (DSCP) rewrite markers, the marker alters the first three bits on the type-of-service (ToS) byte while leaving the last three bits unchanged.

Setting IPv6 DSCP and MPLS EXP Values Independently

On the M120, M320 with Enhanced III FPCs, MX Series Ethernet Services routers, and EX Series switches, you can set the DSCP and MPLS EXP bits independently on IPv6 packets. To enable this feature, include the **protocol mpls** statement at the **[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules dscp-ipv6 rewrite-name]** hierarchy level.

You can set DSCP IPv6 values only at the ingress MPLS node.

The following limitations apply to this feature:

- This feature is supported only on M120, M320 with Enhanced III FPCs, MX Series Ethernet Services routers, and EX Series switches.
- MPLS packets entering another MPLS tunnel at the ingress node may mark only the EXP value if EXP rewrite rules are configured, but not the DSCP value in the IPv6 header.
- This feature does not support MPLS packets generated by the Routing Engine.
- The IP precedence field is not applicable for IPv6, and is not supported.

Related Documentation

- [Configuring DSCP Values for IPv6 Packets Entering the MPLS Tunnel on page 1697](#)

Classifiers and Rewrite Rules at the Global and Physical Interface Levels Overview

On ACX Series Universal Access Routers and EX Series switches, CoS supports classification and rewrite at the global level and physical interface levels.

At a global level, you can define EXP classification.

At a physical interface level, you can define the following features:

- DSCP and inet-precedence classifiers
- DSCP and inet-precedence rewrites
- ieee-802.1 classifiers (inner and outer)
- ieee-802.1 rewrites (outer)

At a logical interface level, you can define the fixed classification and EXP rewrites.

To configure global EXP classifiers, include the **classifiers exp classifier-name** statement at the **[edit class-of-service] system-defaults** hierarchy level.

To configure classifiers or rewrite rules at the physical interface, include either the **classifiers** statement or the **rewrite-rules** statement at the **[edit class-of-service] interfaces interface-name]** hierarchy level.

To display classifiers configured under **system-defaults**, enter the **show class-of-service system-defaults** command.

To display classifiers and rewrite rules bound to physical interfaces, enter the **show class-of-service interfaces interface-name** command.

Related Documentation

- [Configuring Classifiers and Rewrite Rules at the Global and Physical Interface Levels on page 1339](#)

Configuration

- [Configuration Tasks for Applying Rewrite Rules on page 1692](#)
- [Configuration Tasks for Rewriting Packet Header Information on page 1702](#)
- [Example on page 1713](#)
- [Configuration Statements on page 1714](#)

Configuration Tasks for Applying Rewrite Rules

- [Applying Default Rewrite Rules on page 1693](#)
- [Configuring Rewrite Rules on page 1694](#)
- [Applying Rewrite Rules to Output Logical Interfaces on page 1695](#)
- [Configuring DSCP Values for IPv6 Packets Entering the MPLS Tunnel on page 1697](#)
- [Applying IEEE 802.1p Rewrite Rules to Dual VLAN Tags on page 1699](#)
- [Configuring Classifiers and Rewrite Rules at the Global and Physical Interface Levels on page 1701](#)

Applying Default Rewrite Rules

By default, rewrite rules are not usually applied to interfaces. The exceptions are MPLS interfaces: all MPLS-enabled interfaces use the default EXP rewrite rule, even if not configured. Except for MPLS interfaces, if you want to apply a rewrite rule, you can either design your own rule and apply it to an interface, or you can apply a default rewrite rule. To apply default rewrite rules, include one or more of the following statements at the `[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]` hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
dscp default;
dscp-ipv6 default;
exp default;
ieee-802.1 default vlan-tag (outer | outer-and-inner);
inet-precedence default;
```

Table 126 on page 1693 shows the default rewrite rule mappings. These are based on the default bit definitions of DSCP, DSCP IPv6, EXP, IEEE, and IP CoS values, as shown in “Default CoS Values” on page 1372, and the default forwarding classes shown in *Default Forwarding Classes*.

When the software detects packets whose CoS values match the forwarding class and PLP values in the first two columns in Table 126 on page 1693, the software maps the header bits of those packets to the code-point aliases in the last column in Table 126 on page 1693. The code-point aliases in the last column map to the CoS bits shown in “Default CoS Values” on page 1372.

Table 126: Default Packet Header Rewrite Mappings

Map from Forwarding Class	PLP Value	Map to DSCP/DSCP IPv6/ EXP/IEEE/IP
expedited-forwarding	low	ef
expedited-forwarding	high	ef
assured-forwarding	low	af11
assured-forwarding	high	af12 (DSCP/DSCP IPv6/EXP)
best-effort	low	be
best-effort	high	be
network-control	low	nc1/cs6
network-control	high	nc2/cs7

In the following example, the `ge-1/2/3.0` interface is assigned the default DSCP rewrite rule. One result of this configuration is that each packet exiting the interface with the **expedited-forwarding** forwarding class and the **high** or **low** loss priority has its DSCP bits

rewritten to the DSCP **ef** code-point alias. “[Default CoS Values](#)” on page 1372 shows that this code-point alias maps to the **101110** bits.

Another result of this configuration is that all packets exiting the interface with the **best-effort** forwarding class and the **high** or **low** loss priority have their EXP bits rewritten to the EXP **be** code-point alias. “[Default CoS Values](#)” on page 1372 shows that this code-point alias maps to the **000** bits.

To evaluate all the implications of this example, see “[Default CoS Values](#)” on page 1372 and [Table 126](#) on page 1693.

```
class-of-service {
  interfaces {
    ge-1/2/3 {
      unit 0 {
        rewrite-rules {
          dscp default;
        }
      }
    }
  }
}
```

Configuring Rewrite Rules

You define markers in the rewrite rules section of the CoS configuration hierarchy and reference them in the logical interface configuration. This model supports marking on the DSCP, DSCP IPv6, IP precedence, IEEE 802.1, and MPLS EXP CoS values.

To configure a rewrite-rules mapping and associate it with the appropriate forwarding class and code-point alias or bit set, include the **rewrite-rules** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
rewrite-rules {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) rewrite-name {
    import (rewrite-name | default);
    forwarding-class class-name {
      loss-priority level code-point (alias | bits);
    }
  }
}
```



NOTE: The **inet-precedence** statement is not supported on PTX Series Packet Transport Switches.

The rewrite rule sets the code-point aliases and bit patterns for a specific forwarding class and PLP. The inputs for the map are the forwarding class and the PLP. The output of the map is the code-point alias or bit pattern. For more information about how CoS maps work, see “[CoS Inputs and Outputs Overview](#)” on page 1272.

By default, IP precedence rewrite rules alter the first three bits on the type-of-service (ToS) byte while leaving the last three bits unchanged. This default behavior is not

configurable. The default behavior applies to rules you configure by including the **inet-precedence** statement at the **[edit class-of-service rewrite-rules]** hierarchy level. The default behavior also applies to rewrite rules you configure for MPLS packets with IPv4 payloads. You configure these types of rewrite rules by including the **mpls-inet-both** or **mpls-inet-both-non-vpn** option at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules exp *rewrite-rule-name* protocol]** hierarchy level.

On the M320, T1600, and MX960 routers and EX Series switches, if you configure **vlan-vpls** encapsulation and add an IEEE 802.1 header on a Gigabit Ethernet or 10 Gigabit Ethernet interface to output traffic, but do not apply an IEEE 802.1 rewrite rule, then the default IEEE 802.1 rewrite rule is ignored and the IEEE 802.1p bits are set to match the forwarding class queue.



NOTE: The forwarding class is determined by ingress classification.

Related Documentation

- [Applying Rewrite Rules to Output Logical Interfaces](#)
- [Applying Egress Interface Rewrite Rules to the IEEE 802.1p Field for All Host Outbound Traffic on the Interface on page 1711](#)

Applying Rewrite Rules to Output Logical Interfaces

To assign the rewrite-rules configuration to the output logical interface, include the **rewrite-rules** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
rewrite-rules {
  dscp (rewrite-name | <default>) protocol protocol-types;
  dscp-ipv6 (rewrite-name | <default>) protocol protocol-types
  exp (rewrite-name | <default>) protocol protocol-types;
  exp-push-push-push <default>;
  exp-swap-push-push <default>;
  ieee-802.1 (rewrite-name | <default>) inet-prec vlan-tag (outer | outer-and-inner);
  inet-precedence (rewrite-name | <default>) protocol protocol-types;
}
```

On M120, M320 with an Enhanced III FPC, MX Series routers and T 4000 routers with Type 5 FPCs and EX Series switches, you can combine the **dscp** or **inet-prec** and **exp** options to set the DSCP or IP precedence bits and MPLS EXP bits independently on IP packets entering an MPLS tunnel.

For IQ PICs, you can configure only one IEEE 802.1 rewrite rule on a physical port. All logical ports (units) on that physical port should apply the same IEEE 802.1 rewrite rule. If you configure more than one IEEE 802.1 rewrite rule for the IQ PIC, the configuration check fails.

In the following example, the DSCP bits specified in **ss-dscp** are applied to packets entering the MPLS tunnel on **ge-2/1/1**, and the DSCP bits specified in **ss-v6dscp** are applied to IPv6 packets. The EXP bits are set to the bit configuration specified in **ss-exp**:

```
[edit class-of-service interfaces]
ge-2/1/1
  unit 10 {
    rewrite-rules {
      dscp ssf-dscp protocol mpls; # Applies to IPv4 packets entering MPLS tunnel
      dscp-ipv6 ss-v6dscp protocol mpls; # Applies to IPv6 packets entering MPLS tunnel
      exp ss-exp; # Sets label EXP bits independently
    }
  }
}
```

You can use interface wildcards for *interface-name* and *logical-unit-number*. You can also include Layer 2 and Layer 3 rewrite information in the same configuration.



NOTE: On M Series routers only, if you include the `control-word` statement at the `[edit protocols l2circuit neighbor address interface interface-name]` hierarchy level, the software cannot rewrite MPLS EXP bits.

DSCP and DSCP IPv6 rewrite rules are supported on M Series and T Series routers when non-queuing PICs are installed, but are disabled when queuing PICs are installed with the following exceptions:

- On M320 routers, DSCP rewrite is supported on IQ, IQ2, IQE, and IQ2E PICs when used with the Enhanced III FPC.
- On M120 routers, DSCP rewrite is supported on IQ, IQ2, IQE, and IQ2E PICs.

DSCP and DCSP IPv6 rewrite rules are supported on MX Series routers with MPC/MIC interfaces and EX Series switches.

DSCP rewrite rules are not supported on T Series routers when IQ, IQ2, IQE, IQ2E, or PD-5-10XGE-SFPP PICs are installed.

For IQ PICs, you can configure only one IEEE 802.1 rewrite rule on a physical port. All logical ports (units) on that physical port should apply the same IEEE 802.1 rewrite rule.

On M320 and T Series routers (except for T4000 routers with Type 5 FPCs), for a single interface, you cannot enable a rewrite rule on a subset of forwarding classes. You must assign a rewrite rule to either none of the forwarding classes or all of the forwarding classes. When you assign a rewrite rule to a subset of forwarding classes, the commit does not fail, and the subset of forwarding classes works as expected. However, the forwarding classes to which the rewrite rule is not assigned are rewritten to all zeros.

For example, if you configure a Differentiated Services code point (DSCP) rewrite rule, the bits in the forwarding classes to which you do not assign the rewrite rule are rewritten to 000000; if you configure an IP precedence rewrite rule, the bits in the forwarding classes to which you do not assign the rewrite rule are rewritten to 000.

Related Documentation

- [Setting IPv6 DSCP and MPLS EXP Values Independently on page 1691](#)
- [Configuring DSCP Values for IPv6 Packets Entering the MPLS Tunnel on page 1697](#)

Configuring DSCP Values for IPv6 Packets Entering the MPLS Tunnel

The following configuration example explains in detail how to set the DSCP and MPLS EXP bits independently on IPv6 packets.

1. Configure the router device (ingress PE router) to classify (behavior aggregate or multifield) the incoming packets to a particular forwarding class.

```
[edit firewall]
family inet6 {
```

```
filter ss-v6filt {
  term ss-vpn {
    from {
      destination-address {
        ::ffff:192.0.2.128/120;
      }
    }
    then {
      loss-priority low;
      forwarding-class ss-fc;
    }
  }
}
```

In the preceding example, the ingress FPC classifies (MF) incoming IPv6 packets destined for address “::ffff:192.0.2.128/120” to forwarding class “ss-fc” and loss priority “low.”

2. Attach the preceding firewall filter to an interface. Because you are matching on inbound traffic, this would be an input filter. This classifies all traffic on the interface “ge-2/1/0” that matches the filter “ss-v6.”

```
[edit interfaces]
ge-2/1/0 {
  hierarchical-scheduler;
  vlan-tagging;
  unit 300 {
    family inet6 {
      filter {
        input ss-v6filt;
      }
      address ::ffff:192.0.2.100/120;
    }
  }
}
```

3. Configure the DSCP–IPv6 rewrite rule for the forwarding class “ss-fc.” This causes the outgoing IPv6 packets belonging to the forwarding class “ss-fc” and loss priority “low” to have their DSCP value rewritten to 100000.

```
[edit class-of-service rewrite-rules]
dscp-ipv6 ss-v6dscp {
  forwarding-class ss-fc {
    loss-priority low code-point 100000;
  }
}
```

4. Configure the EXP rewrite values for the forwarding class “ss-fc.” This rewrite rule stamps an EXP value of 100 on all outgoing MPLS packets assigned to the forwarding class “ss-fc” and loss priority “low.”

```
[edit class-of-service rewrite-rules]
exp ss-exp {
  forwarding-class ss-fc {
    loss-priority low code-point 100;
  }
}
```



```
}
```

5. Apply the preceding rewrite rule to an egress interface. On the egress FPC, all IPv6 packets in the forwarding class “ss-fc” with loss priority “low” are marked with the DSCP value “100000” and an EXP value of “100” before they enter the MPLS tunnel.

```
[edit class-of-service interfaces]
ge-2/1/1 {
  unit 10 {
    rewrite-rules {
      dscp-ipv6 ss-v6dscp protocol mpls;
      exp ss-exp;
    }
  }
}
```

6. To support IPv4 DSCP and MPLS EXP independent rewrite at the same time, you can define and apply an IPv4 DSCP rewrite rule “ss-dscp” to the same interface.

```
[edit class-of-service interfaces]
ge-2/1/1 {
  unit 10 {
    rewrite-rules {
      dscp ss-dscp protocol mpls;
      dscp-ipv6 ss-v6dscp protocol mpls;
      exp ss-exp;
    }
  }
}
```

Related Documentation

- [Setting IPv6 DSCP and MPLS EXP Values Independently on page 1691](#)

Applying IEEE 802.1p Rewrite Rules to Dual VLAN Tags

By default, when you apply an IEEE 802.1p rewrite rule to an output logical interface, the software rewrites the IEEE bits in the outer VLAN tag only.

For Gigabit Ethernet IQ2 PICs, 10-port 10-Gigabit OSE PICs, and 10-Gigabit Ethernet IQ2 PICs only, you can rewrite the IEEE bits in both the outer and inner VLAN tags of the tagged Ethernet frames. When you enable class of service (CoS) rewrite for both tags, the same IEEE 802.1p rewrite table is used for the inner and outer VLAN tag.

For IQ PICs, you can only configure one IEEE 802.1 rewrite rule on a physical port. All logical ports (units) on that physical port should apply the same IEEE 802.1 rewrite rule.

To rewrite both the outer and inner VLAN tags, include the **vlan-tag outer-and-inner** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules *ieee-802.1* (*rewrite-name* | default)]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules
  ieee-802.1 (rewrite-name | default)]
  vlan-tag outer-and-inner;
```

To explicitly specify the default behavior, include the **vlan-tag outer** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules *ieee-802.1* (*rewrite-name* | default)]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules
  ieee-802.1 (rewrite-name | default)]
  vlan-tag outer;
```

For more information about VLAN tags, see the *Junos® OS Network Interfaces*.

On MX routers and EX Series switches, you can perform IEEE 802.1p and DEI rewriting based on forwarding class and PLP at the VPLS ingress PE. You rewrite (mark) the IEEE 802.1p or DEI bits on frames at the VPLS ingress PE based on the value of the forwarding class and PLP established for the traffic. You can rewrite either the outer tag only or the outer and inner tag. When both tags are rewritten, both get the same value. To configure these rewrite rules, include the **ieee-802.1** statement at the **[edit class-of-services routing-instance *routing-instance-name* rewrite-rules]** hierarchy level.

On routing devices with IQ2 or IQ2-E PICs, you can perform IEEE 802.1p and DEI rewriting based on forwarding-class and packet loss priority (PLP) at the VPLS ingress provider edge (PE) router. You rewrite (mark) the IEEE 802.1p or DEI bits on frames at the VPLS ingress PE based on the value of the forwarding-class and PLP established for the traffic. You can rewrite either the outer tag only or both the outer and inner tags. When both tags are rewritten, both get the same value.



NOTE: The 10-port 10-Gigabit OSE PIC does not support DEI rewriting based on forwarding class and PLP at the VPLS ingress PE.

To configure these rewrite rules, include the **ieee-802.1** statement at the **[edit class-of-services routing-instance *routing-instance-name* rewrite-rules]** hierarchy level.

Example: Applying an IEEE 802.1p Rewrite Rule to Dual VLAN Tags

Apply the **ieee8021p-rwrule1** rewrite rule to both inner and outer VLAN tags of Ethernet-tagged frames exiting the **ge-0/0/0.0** interface:

```
class-of-service {
  interfaces {
    ge-0/0/0 {
      unit 0 {
        rewrite-rules {
          ieee-802.1 ieee8021p-rwrule1 vlan-tag outer-and-inner;
        }
      }
    }
  }
}
```

Configuring Classifiers and Rewrite Rules at the Global and Physical Interface Levels

On ACX Series Universal Access Routers and EX Series switches, CoS supports classification and rewrite at the global and physical interface levels.

To configure the global EXP classifier, include the following statements at the **[edit class-of-service] system-defaults** hierarchy level.

```
[edit class-of-service]
{
  system-defaults
  {
    classifiers exp classifier-name
  }
}
```

CoS supports one global system default classifier of the EXP type, as shown in the following example:

```
[edit class-of-service]
{
  system-defaults {
    classifiers {
      exp exp-classf-core;
    }
  }
}
```

To configure classifiers and rewrite rules at the physical interface level, include the following statements at the **[edit class-of-service] interfaces** hierarchy level.

```
[edit class-of-service]
interfaces {
  interface-name
  classifiers dscp classifier-name
  classifiers inet-precedence classifier-name
  classifiers ieee-802.1 [vlan-tag (outer | inner)] classifier-name
  rewrite-rules dscp rewrite-name
  rewrite-rules inet-prec rewrite-name
  rewrite-rules ieee-802.1 rewrite-name
}
```

The following example shows classifiers and rewrite rules configured on physical interfaces:

```
ge-0/1/0 {
  unit 0 {
    rewrite-rules {
      exp custom-exp;
    }
  }
  classifiers {
    dscp dl;
    ieee-802.1 ci;
  }
  rewrite-rules {
```

```
    dscp default;
  }
}
ge-0/1/2 {
  classifiers {
    ieee-802.1 ci;
  }
  rewrite-rules {
    ieee-802.1 ri;
  }
}
ge-0/1/3 {
  unit 0 {
    rewrite-rules {
      exp custom-exp2;
    }
  }
}
ge-0/1/7 {
  classifiers {
    dscp d1;
  }
}
ge-0/1/8 {
  classifiers {
    dscp d1;
  }
}
```

**Related
Documentation**

- [Classifiers and Rewrite Rules at the Global and Physical Interface Levels Overview on page 1324](#)

Configuration Tasks for Rewriting Packet Header Information

- [Rewriting MPLS and IPv4 Packet Headers on page 1702](#)
- [Rewriting the EXP Bits of All Three Labels of an Outgoing Packet on page 1706](#)
- [Rewriting IEEE 802.1p Packet Headers with an MPLS EXP Value on page 1708](#)
- [Configuring the IEEE 802.1p Field for CoS Host Outbound Traffic on page 1710](#)
- [Configuring a Global Default IEEE 802.1p Value for All Host Outbound Traffic on page 1710](#)
- [Applying Egress Interface Rewrite Rules to the IEEE 802.1p Field for All Host Outbound Traffic on the Interface on page 1711](#)
- [Setting Ingress DSCP Bits for Multicast Traffic over Layer 3 VPNs on page 1712](#)

Rewriting MPLS and IPv4 Packet Headers

You can apply a rewrite rule to MPLS and IPv4 packet headers simultaneously. This allows you to initialize MPLS EXP and IP precedence bits at LSP ingress. You can configure different rewrite rules depending on whether the traffic is VPN or non-VPN.

The default MPLS EXP rewrite table contents are shown in [Table 127 on page 1703](#).

Table 127: Default MPLS EXP Rewrite Table

Forwarding Class	Loss Priority	CoS Value
best-effort	low	000
best-effort	high	001
expedited-forwarding	low	010
expedited-forwarding	high	011
assured-forwarding	low	100
assured-forwarding	high	101
network-control	low	110
network-control	high	111

By default, IP precedence rewrite rules alter the first three bits on the type-of-service (ToS) byte while leaving the last three bits unchanged. This default behavior applies to rewrite rules you configure for MPLS packets with IPv4 payloads.

To override the default MPLS EXP rewrite table and rewrite MPLS and IPv4 packet headers simultaneously, include the **protocol** statement at the **[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules exp rewrite-rule-name]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules exp
rewrite-rule-name]
protocol protocol-types;
```

The **protocol** statement defines the types of MPLS packets and packet headers to which the specified rewrite rule is applied. The MPLS packet can be a standard MPLS packet or an MPLS packet with an IPv4 payload. Specify the type of MPLS packet using the following options:

- **mpls**—Applies the rewrite rule to MPLS packets and writes the CoS value to MPLS headers.
- **mpls-inet-both**—Applies the rewrite rule to VPN MPLS packets with IPv4 payloads. On Juniper Networks M120 Multiservice Edge Routers, M320 Multiservice Edge Routers, and T Series Core Routers (except T4000 routers), writes the CoS value to the MPLS and IPv4 headers. On other M Series Multiservice Edge Router routers, causes all ingress MPLS LSP packets with IPv4 payloads to be initialized with **000** code points for the MPLS EXP value, and the configured rewrite code point for IP precedence.
- **mpls-inet-both-non-vpn**—Applies the rewrite rule to non-VPN MPLS packets with IPv4 payloads. On Juniper Networks M120 Multiservice Edge Routers, M320 Multiservice Edge Routers, and T Series Core Routers, writes the CoS value to the MPLS and IPv4 headers. On other M Series Multiservice Edge Routers, causes all ingress MPLS LSP

packets with IPv4 payloads to be initialized with **000** code points for the MPLS EXP value, and the configured rewrite code point for IP precedence.

On M120 routers, M320 routers with Enhanced-III FPCs, MX Series routers, and EX Series switches, you can perform simultaneous DSCP and EXP rewrite by attaching independent DSCP or IPv4 precedence rewrite rules and EXP rewrite rules to the same core interface. Thus, you can rewrite both code points (DSCP and EXP) when the packet is received by the ingress provider edge (PE) router on the MPLS core.

An alternative to overwriting the default with a rewrite-rules mapping is to configure the default packet header rewrite mappings, as discussed in *Applying Default Rewrite Rules*.

By default, IP precedence rewrite rules alter the first three bits on the ToS byte while leaving the last three bits unchanged. This default behavior is not configurable. The default behavior applies to rules you configure by including the **inet-precedence** statement at the **[edit class-of-service rewrite-rules]** hierarchy level. The default behavior also applies to rewrite rules you configure for MPLS packets with IPv4 payloads. You configure these types of rewrite rules by including the **mpls-inet-both** or **mpls-inet-both-non-vpn** option at the **[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules exp rewrite-rule-name protocol]** hierarchy level.

Example: Rewriting MPLS and IPv4 Packet Headers

On M320 and T Series routers, configure rewrite tables and apply them in various ways to achieve the following results:

- For interface **ge-3/1/0**, the three EXP rewrite tables are applied to packets, depending on the protocol of the payload:
 - IPv4 packets (VPN) that enter the LSPs on interface **ge-3/1/0** are initialized with values from rewrite table **exp-inet-table**. An identical 3-bit value is written into the IP precedence and MPLS EXP bit fields.
 - IPv4 packets (non-VPN) that enter the LSPs on interface **ge-3/1/0** are initialized with values from rewrite table **rule-non-vpn**. An identical 3-bit value is written into the IP precedence and MPLS EXP bit fields.
 - Non-IPv4 packets that enter the LSPs on interface **ge-3/1/0** are initialized with values from rewrite table **rule1**, and written into the MPLS EXP header field only. The statement **exp rule1** has the same result as **exp rule1 protocol mpls**.
- For interface **ge-3/1/0**, IPv4 packets transmitted over a non-LSP layer are initialized with values from IP precedence rewrite table **rule2**.
- For interface **ge-3/1/1**, IPv4 packets that enter the LSPs are initialized with values from EXP rewrite table **exp-inet-table**. An identical 3-bit value is written into the IP precedence and MPLS EXP bit fields.
- For interface **ge-3/1/1**, MPLS packets other than IPv4 Layer 3 types are also initialized with values from table **exp-inet-table**. For VPN MPLS packets with IPv4 payloads, the CoS value is written to MPLS and IPv4 headers. For VPN MPLS packets without IPv4 payloads, the CoS value is written to MPLS headers only.

[edit class-of-service]

```

rewrite-rules {
  exp exp-inet-table {
    forwarding-class best-effort {
      loss-priority low code-point 000;
      loss-priority high code-point 001;
    }
    forwarding-class assured-forwarding {
      loss-priority low code-point 010;
      loss-priority high code-point 011;
    }
    forwarding-class expedited-forwarding {
      loss-priority low code-point 111;
      loss-priority high code-point 110;
    }
    forwarding-class network-control {
      loss-priority low code-point 100;
      loss-priority high code-point 101;
    }
  }
  exp rule1 {
    ...
  }
  inet-precedence rule2 {
    ...
  }
}
exp rule_non_vpn {
  ...
}

interfaces {
  ge-3/1/0 {
    unit 0 {
      rewrite-rules {
        exp rule1;
        inet-precedence rule2;
        exp exp-inet-table protocol mpls-inet-both; # For all VPN traffic.
        exp rule_non_vpn protocol mpls-inet-both-non-vpn; # For all non-VPN
          # traffic.
      }
    }
  }
  ge-3/1/1 {
    unit 0 {
      rewrite-rules {
        exp exp-inet-table protocol [mpls mpls-inet-both];
      }
    }
  }
}

```

Example: Simultaneous DSCP and EXP Rewrite

On M120 routers, M320 routers with Enhanced-III FPCs, MX Series routers, and EX Series switches, configure the simultaneous DSCP and EXP rewrite rules as shown below:

1. Configure CoS.

```
[edit]
user@host# edit class-of-service
```

2. Configure the EXP rewrite rule on the interface.

```
[edit class-of-service]
user@host# set interfaces ge-2/0/3 unit 0 rewrite-rule exp rule1
```

3. Configure the IPv4 rewrite rule on the interface.

```
[edit class-of-service]
user@host# set interfaces ge-2/0/3 unit 0 rewrite-rule inet-precedence rule2
```

4. Configure the IPv4 rewrite rule on the interface and apply it to packets entering the MPLS tunnel.

```
[edit class-of-service]
user@host# set interfaces ge-2/0/3 unit 0 rewrite-rule inet-precedence rule3 protocol
mpls
```

5. Verify the configuration by using the **show interfaces** command.

```
[edit class-of-service]
user@host# show interfaces ge-2/0/3 unit 0
rewrite-rules {
  exp rule1;
  inet-precedence rule2;
  inet-precedence rule3 protocol mpls;
}
```

In the example above, there are two different IPv4 precedence rewrite rules: **rule2** and **rule3**. **rule2** affects the IPv4 to IPv4 traffic and **rule3** affects the IPv4 to MPLS traffic.

Rewriting the EXP Bits of All Three Labels of an Outgoing Packet

In interprovider, carrier-of-carrier, and complex traffic engineering scenarios, it is sometimes necessary to push three labels on the next hop, using a swap-push-push or triple-push operation.

By default, on M Series routers and EX Series switches, the top MPLS EXP label of an outgoing packet is not rewritten when you configure swap-push-push and triple-push operations. On these routing devices, you can rewrite the EXP bits of all three labels of an outgoing packet, thereby maintaining the CoS of an incoming MPLS or non-MPLS packet.

When the software performs a swap-push-push operation and no rewriting is configured, the EXP fields of all three labels are the same as in the old label. If there is EXP rewriting configured, the EXP bits of the bottom two labels are overwritten with the table entry. The EXP setting of the top label is retained even with rewriting.

To push three labels on all incoming MPLS packets, include the **exp-swap-push-push default** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
exp-swap-push-push default;
```


When the software performs a push-push-push operation and if no rewriting is configured, the EXP fields of the bottom two labels are zero. If EXP rewriting is configured, the EXP fields of the bottom two labels are rewritten with the table entry's rewrite value. The EXP field of the top label is inserted with the Qn+PLP value. This Qn reflects the final classification by a multifield classifier if one exists, regardless of whether rewriting is configured.



NOTE: The `exp-push-push-push` and `exp-swap-push-push` configuration on the egress interface does not rewrite the top label's EXP field with the Qn+PLP value on an IQ or IQ2 PIC.

To push three labels on incoming non-MPLS packets, include the `exp-push-push-push default` statement at the `[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]` hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
  exp-push-push-push default;
```

These configurations apply the default MPLS EXP rewrite table, as described in *Rewriting MPLS and IPv4 Packet Headers*. You can configure these operations and override the default MPLS EXP rewrite table with a custom table. For more information about writing and applying a custom rewrite table, see [“Configuring Rewrite Rules” on page 1694](#) and *Applying Rewrite Rules to Output Logical Interfaces*.



NOTE: With a three-label stack, if you do not include the `exp-swap-push-push default` or `exp-push-push-push default` statement in the configuration, the top label's EXP bits are set to zero.

Example: Rewriting the EXP Bits of All Three Labels of an Outgoing Packet

Configure a swap-push-push operation, and override the default rewrite table with a custom table:

```
[edit class-of-service]
forwarding-classes {
  queue 0 be;
  queue 1 ef;
  queue 2 af;
  queue 3 nc;
}
interfaces {
  ge-1/1/3 {
    unit 0 {
      rewrite-rules {
        exp exp_rewrite; # Apply custom rewrite table
        exp-swap-push-push default;
      }
    }
  }
}
```

```
rewrite-rules {  
  exp exp_rew {  
    forwarding-class be {  
      loss-priority low code-point 000;  
      loss-priority high code-point 100;  
    }  
    forwarding-class ef {  
      loss-priority low code-point 001;  
      loss-priority high code-point 101;  
    }  
    forwarding-class af {  
      loss-priority low code-point 010;  
      loss-priority high code-point 110;  
    }  
    forwarding-class nc {  
      loss-priority low code-point 011;  
      loss-priority high code-point 111;  
    }  
  }  
}
```

Rewriting IEEE 802.1p Packet Headers with an MPLS EXP Value

For Ethernet interfaces on Juniper Networks M320 Multiservice Edge Routers, MX Series Ethernet Service Routers, T Series Core Routers, and EX Series switches that have a peer connection to an M Series Multiservice Edge Router, MX Series, T Series router, or EX Series switches, you can rewrite both MPLS EXP and IEEE 802.1p bits to a configured value. This enables you to pass the configured value to the Layer 2 VLAN path. For IQ PICs, you can only configure one IEEE 802.1 rewrite rule on a physical port. All logical ports (units) on that physical port should apply the same IEEE 802.1 rewrite rule.

To rewrite both the MPLS EXP and IEEE 802.1p bits, you must include EXP and IEEE 802.1p rewrite rules in the interface configuration. To configure EXP and IEEE 802.1p rewrite rules, include the **rewrite-rules** statement at the **[edit class-of-service interfaces interface-name unit logical-unit-number]** hierarchy level, specifying the **exp** and **ieee-802.1** options:

```
[edit class-of-service interfaces interface-name unit logical-unit-number]  
rewrite-rules {  
  exp rewrite-rule-name;  
  ieee-802.1 default;  
}
```

When you combine these two rewrite rules, only the EXP rewrite table is used for rewriting packet headers. If you do not configure a VLAN on the interface, only the EXP rewriting is in effect. If you do not configure an LSP on the interface or if the MPLS EXP rewrite rule mapping is removed, the IEEE 802.1p default rewrite rules mapping takes effect.



NOTE: You can also combine other rewrite rules. IP, DSCP, DSCP IPv6, and MPLS EXP are associated with Layer 3 packet headers, and IEEE 802.1p is associated with Layer 2 packet headers.

For IQ PICs, you can only configure one IEEE 802.1 rewrite rule on a physical port. All logical ports (units) on that physical port should apply the same IEEE 802.1 rewrite rule.

If you combine IEEE 802.1p with IP rewrite rules, the Layer 3 packets and Layer 2 headers are rewritten with the IP rewrite rule.

If you combine IEEE 802.1p with DSCP or DSCP IPv6 rewrite rules, three bits of the Layer 2 header and six bits of the Layer 3 packet header are rewritten with the DSCP or DSCP IPv6 rewrite rule.

The following example shows how to configure an EXP rewrite rule and apply it to both MPLS EXP and IEEE 802.1p bits:

```
[edit class-of-service]
rewrite-rules {
  exp exp-ieee-table {
    forwarding-class best-effort {
      loss-priority low code-point 000;
      loss-priority high code-point 001;
    }
    forwarding-class assured-forwarding {
      loss-priority low code-point 010;
      loss-priority high code-point 011;
    }
    forwarding-class expedited-forwarding {
      loss-priority low code-point 111;
      loss-priority high code-point 110;
    }
    forwarding-class network-control {
      loss-priority low code-point 100;
      loss-priority high code-point 101;
    }
  }
}
interfaces {
  ge-3/1/0 {
    unit 0 {
      rewrite-rules {
        exp exp-ieee-table;
        ieee-802.1 default;
      }
    }
  }
}
```

Configuring the IEEE 802.1p Field for CoS Host Outbound Traffic

This topic provides a summary of the configuration for setting the IEEE 802.1p field in the Ethernet frame header for host outbound traffic (control plane traffic). You can set a global value for the priority code point that applies to all host outbound traffic. Additionally, or alternatively, you can specify that rewrite rules are applied to all host outbound traffic on egress logical interfaces. These are rules that have been previously configured to set the IEEE 802.1p field for data traffic on those interfaces.

Configuration of 802.1p bits is supported only on the following hardware and software components:

- EX Series switches
- MX Series 3D Universal Edge Routers
- Enhanced Queuing DPCs
- MPCs
- Junos OS Release 12.3 or later

To configure the IEEE 802.1p field settings:

1. (Optional) Specify a global default value for the IEEE 802.1p field for all host outbound traffic.

See [“Configuring a Global Default IEEE 802.1p Value for All Host Outbound Traffic” on page 1710](#).

2. (Optional) Specify that the IEEE 802.1p rewrite rules for the egress logical interfaces are applied to all host outbound traffic on those interfaces.

See [“Applying Egress Interface Rewrite Rules to the IEEE 802.1p Field for All Host Outbound Traffic on the Interface” on page 1711](#).

Related Documentation

- [Rewriting Packet Header Information Overview on page 1690](#)

Configuring a Global Default IEEE 802.1p Value for All Host Outbound Traffic

This topic describes how to configure a global default value for the IEEE 802.1p field for all host outbound traffic on MX Series routers and EX Series switches.

To configure a global default value for the IEEE 802.1p field:

- Specify the value.

```
[edit class-of-service host-outbound-traffic ieee-802.1]
user@host# set default value
```

For example, specify that a value of 010 is applied to all host outbound traffic:

```
[edit class-of-service host-outbound-traffic ieee-802.1]
user@host# set default 010
```

Related Documentation

- [Configuring the IEEE 802.1p Field for CoS Host Outbound Traffic on page 1710](#)
- [Rewriting Packet Header Information Overview on page 1690](#)

Applying Egress Interface Rewrite Rules to the IEEE 802.1p Field for All Host Outbound Traffic on the Interface

This topic describes how to apply rewrite rules for egress logical interfaces to the IEEE 802.1p field for all host outbound traffic on those interfaces on MX Series routers and EX Series switches.

This task requires separately configured rewrite rules that map packet loss priority information to the code point value in the 802.1p field for data traffic on egress logical interfaces. See *Rewriting Packet Header Information Overview* in the *Junos OS Class of Service Configuration Guide*.

To configure the rewrite rules:

1. Configure the CoS rewrite rules to map the forwarding class to the desired value for the 802.1p field.

See [“Configuring Rewrite Rules” on page 1694](#).

2. Associate the rewrite rules to the desired egress logical interfaces.

See *Applying Rewrite Rules to Output Logical Interfaces*.

3. (Optional) Configure the forwarding class for host outbound traffic. Do not configure this forwarding class if you want to use the default forwarding class assignment (input classification).

See [“Overriding the Input Classification” on page 1560](#).

To configure the rewrite rules to apply to the host outbound traffic IEEE 802.1p field:

- Configure the rewrite rules.

```
[edit class-of-service host-outbound-traffic ieee-802.1]
user@host# set rewrite-rules
```

```
[edit class-of-service]
rewrite-rules {
  ieee-802.1 rewrite_foo {
    forwarding-class network-control {
      loss-priority low code-point 101;
    }
  }
}
interfaces {
  ge-1/0/0 {
    unit 100 {
      rewrite-rules {
        ieee-802.1 rewrite_foo vlan-tag outer-and-inner;
      }
    }
  }
}
```

```

}
host-outbound-traffic {
    forwarding-class network-control;
}
host-outbound-traffic {
    ieee-802.1 {
        rewrite-rules;
    }
}

```

**Related
Documentation**

- [Configuring the IEEE 802.1p Field for CoS Host Outbound Traffic on page 1710](#)
- [Rewriting Packet Header Information Overview on page 1690](#)

Setting Ingress DSCP Bits for Multicast Traffic over Layer 3 VPNs

By default, the DSCP bits on outer IP headers arriving at an ingress PE router using generic routing encapsulation (GRE) are not set for multicast traffic sent over an Layer 3 virtual private network (VPN) provider network. However, you can configure a type-of-service (ToS) rewrite rule so the router sets the DSCP bits of GRE packets to be consistent with the service provider's overall core network CoS policy. The bits are set at the core-facing interface of the ingress provider edge (PE) router. For more information about rewriting IP header bits, see [“Rewriting Packet Header Information Overview” on page 1690](#).

This section describes this configuration from a CoS perspective. The examples are not complete multicast or VPN configurations. For more information about multicast, see the *Multicast Protocols Configuration Guide*. For more information about Layer 3 VPNs, see the *Junos OS VPNs Configuration Guide*.

To configure the rewrite rules on the core-facing interface of the ingress PE, include the **rewrite-rules** statement at the **[edit class-of-service]** hierarchy level. You apply the rule to the proper ingress interface at the **[edit class-of-service interfaces]** hierarchy level to complete the configuration. This ingress DSCP rewrite is independent of classifiers placed on ingress traffic arriving on the customer-facing interface of the PE router.

The rewrite rules are applied to all unicast packets and multicast groups. You cannot configure different rewrite rules for different multicast groups. The use of DSCPv6 bits is not supported because IPv6 multicast is not supported. You can configure another rewrite rule for the EXP bits on MPLS CE-CE unicast traffic.

This example defines a rewrite rule called **dscp-rule** that establishes a value of **000000** for best-effort traffic. The rule is applied to the outgoing, core-facing PE interface **ge-2/3/0**.

```

[edit class-of-service]
rewrite-rules {
    dscp dscp-rule {
        forwarding-class best-effort {
            loss-priority low code-point 000000;
        }
    }
}

```

```

[edit class-of-service interfaces]

```

```

ge-2/3/0 {
  unit 0 {
    rewrite-rules {
      dscp dscp-rule;
    }
  }
}

```

Example

- [Example: Per-Node Rewriting of EXP Bits on page 1713](#)

Example: Per-Node Rewriting of EXP Bits

To configure a custom table to rewrite the EXP bits, also known as CoS bits, on a particular node, the classifier table and the rewrite table must specify exactly the same CoS values.

In addition, the least significant bit of the CoS value itself must represent the PLP value. For example, CoS value **000** must be associated with PLP **low**, **001** must be associated with PLP **high**, and so forth.

This example configures a custom table to rewrite the EXP bits on a particular node:

```

[edit class-of-service]
classifiers {
  exp exp-class {
    forwarding-class be {
      loss-priority low code-points 000;
      loss-priority high code-points 001;
    }
    forwarding-class af {
      loss-priority low code-points 010;
      loss-priority high code-points 011;
    }
    forwarding-class ef {
      loss-priority low code-points 100;
      loss-priority high code-points 101;
    }
    forwarding-class nc {
      loss-priority low code-points 110;
      loss-priority high code-points 111;
    }
  }
}
rewrite-rules {
  exp exp-rw {
    forwarding-class be {
      loss-priority low code-point 000;
      loss-priority high code-point 001;
    }
    forwarding-class af {
      loss-priority low code-point 010;
      loss-priority high code-point 011;
    }
    forwarding-class ef {
      loss-priority low code-point 100;

```

```
        loss-priority high code-point 101;
    }
    forwarding-class nc {
        loss-priority low code-point 110;
        loss-priority high code-point 111;
    }
}
}
```

Configuration Statements

- [\[edit class-of-service\] Hierarchy Level on page 1714](#)

[edit class-of-service] Hierarchy Level

```
class-of-service {
  classifiers {
    type classifier-name {
      forwarding-class class-name {
        loss-priority (high | low | medium-high | medium-low) code-points [ aliases bits ];
      }
      import (classifier-name | default);
    }
  }
  code-point-aliases {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) {
      alias-name bits;
    }
  }
  drop-profiles {
    profile-name {
      fill-level percentage drop-probability percentage;
      interpolate {
        drop-probability value;
        fill-level value;
      }
    }
  }
  fabric {
    scheduler-map {
      priority (high | low) scheduler scheduler-name;
    }
  }
  forwarding-class-map {
    map-name {
      class class-name queue-num queue-number <restricted-queue queue-number>;
    }
  }
  forwarding-classes {
    class class-name policing-priority (normal | premium) queue-num queue-number
      priority (high | low);
    queue queue-number class-name policing-priority (normal | premium) priority (high |
      low);
  }
  forwarding-policy {
    class class-name {
      classification-override {

```



```

        forwarding-class class-name;
    }
}
next-hop-map map-name {
    forwarding-class class-name {
        discard;
        lsp-next-hop [ lsp-regular-expressions ];
        next-hop [ next-hop-names ];
        non-lsp-next-hop;
    }
}
fragmentation-maps {
    map-name {
        forwarding-class class-name {
            drop-timeout milliseconds;
            fragment-threshold bytes;
            multilink-class number;
            no-fragmentation;
        }
    }
}
host-outbound-traffic {
    dscp-code-point value;
    forwarding-class class-name;
    ieee-802.1 {
        default value;
        rewrite-rules;
    }
    tcp {
        raise-internet-control-priority;
    }
}
interfaces {
    ... the interfaces subhierarchy appears after the main [edit class-of-service] hierarchy
    ...
}
}
restricted-queues {
    forwarding-class class-name queue-number;
}
rewrite-rules {
    (dscp | dscp-ipv6 | exp | frame-relay-de | ieee-802.1 | ieee-802.1ad | inet-precedence)
    rewrite-rule {
        forwarding-class class-name {
            loss-priority level code-point (alias | bits);
        }
        import (rewrite-rule | default);
    }
}
routing-instances routing-instance-name {
    classifiers {
        dscp (classifier-name | default);
        dscp-ipv6 (classifier-name | default);
        exp (classifier-name | default);
        ieee-208.1 (classifier-name | default | encapsulated | vlan-tag (inner | outer));
    }
}

```

```

    }
  }
  scheduler-maps {
    map-name {
      forwarding-class class-name scheduler scheduler-name;
    }
  }
  schedulers {
    scheduler-name {
      adjust-minimum value;
      adjust-percent value;
      buffer-size (exact | percent percentage | remainder);
      drop-profile-map loss-priority (any | high | low | medium-high | medium-low)
        protocol any;
      excess-priority (high | low | medium-high | medium-low);
      excess-rate (percent percentage | proportion proportion);
      priority (high | low | medium-high | medium-low | strict-high);
      shaping-rate (bps | percent percentage | burst-size size);
      transmit-rate (bps | percent percentage | remainder) <exact | rate-limit>;
    }
  }
  traceoptions {
    file <files number> <match regular-expression> <size maximum-file-size>
      <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
  traffic-control-profiles {
    profile-name {
      adjust-minimum rate;
      delay-buffer-rate (bps | cps cps | percent percentage);
      excess-rate (percent percentage | proportion value);
      guaranteed-rate (bps | percent percentage) <burst-size bytes>;
      overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
      scheduler-map map-name;
      shaping-rate (bps | percent percentage) <burst-size bytes>;
    }
  }
  tri-color;
}

class-of-service {
  interfaces {
    interface-name {
      excess-bandwidth-share (equal | proportional value);
      input-excess-bandwidth-share (equal | proportional value);
      input-scheduler-map map-name;
      input-shaping-rate bps;
      input-traffic-control-profile profile-name;
      output-forwarding-class-map map-name;
      output-traffic-control-profile profile-name;
      scheduler-map map-name;
      scheduler-map-chassis (map-name | derived);
      shaping-rate bps;
      unit (logical-unit-number | *) {
        classifiers {

```

```

    dscp (classifier-name | default) {
        family [ inet mpls ];
    }
    dscp-ipv6 (classifier-name | default) {
        family [ inet mpls ];
    }
    exp (classifier-name | default);
    ieee-208.1 (classifier-name | default) <vlan-tag (inner | outer)>;
    ieee-208.1ad (classifier-name | default);
    inet-precedence (classifier-name | default);
}
forwarding-class class-name;
input-scheduler-map map-name;
input-shaping-rate bps;
input-traffic-control-profile profile-name shared-instance instance-name;
loss-priority-maps {
    (map-name | default);
}
loss-priority-rewrites {
    (map-name | default);
}
output-forwarding-class-map map-name;
output-traffic-control-profile profile-name shared-instance instance-name;
rewrite-rules {
    dscp (rule-name | default) <protocol mpls>;
    dscp-ipv6 (rule-name | default);
    exp (rule-name | default) <protocol [ mpls-any | mpls-inet-both |
        mpls-inet-both-non-vpn ]>;
    exp-push-push-push default;
    exp-swap-push-push default;
    ieee-802.1 (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
    ieee-802.1ad (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
    inet-precedence (rewrite-name | default) <protocol mpls>;
}
scheduler-map map-name;
shaping-rate bps;
translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 |
    to-exp-from-exp | to-inet-precedence-from-inet-precedence) table-name;
}
}
interface-set interface-set-name {
    excess-bandwidth-share (equal | proportional value);
    input-excess-bandwidth-share (equal | proportional value);
    input-traffic-control-profile profile-name;
    input-traffic-control-profile-remaining profile-name;
    internal-node;
    output-traffic-control-profile profile-name;
    output-traffic-control-profile-remaining profile-name;
}
}
}

```

Related Documentation • *Notational Conventions Used in Junos OS Configuration Hierarchies*

code-point

Syntax	<code>code-point [<i>aliases</i>] [<i>bit-patterns</i>];</code>
Hierarchy Level	[edit class-of-service rewrite-rules <i>type</i> <i>rewrite-name</i> forwarding-class <i>class-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify one or more code-point aliases or bit sets for association with a forwarding class.
Options	<i>aliases</i> —Name of each alias. <i>bit-patterns</i> —Value of the code-point bits, in decimal form.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694

default (CoS Host Outbound Traffic)

Syntax	<code>default <i>value</i>;</code>
Hierarchy Level	[edit class-of-service host-outbound-traffic ieee-802.1p]
Release Information	Statement introduced in Junos OS Release 12.3.
Description	Apply a global default value to the IEEE 802.1p—priority code point (PCP)—field in the Ethernet frame header for all host outbound traffic.
Options	<i>value</i> —Three-bit binary number. Range: 000 through 111
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Global Default IEEE 802.1p Value for All Host Outbound Traffic on page 1710• Rewriting Packet Header Information Overview on page 1690

dscp (Rewrite Rules)

Syntax	<code>dscp (rewrite-name default);</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For IPv4 traffic, apply a Differentiated Services (DiffServ) code point (DSCP) rewrite rule.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules dscp] hierarchy level.</p> <p>default—The default mapping.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Rewrite Rules on page 1694 • dscp-ipv6 (Class-of-Service) on page 1357 • exp on page 1358 • exp-push-push-push on page 1723 • exp-swap-push-push on page 1724 • ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360 • ieee-802.1ad on page 1362 • inet-precedence on page 1364 • rewrite-rules (Definition) on page 1481

dscp (Rewrite Rules on Physical Interface)

Syntax	dscp (<i>rewrite-name</i> default);
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> rewrite-rules
Release Information	Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.
Description	Associate a rewrite-rules configuration or default mapping with a specific interface.
Options	rewrite-name —Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules] hierarchy level. default —The default mapping.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

dscp-ipv6 (Class-of-Service)

Syntax	<code>dscp-ipv6 (<i>rewrite-name</i> <default>) { protocol mpls }</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For IPv6 traffic, apply a DSCP rewrite rule.
Options	<p><i>rewrite-name</i>—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules dscp-ipv6] hierarchy level.</p> <p><i>default</i>— Default mapping.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Rewrite Rules on page 1694 • dscp (Rewrite Rules) on page 1382 • exp on page 1358 • exp-push-push-push on page 1723 • exp-swap-push-push on page 1724 • ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360 • ieee-802.1ad on page 1362 • inet-precedence on page 1364 • rewrite-rules (Definition) on page 1481

exp

Syntax	<code>exp (rewrite-name default) protocol protocol-types;</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS Release 12.2. for ACX series
Description	Apply an MPLS experimental (EXP) rewrite rule.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules exp] hierarchy level.</p> <p>default—The default mapping.</p> <p>By default, IP precedence rewrite rules alter the first three bits on the type-of-service (ToS) byte while leaving the last three bits unchanged. This default behavior applies to rewrite rules you configure for MPLS packets with IPv4 payloads. You configure these types of rewrite rules by including the mpls-inet-both or mpls-inet-both-non-vpn option at the [edit class-of-service interfaces <i>interface</i> <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules exp <i>rewrite-rule-name</i> protocol] hierarchy level. The IP precedence rewrite rules explanation does not apply to ACX Series Universal Access routers.</p> <p>On interfaces configured on Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) on MX Series Ethernet Services Routers and EX Series switches, we highly recommend that you configure the default option when you configure a behavior aggregate (BA) classifier that does not include a specific rewrite rule for MPLS packets. Doing so ensures that MPLS exp value is rewritten according to the BA classifier rules configured for forwarding or packet loss priority. This does not apply to ACX Series Universal Access routers.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Rewriting the EXP Bits of All Three Labels of an Outgoing Packet</i>• dscp (Rewrite Rules) on page 1382• dscp-ipv6 (Class-of-Service) on page 1357• exp-push-push-push on page 1723• exp-swap-push-push on page 1724• ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360• ieee-802.1ad on page 1362• inet-precedence on page 1364

- [rewrite-rules \(Definition\) on page 1481](#)

exp-push-push-push

Syntax	exp-push-push-push default;
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For M Series routers and EX Series switches, rewrite the EXP bits of all three labels of an outgoing packet, thereby maintaining CoS of an incoming non-MPLS packet.
Options	default —Apply the default MPLS EXP rewrite table.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Rewriting the EXP Bits of All Three Labels of an Outgoing Packet</i> • dscp (Rewrite Rules) on page 1382 • dscp-ipv6 (Class-of-Service) on page 1357 • exp on page 1358 • exp-swap-push-push on page 1724 • ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360 • ieee-802.1ad on page 1362 • inet-precedence on page 1364 • rewrite-rules (Definition) on page 1481

exp-swap-push-push

Syntax	exp-swap-push-push default;
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For M Series routers and EX Series switches, rewrite the EXP bits of all three labels of an outgoing packet, thereby maintaining CoS of an incoming MPLS packet.
Options	default —Apply the default MPLS EXP rewrite table.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Rewriting the EXP Bits of All Three Labels of an Outgoing Packet</i>• dscp (Rewrite Rules) on page 1382• dscp-ipv6 (Class-of-Service) on page 1357• exp on page 1358• exp-push-push-push on page 1723• ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360• ieee-802.1ad on page 1362• inet-precedence on page 1364• rewrite-rules (Definition) on page 1481

forwarding-class (BA Classifiers)

Syntax	<code>forwarding-class <i>class-name</i> { <code>loss-priority level</code> <code>code-points</code> [<i>aliases</i>] [<i>bit-patterns</i>]; }</code>
Hierarchy Level	[edit class-of-service classifiers <i>type classifier-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define forwarding class name and option values.
Options	<p><i>class-name</i>—Name of the forwarding class.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Defining Classifiers on page 1330 • <i>Example: Configuring CoS for a PBB Network</i>

frame-relay-de (Defining Loss Priority Maps)

Syntax	<pre>frame-relay-de <i>name</i> { loss-priority <i>level</i> <i>code-points</i> [<i>alias</i> <i>bits</i>]; }</pre>
Hierarchy Level	[edit class-of-service loss-priority-maps]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Define a Frame Relay discard eligibility (DE) bit loss priority map.
Options	<p><i>name</i>—Name of the loss priority map.</p> <p><i>loss-priority level</i>—Level of the loss priority to be applied based on the specified CoS values. The loss priority level can be one of the following:</p> <ul style="list-style-type: none">• high—Packet has high loss priority.• low—Packet has low loss priority.• medium-high—Packet has medium-high loss priority.• medium-low—Packet has medium-low loss priority. <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Defining a Custom Frame Relay Loss Priority Map</i>

host-outbound-traffic

Syntax	<pre> host-outbound-traffic { forwarding-class class-name; dscp-code-point value; ieee-802.1 { default value; rewrite-rules; } } </pre>
Hierarchy Level	[edit class-of-service]
Release Information	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Statement introduced before Junos OS 11.4 for EX Series switches.</p>
Description	<p>Allow queue selection for all traffic generated by the Routing Engine (host). The selected queue must be configured properly. The configuration of specific DSCP code point bits for the ToS field of the generated packets is also allowed. Transit packets are not affected; only packets originating on the Routing Engine are affected. By default, the forwarding class (queue) and DSCP bits are set according to those given in <i>Default Routing Engine Protocol Queue Assignments</i>. This feature is not available on J Series routers.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • See Changing the Routing Engine Outbound Traffic Defaults on page 1747. • Configuring a Global Default IEEE 802.1p Value for All Host Outbound Traffic on page 1710 • Applying Egress Interface Rewrite Rules to the IEEE 802.1p Field for All Host Outbound Traffic on the Interface on page 1711 • <i>Understanding Junos OS CoS Components for EX Series Switches</i>

ieee-802.1 (Rewrite Rules on Logical Interface)

Syntax	ieee-802.1 (<i>rewrite-name</i> default) vlan-tag (outer outer-and-inner);
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4. vlan-tag statement introduced in Junos OS Release 8.1.
Description	Apply an IEEE-802.1 rewrite rule. For IQ PICs, you can only configure one IEEE 802.1 rewrite rule on a physical port. All logical ports (units) on that physical port should apply the same IEEE 802.1 rewrite rule.
Options	rewrite-name —Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules ieee-802.1] hierarchy level. default —The default mapping.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694• <i>Example: Configuring CoS for a PBB Network</i>• dscp (Rewrite Rules) on page 1382• dscp-ipv6 (Class-of-Service) on page 1357• exp on page 1358• exp-push-push-push on page 1723• exp-swap-push-push on page 1724• ieee-802.1ad on page 1362• inet-precedence on page 1364• rewrite-rules (Definition) on page 1481

ieee-802.1 (Host Outbound Traffic)

Syntax	<pre>ieee-802.1 { default <i>value</i>; rewrite-rules; }</pre>
Hierarchy Level	[edit class-of-service host-outbound-traffic]
Release Information	Statement introduced in Junos OS Release 12.3.
Description	<p>Apply the IEEE 802.1p rewrite rules associated with the egress logical interface to the IEEE 802.1p PCP field for all host outbound traffic on that interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Global Default IEEE 802.1p Value for All Host Outbound Traffic on page 1710 • Applying Egress Interface Rewrite Rules to the IEEE 802.1p Field for All Host Outbound Traffic on the Interface on page 1711 • Rewriting Packet Header Information Overview on page 1690 • Configuring Rewrite Rules on page 1694

ieee-802.1 (Rewrite Rules on Physical Interface)

Syntax	<pre>ieee-802.1 (<i>rewrite-name</i> default) ;</pre>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i>] rewrite-rules
Release Information	Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.
Description	Apply an IEEE-802.1 rewrite rule.
Options	<p><i>rewrite-name</i>—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules ieee-802.1] hierarchy level.</p> <p>default—The default mapping.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

ieee-802.1ad

Syntax	ieee-802.1ad (<i>rewrite-name</i> default) vlan-tag (outer outer-and-inner);
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Apply a IEEE-802.1ad rewrite rule.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules ieee-802.1ad] hierarchy level.</p> <p>default—The default rewrite bit mapping.</p> <p>vlan-tag—The rewrite rule is applied to the outer or outer-and-inner VLAN tag.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694• <i>Example: Configuring CoS for a PBB Network</i>• dscp (Rewrite Rules) on page 1382• dscp-ipv6 (Class-of-Service) on page 1357• exp on page 1358• exp-push-push-push on page 1723• exp-swap-push-push on page 1724• ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360• inet-precedence on page 1364• rewrite-rules (Definition) on page 1481

import (Rewrite Rules)

Syntax	<code>import (rewrite-name default);</code>
Hierarchy Level	<code>[edit class-of-service rewrite-rules type rewrite-name]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify a default or previously defined rewrite-rules mapping to import.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the <code>[edit class-of-service rewrite-rules]</code> hierarchy level.</p> <p>default—The default rewrite-rules mapping.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694

inet-precedence

Syntax	<code>inet-precedence (rewrite-name default);</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Apply a IPv4 precedence rewrite rule.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules inet-precedence] hierarchy level.</p> <p>default—The default mapping. By default, IP precedence rewrite rules alter the first three bits on the type of service (ToS) byte while leaving the last three bits unchanged.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694• dscp (Rewrite Rules) on page 1382• dscp-ipv6 (Class-of-Service) on page 1357• exp on page 1358• exp-push-push-push on page 1723• exp-swap-push-push on page 1724• ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360• ieee-802.1ad on page 1362• rewrite-rules (Definition) on page 1481

inet-precedence (Rewrite Rules on Physical Interface)

Syntax	<code>inet-precedence (<i>rewrite-name</i> default);</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> rewrite-rules]
Release Information	Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.
Description	Apply a IPv4 precedence rewrite rule.
Options	<p><i>rewrite-name</i>—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules inet-precedence] hierarchy level.</p> <p>default—The default mapping. By default, IP precedence rewrite rules alter the first three bits on the type of service (ToS) byte while leaving the last three bits unchanged.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

interfaces

```
Syntax  interfaces {
    interface-name {
        classifiers{
            dscp(classifier-name | default) {
            }
            ieee-802.1 (classifier-name | default) vlan-tag (inner | outer | classifier-name);
            inet-precedence (rewrite-name | default);
        }
        input-scheduler-map map-name;
        input-shaping-rate rate;
        irb {
            unit logical-unit-number {
                classifiers {
                    type (classifier-name | default);
                }
                rewrite-rules {
                    dscp (rewrite-name | default);
                    dscp-ipv6 (rewrite-name | default);
                    exp (rewrite-name | default) protocol protocol-types;
                    ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                    inet-precedence (rewrite-name | default);
                }
            }
        }
        member-link-scheduler (replicate | scale);
        rewrite-rules {
            dscp (rewrite-name | default);
            ieee-802.1 (rewrite-name | default) vlan-tag (outer);
            inet-precedence (rewrite-name | default);
        }
        scheduler-map map-name;
        scheduler-map-chassis map-name;
        shaping-rate rate;
        unit logical-unit-number {
            classifiers {
                type (classifier-name | default) family (mpls | inet);
            }
            forwarding-class class-name;
            fragmentation-map map-name;
            input-shaping-rate (percent percentage | rate);
            input-traffic-control-profile profile-name shared-instance instance-name;
            output-traffic-control-profile profile-name shared-instance instance-name;
            per-session-scheduler;
            rewrite-rules {
                dscp (rewrite-name | default);
                dscp-ipv6 (rewrite-name | default);
                exp (rewrite-name | default) protocol protocol-types;
                exp-push-push-push default;
                exp-swap-push-push default;
                ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                inet-precedence (rewrite-name | default);
            }
        }
    }
}
```

```

    }
    scheduler-map map-name;
    shaping-rate rate;
    translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 | to-exp-from-exp
    | to-inet-precedence-from-inet-precedence) table-name;
  }
}
interface-set interface-set-name {
  excess-bandwidth-share;
  internal-node;
  output-traffic-control-profile profile-name;
  output-traffic-control-profile-remaining profile-name;
}
}

```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 7.4.
Interface-set level added in Junos OS Release 8.5.

Description Configure interface-specific CoS properties for incoming packets.



NOTE: The dscp-ipv6 and ieee-802.1ad classifier types are not supported on ACX Series routers. For further information about support on ACX Series routers, see *Understanding CoS CLI Configuration Statements on ACX Series Universal Access Routers*.

Options The remaining statements are explained separately.

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

Related Documentation

- [Overview of BA Classifier Types on page 1321](#)
- [Configuring Rewrite Rules on page 1694](#)
- [Understanding CoS CLI Configuration Statements on ACX Series Universal Access Routers](#)

loss-priority (BA Classifiers)

Syntax	loss-priority <i>level</i> ;
Hierarchy Level	[edit class-of-service classifiers <i>type classifier-name</i> forwarding-class <i>class-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify packet loss priority value for a specific set of code-point aliases and bit patterns.
Options	<i>level</i> can be one of the following: <ul style="list-style-type: none">• high—Packet has high loss priority.• medium-high—Packet has medium-high loss priority.• medium-low—Packet has medium-low loss priority.• low—Packet has low loss priority.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Overview of BA Classifier Types on page 1321• Example: Configuring CoS for a PBB Network• Configuring Tricolor Marking on page 1447

loss-priority-maps

Syntax	<pre> loss-priority-maps { frame-relay-de <i>rewrite-name</i> { loss-priority <i>level</i> { <i>code-points</i> [<i>aliases</i>] [<i>bit-patterns</i>]; } } } </pre>
Hierarchy Level	[edit class-of-service]
Release Information	Statement introduced in JUNOS Release 11.4.
Description	Map the loss priority of incoming packets based on the CoS values.
Options	<p>frame-relay-de <i>rewrite-name</i>—Name of the Frame Relay DE bit loss priority map.</p> <p>loss-priority <i>level</i>—The loss priority level can be one of the following:</p> <ul style="list-style-type: none"> • high—Packet has high loss priority. • low—Packet has low loss priority. • medium-high—Packet has medium-high loss priority. • medium-low—Packet has medium-low loss priority. <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Assigning the Default Frame Relay DE Loss Priority Map to an Interface</i>

loss-priority-maps (Assigning to an Interface)

Syntax	<pre>loss-priority-maps { frame-relay-de (<i>loss-priority-rewrite-name</i> default); }</pre>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 11.4.
Description	Assign the loss priority map to a logical interface.
Options	<p>default—Apply the default loss priority map. The default map includes the following configuration:</p> <pre>loss-priority low code-point 0; loss-priority high code-point 1;</pre> <p>map-name—Name of loss priority map to be applied.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Assigning the Default Frame Relay DE Loss Priority Map to an Interface</i>• unit on page 1371

protocol (Rewrite Rules)

Syntax	<code>protocol <i>protocol-types</i>;</code>
Hierarchy Level	<p>[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules exp <i>rewrite-name</i>],</p> <p>[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules dscp <i>rewrite-name</i>],</p> <p>[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules inet-prec <i>rewrite-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Option for dscp and inet-prec introduced in Junos OS Release 8.4.</p>
Description	Apply a rewrite rule to MPLS packets only, and write the CoS value to MPLS headers only; or apply a rewrite rule to MPLS and IPv4 packets, and write the CoS value to MPLS and IPv4 headers.
Options	<p><i>protocol-types</i> can be one of the following:</p> <ul style="list-style-type: none"> • mpls—Apply a rewrite rule to MPLS packets and write the CoS value to MPLS headers. • mpls-inet-both—Apply a rewrite rule to VPN MPLS packets with IPv4 payloads. On M120, M320, MX Series, T Series routers (except T4000 routers), and EX Series switches, write the CoS value to the MPLS and IPv4 headers. On M Series routers, initialize all ingress MPLS LSP packets with IPv4 payloads with 000 code points for the MPLS EXP value, and the configured rewrite code point for IP precedence. • mpls-inet-both-non-vpn—Apply a rewrite rule to non-VPN MPLS packets with IPv4 payloads. On M120, M320, MX Series, T Series routers, and EX Series switches write the CoS value to the MPLS and IPv4 headers. On M Series routers, initialize all ingress MPLS LSP packets with IPv4 payloads with 000 code points for the MPLS EXP value, and the configured rewrite code point for IP precedence.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Rewriting MPLS and IPv4 Packet Headers</i>

rewrite-rules (CoS Host Outbound Traffic)

Syntax	rewrite-rules;
Hierarchy Level	[edit class-of-service host-outbound-traffic ieee-802.1]
Release Information	Statement introduced in Junos OS Release 12.3.
Description	Apply the IEEE 802.1p rewrite rules associated with the egress logical interface to the IEEE 802.1p PCP field for all host outbound traffic on that interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Applying Egress Interface Rewrite Rules to the IEEE 802.1p Field for All Host Outbound Traffic on the Interface on page 1711• Rewriting Packet Header Information Overview on page 1690• Configuring Rewrite Rules on page 1694

rewrite-rules (Definition)

Syntax	<pre>rewrite-rules { type <i>rewrite-name</i>{ import (<i>rewrite-name</i> default); forwarding-class <i>class-name</i> { loss-priority <i>level</i> <i>code-point</i> [<i>aliases</i>] [<i>bit-patterns</i>]; } } }</pre>
Hierarchy Level	[edit class-of-service]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>ieee-802.1ad option introduced in Junos OS Release 9.2.</p>
Description	Specify a rewrite-rules mapping for the traffic that passes through all queues on the interface.
Options	<p><i>rewrite-name</i>—Name of a rewrite-rules mapping.</p> <p><i>type</i>—Traffic type.</p> <p>Values: dscp, dscp-ipv6, exp, frame-relay-de (J Series routers only), ieee-802.1, ieee-802.1ad, inet-precedence</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Rewrite Rules on page 1694 • <i>Example: Configuring CoS for a PBB Network</i> • J Series router documentation

rewrite-rules (Interfaces)

Syntax	<pre>rewrite-rules { dscp (rewrite-name default); dscp-ipv6 (rewrite-name default); exp (rewrite-name default) protocol protocol-types; exp-push-push-push default; exp-swap-push-push default; ieee-802.1 (rewrite-name default) vlan-tag (outer outer-and-inner); ieee-802.1ad (rewrite-name default) vlan-tag (outer outer-and-inner); inet-precedence (rewrite-name default); }</pre>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i>], [edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Associate a rewrite-rules configuration or default mapping with a specific interface.</p> <p>The [edit class-of-service interfaces <i>interface-name</i>] hierarchy level is not supported on M Series routers.</p> <p>The [edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level is not supported on ACX Series routers.</p> <p>On an MX Series router and on an EX Series switch, exp-push-push-push, exp-swap-push-push, and frame-relay-de are not supported on an integrated routing and bridging (IRB) interface.</p> <p>On an ACX Series router, only the outer tag is supported for dscp, inet-precedence, and ieee802.1.</p>
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules] hierarchy level.</p> <p>default—The default mapping.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694

rewrite-rules (Physical Interfaces)

Syntax	<pre>rewrite-rules { dscp (rewrite-name default); ieee-802.1 (rewrite-name default); inet-precedence (rewrite-name default); }</pre>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.
Description	Associate a rewrite-rules configuration or default mapping with a specific interface.
Options	<p><i>rewrite-name</i>—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules] hierarchy level.</p> <p>default—The default mapping.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

unit

Syntax	<pre>unit logical-unit-number { classifiers { type (classifier-name default) family (mpls all); } forwarding-class class-name; fragmentation-map map-name; input-traffic-control-profile profile-name shared-instance instance-name; output-traffic-control-profile profile-name shared-instance instance-name; per-session-scheduler; rewrite-rules { dscp (rewrite-name default); dscp-ipv6 (rewrite-name default); exp (rewrite-name default) protocol protocol-types; exp-push-push-push default; exp-swap-push-push default; ieee-802.1 (rewrite-name default) vlan-tag (outer outer-and-inner); inet-precedence (rewrite-name default); } scheduler-map map-name; shaping-rate rate; }</pre>
Hierarchy Level	[edit class-of-service interfaces interface-name]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.
Options	<p>logical-unit-number—Number of the logical unit.</p> <p>Range: 0 through 16,384</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Overview of BA Classifier Types on page 1321• Configuring Rewrite Rules on page 1694

vlan-tag

Syntax	vlan-tag (outer outer-and-inner);
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules ieee-802.1 (<i>rewrite-name</i> default)]
Release Information	Statement introduced in Junos OS Release 8.1.
Description	For Gigabit Ethernet IQ2 PICs only, apply this IEEE-802.1 rewrite rule to the outer or outer and inner VLAN tags.
Default	If you do not include this statement, the rewrite rule applies to the outer VLAN tag only.
Options	outer —Apply the rewrite rule to the outer VLAN tag only. outer-and-inner —Apply the rewrite rule to both the outer and inner VLAN tags.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Applying IEEE 802.1p Rewrite Rules to Dual VLAN Tags on page 1699

Routing Engine Protocol Queue Assignments

- [Overview on page 1745](#)
- [Configuration on page 1748](#)

Overview

- [Routing Engine Protocol Queue Assignments on page 1745](#)

Routing Engine Protocol Queue Assignments

- [Default Routing Engine Protocol Queue Assignments on page 1745](#)
- [Changing the Routing Engine Outbound Traffic Defaults on page 1747](#)

Default Routing Engine Protocol Queue Assignments

[Table 128 on page 1746](#) lists (in alphabetical order) how Routing Engine-sourced traffic is mapped to output queues. The following caveats apply to [Table 128 on page 1746](#):

- For all packets sent to queue 3 over a VLAN-tagged interface, the software sets the 802.1p bit to 110.
- For IPv4 and IPv6 packets, the software copies the IP type-of-service (ToS) value into the 802.1p field independently of which queue the packets are sent out.
- For MPLS packets, the software copies the EXP bits into the 802.1p field.

Table 128: Routing Engine Protocol Queue Assignments

Routing Engine Protocol	Queue Assignment
Adaptive Services PIC	TCP tickle (keepalive packets for idle session generated with stateful firewall to probe idle TCP sessions) are sent from queue 0.
Bidirectional Forwarding Detection (BFD) Protocol	Queue 3
Border Gateway Protocol (BGP)	Queue 0
BGP TCP Retransmission	Queue 3
Cisco High-Level Data Link Control (HDLC)	Queue 3
Distance Vector Multicast Routing Protocol (DVMRP)	Queue 3
Frame Relay Local Management Interface (LMI)	Queue 3
Frame Relay Asynchronization permanent virtual circuit (PVC)/data link connection identifier (DLCI) status messages	Queue 3
FTP	Queue 0
Intermediate System-to-Intermediate System (IS-IS) Open Systems Interconnection (OSI)	Queue 3
Internet Group Management Protocol (IGMP) query	Queue 3
IGMP Report	Queue 0
IP version 6 (IPv6) Neighbor Solicitation	Queue 3
IPv6 Neighbor Advertisement	Queue 3
IPv6 Router Advertisement	Queue 0
Label Distribution Protocol (LDP) User Datagram Protocol (UDP) hello	Queue 3
LDP keepalive and Session data	Queue 0
LDP TCP Retransmission	Queue 3
Link Aggregation Control Protocol (LACP)	Queue 3

Table 128: Routing Engine Protocol Queue Assignments (*continued*)

Routing Engine Protocol	Queue Assignment
Link Services (LS) PIC	If link fragmentation and interleaving (LFI) is enabled, all routing protocol packets larger than 128 bytes are transmitted from queue 0. This ensures that VoIP traffic is not affected. Fragmentation is supported on queue 0 only.
Multicast listener discovery (MLD)	Queue 0
Multicast Source Discovery Protocol (MSDP)	Queue 0
MSDP TCP Retransmission	Queue 3
Multilink Frame Relay Link Integrity Protocol (LIP)	Queue 3
Open Shortest Path First (OSPF) protocol data unit (PDU)	Queue 3
Protocol Independent Multicast (PIM)	Queue 3
Real-time performance monitoring (RPM) probe packets	Queue 3
Resource Reservation Protocol (RSVP)	Queue 3
Routing Information Protocol (RIP)	Queue 3
Simple Network Management Protocol (SNMP)	Queue 0
SSH	Queue 0
Telnet	Queue 0
Virtual Router Redundancy Protocol (VRRP)	Queue 3
xnm-clear-text	Queue 0
xnm-ssl	Queue 0

Changing the Routing Engine Outbound Traffic Defaults

You can modify the default queue assignment (forwarding class) and DSCP bits used in the ToS field of packets generated by the Routing Engine. By default, the forwarding class (queue) and packet loss priority (PLP) bits are set according to the values given in [“Default DSCP and DSCP IPv6 Classifier” on page 1326](#).

TCP-related packets, such as BGP or LDP, use queue 3 (network control) for retransmitted traffic. Changing the defaults for Routing Engine-sourced traffic does not affect transit

or incoming traffic. The changes apply to all packets relating to Layer 3 and Layer 2 protocols, but not MPLS EXP bits or IEEE 802.1p bits. This feature applies to all application-level traffic such as FTP or ping operations as well.

This feature is not available on Juniper Networks J Series Services Routers.

The queue selected is global to the routing device. That is, the traffic is placed in the selected queue on all egress interfaces. In the case of a restricted interface, the Routing Engine-sourced traffic flows through the restricted queue.

The queue selected must be properly configured on all interfaces. For more information about configuring queues and forwarding classes, see [“Overview of Forwarding Classes” on page 1526](#).

To change the default queue and DSCP bits for Routing Engine-sourced traffic, include the **host-outbound-traffic** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
host-outbound-traffic {
  forwarding-class class-name;
  dscp-code-point value;
}
```

The following example places all Routing Engine-sourced traffic into queue 3 (network control) with a DSCP code point value of 101010:

```
[edit class-of-service]
host-outbound-traffic {
  forwarding-class network-control;
  dscp-code-point 101010;
}
```

Configuration

- [Configuration Statements on page 1748](#)

Configuration Statements

- [\[edit class-of-service\] Hierarchy Level on page 1748](#)

[edit class-of-service] Hierarchy Level

```
class-of-service {
  classifiers {
    type classifier-name {
      forwarding-class class-name {
        loss-priority (high | low | medium-high | medium-low) code-points [ aliases bits ];
      }
      import (classifier-name | default);
    }
  }
  code-point-aliases {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) {
      alias-name bits;
    }
  }
}
```

```

drop-profiles {
  profile-name {
    fill-level percentage drop-probability percentage;
    interpolate {
      drop-probability value;
      fill-level value;
    }
  }
}
fabric {
  scheduler-map {
    priority (high | low) scheduler scheduler-name;
  }
}
forwarding-class-map {
  map-name {
    class class-name queue-num queue-number <restricted-queue queue-number>;
  }
}
forwarding-classes {
  class class-name policing-priority (normal | premium) queue-num queue-number
  priority (high | low);
  queue queue-number class-name policing-priority (normal | premium) priority (high |
  low);
}
forwarding-policy {
  class class-name {
    classification-override {
      forwarding-class class-name;
    }
  }
  next-hop-map map-name {
    forwarding-class class-name {
      discard;
      lsp-next-hop [ lsp-regular-expressions ];
      next-hop [ next-hop-names ];
      non-lsp-next-hop;
    }
  }
}
fragmentation-maps {
  map-name {
    forwarding-class class-name {
      drop-timeout milliseconds;
      fragment-threshold bytes;
      multilink-class number;
      no-fragmentation;
    }
  }
}
host-outbound-traffic {
  dscp-code-point value;
  forwarding-class class-name;
  ieee-802.1 {
    default value;
    rewrite-rules;
  }
}

```

```

    }
    tcp {
        raise-internet-control-priority;
    }
}
interfaces {
    ... the interfaces subhierarchy appears after the main [edit class-of-service] hierarchy
    ...
}
}
restricted-queues {
    forwarding-class class-name queue-number;
}
rewrite-rules {
    (dscp | dscp-ipv6 | exp | frame-relay-de | ieee-802.1 | ieee-802.1ad | inet-precedence)
    rewrite-rule {
        forwarding-class class-name {
            loss-priority level code-point (alias | bits);
        }
        import (rewrite-rule | default);
    }
}
}
routing-instances routing-instance-name {
    classifiers {
        dscp (classifier-name | default);
        dscp-ipv6 (classifier-name | default);
        exp (classifier-name | default);
        ieee-208.1 (classifier-name | default | encapsulated | vlan-tag (inner | outer));
    }
}
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
schedulers {
    scheduler-name {
        adjust-minimum value;
        adjust-percent value;
        buffer-size (exact | percent percentage | remainder);
        drop-profile-map loss-priority (any | high | low | medium-high | medium-low)
            protocol any;
        excess-priority (high | low | medium-high | medium-low);
        excess-rate (percent percentage | proportion proportion);
        priority (high | low | medium-high | medium-low | strict-high);
        shaping-rate (bps | percent percentage | burst-size size);
        transmit-rate (bps | percent percentage | remainder) <exact | rate-limit>;
    }
}
traceoptions {
    file <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
traffic-control-profiles {

```

```

    profile-name {
        adjust-minimum rate;
        delay-buffer-rate (bps | cps cps | percent percentage);
        excess-rate (percent percentage | proportion value);
        guaranteed-rate (bps | percent percentage) <burst-size bytes>;
        overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
        scheduler-map map-name;
        shaping-rate (bps | percent percentage) <burst-size bytes>;
    }
}
tri-color;
}

class-of-service {
    interfaces {
        interface-name {
            excess-bandwidth-share (equal | proportional value);
            input-excess-bandwidth-share (equal | proportional value);
            input-scheduler-map map-name;
            input-shaping-rate bps;
            input-traffic-control-profile profile-name;
            output-forwarding-class-map map-name;
            output-traffic-control-profile profile-name;
            scheduler-map map-name;
            scheduler-map-chassis (map-name | derived);
            shaping-rate bps;
            unit (logical-unit-number | *) {
                classifiers {
                    dscp (classifier-name | default) {
                        family [ inet mpls ];
                    }
                    dscp-ipv6 (classifier-name | default) {
                        family [ inet mpls ];
                    }
                    exp (classifier-name | default);
                    ieee-208.1 (classifier-name | default) <vlan-tag (inner | outer)>;
                    ieee-208.1ad (classifier-name | default);
                    inet-precedence (classifier-name | default);
                }
                forwarding-class class-name;
                input-scheduler-map map-name;
                input-shaping-rate bps;
                input-traffic-control-profile profile-name shared-instance instance-name;
                loss-priority-maps {
                    (map-name | default);
                }
                loss-priority-rewrites {
                    (map-name | default);
                }
                output-forwarding-class-map map-name;
                output-traffic-control-profile profile-name shared-instance instance-name;
                rewrite-rules {
                    dscp (rule-name | default) <protocol mpls>;
                    dscp-ipv6 (rule-name | default);
                    exp (rule-name | default) <protocol [ mpls-any | mpls-inet-both |
                        mpls-inet-both-non-vpn ]>;
                }
            }
        }
    }
}


```

```
exp-push-push-push default;
exp-swap-push-push default;
ieee-802.1 (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
ieee-802.1ad (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
inet-precedence (rewrite-name | default) <protocol mpls>;
}
scheduler-map map-name;
shaping-rate bps;
translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 |
to-exp-from-exp | to-inet-precedence-from-inet-precedence) table-name;
}
}
interface-set interface-set-name {
excess-bandwidth-share (equal | proportional value);
input-excess-bandwidth-share (equal | proportional value);
input-traffic-control-profile profile-name;
input-traffic-control-profile-remaining profile-name;
internal-node;
output-traffic-control-profile profile-name;
output-traffic-control-profile-remaining profile-name;
}
}
}
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

classifiers (Application)

Syntax	<pre> classifiers { type (classifier-name default) family (mpls inet); } </pre>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Apply a CoS aggregate behavior classifier to a logical interface. You can apply a default classifier or one that is previously defined.
Options	<p>classifier-name—Name of the aggregate behavior classifier.</p> <p>type—Traffic type.</p> <p>Values: dscp, dscp-ipv6, exp, ieee-802.1, inet-precedence</p>
	<div>  <p>NOTE: You can only specify a family for the dscp and dscp-ipv6 types.</p> </div>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Applying Classifiers to Logical Interfaces on page 1331

dscp (Rewrite Rules)

Syntax	<code>dscp (rewrite-name default);</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For IPv4 traffic, apply a Differentiated Services (DiffServ) code point (DSCP) rewrite rule.
Options	rewrite-name —Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules dscp] hierarchy level. default —The default mapping.
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694• dscp-ipv6 (Class-of-Service) on page 1357• exp on page 1358• exp-push-push-push on page 1723• exp-swap-push-push on page 1724• ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360• ieee-802.1ad on page 1362• inet-precedence on page 1364• rewrite-rules (Definition) on page 1481

dscp-code-point (Class-of-Service)

Syntax	<code>dscp-code-point value;</code>
Hierarchy Level	[edit class-of-service host-outbound-traffic]
Release Information	Statement introduced in Junos OS Release 8.4. Statement introduced before Junos OS 11.4 for EX Series switches.
Description	Set the value of the DSCP code point in the ToS field of the packet generated by the Routing Engine (host).
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Changing the Routing Engine Outbound Traffic Defaults on page 1747.

dscp-ipv6 (Class-of-Service)

Syntax	<code>dscp-ipv6 (<i>rewrite-name</i> <default>) { protocol mpls }</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For IPv6 traffic, apply a DSCP rewrite rule.
Options	<p><i>rewrite-name</i>—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules dscp-ipv6] hierarchy level.</p> <p><i>default</i>— Default mapping.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Rewrite Rules on page 1694 • dscp (Rewrite Rules) on page 1382 • exp on page 1358 • exp-push-push-push on page 1723 • exp-swap-push-push on page 1724 • ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360 • ieee-802.1ad on page 1362 • inet-precedence on page 1364 • rewrite-rules (Definition) on page 1481

exp

Syntax	exp (<i>rewrite-name</i> default) protocol <i>protocol-types</i> ;
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS Release 12.2. for ACX series
Description	Apply an MPLS experimental (EXP) rewrite rule.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules exp] hierarchy level.</p> <p>default—The default mapping.</p> <p>By default, IP precedence rewrite rules alter the first three bits on the type-of-service (ToS) byte while leaving the last three bits unchanged. This default behavior applies to rewrite rules you configure for MPLS packets with IPv4 payloads. You configure these types of rewrite rules by including the mpls-inet-both or mpls-inet-both-non-vpn option at the [edit class-of-service interfaces <i>interface</i> <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules exp <i>rewrite-rule-name</i> protocol] hierarchy level. The IP precedence rewrite rules explanation does not apply to ACX Series Universal Access routers.</p> <p>On interfaces configured on Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) on MX Series Ethernet Services Routers and EX Series switches, we highly recommend that you configure the default option when you configure a behavior aggregate (BA) classifier that does not include a specific rewrite rule for MPLS packets. Doing so ensures that MPLS exp value is rewritten according to the BA classifier rules configured for forwarding or packet loss priority. This does not apply to ACX Series Universal Access routers.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Rewriting the EXP Bits of All Three Labels of an Outgoing Packet</i> • dscp (Rewrite Rules) on page 1382 • dscp-ipv6 (Class-of-Service) on page 1357 • exp-push-push-push on page 1723 • exp-swap-push-push on page 1724 • ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360 • ieee-802.1ad on page 1362 • inet-precedence on page 1364

- [rewrite-rules \(Definition\)](#) on page 1481

forwarding-class (Forwarding Policy)

Syntax	<pre>forwarding-class <i>class-name</i> { next-hop [<i>next-hop-name</i>]; lsp-next-hop [<i>lsp-regular-expression</i>]; non-lsp-next-hop; discard; }</pre>
Hierarchy Level	[edit class-of-service forwarding-policy next-hop-map <i>map-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define forwarding class name and associated next hops.
Options	<p><i>class-name</i>—Name of the forwarding class.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Overriding the Input Classification on page 1560

host-outbound-traffic

Syntax host-outbound-traffic {
 forwarding-class *class-name*;
 dscp-code-point *value*;
 ieee-802.1 {
 default *value*;
 rewrite-rules;
 }
 }

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 8.4.
 Statement introduced before Junos OS 11.4 for EX Series switches.

Description Allow queue selection for all traffic generated by the Routing Engine (host). The selected queue must be configured properly. The configuration of specific DSCP code point bits for the ToS field of the generated packets is also allowed. Transit packets are not affected; only packets originating on the Routing Engine are affected. By default, the forwarding class (queue) and DSCP bits are set according to those given in *Default Routing Engine Protocol Queue Assignments*. This feature is not available on J Series routers.

The remaining statements are explained separately.

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

Related Documentation

- See [Changing the Routing Engine Outbound Traffic Defaults on page 1747](#).
- [Configuring a Global Default IEEE 802.1p Value for All Host Outbound Traffic on page 1710](#)
- [Applying Egress Interface Rewrite Rules to the IEEE 802.1p Field for All Host Outbound Traffic on the Interface on page 1711](#)
- *Understanding Junos OS CoS Components for EX Series Switches*

ieee-802.1 (Rewrite Rules on Logical Interface)

Syntax	<code>ieee-802.1</code> (<i>rewrite-name</i> default) vlan-tag (outer outer-and-inner);
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4. vlan-tag statement introduced in Junos OS Release 8.1.
Description	Apply an IEEE-802.1 rewrite rule. For IQ PICs, you can only configure one IEEE 802.1 rewrite rule on a physical port. All logical ports (units) on that physical port should apply the same IEEE 802.1 rewrite rule.
Options	rewrite-name —Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules ieee-802.1] hierarchy level. default —The default mapping.
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Rewrite Rules on page 1694 • <i>Example: Configuring CoS for a PBB Network</i> • dscp (Rewrite Rules) on page 1382 • dscp-ipv6 (Class-of-Service) on page 1357 • exp on page 1358 • exp-push-push-push on page 1723 • exp-swap-push-push on page 1724 • ieee-802.1ad on page 1362 • inet-precedence on page 1364 • rewrite-rules (Definition) on page 1481

inet-precedence

Syntax	<code>inet-precedence (<i>rewrite-name</i> default);</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Apply a IPv4 precedence rewrite rule.
Options	<p><i>rewrite-name</i>—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules inet-precedence] hierarchy level.</p> <p>default—The default mapping. By default, IP precedence rewrite rules alter the first three bits on the type of service (ToS) byte while leaving the last three bits unchanged.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694• dscp (Rewrite Rules) on page 1382• dscp-ipv6 (Class-of-Service) on page 1357• exp on page 1358• exp-push-push-push on page 1723• exp-swap-push-push on page 1724• ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360• ieee-802.1ad on page 1362• rewrite-rules (Definition) on page 1481

irb

```
Syntax  irb {
        unit logical-unit-number {
            classifiers {
                type (classifier-name | default);
            }
            rewrite-rules {
                dscp (rewrite-name | default);
                dscp-ipv6 (rewrite-name | default);
                exp (rewrite-name | default) protocol protocol-types;
                ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                inet-precedence (rewrite-name | default);
            }
        }
    }
```

Hierarchy Level [edit class-of-service [interfaces](#)]

Release Information Statement introduced in Junos OS Release 8.4.

Description On the MX Series routers and EX Series switches, you can apply classifiers or rewrite rules to an integrated bridging and routing (IRB) interface. All types of classifiers and rewrite rules are allowed. These classifiers and rewrite rules are independent of others configured on the MX Series router and on EX Series switches.

The statements are explained separately.

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

Related Documentation

- *MX Series Router CoS Hardware Capabilities and Limitations*

protocol (Rewrite Rules)

Syntax	<code>protocol protocol-types;</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules exp <i>rewrite-name</i>], [edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules dscp <i>rewrite-name</i>], [edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules inet-prec <i>rewrite-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Option for dscp and inet-prec introduced in Junos OS Release 8.4.
Description	Apply a rewrite rule to MPLS packets only, and write the CoS value to MPLS headers only; or apply a rewrite rule to MPLS and IPv4 packets, and write the CoS value to MPLS and IPv4 headers.
Options	<i>protocol-types</i> can be one of the following: <ul style="list-style-type: none">• mpls—Apply a rewrite rule to MPLS packets and write the CoS value to MPLS headers.• mpls-inet-both—Apply a rewrite rule to VPN MPLS packets with IPv4 payloads. On M120, M320, MX Series, T Series routers (except T4000 routers), and EX Series switches, write the CoS value to the MPLS and IPv4 headers. On M Series routers, initialize all ingress MPLS LSP packets with IPv4 payloads with 000 code points for the MPLS EXP value, and the configured rewrite code point for IP precedence.• mpls-inet-both-non-vpn—Apply a rewrite rule to non-VPN MPLS packets with IPv4 payloads. On M120, M320, MX Series, T Series routers, and EX Series switches write the CoS value to the MPLS and IPv4 headers. On M Series routers, initialize all ingress MPLS LSP packets with IPv4 payloads with 000 code points for the MPLS EXP value, and the configured rewrite code point for IP precedence.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Rewriting MPLS and IPv4 Packet Headers</i>

rewrite-rules (Interfaces)

Syntax	<pre>rewrite-rules { dscp (rewrite-name default); dscp-ipv6 (rewrite-name default); exp (rewrite-name default) protocol protocol-types; exp-push-push-push default; exp-swap-push-push default; ieee-802.1 (rewrite-name default) vlan-tag (outer outer-and-inner); ieee-802.1ad (rewrite-name default) vlan-tag (outer outer-and-inner); inet-precedence (rewrite-name default); }</pre>
Hierarchy Level	<p>[edit class-of-service interfaces <i>interface-name</i>], [edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Associate a rewrite-rules configuration or default mapping with a specific interface.</p> <p>The [edit class-of-service interfaces <i>interface-name</i>] hierarchy level is not supported on M Series routers.</p> <p>The [edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level is not supported on ACX Series routers.</p> <p>On an MX Series router and on an EX Series switch, exp-push-push-push, exp-swap-push-push, and frame-relay-de are not supported on an integrated routing and bridging (IRB) interface.</p> <p>On an ACX Series router, only the outer tag is supported for dscp, inet-precedence, and ieee802.1.</p>
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules] hierarchy level.</p> <p>default—The default mapping.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Rewrite Rules on page 1694

unit

Syntax `unit logical-unit-number {
 classifiers {
 type (classifier-name | default) family (mpls | all);
 }
 forwarding-class class-name;
 fragmentation-map map-name;
 input-traffic-control-profile profile-name shared-instance instance-name;
 output-traffic-control-profile profile-name shared-instance instance-name;
 per-session-scheduler;
 rewrite-rules {
 dscp (rewrite-name | default);
 dscp-ipv6 (rewrite-name | default);
 exp (rewrite-name | default) protocol protocol-types;
 exp-push-push-push default;
 exp-swap-push-push default;
 ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
 inet-precedence (rewrite-name | default);
 }
 scheduler-map map-name;
 shaping-rate rate;
 }`

Hierarchy Level [edit class-of-service **interfaces** interface-name]

Release Information Statement introduced before Junos OS Release 7.4.

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

Options **logical-unit-number**—Number of the logical unit.

Range: 0 through 16,384

The remaining statements are explained separately.

Required Privilege interface—To view this statement in the configuration.

Level interface-control—To add this statement to the configuration.

Related Documentation

- [Overview of BA Classifier Types on page 1321](#)
- [Configuring Rewrite Rules on page 1694](#)

vlan-tag

Syntax	<code>vlan-tag (outer outer-and-inner);</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules ieee-802.1 (<i>rewrite-name</i> default)]
Release Information	Statement introduced in Junos OS Release 8.1.
Description	For Gigabit Ethernet IQ2 PICs only, apply this IEEE-802.1 rewrite rule to the outer or outer and inner VLAN tags.
Default	If you do not include this statement, the rewrite rule applies to the outer VLAN tag only.
Options	<p>outer—Apply the rewrite rule to the outer VLAN tag only.</p> <p>outer-and-inner—Apply the rewrite rule to both the outer and inner VLAN tags.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Applying IEEE 802.1p Rewrite Rules to Dual VLAN Tags on page 1699

CoS for Tunnels

- [Overview on page 1765](#)
- [Configuration on page 1766](#)

Overview

- [CoS for Tunnels on page 1765](#)

CoS for Tunnels

- [CoS for Tunnels Overview on page 1765](#)

CoS for Tunnels Overview

For Adaptive Services, Link Services, and Tunnel PICs installed on Juniper Networks M Series Multiservice Edge Routers and T Series Core Routers with enhanced Flexible PIC Concentrators (FPCs), class-of-service (CoS) information is preserved inside generic routing encapsulation (GRE) and IP-IP tunnels.

For the ES PIC installed on M Series and T Series routers with enhanced FPCs, class-of-service information is preserved inside IP Security (IPsec) tunnels. For IPsec tunnels, you do not need to configure CoS, because the ES PIC copies the type-of-service (ToS) byte from the inner IP header to the GRE or IP-IP header.

For IPsec tunnels, the IP header type-of-service (ToS) bits are copied to the outer IPsec header at encryption side of the tunnel. You can rewrite the outer ToS bits in the IPsec

header using a rewrite rule. On the decryption side of the IPsec tunnel, the ToS bits in the IPsec header are not written back to the original IP header field. You can still apply a firewall filter to the ToS bits to apply a packet action on egress. For more information about ToS bits and the Multiservices PICs, see *Multiservices PIC ToS Translation*. For more information about IPsec and Multiservices PICs, see the *Junos Services Interfaces Configuration Release 11.2*.

To configure CoS for tunnels, include the following statements at the **[edit class-of-service]** and **[edit interfaces]** hierarchy level:

```
[edit class-of-service]
interfaces {
  interface-name {
    unit logical-unit-number {
      rewrite-rules {
        dscp (rewrite-name | default);
        dscp-ipv6 (rewrite-name | default);
        exp (rewrite-name | default) protocol protocol-types;
        exp-push-push-push default;
        exp-swap-push-push default;
        ieee-802.1 (rewrite-name | default);
        inet-precedence (rewrite-name | default);
      }
    }
  }
}
rewrite-rules {
  (dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence) rewrite-name {
    import (rewrite-name | default);
    forwarding-class class-name {
      loss-priority level code-point (alias | bits);
    }
  }
}
[edit interfaces]
gre-interface-name {
  unit logical-unit-number;
  copy-tos-to-outer-ip-header;
}
```

Configuration

- [Configuration Task on page 1766](#)
- [Examples on page 1767](#)
- [Configuration Statements on page 1770](#)

Configuration Task

- [Configuring CoS for Tunnels on page 1766](#)

Configuring CoS for Tunnels

To configure CoS for GRE and IP-IP tunnels, perform the following configuration tasks:

1. To configure the tunnel, include the **tunnel** statement at the **[edit interfaces ip-fpc/pic/port unit logical-unit-number]** or **[edit interfaces gr-fpc/pic/port unit logical-unit-number]** hierarchy level.
2. To rewrite traffic on the outbound interface, include the **rewrite-rules** statement at the **[edit class-of-service]** and **[edit class-of-service interfaces interface-name unit logical-unit-number]** hierarchy levels. For GRE and IP-IP tunnels, you can configure IP precedence and DSCP rewrite rules.
3. To classify traffic on the inbound interface, you can configure a behavior aggregate (BA) classifier or firewall filter. Include the **loss-priority** and **forwarding-class** statements at the **[edit firewall filter filter-name term term-name then]** hierarchy level, or the **classifiers** statement at the **[edit class-of-service]** hierarchy level.
4. For a GRE tunnel, the default is to set the ToS bits in the outer IP header to all 0s. To copy the ToS bits from the inner IP header to the outer, include the **copy-tos-to-outer-ip-header** statement at the **[edit interfaces gr-fpc/pic/port unit logical-unit-number]** hierarchy level. (This inner-to-outer ToS bits copying is already the default behavior for IP-IP tunnels.)


Examples

- [Example: Configuring CoS for Tunnels on page 1767](#)
- [Example: Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header on page 1770](#)

Example: Configuring CoS for Tunnels

In [Figure 16 on page 1767](#), Router A acts as a tunnel ingress device. The link between interfaces **ge-1/0/0** in Router A and **ge-1/3/0** in Router B is the GRE or IP-IP tunnel. Router A monitors the traffic received from interface **ge-1/3/0**. By way of interface **ge-1/0/0**, Router C generates traffic to Router B.

Figure 16: CoS with a Tunnel Configuration



```
Router A [edit interfaces]
ge-1/0/0 {
  unit 0 {
    family inet {
      address 10.80.0.2/24;
    }
  }
}
ge-1/0/1 {
  unit 0 {
    family inet {
```

```
        filter {
            input zf-catch-all;
        }
        address 10.90.0.2/24;
    }
}
gr-2/1/0 {
    unit 0 {
        tunnel {
            source 11.11.11.11;
            destination 10.255.245.46;
        }
        family inet {
            address 21.21.21.21/24;
        }
    }
}
ip-2/1/0 {
    unit 0 {
        tunnel {
            source 12.12.12.12;
            destination 10.255.245.46;
        }
        family inet {
            address 22.22.22.22/24;
        }
    }
}

[edit routing-options]
static {
    route 1.1.1.1/32 next-hop gr-2/1/0.0;
    route 2.2.2.2/32 next-hop ip-2/1/0.0;
}

[edit class-of-service]
interfaces {
    ge-1/0/0 {
        unit 0 {
            rewrite-rules {
                inet-precedence zf-tun-rw-ipprec-00;
            }
        }
    }
}
rewrite-rules {
    inet-precedence zf-tun-rw-ipprec-00 {
        forwarding-class best-effort {
            loss-priority low code-point 000;
            loss-priority high code-point 001;
        }
        forwarding-class expedited-forwarding {
            loss-priority low code-point 010;
            loss-priority high code-point 011;
        }
    }
}
```

```

        forwarding-class assured-forwarding {
            loss-priority low code-point 100;
            loss-priority high code-point 101;
        }
        forwarding-class network-control {
            loss-priority low code-point 110;
            loss-priority high code-point 111;
        }
    }
}
dscp zf-tun-rw-dscp-00 {
    forwarding-class best-effort {
        loss-priority low code-point 000000;
        loss-priority high code-point 001001;
    }
    forwarding-class expedited-forwarding {
        loss-priority low code-point 010010;
        loss-priority high code-point 011011;
    }
    forwarding-class assured-forwarding {
        loss-priority low code-point 100100;
        loss-priority high code-point 101101;
    }
    forwarding-class network-control {
        loss-priority low code-point 110110;
        loss-priority high code-point 111111;
    }
}

[edit firewall]
filter zf-catch-all {
    term term1 {
        then {
            loss-priority high;
            forwarding-class network-control;
        }
    }
}

```

```

Router B [edit interfaces]
ge-1/3/0 {
    unit 0 {
        family inet {
            address 10.80.0.1/24;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.245.46/32;
        }
    }
}

```

```
Router C    [edit interfaces]
            ge-1/0/0 {
              unit 0 {
                family inet {
                  address 10.90.0.1/24;
                }
              }
            }

            [edit routing-options]
            static {
              route 1.1.1.1/32 next-hop 10.90.0.2;
              route 2.2.2.2/32 next-hop 10.90.0.2;
            }
```

Example: Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header

Unlike IP-IP tunnels, GRE tunnels do not copy the ToS bits to the outer IP header by default. To copy the inner ToS bits to the outer IP header (which is required for some tunneled routing protocols) on packets sent by the Routing Engine, include the **copy-tos-to-outer-ip-header** statement at the logical unit hierarchy level of a GRE interface. This example copies the inner ToS bits to the outer IP header on a GRE tunnel:

```
[edit interfaces]
gr-0/0/0 {
  unit 0 {
    copy-tos-to-outer-ip-header;
    family inet;
  }
}
```

Configuration Statements

- [\[edit class-of-service\] Hierarchy Level on page 1770](#)
- [\[edit interfaces\] Hierarchy Level on page 1790](#)

[edit class-of-service] Hierarchy Level

```
class-of-service {
  classifiers {
    type classifier-name {
      forwarding-class class-name {
        loss-priority (high | low | medium-high | medium-low) code-points [ aliases bits ];
      }
      import (classifier-name | default);
    }
  }
  code-point-aliases {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) {
      alias-name bits;
    }
  }
  drop-profiles {
    profile-name {
      fill-level percentage drop-probability percentage;
      interpolate {
```



```

        drop-probability value;
        fill-level value;
    }
}
fabric {
    scheduler-map {
        priority (high | low) scheduler scheduler-name;
    }
}
forwarding-class-map {
    map-name {
        class class-name queue-num queue-number <restricted-queue queue-number>;
    }
}
forwarding-classes {
    class class-name policing-priority (normal | premium) queue-num queue-number
    priority (high | low);
    queue queue-number class-name policing-priority (normal | premium) priority (high |
    low);
}
forwarding-policy {
    class class-name {
        classification-override {
            forwarding-class class-name;
        }
    }
}
next-hop-map map-name {
    forwarding-class class-name {
        discard;
        lsp-next-hop [ lsp-regular-expressions ];
        next-hop [ next-hop-names ];
        non-lsp-next-hop;
    }
}
fragmentation-maps {
    map-name {
        forwarding-class class-name {
            drop-timeout milliseconds;
            fragment-threshold bytes;
            multilink-class number;
            no-fragmentation;
        }
    }
}
host-outbound-traffic {
    dscp-code-point value;
    forwarding-class class-name;
    ieee-802.1 {
        default value;
        rewrite-rules;
    }
    tcp {
        raise-internet-control-priority;
    }
}

```

```

}
interfaces {
  ... the interfaces subhierarchy appears after the main [edit class-of-service] hierarchy
  ...
}
restricted-queues {
  forwarding-class class-name queue-number;
}
rewrite-rules {
  (dscp | dscp-ipv6 | exp | frame-relay-de | ieee-802.1 | ieee-802.1ad | inet-precedence)
  rewrite-rule {
    forwarding-class class-name {
      loss-priority level code-point (alias | bits);
    }
    import (rewrite-rule | default);
  }
}
routing-instances routing-instance-name {
  classifiers {
    dscp (classifier-name | default);
    dscp-ipv6 (classifier-name | default);
    exp (classifier-name | default);
    ieee-208.1 (classifier-name | default | encapsulated | vlan-tag (inner | outer));
  }
}
scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
schedulers {
  scheduler-name {
    adjust-minimum value;
    adjust-percent value;
    buffer-size (exact | percent percentage | remainder);
    drop-profile-map loss-priority (any | high | low | medium-high | medium-low)
      protocol any;
    excess-priority (high | low | medium-high | medium-low);
    excess-rate (percent percentage | proportion proportion);
    priority (high | low | medium-high | medium-low | strict-high);
    shaping-rate (bps | percent percentage | burst-size size);
    transmit-rate (bps | percent percentage | remainder) <exact | rate-limit>;
  }
}
traceoptions {
  file <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
traffic-control-profiles {
  profile-name {
    adjust-minimum rate;
    delay-buffer-rate (bps | cps cps | percent percentage);
    excess-rate (percent percentage | proportion value);
  }
}

```

```

    guaranteed-rate (bps | percent percentage) <burst-size bytes>;
    overhead-accounting (frame-mode | cell-mode) <bytes byte-value>;
    scheduler-map map-name;
    shaping-rate (bps | percent percentage) <burst-size bytes>;
  }
}
tri-color;
}

class-of-service {
  interfaces {
    interface-name {
      excess-bandwidth-share (equal | proportional value);
      input-excess-bandwidth-share (equal | proportional value);
      input-scheduler-map map-name;
      input-shaping-rate bps;
      input-traffic-control-profile profile-name;
      output-forwarding-class-map map-name;
      output-traffic-control-profile profile-name;
      scheduler-map map-name;
      scheduler-map-chassis (map-name | derived);
      shaping-rate bps;
      unit (logical-unit-number | *) {
        classifiers {
          dscp (classifier-name | default) {
            family [ inet mpls ];
          }
          dscp-ipv6 (classifier-name | default) {
            family [ inet mpls ];
          }
          exp (classifier-name | default);
          ieee-208.1 (classifier-name | default) <vlan-tag (inner | outer)>;
          ieee-208.1ad (classifier-name | default);
          inet-precedence (classifier-name | default);
        }
        forwarding-class class-name;
        input-scheduler-map map-name;
        input-shaping-rate bps;
        input-traffic-control-profile profile-name shared-instance instance-name;
        loss-priority-maps {
          (map-name | default);
        }
        loss-priority-rewrites {
          (map-name | default);
        }
        output-forwarding-class-map map-name;
        output-traffic-control-profile profile-name shared-instance instance-name;
        rewrite-rules {
          dscp (rule-name | default) <protocol mpls>;
          dscp-ipv6 (rule-name | default);
          exp (rule-name | default) <protocol [ mpls-any | mpls-inet-both |
            mpls-inet-both-non-vpn ]>;
          exp-push-push-push default;
          exp-swap-push-push default;
          ieee-802.1 (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
          ieee-802.1ad (rewrite-name | default) <vlan-tag (outer | outer-and-inner)>;
        }
      }
    }
  }
}

```

```
        inet-precedence (rewrite-name | default) <protocol mpls>;
    }
    scheduler-map map-name;
    shaping-rate bps;
    translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 |
        to-exp-from-exp | to-inet-precedence-from-inet-precedence) table-name;
    }
}
interface-set interface-set-name {
    excess-bandwidth-share (equal | proportional value);
    input-excess-bandwidth-share (equal | proportional value);
    input-traffic-control-profile profile-name;
    input-traffic-control-profile-remaining profile-name;
    internal-node;
    output-traffic-control-profile profile-name;
    output-traffic-control-profile-remaining profile-name;
}
}
```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

code-point

Syntax	<code>code-point [<i>aliases</i>] [<i>bit-patterns</i>];</code>
Hierarchy Level	[edit class-of-service rewrite-rules type <i>rewrite-name</i> forwarding-class <i>class-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify one or more code-point aliases or bit sets for association with a forwarding class.
Options	<i>aliases</i> —Name of each alias. <i>bit-patterns</i> —Value of the code-point bits, in decimal form.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	• Configuring Rewrite Rules on page 1694

dscp (Rewrite Rules)

Syntax	<code>dscp (rewrite-name default);</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For IPv4 traffic, apply a Differentiated Services (DiffServ) code point (DSCP) rewrite rule.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules dscp] hierarchy level.</p> <p>default—The default mapping.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Rewrite Rules on page 1694 • dscp-ipv6 (Class-of-Service) on page 1357 • exp on page 1358 • exp-push-push-push on page 1723 • exp-swap-push-push on page 1724 • ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360 • ieee-802.1ad on page 1362 • inet-precedence on page 1364 • rewrite-rules (Definition) on page 1481

dscp-ipv6 (Class-of-Service)

Syntax	<code>dscp-ipv6 (<i>rewrite-name</i> <default>) { protocol mpls }</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For IPv6 traffic, apply a DSCP rewrite rule.
Options	<p><i>rewrite-name</i>—Name of a <i>rewrite-rules</i> mapping configured at the [edit class-of-service <i>rewrite-rules dscp-ipv6</i>] hierarchy level.</p> <p><i>default</i>— Default mapping.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694• dscp (Rewrite Rules) on page 1382• exp on page 1358• exp-push-push-push on page 1723• exp-swap-push-push on page 1724• ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360• ieee-802.1ad on page 1362• inet-precedence on page 1364• rewrite-rules (Definition) on page 1481

exp

Syntax	<code>exp (rewrite-name default) protocol protocol-types;</code>
Hierarchy Level	<code>[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced before Junos OS Release 12.2. for ACX series
Description	Apply an MPLS experimental (EXP) rewrite rule.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the <code>[edit class-of-service rewrite-rules exp]</code> hierarchy level.</p> <p>default—The default mapping.</p> <p>By default, IP precedence rewrite rules alter the first three bits on the type-of-service (ToS) byte while leaving the last three bits unchanged. This default behavior applies to rewrite rules you configure for MPLS packets with IPv4 payloads. You configure these types of rewrite rules by including the mpls-inet-both or mpls-inet-both-non-vpn option at the <code>[edit class-of-service interfaces interface interface-name unit logical-unit-number rewrite-rules exp rewrite-rule-name protocol]</code> hierarchy level. The IP precedence rewrite rules explanation does not apply to ACX Series Universal Access routers.</p> <p>On interfaces configured on Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) on MX Series Ethernet Services Routers and EX Series switches, we highly recommend that you configure the default option when you configure a behavior aggregate (BA) classifier that does not include a specific rewrite rule for MPLS packets. Doing so ensures that MPLS exp value is rewritten according to the BA classifier rules configured for forwarding or packet loss priority. This does not apply to ACX Series Universal Access routers.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Rewriting the EXP Bits of All Three Labels of an Outgoing Packet</i> • dscp (Rewrite Rules) on page 1382 • dscp-ipv6 (Class-of-Service) on page 1357 • exp-push-push-push on page 1723 • exp-swap-push-push on page 1724 • ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360 • ieee-802.1ad on page 1362 • inet-precedence on page 1364

- [rewrite-rules \(Definition\) on page 1481](#)

exp-push-push-push

Syntax	exp-push-push-push default;
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For M Series routers and EX Series switches, rewrite the EXP bits of all three labels of an outgoing packet, thereby maintaining CoS of an incoming non-MPLS packet.
Options	default —Apply the default MPLS EXP rewrite table.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Rewriting the EXP Bits of All Three Labels of an Outgoing Packet</i>• dscp (Rewrite Rules) on page 1382• dscp-ipv6 (Class-of-Service) on page 1357• exp on page 1358• exp-swap-push-push on page 1724• ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360• ieee-802.1ad on page 1362• inet-precedence on page 1364• rewrite-rules (Definition) on page 1481

exp-swap-push-push

Syntax	exp-swap-push-push default;
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For M Series routers and EX Series switches, rewrite the EXP bits of all three labels of an outgoing packet, thereby maintaining CoS of an incoming MPLS packet.
Options	default —Apply the default MPLS EXP rewrite table.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Rewriting the EXP Bits of All Three Labels of an Outgoing Packet</i> • dscp (Rewrite Rules) on page 1382 • dscp-ipv6 (Class-of-Service) on page 1357 • exp on page 1358 • exp-push-push-push on page 1723 • ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360 • ieee-802.1ad on page 1362 • inet-precedence on page 1364 • rewrite-rules (Definition) on page 1481

forwarding-class (BA Classifiers)

Syntax	<code>forwarding-class <i>class-name</i> { <code>loss-priority level code-points</code> [<i>aliases</i>] [<i>bit-patterns</i>]; }</code>
Hierarchy Level	[edit class-of-service classifiers <i>type classifier-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define forwarding class name and option values.
Options	<i>class-name</i> —Name of the forwarding class. The remaining statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Defining Classifiers on page 1330• <i>Example: Configuring CoS for a PBB Network</i>

ieee-802.1 (Rewrite Rules on Logical Interface)

Syntax	<code>ieee-802.1</code> (<i>rewrite-name</i> default) vlan-tag (outer outer-and-inner);
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4. vlan-tag statement introduced in Junos OS Release 8.1.
Description	Apply an IEEE-802.1 rewrite rule. For IQ PICs, you can only configure one IEEE 802.1 rewrite rule on a physical port. All logical ports (units) on that physical port should apply the same IEEE 802.1 rewrite rule.
Options	rewrite-name —Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules ieee-802.1] hierarchy level. default —The default mapping.
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Rewrite Rules on page 1694 • <i>Example: Configuring CoS for a PBB Network</i> • dscp (Rewrite Rules) on page 1382 • dscp-ipv6 (Class-of-Service) on page 1357 • exp on page 1358 • exp-push-push-push on page 1723 • exp-swap-push-push on page 1724 • ieee-802.1ad on page 1362 • inet-precedence on page 1364 • rewrite-rules (Definition) on page 1481

import (Rewrite Rules)

Syntax	<code>import (rewrite-name default);</code>
Hierarchy Level	<code>[edit class-of-service rewrite-rules type rewrite-name]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify a default or previously defined rewrite-rules mapping to import.
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the <code>[edit class-of-service rewrite-rules]</code> hierarchy level.</p> <p>default—The default rewrite-rules mapping.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694

inet-precedence

Syntax	<code>inet-precedence (<i>rewrite-name</i> default);</code>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Apply a IPv4 precedence rewrite rule.
Options	<p><i>rewrite-name</i>—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules inet-precedence] hierarchy level.</p> <p>default—The default mapping. By default, IP precedence rewrite rules alter the first three bits on the type of service (ToS) byte while leaving the last three bits unchanged.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Rewrite Rules on page 1694 • dscp (Rewrite Rules) on page 1382 • dscp-ipv6 (Class-of-Service) on page 1357 • exp on page 1358 • exp-push-push-push on page 1723 • exp-swap-push-push on page 1724 • ieee-802.1 (Rewrite Rules on Logical Interface) on page 1360 • ieee-802.1ad on page 1362 • rewrite-rules (Definition) on page 1481

interfaces

```
Syntax  interfaces {
    interface-name {
        classifiers{
            dscp(classifier-name | default) {
            }
            ieee-802.1 (classifier-name | default) vlan-tag (inner | outer | classifier-name);
            inet-precedence (rewrite-name | default);
        }
        input-scheduler-map map-name;
        input-shaping-rate rate;
        irb {
            unit logical-unit-number {
                classifiers {
                    type (classifier-name | default);
                }
                rewrite-rules {
                    dscp (rewrite-name | default);
                    dscp-ipv6 (rewrite-name | default);
                    exp (rewrite-name | default) protocol protocol-types;
                    ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                    inet-precedence (rewrite-name | default);
                }
            }
        }
        member-link-scheduler (replicate | scale);
        rewrite-rules {
            dscp (rewrite-name | default);
            ieee-802.1 (rewrite-name | default) vlan-tag (outer);
            inet-precedence (rewrite-name | default);
        }
        scheduler-map map-name;
        scheduler-map-chassis map-name;
        shaping-rate rate;
        unit logical-unit-number {
            classifiers {
                type (classifier-name | default) family (mpls | inet);
            }
            forwarding-class class-name;
            fragmentation-map map-name;
            input-shaping-rate (percent percentage | rate);
            input-traffic-control-profile profile-name shared-instance instance-name;
            output-traffic-control-profile profile-name shared-instance instance-name;
            per-session-scheduler;
            rewrite-rules {
                dscp (rewrite-name | default);
                dscp-ipv6 (rewrite-name | default);
                exp (rewrite-name | default) protocol protocol-types;
                exp-push-push-push default;
                exp-swap-push-push default;
                ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                inet-precedence (rewrite-name | default);
            }
        }
    }
}
```

```

    }
    scheduler-map map-name;
    shaping-rate rate;
    translation-table (to-dscp-from-dscp | to-dscp-ipv6-from-dscp-ipv6 | to-exp-from-exp
    | to-inet-precedence-from-inet-precedence) table-name;
  }
}
interface-set interface-set-name {
  excess-bandwidth-share;
  internal-node;
  output-traffic-control-profile profile-name;
  output-traffic-control-profile-remaining profile-name;
}
}

```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 7.4.
Interface-set level added in Junos OS Release 8.5.

Description Configure interface-specific CoS properties for incoming packets.



NOTE: The dscp-ipv6 and ieee-802.1ad classifier types are not supported on ACX Series routers. For further information about support on ACX Series routers, see *Understanding CoS CLI Configuration Statements on ACX Series Universal Access Routers*.

Options The remaining statements are explained separately.

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

Related Documentation

- [Overview of BA Classifier Types on page 1321](#)
- [Configuring Rewrite Rules on page 1694](#)
- [Understanding CoS CLI Configuration Statements on ACX Series Universal Access Routers](#)

loss-priority (BA Classifiers)

Syntax	loss-priority <i>level</i> ;
Hierarchy Level	[edit class-of-service classifiers <i>type classifier-name</i> forwarding-class <i>class-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify packet loss priority value for a specific set of code-point aliases and bit patterns.
Options	<i>level</i> can be one of the following: <ul style="list-style-type: none">• high—Packet has high loss priority.• medium-high—Packet has medium-high loss priority.• medium-low—Packet has medium-low loss priority.• low—Packet has low loss priority.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Overview of BA Classifier Types on page 1321• Example: Configuring CoS for a PBB Network• Configuring Tricolor Marking on page 1447

protocol (Rewrite Rules)

Syntax	<code>protocol protocol-types;</code>
Hierarchy Level	<p>[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules exp <i>rewrite-name</i>],</p> <p>[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules dscp <i>rewrite-name</i>],</p> <p>[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules inet-prec <i>rewrite-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Option for dscp and inet-prec introduced in Junos OS Release 8.4.</p>
Description	Apply a rewrite rule to MPLS packets only, and write the CoS value to MPLS headers only; or apply a rewrite rule to MPLS and IPv4 packets, and write the CoS value to MPLS and IPv4 headers.
Options	<p>protocol-types can be one of the following:</p> <ul style="list-style-type: none"> • mpls—Apply a rewrite rule to MPLS packets and write the CoS value to MPLS headers. • mpls-inet-both—Apply a rewrite rule to VPN MPLS packets with IPv4 payloads. On M120, M320, MX Series, T Series routers (except T4000 routers), and EX Series switches, write the CoS value to the MPLS and IPv4 headers. On M Series routers, initialize all ingress MPLS LSP packets with IPv4 payloads with 000 code points for the MPLS EXP value, and the configured rewrite code point for IP precedence. • mpls-inet-both-non-vpn—Apply a rewrite rule to non-VPN MPLS packets with IPv4 payloads. On M120, M320, MX Series, T Series routers, and EX Series switches write the CoS value to the MPLS and IPv4 headers. On M Series routers, initialize all ingress MPLS LSP packets with IPv4 payloads with 000 code points for the MPLS EXP value, and the configured rewrite code point for IP precedence.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Rewriting MPLS and IPv4 Packet Headers</i>

rewrite-rules (Definition)

Syntax	<pre>rewrite-rules { type <i>rewrite-name</i>{ import (<i>rewrite-name</i> default); forwarding-class <i>class-name</i> { loss-priority <i>level</i> <i>code-point</i> [<i>aliases</i>] [<i>bit-patterns</i>]; } } }</pre>
Hierarchy Level	[edit class-of-service]
Release Information	Statement introduced before Junos OS Release 7.4. <i>ieee-802.1ad</i> option introduced in Junos OS Release 9.2.
Description	Specify a rewrite-rules mapping for the traffic that passes through all queues on the interface.
Options	<p><i>rewrite-name</i>—Name of a <i>rewrite-rules</i> mapping.</p> <p><i>type</i>—Traffic type.</p> <p>Values: <i>dscp</i>, <i>dscp-ipv6</i>, <i>exp</i>, <i>frame-relay-de</i> (J Series routers only), <i>ieee-802.1</i>, <i>ieee-802.1ad</i>, <i>inet-precedence</i></p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<i>interface</i> —To view this statement in the configuration. <i>interface-control</i> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Rewrite Rules on page 1694• <i>Example: Configuring CoS for a PBB Network</i>• J Series router documentation

rewrite-rules (Interfaces)

Syntax	<pre>rewrite-rules { dscp (rewrite-name default); dscp-ipv6 (rewrite-name default); exp (rewrite-name default) protocol protocol-types; exp-push-push-push default; exp-swap-push-push default; ieee-802.1 (rewrite-name default) vlan-tag (outer outer-and-inner); ieee-802.1ad (rewrite-name default) vlan-tag (outer outer-and-inner); inet-precedence (rewrite-name default); }</pre>
Hierarchy Level	<p>[edit class-of-service interfaces <i>interface-name</i>], [edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Associate a rewrite-rules configuration or default mapping with a specific interface.</p> <p>The [edit class-of-service interfaces <i>interface-name</i>] hierarchy level is not supported on M Series routers.</p> <p>The [edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level is not supported on ACX Series routers.</p> <p>On an MX Series router and on an EX Series switch, exp-push-push-push, exp-swap-push-push, and frame-relay-de are not supported on an integrated routing and bridging (IRB) interface.</p> <p>On an ACX Series router, only the outer tag is supported for dscp, inet-precedence, and ieee802.1.</p>
Options	<p>rewrite-name—Name of a rewrite-rules mapping configured at the [edit class-of-service rewrite-rules] hierarchy level.</p> <p>default—The default mapping.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Rewrite Rules on page 1694

unit

Syntax	<pre> unit <i>logical-unit-number</i> { classifiers { type (<i>classifier-name</i> default) family (mpls all); } forwarding-class <i>class-name</i>; fragmentation-map <i>map-name</i>; input-traffic-control-profile <i>profile-name</i> shared-instance <i>instance-name</i>; output-traffic-control-profile <i>profile-name</i> shared-instance <i>instance-name</i>; per-session-scheduler; rewrite-rules { dscp (<i>rewrite-name</i> default); dscp-ipv6 (<i>rewrite-name</i> default); exp (<i>rewrite-name</i> default) <i>protocol</i> <i>protocol-types</i>; exp-push-push-push default; exp-swap-push-push default; ieee-802.1 (<i>rewrite-name</i> default) <i>vlan-tag</i> (outer outer-and-inner); inet-precedence (<i>rewrite-name</i> default); } scheduler-map <i>map-name</i>; shaping-rate <i>rate</i>; } </pre>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.
Options	<p><i>logical-unit-number</i>—Number of the logical unit.</p> <p>Range: 0 through 16,384</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Overview of BA Classifier Types on page 1321 • Configuring Rewrite Rules on page 1694

[edit interfaces] Hierarchy Level

The following statement hierarchy can also be included at the [edit logical-systems *logical-system-name*] hierarchy level.

```

interfaces {
  interface-name {
    ... the "interface-name" subhierarchy appears after the main [edit interfaces] hierarchy level ...
  }
}

```

```

interface-set interface-set-name {
  interface interface-name {
    (unit unit-number | vlan-tags-outer vlan-tag);
  }
}
irb (Interfaces) {
  accounting-profile name;
  description text;
  disable;

  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  hold-time up milliseconds down milliseconds;
  mtu bytes;
  no-gratuitous-arp-request;

  traceoptions {
    flag flag;
  }
  (traps | no-traps);
  unit logical-unit-number {
    accounting-profile name;
    bandwidth rate;
    description text;
    disable;
    encapsulation type;
    family inet {
      accounting {
        destination-class-usage;
        source-class-usage {
          input;
          output;
        }
      }
    }
    address ipv4-address {
      arp ip-address (mac | multicast-mac) mac-address <publish>;
      broadcast address;
      preferred;
      primary;
      vrrp-group group-id {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        (preempt | no-preempt) {
          hold-time seconds;
        }
        priority number;
        track {
          interface interface-name {
            bandwidth-threshold bits-per-second priority-cost priority;
            priority-cost priority;
          }
          priority-hold-time seconds;
        }
      }
    }
  }
}

```

```

        route prefix/prefix-length routing-instance instance-name priority-cost priority;
    }
    virtual-address [ addresses ];
    vrrp-inherit-from vrrp-group;
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
    }
    priority number;
    track {
        interface interface-name {
            bandwidth-threshold bandwidth priority-cost number;
            priority-cost number;
        }
    }
    priority-hold-time seconds;
    route ip-address/mask routing-instance instance-name priority-cost cost;
}

```

```

    }
    virtual-inet6-address [addresses];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
    }
}
}
(dad-disable | no-dad-disable);
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
no-redirects;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    filter {
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    policer {
        input policer-name;
        output policer-name;
    }
}
native-inner-vlan-id vlan-id;
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id-list [vlan-id's];
vlan-id-range [vlan-id-range];
}
}
traceoptions {
    file <filename> <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag <disable>;
    no-remote-trace;
}

```

```
    }
  }

interfaces {
  interface-name {
    disable;
    accounting-profile name;
    aggregated-ether-options {
      ethernet-switch-profile {
        tag-protocol-id [ hexadecimal-identifiers ];
      }
      (flow-control | no-flow-control);
      lacp {
        (active | passive);
        admin-key key;
        fast-failover;
        link-protection {
          disable;
          (revertive | non-revertive);
        }
        periodic (fast | slow);
        system-id mac-address;
        system-priority priority;
      }
      (link-protection | no-link-protection);
      link-speed (100m | 1g | 8g | 10g | 40g | 50g | 80g | 100g | oc192);
      logical-interface-fpc-redundancy;
      (loopback | no-loopback);
      mc-ae {
        chassis-id chassis-id;
        events {
          iccp-peer-down {
            force-icl-down;
            prefer-status-control-active;
          }
        }
        mc-ae-id mc-ae-id;
        mode (active-active | active-standby);
        redundancy-group group-id;
        status-control (active | standby);
      }
      minimum-links number;
      rebalance-periodic {
        start-time time;
        interval number;
      }
      source-address-filter {
        mac-address;
      }
      (source-filtering | no-source-filtering);
    }
  }
  auto-configure {
    remove-when-no-subscribers;
    stacked-vlan-ranges {
      access-profile profile-name;
      authentication {
```



```

password password-string;
username-include {
    circuit-type;
    delimiter delimiter-character;
    domain-name domain-name-string;
    interface-name;
    mac-address;
    option-82 ( circuit-id | remote-id);
    radius-realm radius-realm-string;
    user-prefix user-prefix-string;
}
}
dynamic-profile profile-name {
    accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
    ranges (any | low-tag-high-tag), (any | low-tag-high-tag);
}
}
vlan-ranges {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-name;
            mac-address;
            option-82;
            radius-realm radius-realm-string;
            user-prefix user-prefix-string;
        }
    }
    dynamic-profile profile-name {
        accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
        ranges (any | low-tag)—(any | high-tag);
    }
}
override tag vlan-tag dynamic-profile profile name;
}
encapsulation (ethernet-bridge | ethernet-vpls | extended-vlan-bridge |
    extended-vlan-vpls | flexible-ethernet-services | vlan-vpls);
ether-options {
    802.3ad {
        aex;
        (backup | primary);
        lacp {
            force-up;
            port-priority
        }
    }
}
asynchronous-notification;
(auto-negotiation | no-auto-negotiation);
ethernet-switch-profile {
    ethernet-policer-profile {
        input-priority-map {
            ieee802.1p premium [ values ];

```

```

    }
    output-priority-map {
        classifier {
            premium {
                forwarding-class class-name {
                    loss-priority (high | low);
                }
            }
        }
    }
    policer cos-policer-name {
        aggregate {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
        premium {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
    }
    tag-protocol-id;
}
(mac-learn-enable | no-mac-learn-enable);
}
(flow-control | no-flow-control);
ignore-l3-incompletes;
link-mode (automatic | full-duplex | half-duplex);
(lloopback | no-loopback);
keepalives <interval seconds> <down-count number> <up-count number>;
speed (1g | 10m | 100m | 10m-100m | auto-negotiation);
source-address-filter {
    mac-address;
}
source-filtering | no-source-filtering;
}
flexible-vlan-tagging;
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time (up milliseconds | down milliseconds);
interface-transmit-statistics;
(keepalives <down-count number> <interval seconds> <up-count number> |
no-keepalives);
layer2-policer {
    apply-groups [ group-names ];
    apply-groups-except [ group-names ];
}
link-mode (automatic | full-duplex);
mac mac-address;
mtu bytes;
multi-chassis-protection peer-ip-address {
    interface interface-name;
}
native-vlan-id number;
no-gratuitous-arp-request;
optics-options {
    alarm low-light-alarm {
        (link-down | syslog);
    }
}

```

```

    }
    warning low-light-warning {
        (link-down | syslog);
    }
    wavelength nm;
}
passive-monitor-mode;
per-unit-scheduler;
speed (10m | 100m | 1g | auto | oc3 | oc12 | oc48);
stacked-vlan-tagging;
traceoptions {
    flag flag;
}
transmit-bucket {
    overflow discard;
    rate percentage;
    threshold bytes;
}
(traps | no-traps);
unidirectional;
vlan-tagging;
}

interface-name {
    unit logical-unit-number {
        disable;
        accept-source-mac {
            mac-address mac-address {
                policer {
                    input policer-name;
                    output policer-name;
                }
            }
        }
    }
    account-layer2-overhead (Interface Level) {
        value;
        egress bytes;
        ingress bytes;
    }
    accounting-profile name;
    advisory-options {
        downstream-rate rate;
        upstream-rate rate;
    }
    arp-resp (restricted|unrestricted);
    bandwidth rate;
    clear-dont-fragment-bit;
    copy-tos-to-outer-ip-header;
    demux-destination family;
    encapsulation (vlan-bridge | vlan-vpls);
    epd-threshold cells plp1 cells;
    filter filter-name;
    inner-vlan-id-range start start-id end end-id;
    input-vlan-map {
        (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    }
}

```

```

    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
interface-shared-with psd numerical-index;
layer2-policer {
    input-hierarchical-policer policer-name;
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
}
multi-chassis-protection peer-ip-address {
    interface interface-name;
}
native-inner-vlan-id number;
output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
peer-interface interface-name;
peer-unit unit-number;
plp-to-clp;
proxy-arp <restricted | unrestricted>;
rpm {
    (client | server);
    twamp-server;
}
swap-by-poppush;
vlan-id number;
vlan-id-list [ vlan-id vlan-id-vlan-id ];
vlan-id-range number-number;
vlan-tags (inner <tpid.>vlan-id | inner-list [ vlan-id vlan-id-vlan-id ] |
    inner-range <tpid.>vlan-id-vlan-id) outer <tpid.>vlan-id;
}

unit logical-unit-number {
    family ethernet-switching {
        filter {
            group filter-group-number;
            (input filter-name | input-list [ filter-names ]);
            (output filter-name | output-list [ filter-names ]);
            (inner-vlan-id-list [ vlan-ids ] | vlan-id number | vlan-id-list [ number
                number-number ]);
            interface-mode (access | trunk);
            policer {
                input policer-name;
                output policer-name;
            }
            vlan-rewrite {
                translate old-vlan-id new-vlan-id;
            }
        }
    }
}

```

```

    vlan {
        members [ all vlan-identifiers ];
    }
}
family inet {
    filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
    input-hierarchical-policer policer-name;
    mac-validate (loose | strict);
    mtu bytes;
    no-neighbor-learn;
    no-redirects;
    policer {
        arp policer-template-name;
        input policer-name;
        output policer-name;
    }
    primary;
    receive-options-packets;
    receive-ttl-exceeded;
    rpf-check {
        fail-filter filter-name;
        mode loose;
    }
    sampling {
        (input | output | input output);
    }
    simple-filter {
        input filter-name;
    }
    targeted-broadcast {
        forward-and-send-to-re;
        forward-only;
    }
    unnumbered-address interface-name <destination address>
        <destination-profile profile-name> <preferred-source-address address>;
}

family inet6 {
    address ipv6-address {
        destination destination-address;
        eui-64;
        ndp ipv6-address <l2-interface interface-name> <(mac mac-address |
            multicast-mac multicast-mac-address) <publish>>;
        preferred;
        primary;
        vrrp-inet6-group group-number {
            (accept-data | no-accept-data);
            fast-interval milliseconds;
            inet6-advertise-interval seconds;
            (no-preempt; | ... the following preempt statement ...)
            preempt {

```

```

        hold-time seconds;
    }
    priority number;
    track {
        interface interface-name {
            bandwidth-threshold bits-per-second priority-cost priority;
            priority-cost priority;
        }
        priority-hold-time seconds;
        route ip-address-prefix/prefix-length routing-instance instance-name
            priority-cost priority;
    }
    virtual-inet6-address [ addresses ];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
    }
}
(dad-disable | no-dad-disable);
filter {
    group filter-group-number;
    (input filter-name | input-list [ filter-names ]);
    (output filter-name | output-list [ filter-names ]);
}
input-hierarchical-policer policer-name;
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode loose;
}
sampling {
    (input | output | input output);
}
unnumbered-address interface-name preferred-source-address address;
}

family iso {
    address iso-address;
    mtu bytes;
}

family mlfrr-end-to-end {
    bundle logical-interface-name;
}

```

```

family mpls {
  filter {
    group filter-group-number;
    (input filter-name | input-list [ filter-names ]);
    (output filter-name | output-list [ filter-names ]);
  }
  input-hierarchical-policer policer-name;
  maximum-labels maximum-labels;
  mtu bytes;
  policer {
    input policer-name;
    output policer-name;
  }
}

family vpls {
  core-facing;
  filter {
    group filter-group-number;
    (input filter-name | input-list [ filter-names ]);
    (output filter-name | output-list [ filter-names ]);
  }
  policer {
    input policer-name;
    output policer-name;
  }
}
}
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

copy-tos-to-outer-ip-header

Syntax	copy-tos-to-outer-ip-header;
Hierarchy Level	[edit interfaces at- <i>fpc/pic/port</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces at- <i>fpc/pic/port</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 8.2.
Description	For GRE tunnel interfaces only, enables the inner IP header's ToS bits to be copied to the outer IP packet header.
Default	If you omit this statement, the ToS bits in the outer IP header are set to 0.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header on page 1770

CHAPTER 7

Device Security

- Overview on page 1803
- Configuration on page 1804

Overview

- Rate Limiting on page 1803

Rate Limiting

- Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages on page 1803
- Configuring the Junos OS ICMPv6 Rate Limit for ICMPv6 Routing Engine Messages on page 1803

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

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

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

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

Related Documentation

- Configuring the Junos OS ICMPv6 Rate Limit for ICMPv6 Routing Engine Messages on page 1803

Configuring the Junos OS ICMPv6 Rate Limit for ICMPv6 Routing Engine Messages

To limit the rate at which ICMPv6 messages are sent, include the **icmpv6-rate-limit** statement at the **[edit system internet-options]** hierarchy level:

icmpv6-rate-limit bucket-size *bucket-size* packet-rate *packet-rate*;

The bucket size is the the number of seconds in the rate-limiting bucket. The packet rate is the rate-limiting packets earned per second. Specify a **bucket-size** from 0 through

4294967295 seconds. The default value is 5 seconds. Specify a **packet-rate** from 0 through 4294967295. The default value is 1000.

**Related
Documentation**

- [Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages on page 1803](#)

Configuration

- [Configuration Statements on page 1804](#)

Configuration Statements

- [\[edit system\] Hierarchy Level on page 1804](#)

[edit system] Hierarchy Level

```
system {
  accounting {
    destination {
      radius {
        server {
          server-address {
            accounting-port port-number;
            max-outstanding-requests
            port port-number;
            retry number;
            secret password;
            source-address address;
            timeout seconds;
          }
        }
      }
    }
    tacplus {
      server {
        server-address {
          port port-number;
          secret password;
          single-connection;
          source-address address;
          timeout seconds;
        }
      }
    }
  }
  events [ change-log interactive-commands login ];
}
allow-6pe-traceroute;
allow-v4mapped-packets;
archival {
  configuration {
    archive-sites {
      ftp://<username>:<password>@<host>:<port>/<url-path>;
      scp://<username>:<password>@<host>:<port>/<url-path>;
    }
  }
}
```

```

        transfer-interval interval;
        transfer-on-commit;
    }
}
arp {
    aging-timer minutes;
    gratuitous-arp-delay;
    gratuitous-arp-on-ifup;
    interfaces {
        logical-interface-name {
            aging-timer minutes;
        }
    }
    passive-learning;
    purging;
}
authentication-order [ authentication-methods ];
auto-configuration {
    traceoptions {
        file <filename> <files number> <match regular-expression> <size size>
            <world-readable | no-world-readable>;
        flag <all | auth | configuration | ;interfaces | io | rtsock | ui>
        level level;
        no-remote-trace;
    }
}
backup-router address <destination [ destination-addresses ]>;
commit {
    fast-synchronize;
    synchronize;
    server {
        commit-interval number;
        days-to-keep-error-logs number;
        maximum-aggregate-pool number;
        maximum-entries number;
        traceoptions {
            file <filename> <files number> <match regular-expression> <size size>
                <world-readable | no-world-readable>;
            flag <all | auth | configuration | ;interfaces | io | rtsock | ui>
            level level;
            no-remote-trace;
        }
    }
}
(compress-configuration-files | no-compress-configuration-files);
ddos-protection {
    global {
        disable-fpc;
        disable-logging;
        disable-routing-engine;
        flow-detection;
        flow-report-rate;
        violation-report-rate;
    }
    protocols protocol-group (aggregate | packet-type) {
        bandwidth packets-per-second;
    }
}

```

```
burst size;
disable-fpc;
disable-logging;
disable-routing-engine;
fpc {
    bandwidth-scale percentage;
    burst-scale percentage;
    disable-fpc;
}
priority level;
recover-time seconds;
flow-detection {
    flow-detect-time detect-period;
    no-flow-logging;
    timeout-active-flows enable-period;
    flow-level-bandwidth;
    flow-level-control (all | keep-all | police);
    flow-detection-mode (always-on | automatic | disabled);
    physical-interface;
    flow-recover-time recover-period;
    flow-timeout-time timeout-period;
    subscriber;
}
}
traceoptions{
    file filename <files number> <match regular-expression > <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag;
    level (all | error | info | notice | verbose | warning);
    no-remote-trace;
}
}
default-address-selection;
diag-port-authentication (encrypted-password "password" | plain-text-password);
dynamic-profile-options {
    versioning;
}
domain-name domain-name;
domain-search [ domain-list ];
donot-disable-ip6op-ondad;
extensions {
    providers {
        provider-id {
            license-type license deployment-scope [ deployments ];
        }
    }
}
resource-limits {
    package package-name {
        resources {
            cpu {
                priority number;
                time seconds;
            }
            file {
                core-size bytes;
                open number;
            }
        }
    }
}
```

```

        size bytes;
    }
    memory {
        data-size bytes;
        locked-in bytes;
        resident-set-size bytes;
        socket-buffers bytes;
        stack-size bytes;
    }
}
}
process process-ui-name {
    resources {
        cpu {
            priority number;
            time seconds;
        }
        file {
            core-size bytes;
            open number;
            size bytes;
        }
        memory {
            data-size bytes;
            locked-in bytes;
            resident-set-size bytes;
            socket-buffers bytes;
            stack-size bytes;
        }
    }
}
}
}
fips {
    level level;
}
host-name hostname;
inet6-backup-router ipv6-address <destination address>;
internet-options {
    (gre-path-mtu-discovery | no-gre-path-mtu-discovery);
    icmpv4-rate-limit bucket-size number packet-rate rate;
    icmpv6-rate-limit bucket-size number packet-rate rate;
    (ipip-path-mtu-discovery | no-ipip-path-mtu-discovery);
    (ipv6-path-mtu-discovery | noipv6-path-mtu-discovery);
    ipv6-path-mtu-discovery-timeout;
    no-tcp-rfc1323-paws;
    no-tcp-rfc1323;
    (path-mtu-discovery | no-path-mtu-discovery);
    source-port upper-limit port-number;
    (source-quench | no-source-quench);
    tcp-drop-synfin-set;
}
kernel-replication;
license {
    autoupdate {
        url URL;
    }
}

```

```
    password password;
  }
  renew before-expiration number;
  interval number
  traceoptions {
    file <filename> <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
location {
  altitude feet;
  building name;
  country-code code;
  floor number;
  hcoord horizontal-coordinate;
  lata service-area;
  latitude degrees;
  longitude degrees;
  npa-nxx number;
  postal-code postal-code;
  rack number;
  vcoord vertical-coordinate;
}
login {
  announcement "text";
  class class-name {
    access-end "hh<:mm:<ss>>";
    access-start "hh<:mm:<ss>>";
    allow-commands "regular-expression";
    ( allow-configuration | allow-configuration-regexps ) "regular expression 1" "regular
      expression 2";
    allowed-days [ sunday monday tuesday wednesday thursday friday saturday ];
    configuration-breadcrumbs;
    deny-commands "regular-expression";
    ( deny-configuration | deny-configuration-regexps ) "regular expression 1" "regular
      expression 2";
    idle-timeout minutes;
    logical-system logical-system-name;
    login-alarms;
    login-script filename;
    login-tip;
    permissions [ permissions ];
    security-role [ security-role ] ;
  }
  deny-sources (address address | apply-groups | apply-groups-except) ;
  message "text";
  password {
    change-type (character-sets | set-transitions);
    format (des | md5 | sha1);
    maximum-length length;
    minimum-changes number;
    minimum-length length;
    minimum-lower-cases number;
    minimum-numeric number;
```

```

        minimum-punctuations number;
        minimum-upper-cases number;
    }
    retry-options {
        backoff-factor number;
        backoff-threshold number;
        maximum-time number;
        minimum-time number;
        tries-before-disconnect number;
    }
    user username {
        authentication {
            (encrypted-password "password" | plain-text-password);
            load-key-file filename;
            ssh-dsa "public-key" <from hostname>;
            ssh-ecdsa "public-key" <from hostname>;
            ssh-rsa "public-key" <from hostname>;
        }
        class class-name;
        full-name "complete-name";
        uid uid-value;
    }
}
max-configurations-on-flash number;
mirror-flash-on-disk;
name-server {
    address;
}
nd-maxmcast-solicit
nd-retransmit-timer
no-multicast-echo;
no-neighbor-learn;;
no-ping-record-route;
no-ping-time-stamp;
no-redirects;
no-redirects-ipv6;
ntp {
    authentication-key key-number type md5 value password;
    boot-server address;
    broadcast <address> <key key-number> <ttl value> <version value>;
    broadcast-client;
    multicast-client <address>;
    peer address <key key-number> <prefer> <version value>;
    server address <key key-number> <prefer> <version value>;
    source-address source-address;
    trusted-key [ key-numbers ];
}
pic-console-authentication {
    (encrypted-password "encrypted-password" | plain-text-password);
}
ports {
    auxiliary {
        disable;
        insecure;
        type (ansi | small-xterm | vt100 | xterm);
        port-type (mini-usb | rj45) ;
    }
}

```

```
    }
  }
  console {
    disable;
    insecure;
    log-out-on-disconnect;
    type (ansi | small-xterm | vt100 | xterm);
  }
}
processes {
  process-name (enable | disable) failover (alternate-media | other-routing-engine);
  command path;
  timeout seconds;
}
proxy {
  password password;
  port port-number;
  server (hostname | ip-address);
  username username;
}
radius-options {
  attributes {
    nas-ip-address address;
  }
  password-protocol mschap-v2;
}
radius-server {
  server-address {
    accounting-port port-number;
    max-outstanding-requests number;
    port port-number;
    retry number;
    secret password;
    source-address source-address;
    timeout seconds;
  }
}
root-authentication {
  (encrypted-password "password" | plain-text-password);
  load-key-file filename;
  ssh-dsa "public-key" <from hostname>;
  ssh-ecdsa "public-key" <from hostname>;
  ssh-rsa "public-key" <from hostname>;
}
(saved-core-context | no-saved-core-context);
saved-core-files number;
scripts {
  load-scripts-from-flash;
  commit {
    allow-transients;
    direct-access;
    file filename.xml {
      checksum (md5 | sha-256 | sha1) hash;
      optional;
      refresh;
      refresh-from url;
    }
  }
}
```



```

        source url;
    }
    max-datasize
    refresh;
    refresh-from url;
    traceoptions {
        file <filename> <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
op {
    file filename.xml {
        arguments {
            argument-name {
                description descriptive-text;
            }
        }
        checksum (md5 | sha-256 | sha1) hash;
        command filename-alias;
        description descriptive-text;
        refresh;
        refresh-from url;
        source url;
    }
    max-datasize
    no-allow-url
    refresh;
    refresh-from url;
    traceoptions {
        file <filename> <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
}
static-host-mapping {
    hostname {
        alias [ aliases ];
        inet [ addresses ];
        inet6 [ addresses ];
        sysid system-identifier;
    }
}
syslog {
    allow-duplicates;
    archive <binary-data | no-binary-data> <files number> <size size> <world-readable |
    no-world-readable>;
    console {
        any | authorization | change-log | conflict-log | daemon | dfc | external | firewall | ftp
        | interactive-commands | kernel | ntp | pfe | security | user (alert | any | critical |
        emergency | error | info | none | notice | warning);
    }
    file filename {

```

```

    facility severity;
    allow-duplicates;
    any (alert | any | critical | emergency | error | info | none | notice | warning);
    archive <archive-sites {ftp-url <password password>} > <files number> <size size>
        <start-time "YYYY-MM-DD.hh:mm"> <transfer-interval minutes> <world-readable |
        no-world-readable>;
    authorization (alert | any | critical | emergency | error | info | none | notice | warning);
    change-log (alert | any | critical | emergency | error | info | none | notice | warning);
    conflict-log (alert | any | critical | emergency | error | info | none | notice | warning);
    daemon (alert | any | critical | emergency | error | info | none | notice | warning);
    dfc (alert | any | critical | emergency | error | info | none | notice | warning);
    explicit-priority;
    external (alert | any | critical | emergency | error | info | none | notice | warning);
    firewall (alert | any | critical | emergency | error | info | none | notice | warning);
    ftp (alert | any | critical | emergency | error | info | none | notice | warning);
    interactive-commands (alert | any | critical | emergency | error | info | none | notice
        | warning);
    kernel (alert | any | critical | emergency | error | info | none | notice | warning);
    match "regular-expression";
    ntp (alert | any | critical | emergency | error | info | none | notice | warning);
    pfe (alert | any | critical | emergency | error | info | none | notice | warning);
    security (alert | any | critical | emergency | error | info | none | notice | warning);
    structured-data {
        brief
    }
}
host (hostname | other-routing-engine | scc-master) {
    facility severity;
    authorization (alert | any | critical | emergency | error | info | none | notice | warning);
    change-log (alert | any | critical | emergency | error | info | none | notice | warning);
    conflict-log (alert | any | critical | emergency | error | info | none | notice | warning);
    daemon (alert | any | critical | emergency | error | info | none | notice | warning);
    dfc (alert | any | critical | emergency | error | info | none | notice | warning);
    explicit-priority;
    external (alert | any | critical | emergency | error | info | none | notice | warning);
    facility-override facility;
    firewall (alert | any | critical | emergency | error | info | none | notice | warning);
    ftp (alert | any | critical | emergency | error | info | none | notice | warning);
    interactive-commands (alert | any | critical | emergency | error | info | none | notice
        | warning);
    kernel (alert | any | critical | emergency | error | info | none | notice | warning);
    log-prefix string;
    match "regular-expression";
    ntp (alert | any | critical | emergency | error | info | none | notice | warning);
    pfe (alert | any | critical | emergency | error | info | none | notice | warning);
    security (alert | any | critical | emergency | error | info | none | notice | warning);
    source-address source-address;
    structured-data {
        brief
        user (username | *) {
    }
}
log-rotate-frequency minutes;
server;
source-address address;
time-format (year | millisecond | year millisecond);
user (username | *) {
    facility severity;

```

```

        match "regular-expression";
    }
}
tacplus-options {
    (exclude-cmd-attribute | no-cmd-attribute-value);
    service-name service-name;
}
tacplus-server {
    server-address {
        port port-number;
        secret password;
        single-connection;
        source-address source-address;
        timeout seconds;
    }
}
time-zone (GMT | GMT+hour-offset | GMT-hour-offset | zone-name);
tracing destination-override syslog host address;
use-imported-time-zones;
}
}
system {
    services {
        database-replication {
            traceoptions {
                file <filename> <files number> <match regular-expression>
                <size maximum-file-size> <world-readable | no-world-readable>;
                flag flag;
                no-remote-trace;
            }
        }
    }
    dhcp-local-server {
        authentication {
            password password;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name;
                logical-system-name;
                mac-address;
                option-60;
                option-82 <circuit-id> <remote-id>;
                routing-instance-name;
                user-prefix user-prefix;
            }
        }
    }
    duplicate-clients-on-interface;
    dynamic-profile (profile-name | junos-default-profile) <aggregate-clients <merge |
    replace> | use-primary primary-profile-name>;
    forward-snooped-clients (all-interfaces | configured-interfaces |
    non-configured-interfaces);
    group group-name {
        dynamic-profile (profile-name | junos-default-profile) <aggregate-clients <merge |
        replace> | use-primary primary-profile-name>;
        interface interface-name {
            exclude;

```

```

    overrides {
        ...same statements as at the [edit system services dhcp-local-server overrides]
        hierarchy level ...
    }
    trace;
    upto upto-interface-name;
}
}
overrides {
    client-discover-match <option60-and-option82>;
    interface-client-limit number;
    no-arp;
    process-inform {
        pool pool-name;
    }
}
pool-match-order {
    external-authority;
    ip-address-first;
    option-82;
}
reconfigure {
    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {
        radius-disconnect;
    }
}
traceoptions {
    file <filename> <files number> <match regular-expression>
    <size maximum-file-size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
}
dhcpv4-profiles profile-name {
    bind-interface interface-name;
    dead-server-retry-interval interval-in-seconds;
    dead-server-successive-retry-attempt number-of-attempts;
    dhcp-server-selection-algorithm (highest-priority-server | round-robin);
    lease-time time-in-seconds;
    pool-name pool-name;
    retransmission-attempt number-of-attempts;
    retransmission-interval interval-in-seconds;
    servers ip-address {
        priority value;
    }
}
}
dhcpv6-profiles profile-name {
    bind-interface interface-name;
    lease-time time-in-seconds;
    pool-name pool-name;
    retransmission-attempt number-of-attempts;
}

```

```

        retransmission-interval interval-in-seconds;
    }
    traceoptions {
        file <filename> <files number> <match regular-expression>
            <size maximum-file-size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
finger {
    connection-limit limit;
    rate-limit limit;
}
flow-tap-dtcp {
    ssh {
        connection-limit limit;
        rate-limit limit;
    }
}
ftp {
    connection-limit limit;
    rate-limit limit;
}
local-policy-decision-function {
    statistics {
        aacl-statistics-profile profile-name {
            aacl-fields {
                address;
                all-fields;
                application;
                application-group;
                input-bytes;
                input-interface;
                input-packets;
                ipv6-address
                ipv6-prefix-length
                mask;
                output-bytes;
                output-packets;
                subscriber-name;
                timestamp;
                vrf-name;
            }
            file filename;
            record-type (delta | interim);
        }
    }
    file filename {
        archive-sites {
            url;
        }
        files number;
        size bytes;
        transfer-interval minutes;
    }
    record-type (data | interim);
}

```

```
    traceoptions {
      file <filename> <files number> <match regular-expression>
        <size maximum-file-size> <world-readable | no-world-readable>;
      flag flag;
      no-remote-trace;
    }
  }
  netconf {
    ssh {
      connection-limit limit;
      port port;
      rate-limit limit;
    }
    traceoptions {
      file <filename> <files number> <match regular-expression> <size size>
        <world-readable | no-world-readable>;
      flag flag;
      no-remote-trace;
      on-demand;
    }
  }
  outbound-ssh {
    client client-id {
      address {
        port port-number;
        retry number;
        timeout seconds;
      }
      device-id device-id;
      keep-alive {
        retry number;
        timeout seconds;
      }
      reconnect-strategy (in-order | sticky);
      secret secret;
      services netconf;
    }
    traceoptions {
      file <filename> <files number> <match regular-expression>
        <size maximum-file-size> <world-readable | no-world-readable>;
      flag flag;
      no-remote-trace;
    }
  }
  resource-monitor {
    resource-category jtree {
      resource-type free-dwords {
        low-watermark number;
        high-watermark number;
      }
      resource-type free-pages {
        low-watermark number;
        high-watermark number;
      }
    }
  }
  no-throttle;
```

```

no-logging;
high-threshold number;
traceoptions {
  file filename <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
}
service-deployment {
  local-certificate certificate-name;
  servers {
    server-address {
      port port-number;
      security-options {
        (ssl3 | tls);
      }
      user username;
    }
  }
  source-address source-address;
  traceoptions {
    flag flag;
  }
}
ssh {
  ciphers [ cipher-1 cipher-2 cipher-3 ... ]
  client-alive-count-max seconds;
  client-alive-interval seconds;
  connection-limit limit;
  hostkey-algorithm limit;
  key-exchange limit;
  macs limit;
  max-sessions-per-connection number;
  no-tcp-forwarding;
  protocol-version [v1 v2];
  rate-limit limit;
  root-login (allow | deny | deny-password);
}
subscriber-management {
  enforce-strict-scale-limit-license;
  gres-route-flush-delay;
  maintain-subscriber {
    interface-delete;
  }
  traceoptions {
    file filename <files number> <match regular-expression> <size maximum-file-size>
      <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
}
traceoptions {
  file filename <files number> <match regular-expression> <size maximum-file-size>
    <world-readable | no-world-readable>;
  flag flag;
}

```

```
        no-remote-trace;
    }
    telnet {
        connection-limit limit;
        rate-limit limit;
    }
    tftp-server {
        connection-limit limit;
        rate-limit limit;
    }
    xnm-clear-text {
        connection-limit limit;
        rate-limit limit;
    }
    xnm-ssl {
        connection-limit limit;
        local-certificate certificate-name;
        rate-limit limit;
    }
}
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

icmpv4-rate-limit

Syntax	<pre>icmpv4-rate-limit { bucket-size <i>seconds</i>; packet-rate <i>pps</i>; }</pre>
Hierarchy Level	[edit system internet-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure rate-limiting parameters for ICMPv4 messages sent.
Options	<p>bucket-size <i>seconds</i>—Number of seconds in the rate-limiting bucket. Range: 0 through 4294967295 seconds Default: 5</p> <p>packet-rate <i>pps</i>—Rate-limiting packets earned per second. Range: 0 through 4294967295 pps Default: 1000</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages on page 1803

icmpv6-rate-limit

Syntax	icmpv6-rate-limit { bucket-size <i>seconds</i> ; packet-rate <i>packet-rate</i> ; }
Hierarchy Level	[edit system internet-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure rate-limiting parameters for ICMPv6 messages sent.
Options	<p>bucket-size <i>seconds</i>—Number of seconds in the rate-limiting bucket. Range: 0 through 4294967295 seconds Default: 5</p> <p>packet-rate <i>pps</i>—Rate-limiting packets earned per second. Range: 0 through 4294967295 pps Default: 1000</p>
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Junos OS ICMPv6 Rate Limit for ICMPv6 Routing Engine Messages on page 1803

CHAPTER 8

Ethernet Switching

- [Overview on page 1821](#)
- [Configuration on page 1829](#)
- [Administration on page 1937](#)

Overview

- [Bridging on page 1821](#)
- [MVRP on page 1824](#)
- [Proxy ARP on page 1826](#)

Bridging

- [Layer 2 VLANs Overview on page 1821](#)
- [Layer 2 Learning and Forwarding Overview on page 1822](#)
- [Layer 2 Learning and Forwarding for VLANs Overview on page 1822](#)
- [Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port on page 1823](#)
- [Guidelines for Configuring VLAN Identifiers for VLANs and VPLS Routing Instances on page 1823](#)

Layer 2 VLANs Overview

You can configure one or more VLANs to perform Layer 2 bridging. The Layer 2 bridging functions include integrated routing and bridging (IRB) for support for Layer 2 bridging and Layer 3 IP routing on the same interface, and virtual switches that isolate a LAN segment with its spanning-tree protocol instance and separate its VLAN ID space.

A VLAN is a set of logical ports that share the same flooding or broadcast characteristics and span one or more ports of multiple devices.

You can configure one or more VLANs to perform Layer 2 bridging. Thus, MX Series routers or EX Series switches can function as Layer 2 switches, each with multiple bridging, or broadcast, domains that participate in the same Layer 2 network. You can also configure Layer 3 routing support for a VLAN. Integrated routing and bridging (IRB) provides support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route

packets to another routed interface or to another VLAN that has a Layer 3 protocol configured.

You can also group one or more VLANs within a single instance, or virtual switch. Multiple virtual switches, each of which operates independently of other virtual switches on the device, are supported. Virtual switches isolate a LAN segment with its spanning-tree protocol instance and separate its VLAN ID space. Thus, each virtual switch can participate in a different Layer 2 network.

VLANs provide support for a Layer 2 trunk port. A Layer 2 trunk interface enables you to configure a single logical interface to represent multiple VLANs on a physical interface. You can configure a set of VLANs and VLAN identifiers that are automatically associated with one or more Layer 2 trunk interfaces. Packets received on a trunk interface are forwarded within a VLAN that has the same VLAN identifier. A Layer 2 trunk interface also supports IRB within a VLAN. In addition, you can configure Layer 2 learning and forwarding properties that apply to the entire set of VLANs.

You can configure VPLS ports in a virtual switch instead of a dedicated routing instance of type **vpls** so that the logical interfaces of the Layer 2 VLANs in the virtual switch can handle VPLS routing instance traffic. Packets received on a Layer 2 trunk interface are forwarded within a VLAN that has the same VLAN identifier.

Layer 2 Learning and Forwarding Overview

You can configure Layer 2 MAC address and VLAN learning and forwarding properties in support of Layer 2 bridging. Unicast media access control (MAC) addresses are learned to avoid flooding the packets to all the ports in a VLAN. A source MAC entry is created in its source and destination MAC tables for each MAC address learned from packets received on ports that belong to the VLAN.

By default, Layer 2 address learning is enabled. You can disable MAC learning for a device or for a specific VLAN or logical interfaces. You can also configure the following Layer 2 forwarding properties:

- Timeout interval for MAC entries
- MAC accounting
- A limit to the number of MAC addresses learned from the logical interfaces

For more information about how to configure VLANs and virtual switches, see [“Configuring a VLAN” on page 1829](#) and [“Configuring a Layer 2 Virtual Switch” on page 1847](#).

Layer 2 Learning and Forwarding for VLANs Overview

When you configure a VLAN, Layer 2 address learning is enabled by default. The VLAN learns unicast media access control (MAC) addresses to avoid flooding the packets to all the ports in the VLAN. Each VLAN creates a source MAC entry in its source and destination MAC tables for each source MAC address learned from packets received on the ports that belong to the VLAN.



NOTE: Traffic is not flooded back onto the interface on which it was received. However, because this “split horizon” occurs at a late stage, the packet statistics displayed by commands such as `show interfaces queue` will include flood traffic.

You can optionally disable MAC learning either for the entire device or for a specific VLAN or logical interface. You can also configure the following Layer 2 learning and forwarding properties:

- Static MAC entries for logical interfaces only
- Limit to the number of MAC addresses learned from a specific logical interface or from all the logical interfaces in a VLAN
- Size of the MAC address table for the VLAN
- MAC accounting for a VLAN

Related Documentation

- [Layer 2 Learning and Forwarding Overview on page 1822](#)

Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port

Layer 2 learning is enabled by default. A set of VLANs, configured to function as a switch with a Layer 2 trunk port, learns unicast media access control (MAC) addresses to avoid flooding packets to the trunk port.



NOTE: Traffic is not flooded back onto the interface on which it was received. However, because this “split horizon” occurs at a late stage, the packet statistics displayed by commands such as `show interfaces queue` will include flood traffic.

You can optionally disable Layer 2 learning for the entire set of VLANs as well as modify the following Layer 2 learning and forwarding properties:

- Limit the number of MAC addresses learned from the Layer 2 trunk port associated with the set of VLANs
- Modify the size of the MAC address table for the set of VLANs
- Enable MAC accounting for the set of VLANs

Related Documentation

- [Layer 2 Learning and Forwarding Overview on page 1822](#)

Guidelines for Configuring VLAN Identifiers for VLANs and VPLS Routing Instances

For a VLAN that is performing Layer 2 switching only, you do not have to specify a VLAN identifier.

For a VLAN that is performing Layer 3 IP routing, you must specify either a VLAN identifier or dual VLAN identifier tags.

For a VPLS routing instance, you must specify either a VLAN identifier or dual VLAN identifier tags.

Related Documentation

- [Layer 2 Learning and Forwarding Overview on page 1822](#)

MVRP

- [Understanding Multiple VLAN Registration Protocol \(MVRP\) on page 1824](#)

Understanding Multiple VLAN Registration Protocol (MVRP)

You can configure Multiple VLAN Registration Protocol (MVRP) on Juniper Networks MX Series routers and EX Series switches. The primary purpose of MVRP is to manage dynamic VLAN registration in switching networks. In managing dynamic VLAN registration, MVRP also prunes VLAN information.

MVRP is an Layer 2 application protocol of the Multiple Registration Protocol (MRP) and is defined in the IEEE 802.1ak standard. MRP and MVRP were designed by IEEE to perform the same functions as Generic Attribute Registration Protocol (GARP) and GARP VLAN Registration Protocol (GVRP) while overcoming some GARP and GVRP limitations, in particular limitations involving bandwidth usage and convergence time in large networks with large numbers of VLANs.

MVRP was created by IEEE as a replacement application for GVRP. MVRP and GVRP cannot be run concurrently to share VLAN information in a switching network.

This topic describes:

- [How MVRP Works on page 1824](#)
- [Basics of MVRP on page 1825](#)
- [MVRP Registration Modes on page 1825](#)
- [MRP Timers on page 1825](#)
- [MRP VLAN Messages on page 1826](#)
- [MVRP Limitations on page 1826](#)

How MVRP Works

The VLAN registration information sent by MVRP protocol data units (PDUs) includes the current VLANs membership—that is, which routers are members of which VLANs—and which router interfaces are in which VLAN. MVRP shares all information in the PDU with all routers participating in MVRP in the switching network.

MVRP stays synchronized using these PDUs. The routers in the network participating in MVRP receive these PDUs during state changes and update their MVRP states accordingly. MVRP timers dictate when PDUs can be sent and when routers receiving MVRP PDUs can update their MVRP information.

VLAN information is distributed as part of the MVRP message exchange process and can be used to dynamically create VLANs, which are VLANs created on one switch and propagated to other routers as part of the MVRP message exchange process. Dynamic VLAN creation using MVRP is enabled by default but can be disabled.

As part of ensuring that VLAN membership information is current, MVRP removes routers and interfaces from the VLAN information when they become unavailable. Pruning VLAN information has these benefits:

- Limits the network VLAN configuration to active participants only, reducing network overhead.
- Targets the scope of broadcast, unknown unicast, and multicast (BUM) traffic to interested devices only.

Basics of MVRP

MVRP is disabled by default. You can configure MVRP router interfaces to participate in MVRP for the switching network. MVRP can only be enabled on trunk interfaces, and dynamic VLAN configuration through MVRP is enabled by default when MVRP is enabled.

MVRP Registration Modes

The MVRP registration mode defines whether an interface does or does not participate in MVRP.

The following MVRP registration modes are configurable:

- forbidden—The interface does not register or declare VLANs (except statically configured VLANs).
- normal—The interface accepts MVRP messages and participates in MVRP. This is the default registration mode setting.
- restricted—The interface—The interface ignores all MVRP JOIN messages received for VLANs that are not statically configured on the interface.

MRP Timers

MVRP registration and updates are controlled by timers that are part of the MRP protocol. These timers are set on a per-interface basis and define when MVRP PDUs can be sent and when MVRP information can be updated on a switch.

The following timers are used to control the operation of MVRP:

- Join timer—Controls the interval for the next MVRP PDU transmit opportunity.
- Leave timer—Controls the period of time that an interface on the switch waits in the Leave state before changing to the unregistered state.
- LeaveAll timer—Controls the frequency with which the interface generates LeaveAll messages.



BEST PRACTICE: Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

MRP VLAN Messages

MVRP uses MRP messages to register and declare MVRP states for a switch and to inform the switching network that a switch is leaving MVRP. These messages are communicated as part of the PDU to communicate the state of a particular switch interface on the switching network to the other switches in the network.

The following messages are communicated for MVRP:

- Empty—VLAN information is not being declared and is not registered.
- In—VLAN information is not being declared but is registered.
- JoinEmpty—VLAN information is being declared but not registered.
- JoinIn—VLAN information is being declared and is registered.
- Leave—VLAN information that was previously registered is being withdrawn.
- LeaveAll—All registrations will be de-registered. Participants that want to participate in MVRP will need to re-register.
- New—VLAN information is new and possibly not previously registered.

MVRP Limitations

The following limitations apply when configuring MVRP:

- MVRP works with Rapid Spanning Tree Protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP), but not with VLAN Spanning Tree Protocol (VSTP).
- MVRP is allowed only on single tagged trunk ports.
- MVRP is not allowed if a physical interface has more than one logical interface.
- MVRP is only allowed if a logical has one trunk interface (unit 0).

Related Documentation

- [Example: Configuring Automatic VLAN Administration Using MVRP](#)
- [Configuring Multiple VLAN Registration Protocol \(MVRP\) on page 1845](#)
- [Controlling the Management State of a VLAN in MVRP Configurations \(CLI Procedure\)](#)
- [Verifying That MVRP Is Working Correctly on page 1937](#)

Proxy ARP

- [Restricted and Unrestricted Proxy ARP Overview on page 1827](#)

Restricted and Unrestricted Proxy ARP Overview

By default, the Junos OS responds to an Address Resolution Protocol (ARP) request only if the destination address of the ARP request is local to the incoming interface.

For Ethernet Interfaces, you can configure the router or switches to proxy-reply to the ARP requests using the restricted or unrestricted proxy ARP configuration.

You might want to configure restricted or unrestricted proxy ARP for routers that act as provider edge (PE) devices in Ethernet Layer 2 LAN switching domains.



NOTE: From Junos OS Release 10.0 onward, Junos OS does not respond to proxy ARP requests with the default route 0.0.0.0. This behavior is in compliance with RFC 1027.

Restricted Proxy ARP

Restricted proxy ARP enables the router or switch to respond to the ARP requests in which the physical networks of the source and target are not the same and the router or switch has an active route to the target address in the ARP request. The router does not reply if the target address is on the same subnet and the same interface as the ARP requestor.

Unrestricted Proxy ARP

Unrestricted proxy ARP enables the router or switch to respond to any ARP request, on condition that the router has an active route to the destination address of the ARP request. The route is not limited to the incoming interface of the request, nor is it required to be a direct route.



WARNING: If you configure unrestricted proxy ARP, the proxy router replies to ARP requests for the target IP address on the same interface as the incoming ARP request. This behavior is appropriate for cable modem termination system (CMTS) environments, but might cause Layer 2 reachability problems if you enable unrestricted proxy ARP in other environments.

When an IP client broadcasts the ARP request across the Ethernet wire, the end node with the correct IP address responds to the ARP request and provides the correct MAC address. If the unrestricted proxy ARP feature is enabled, the router response is redundant and might fool the IP client into determining that the destination MAC address within its own subnet is the same as the address of the router.



NOTE: While the destination address can be remote, the source address of the ARP request must be on the same subnet as the interface upon which the ARP request is received. For security reasons, this rule applies to both unrestricted and restricted proxy ARP.

Topology Considerations for Unrestricted Proxy ARP

In most situations, you should not configure the router or switch to perform unrestricted proxy ARP. Do so only for special situations, such as when cable modems are used.

[Figure 17 on page 1828](#) and [Figure 18 on page 1828](#) show examples of situations in which you might want to configure unrestricted proxy ARP.

In [Figure 17 on page 1828](#), the edge device is not running any IP protocols. In this case, you configure the core router to perform unrestricted proxy ARP. The edge device is the client of the proxy.

In [Figure 18 on page 1828](#), the Broadband Remote Access Server (B-RAS) routers are not running any IP protocols. In this case, you configure unrestricted proxy ARP on the B-RAS interfaces. This allows the core device to behave as though it is directly connected to the end users.

Figure 17: Edge Device Case for Unrestricted Proxy ARP

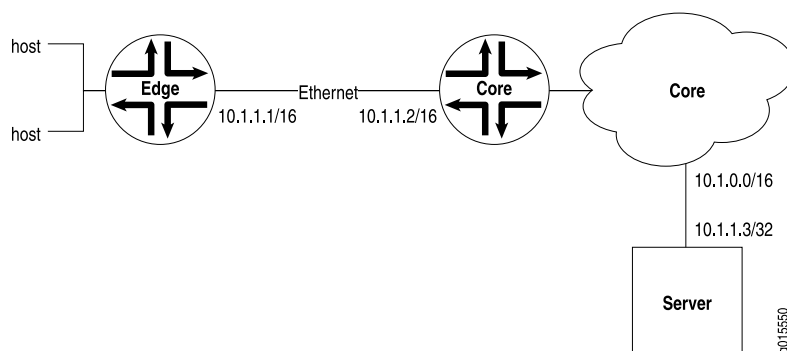



Figure 18: Core Device Case for Unrestricted Proxy ARP



Related Documentation

- [Configuring Restricted and Unrestricted Proxy ARP on page 1850](#)
- [Junos® OS Ethernet Interfaces](#)

Configuration

- [Configuration Tasks on page 1829](#)
- [Configuration Statements on page 1853](#)

Configuration Tasks

- [Configuring a VLAN on page 1829](#)
- [Configuring VLAN Identifiers for VLANs and VPLS Routing Instances on page 1830](#)
- [Configuring Integrated Routing and Bridging for VLANs on page 1835](#)
- [Configuring a Set of VLANs to Act as a Switch for a Layer 2 Trunk Port on page 1836](#)
- [Disabling MAC Learning for a VLAN or Logical Interface on page 1837](#)
- [Disabling MAC Learning for a Set of VLANs on page 1838](#)
- [Enabling MAC Accounting on page 1838](#)
- [Enabling MAC Accounting for a VLAN on page 1839](#)
- [Enabling MAC Accounting for a Set of VLANs on page 1839](#)
- [Configuring Inner and Outer TPIDs and VLAN IDs on page 1839](#)
- [Stacking a VLAN Tag on page 1843](#)
- [Configuring the Size of the MAC Address Table on page 1843](#)
- [Configuring Static MAC Addresses for Logical Interfaces in a VLAN on page 1844](#)
- [Disabling Layer 2 Learning and Forwarding on page 1845](#)
- [Configuring Multiple VLAN Registration Protocol \(MVRP\) on page 1845](#)
- [Configuring a Layer 2 Virtual Switch on page 1847](#)
- [Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port on page 1848](#)
- [Configuring VLAN Encapsulation on page 1849](#)
- [Configuring Restricted and Unrestricted Proxy ARP on page 1850](#)
- [Rewriting a VLAN Tag and Adding a New Tag on page 1851](#)
- [Configuring VLAN Translation with a VLAN ID List on page 1852](#)
- [Configuring a Logical Interface for Access Mode on page 1853](#)

Configuring a VLAN

A VLAN must include a set of logical interfaces that participate in Layer 2 learning and forwarding. You can optionally configure a VLAN identifier and a Layer 3 interface for the VLAN to also support Layer 3 IP routing.

To enable a VLAN, include the following statements:

```
[edit]
vans {
  vlan-name {
    domain-type bridge;
    interface interface-name;
    l3-interface interface-name;
```

```
    vlan-id (none | all | number);
    vlan-id-list [ vlan-id-numbers ];
    vlan-tags outer number inner number);
  }
}
```

You cannot use the slash (/) character in VLAN names. If you do, the configuration does not commit and an error is generated.

For the **vlan-id** statement, you can specify either a valid VLAN identifier or the **none** or **all** options.

To include one or more logical interfaces in the VLAN, specify an **interface-name** for an Ethernet interface you configured at the **[edit interfaces]** hierarchy level.



NOTE: A maximum of 4096 active logical interfaces are supported for a VLAN or on each mesh group in a virtual private LAN service (VPLS) instance configured for Layer 2 bridging.

By default, each VLAN maintains a Layer 2 forwarding database that contains media access control (MAC) addresses learned from packets received on the ports that belong to the VLAN. You can modify Layer 2 forwarding properties, for example, disabling MAC learning for the entire system or a VLAN, adding static MAC addresses for specific logical interfaces, and limiting the number of MAC addresses learned by the entire system, the VLAN, or a logical interface.

You can also configure spanning tree protocols to prevent forwarding loops.

Configuring VLAN Identifiers for VLANs and VPLS Routing Instances

You can configure VLAN identifiers for a VLAN or a VPLS routing instance in the following ways:

- By using either the **vlan-id** statement or the **vlan-tags** statement to configure a normalizing VLAN identifier. This topic describes how normalizing VLAN identifiers are processed and translated in a VLAN or a VPLS routing instance.
- By using the **input-vlan-map** and the **output-vlan-map** statements at the **[edit interfaces interface-name unit logic-unit-number]** or **[edit logical-systems logical-system-name interfaces interface-name unit logic-unit-number]** hierarchy level to configure VLAN mapping.

The **vlan-id** and **vlan-tags** statements are used to specify the normalizing VLAN identifier under the VLAN or VPLS routing instance. The normalizing VLAN identifier is used to perform the following functions:

- Translate, or normalize, the VLAN tags of packets received into a learn VLAN identifier.
- Create multiple learning domains that each contain a learn VLAN identifier. A learning domain is a MAC address database to which MAC addresses are added based on the learn VLAN identifier.



NOTE: You cannot configure VLAN mapping using the `input-vlan-map` and `output-vlan-map` statements if you configure a normalizing VLAN identifier for a VLAN or VPLS routing instance using the `vlan-id` or `vlan-tags` statements.

To configure a VLAN identifier for a VLAN, include either the `vlan-id` or the `vlan-tags` statement at the `[edit interfaces interface-name unit logic-unit-number]` or `[edit logical-systems logical-system-name interfaces interface-name unit logic-unit-number]` hierarchy level, and then include that logical interface in the VLAN configuration.

For a VPLS routing instance, include either the `vlan-id` or `vlan-tags` statement at the `[edit interfaces interface-name unit logic-unit-number]` or `[edit logical-systems logical-system-name interfaces interface-name unit logic-unit-number]` hierarchy level, and then include that logical interface in the VPLS routing instance configuration.



NOTE: For a single VLAN or VPLS routing instance, you can include either the `vlan-id` or the `vlan-tags` statement, but not both. If you do not configure a `vlan-id`, `vlan-tags`, or `vlan-id-list` [`vlan-id-numbers`] for the VLAN or the VPLS routing instance, the Layer 2 packets received are forwarded to the outbound Layer 2 interface without having the VLAN tag modified unless an `output-vlan-map` is configured on the Layer 2 interface. This results in a frame being forwarded to a Layer 2 interface with a VLAN tag that is different from what is configured for the Layer 2 interface. Note that a frame received from the Layer 2 interface is still required to match the VLAN tag(s) specified in the interface configuration. The invalid configuration may cause a Layer 2 loop to occur.

The VLAN tags associated with the inbound logical interface are compared with the normalizing VLAN identifier. If the tags are different, they are rewritten as described in [Table 129 on page 1834](#). The source MAC address of a received packet is learned based on the normalizing VLAN identifier.



NOTE: You do not have to specify a VLAN identifier for a VLAN that is performing Layer 2 switching only. To support Layer 3 IP routing, you must specify either a VLAN identifier or a pair of VLAN tags. However, you cannot specify the same VLAN identifier for more than one VLAN within a routing instance. Each VLAN must have a unique VLAN identifier.

If the VLAN tags associated with the outbound logical interface and the normalizing VLAN identifier are different, the normalizing VLAN identifier is rewritten to match the VLAN tags of the outbound logical interface, as described in [Table 130 on page 1835](#).

For the packets sent over the VPLS routing instance to be tagged by the normalizing VLAN identifier, include one of the following configuration statements:

- **vlan-id *number*** to tag all packets that are sent over the VPLS virtual tunnel (VT) interfaces with the VLAN identifier.
- **vlan-tags outer *number* inner *number*** to tag all packets sent over the VPLS VT interfaces with dual outer and inner VLAN tags.

Use the **vlan-id none** statement to have the VLAN tags removed from packets associated with an inbound logical interface when those packets are sent over VPLS VT interfaces. Note that those packets might still be sent with other customer VLAN tags.

The **vlan-id all** statement enables you to configure bridging for several VLANs with a minimum amount of configuration. Configuring this statement creates a learning domain for:

- Each inner VLAN, or learn VLAN, identifier of a logical interface configured with two VLAN tags
- Each VLAN, or learn VLAN, identifier of a logical interface configured with one VLAN tag

We recommend that you do not use customer VLAN IDs in a VPLS routing instance because customer VLAN IDs are used for learning only.

You should use the service VLAN ID in a VPLS routing instance, as in the following configuration:

```
[edit]
interface ge-1/1/1 {
  vlan-tagging;
  unit 1 {
    vlan-id s1; /* Service vlan */
    encapsulation vlan-vpls;
    input-vlan-map pop; /* Pop the service vlan on input */
    output-vlan-map push; /* Push the service vlan on output */
  }
}
interface ge-1/1/2 {
  encapsulation ethernet-vpls;
  unit 0;
}
routing-instances {
  V1 {
    instance-type vpls;
    vlan-id all;
    interface ge-1/1/1.1;
    interface ge-1/1/2.0;
  }
}
```



NOTE: If you configure the `vlan-id all` statement in a VPLS routing instance, we recommend using the `input-vlan-map pop` and `output-vlan-map push` statements on the logical interface to pop the service VLAN ID on input and push the service VLAN ID on output and in this way, limit the impact of double-tagged frames on scaling. You cannot use the native `vlan-id` statement when the `vlan-id all` statement is included in the configuration.

The `vlan-id-list [vlan-id-numbers]` statement enables you to configure bridging for multiple VLANs on a trunk interface. Configuring this statement creates a learning domain for:

- Each VLAN listed: `vlan-id-list [100 200 300]`
- Each VLAN in a range: `vlan-id-list [100-200]`
- Each VLAN in a list and range combination: `vlan-id-list [50, 100-200, 300]`

The following steps outline the process for bridging a packet received over a Layer 2 logical interface when you specify a normalizing VLAN identifier using either the `vlan-id number` or `vlan-tags` statement for a VLAN or a VPLS routing instance:

1. When a packet is received on a physical port, it is accepted only if the VLAN identifier of the packet matches the VLAN identifier of one of the logical interfaces configured on that port.
2. The VLAN tags of the received packet are then compared with the normalizing VLAN identifier. If the VLAN tags of the packet are different from the normalizing VLAN identifier, the VLAN tags are rewritten as described in [Table 129 on page 1834](#).
3. If the source MAC address of the received packet is not present in the source MAC table, it is learned based on the normalizing VLAN identifier.
4. The packet is then forwarded toward one or more outbound Layer 2 logical interfaces based on the destination MAC address. A packet with a known unicast destination MAC address is forwarded only to one outbound logical interface. For each outbound Layer 2 logical interface, the normalizing VLAN identifier configured for the VLAN or VPLS routing instance is compared with the VLAN tags configured on that logical interface. If the VLAN tags associated with an outbound logical interface do not match the normalizing VLAN identifier configured for the VLAN or VPLS routing instance, the VLAN tags are rewritten as described in [Table 130 on page 1835](#).

The tables below show how VLAN tags are applied for traffic sent to and from the VLAN, depending on how the `vlan-id` and `vlan-tags` statements are configured for the VLAN and on how identifiers are configured for the logical interfaces in a VLAN or VPLS routing instance. Depending on your configuration, the following rewrite operations are performed on VLAN tags:

- **pop**—Remove a VLAN tag from the top of the VLAN tag stack.
- **pop-pop**—Remove both the outer and inner VLAN tags of the frame.
- **pop-swap**—Remove the outer VLAN tag of the frame and replace the inner VLAN tag of the frame.

- **swap**—Replace the VLAN tag of the frame.
- **push**—Add a new VLAN tag to the top of the VLAN stack.
- **push-push**—Push two VLAN tags in front of the frame.
- **swap-push**—Replace the VLAN tag of the frame and add a new VLAN tag to the top of the VLAN stack.
- **swap-swap**—Replace both the outer and inner VLAN tags of the frame.

Table 129 on page 1834 shows specific examples of how the VLAN tags for packets sent to the VLAN are processed and translated, depending on your configuration. “—” means that the statement is not supported for the specified logical interface VLAN identifier. “No operation” means that the VLAN tags of the received packet are not translated for the specified input logical interface.

Table 129: Statement Usage and Input Rewrite Operations for VLAN Identifiers for a VLAN

VLAN Identifier of Logical Interface	VLAN Configurations for a VLAN			
	vlan-id none	vlan-id 200	vlan-id all	vlan tags outer 100 inner 300
none	No operation	push 200	—	push 100, push 300
200	pop 200	No operation	No operation	swap 200 to 300, push 100
1000	pop 1000	swap 1000 to 200	No operation	swap 1000 to 300, push 100
vlan-tags outer 2000 inner 300	pop 2000, pop 300	pop 2000, swap 300 to 200	pop 2000	swap 2000 to 100
vlan-tags outer 100 inner 400	pop 100, pop 400	pop 100, swap 400 to 200	pop 100	swap 400 to 300
vlan-id-range 10-100	—	—	No operation	—
vlan-tags outer 200 inner-range 10-100	—	—	pop 200	—

Table 130 on page 1835 shows specific examples of how the VLAN tags for packets sent from the VLAN are processed and translated, depending on your configuration. “—” means that the statement is not supported for the specified logical interface VLAN identifier. “No operation” means that the VLAN tags of the outbound packet are not translated for the specified output logical interface.

Table 130: Statement Usage and Output Rewrite Operations for VLAN Identifiers for a VLAN

VLAN Identifier of Logical Interface	VLAN Configurations for a VLAN			
	vlan-id none	vlan-id 200	vlan-id all	vlan tags outer 100 inner 300
none	no operation	pop 200	–	pop 100, pop 300
200	push 200	No operation	No operation	pop 100, swap 300 to 200
1000	push 1000	swap 200 to 1000	No operation	pop 100, swap 300 to 1000
vlan-tags outer 2000 inner 300	push 2000, push 300	swap 200 to 300, push 2000	push 2000	swap 100 to 2000
vlan-tags outer 100 inner 400	push 100, push 400	swap 200 to 400, push 100	push 100	swap 300 to 400
vlan-id-range 10-100	–	–	No operation	–
vlan-tags outer 200 inner-range 10-100	–	–	push 200	–

Configuring Integrated Routing and Bridging for VLANs

Integrated routing and bridging (IRB) provides simultaneous support for Layer 2 bridging and Layer 3 routing on the same interface. IRB enables you to route packets to another routed interface or to another VLAN that has an IRB interface configured. You configure a logical routing interface by specifying **irb** as an interface name at the **[edit interfaces]** hierarchy level and including that interface in the VLAN.



NOTE: You can include only one Layer 3 interface in a VLAN.

To configure a VLAN with IRB support, include the following statements:

```
[edit]
vlands {
  vlan-name {
    domain-type bridge;
    interface interface-name;
    l3-interface interface-name;
    vlan-id (none | number);
    vlan-tags outer number inner number;
  }
}
```

For each VLAN that you configure, specify a **vlan-name**. You must also specify the value **bridge** for the **domain-type** statement.

For the **vlan-id** statement, you can specify either a valid VLAN identifier or the **none** option.



NOTE: If you configure a Layer 3 interface to support IRB in a VLAN, you cannot use the **all** option for the **vlan-id** statement.

The **vlan-tags** statement enables you to specify a pair of VLAN identifiers; an **outer** tag and an **inner** tag.



NOTE: For a single VLAN, you can include either the **vlan-id** statement or the **vlan-tags** statement, but not both.

To include one or more logical interfaces in the VLAN, specify the **interface-name** for each Ethernet interface to include that you configured at the **[edit interfaces]** hierarchy level.



NOTE: A maximum of 4096 active logical interfaces are supported for a VLAN or on each mesh group in a VPLS routing instance configured for Layer 2 bridging.

To associate a Layer 3 interface with a VLAN, include the **l3-interface interface-name** statement and specify an **interface-name** you configured at the **[edit interfaces irb]** hierarchy level. You can configure only one Layer 3 interface for each VLAN.

IRB interfaces are supported for multicast snooping.

In multihomed VPLS configurations, you can configure VPLS to keep a VPLS connection up if only an IRB interface is available by configuring the **irb** option for the **connectivity-type** statement at the **[edit routing-instances routing-instance-name protocols vpls]** hierarchy level. The **connectivity-type** statement has the **ce** and **irb** options. The **ce** option is the default and specifies that a CE interface is required to maintain the VPLS connection. By default, if only an IRB interface is available, the VPLS connection is brought down.



NOTE: When you configure IRB interfaces in more than one logical system on a device, all of the IRB logical interfaces share the same MAC address.

Configuring a Set of VLANs to Act as a Switch for a Layer 2 Trunk Port

You can configure a set of VLANs that are associated with a Layer 2 trunk port. The set of VLANs function as a switch. Packets received on a trunk interface are forwarded within a VLAN that has the same VLAN identifier. A trunk interface also provides support for IRB, which provides support for Layer 2 bridging and Layer 3 IP routing on the same interface.

To configure a Layer 2 trunk port and set of VLANs, include the following statements:

[edit interfaces]

```

interface-name {
  unit number {
    family ethernet-switching {
      interface-mode access;
      vlan-members (vlan-name | vlan-tag);
    }
  }
}
interface-name {
  native-vlan-id number;
  unit number {
    family ethernet-switching {
      interface-mode trunk;
      vlan-members (vlan-name | vlan-tag);
    }
  }
}
[edit vlans ]
vlan-name {
  vlan-id number;
  vlan-id-list [ vlan-id-numbers ];
  ....
}

```

You must configure a VLAN and VLAN identifier for each VLAN associated with the trunk interface. You can configure one or more trunk or access interfaces at the **[edit interfaces]** hierarchy level. An access interface enables you to accept packets with no VLAN identifier.

Disabling MAC Learning for a VLAN or Logical Interface

You can disable MAC learning for all logical interfaces in a specified VLAN, or for a specific logical interface in a VLAN. Disabling dynamic MAC learning prevents the specified interfaces from learning source MAC addresses.

To disable MAC learning for all logical interfaces in a VLAN in a virtual switch, include the **no-mac-learning** statement at the **[edit vlans vlan-name switch-options]** hierarchy level:

```

[edit]
vlans {
  vlan-name {
    domain-type bridge;
    interface interface-name;
    switch-options {
      no-mac-learning;
    }
  }
}

```

To disable MAC learning for a specific logical interface in a VLAN, include the **no-mac-learning** statement at the **[edit vlans vlan-name switch-options interface interface-name]** hierarchy level.

```

[edit]
vlans {
  vlan-name {
    domain-type bridge;
    interface interface-name;
  }
}

```

```
switch-options {  
  interface interface-name {  
    no-mac-learning;  
  }  
}
```



NOTE: When you disable MAC learning, source MAC addresses are not dynamically learned, and any packets sent to these source addresses are flooded into the VLAN.



NOTE: When you gather interfaces into a VLAN, the `no-mac-learn-enable` statement at the `[edit interfaces interface-name ether-options ethernet-switch-profile]` hierarchy level is not supported. You must use the `no-mac-learning` statement at the `[edit vlans vlan-name switch-options interface interface-name]` hierarchy level to disable MAC learning on an interface in a VLAN.



NOTE: When MAC learning is disabled for a VPLS routing instance, traffic is not load balanced and only one of the equal-cost next hops is used.

Disabling MAC Learning for a Set of VLANs

You can disable MAC learning for a set of VLANs. Disabling dynamic MAC learning prevents the Layer 2 trunk port associated with the set of VLANs from learning source and destination MAC addresses. When you disable MAC learning, source MAC addresses are not dynamically learned, and any packets sent to these source addresses are flooded into the switch.

To disable MAC learning for a set of VLANs, include the **no-mac-learning** statement at the `[edit switch-options]` hierarchy level:

```
[edit switch-options]  
no-mac-learning;
```

Enabling MAC Accounting

By default, MAC accounting is disabled. You can enable packet accounting either for a router or switch as a whole or for a specific VLAN. After you enable packet accounting, the Junos OS maintains packet counters for each MAC address learned.

To enable MAC accounting, include the **global-mac-statistics** statement at the `[edit protocols l2-learning]` hierarchy level:

```
[edit protocols l2-learning]  
global-mac-statistics;
```

Enabling MAC Accounting for a VLAN

By default, MAC accounting is disabled. You can enable packet counting for a VLAN. When you enable packet accounting, the Junos OS maintains packet counters for each MAC address learned on the interfaces in the VLAN.

To enable MAC accounting for a VLAN, include the **mac-statistics** statement at the **[edit vlans *vlan-name* switch-options]** hierarchy level:

```
[edit vlans vlan-name switch-options]
mac-statistics;
```

Enabling MAC Accounting for a Set of VLANs

By default, MAC accounting is disabled. You can enable packet counting for a set of VLANs. After you enable packet accounting, the Junos OS maintains packet counters for each MAC address learned on the trunk port associated with the set of VLANs.

To enable MAC accounting for a set of VLANs, include the **mac-statistics** statement at the **[edit switch-options]** hierarchy level:

```
[edit switch-options on page 430]
mac-statistics;
```

Configuring Inner and Outer TPIDs and VLAN IDs

For some rewrite operations, you must configure the inner or outer TPID values and inner or outer VLAN ID values. These values can be applied to either the input VLAN map or the output VLAN map.

On Ethernet IQ, IQ2, and IQ2-E interfaces; on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces; and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to configure the inner TPID, include the **inner-tag-protocol-id** statement:

```
inner-tag-protocol-id tpid;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]**

For the inner VLAN ID, include the **inner-vlan-id** statement. For the outer TPID, include the **tag-protocol-id** statement. For the outer VLAN ID, include the **vlan-id** statement:

```
input-vlan-map {
  (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
  inner-tag-protocol-id tpid;
  inner-vlan-id number;
  tag-protocol-id tpid;
```

```

    vlan-id number;
  }
  output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
  }

```

For aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces, include the **tag-protocol-id** statement for the outer TPID. For the outer VLAN ID, include the **vlan-id** statement:

```

input-vlan-map {
  (pop | push | swap);
  tag-protocol-id tpid;
  vlan-id number;
}
output-vlan-map {
  (pop | push | swap);
  tag-protocol-id tpid;
  vlan-id number;
}

```

You can include these statements at the following hierarchy levels:

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

The VLAN IDs you define in the input VLAN maps are stacked on top of the VLAN ID bound to the logical interface. For more information about binding a VLAN ID to the logical interface, see *802.1Q VLANs Overview*.

All TPIDs you include in input and output VLAN maps must be among those you specify at the [edit interfaces *interface-name* ether-options ethernet-switch-profile tag-protocol-id [*tpids*]] hierarchy level.

Table 131 on page 1840 and Table 132 on page 1841 specify when these statements are required. Table 131 on page 1840 indicates valid statement combinations for rewrite operations for the input VLAN map. “No” means the statement must not be included in the input VLAN map for the rewrite operation. “Optional” means the statement may be optionally specified for the rewrite operation in the input VLAN map. “Any” means that you must include the **vlan-id** statement, **tag-protocol-id** statement, **inner-vlan-id** statement, or **inner-tag-protocol-id** statement.

Table 131: Rewrite Operations and Statement Usage for Input VLAN Maps

Rewrite Operation	Input VLAN Map Statements			
	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
push	Optional	Optional	No	No

Table 131: Rewrite Operations and Statement Usage for Input VLAN Maps (*continued*)

	Input VLAN Map Statements			
pop	No	No	No	No
swap	Any	Any	No	No
push-push	Optional	Optional	Optional	optional
swap-push	Optional	Optional	Any	Any
swap-swap	Optional	Optional	Any	Any
pop-swap	No	No	Any	Any
pop-pop	No	No	No	No

Table 132 on page 1841 indicates valid statement combinations for rewrite operations for the output VLAN map. “No” means the statement must not be included in the output VLAN map for the rewrite operation. “Optional” means the statement may be optionally specified for the rewrite operation in the output VLAN map.

Table 132: Rewrite Operations and Statement Usage for Output VLAN Maps

	Output VLAN Map Statements			
Rewrite Operation	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
push	No	Optional	No	No
pop	No	No	No	No
swap	No	Optional	No	No
push-push	No	Optional	No	Optional
swap-push	No	Optional	No	Optional
swap-swap	No	Optional	No	Optional
pop-swap	No	No	No	Optional
pop-pop	No	No	No	No

The following examples use Table 131 on page 1840 and Table 132 on page 1841 and show how the **pop-swap** operation can be configured in an input VLAN map and an output VLAN map:

Input VLAN Map with inner-vlan-id Statement, Output VLAN Map with Optional inner-tag-protocol-id Statement	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] <code>input-vlan-map { pop-swap; inner-vlan-id <i>number</i>; }</code> <code>output-vlan-map { pop-swap; inner-tag-protocol-id <i>tpid</i>; }</code>
--	---

Input VLAN Map with inner-tag-protocol-id Statement, Output VLAN Map with Optional inner-tag-protocol-id Statement	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] <code>input-vlan-map { pop-swap; inner-tag-protocol-id <i>tpid</i>; }</code> <code>output-vlan-map { pop-swap; inner-tag-protocol-id <i>tpid</i>; }</code>
--	---

Input VLAN Map with inner-tag-protocol-id and inner-vlan-id Statements	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] <code>input-vlan-map { pop-swap; inner-vlan-id <i>number</i>; inner-tag-protocol-id <i>tpid</i>; }</code>
---	---

Stacking a VLAN Tag

To stack a VLAN tag on all tagged frames entering or exiting the interface, include the **push**, **vlan-id**, and **tag-protocol-id** statements in the input VLAN map or the output VLAN map:

```
input-vlan-map input-vlan-map {
  push;
  vlan-id number;
  tag-protocol-id tpid;
}
output-vlan-map {
  push;
  tag-protocol-id tpid;
}
```

You can include these statements at the following hierarchy levels:

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

The VLAN IDs you define in the input VLAN maps are stacked on top of the VLAN ID bound to the logical interface.

All TPIDs you include in input and output VLAN maps must be among those you specify at the [edit interfaces *interface-name* ether-options ethernet-switch-profile tag-protocol-id [*tpids*]] hierarchy level.

Configuring the Size of the MAC Address Table

You can modify the size of the MAC address table for each VLAN. The default table size is 5120 addresses. The minimum you can configure is 16 addresses, and the maximum is 1,048,575 addresses.

If the MAC table limit is reached, new addresses can no longer be added to the table. Unused MAC addresses are removed from the MAC address table automatically. This frees space in the table, allowing new entries to be added.

To modify the size of the MAC table, include the **mac-table-size** *limit* statement at the [edit vlans *vlan-name* switch-options] hierarchy level:

```
[edit]
vlans {
  vlan-name {
    domain-type bridge;
    switch-options {
      mac-table-size limit {
        packet-action drop;
```

```

    }
  }
}

```

Related Documentation

- [Disabling MAC Learning for a Bridge Domain or Logical Interface](#)
- [Configuring Static MAC Addresses for Logical Interfaces in a Bridge Domain](#)
- [Limiting MAC Addresses Learned from an Interface in a Bridge Domain](#)
- [Enabling MAC Accounting for a Bridge Domain](#)

Configuring Static MAC Addresses for Logical Interfaces in a VLAN

You can manually add static MAC entries for the logical interfaces in a VLAN. You can specify one or more static MAC addresses for each logical interface.

To add a static MAC address for a logical interface in a VLAN, include the **static-mac mac-address** statement at the **[edit vlans vlan-name switch-options interface interface-name]** hierarchy level.

```

[edit]
vlans {
  vlan-name {
    domain-type bridge;
    switch-options {
      interface interface-name {
        static-mac mac-address {
          <vlan-id number>;
        }
      }
    }
  }
}

```

You can optionally specify a VLAN identifier for the static MAC address by using the **vlan-id** statement. To specify a VLAN identifier for a static MAC address, you must use the **all** option when configuring a VLAN identifier for the VLAN.



NOTE: If a static MAC address you configure for a logical interface appears on a different logical interface, packets sent to that interface are dropped.

Related Documentation

- [Disabling MAC Learning for a VLAN or Logical Interface on page 1837](#)
- [Configuring the Size of the MAC Address Table on page 1843](#)
- [Enabling MAC Accounting for a VLAN on page 1839](#)

Disabling Layer 2 Learning and Forwarding

Disabling dynamic MAC learning on an MX Series router or an EX Series switch prevents all the logical interfaces on the router from learning source and destination MAC addresses.

To disable MAC learning for an MX Series router or an EX Series switch, include the **global-no-mac-learning** statement at the **[edit protocols l2-learning]** hierarchy level:

```
[edit protocols l2-learning]
global-no-mac-learning;
```

For information about how to configure a virtual switch, see *Configuring a Layer 2 Virtual Switch*.

Related Documentation

- [Layer 2 Learning and Forwarding Overview](#)
- [Configuring the MAC Table Timeout Interval](#)
- [Enabling MAC Accounting](#)
- [Limiting the Number of MAC Addresses Learned from Each Logical Interface](#)

Configuring Multiple VLAN Registration Protocol (MVRP)

Multiple VLAN Registration Protocol (MVRP) is used to manage dynamic VLAN registration in Carrier Ethernet network. You can use MVRP on MX Series routers or on EX Series switches.

For information about using MVRP on EX Series switches, see *Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches*.

MVRP is disabled by default on MX Series routers and EX Series switches.

To enable MVRP or set MVRP options, follow these instructions:

- [Enabling MVRP on page 1845](#)
- [Disabling MVRP on page 1846](#)
- [Changing the Registration Mode to Disable Dynamic VLANs on page 1846](#)
- [Configuring Timer Values on page 1846](#)
- [Configuring the Multicast MAC address for MVRP on page 1847](#)
- [Configuring an MVRP Interface as a Point-to-Point Interface on page 1847](#)
- [Configuring MVRP Tracing Options on page 1847](#)

Enabling MVRP

MVRP can only be enabled on trunk interfaces.

To enable MVRP on a specific trunk interface (here, interface **ge-3/0/5**):

```
[edit protocols mvrp]
user@host# set interfaces ge-3/0/5
```

Disabling MVRP

MVRP is disabled by default. You only need to perform this procedure if you have previously enabled MVRP.

To disable MVRP on all trunk interfaces, use one of the following:

```
[edit protocols mvrp]
user@host# deactivate protocols mvrp
user@host# delete protocols mvrp
```

Changing the Registration Mode to Disable Dynamic VLANs

When the registration mode for an interface is set to **normal** (the default), dynamic VLANs are created on interfaces participating in MVRP. The dynamic VLANs created on one router device are then propagated by means of MVRP to other router devices in a topology.

However, Dynamic VLAN creation through MVRP can be disabled for all trunk interfaces on a router device or for individual trunk interfaces.

For information about disabling dynamic VLAN creation on an interface so that the interface does not register and does not participate in MVRP, see *Controlling the Management State of a VLAN in MVRP Configurations (CLI Procedure)*.

Configuring Timer Values

The timers in MVRP define the amount of time an interface waits to join or leave MVRP or to send or process the MVRP information for the router or switch after receiving an MVRP PDU:

- The join timer controls the amount of time the router waits to accept a registration request.
- The leave timer controls the period of time that the router waits in the Leave state before changing to the unregistered state.
- The leaveall timer controls the frequency with which the LeaveAll messages are communicated.

The default MVRP timer values are 200 ms for the join timer, 1000 ms for the leave timer, and 10000 ms for the leaveall timer.



BEST PRACTICE: Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

To set the join timer for a specific interface:

```
[edit protocols mvrp]
user@host# set interfaces ge-3/0/5 join-timer 300
```

To set the leave timer for a specific interface:

```
[edit protocols mvrp]
```

```
user@host# set interfaces ge-3/0/5 leave-timer 1200
```

To set the leaveall timer for a specific interface:

```
[edit protocols mvrp]
user@host# set interface ge-3/0/5 leaveall-timer 12000
```

Configuring the Multicast MAC address for MVRP

MVRP uses the customer MVRP multicast MAC address when MVRP is enabled. However, you can configure MVRP to instead use the provider MVRP multicast MAC address.

To configure MVRP to use the provider MVRP multicast MAC address:

```
[edit protocols mvrp]
user@host# set bpdu-destination-mac-address provider-bridge-group;
```

Configuring an MVRP Interface as a Point-to-Point Interface

Specify that a configured interface is connected point-to-point. If specified, a point-to-point subset of the MRP state machine provides a simpler and more efficient method to accelerate convergence on the network.

To specify that an MVRP interface is point-to-point (here, interface **ge-3/0/5**):

```
[edit protocols mvrp]
user@host# set interfaces ge-3/0/5 point-to-point;
```

Configuring MVRP Tracing Options

Set MVRP protocol-level tracing options.

To specify MVRP protocol tracing (here, the file is **/var/log/mvrp-log**, size is **2m**, number of files is **28**, the option **world-readable** indicates the log can be read by user, and MVRP is flagging **events**):

```
[edit protocols mvrp]
user@host# edit protocols mvrp traceoptions file /var/log/mvrp-log size 2m files 28
world-readable flag events
```

Related Documentation

- *Example: Configuring Automatic VLAN Administration Using MVRP*
- *Verifying That MVRP Is Working Correctly*

Configuring a Layer 2 Virtual Switch

A Layer 2 virtual switch, which isolates a LAN segment with its spanning-tree protocol instance and separates its VLAN ID space, filters and forwards traffic only at the data link layer. Each VLAN consists of a set of logical ports that participate in Layer 2 learning and forwarding. A virtual switch represents a Layer 2 network.

Two main types of interfaces are used in virtual switch hierarchies:

- Layer 2 logical interface—This type of interface uses the VLAN-ID as a virtual circuit identifier and the scope of the VLAN-ID is local to the interface port. This type of interface is often used in service-provider-centric applications.
- Access or trunk interface—This type of interface uses a VLAN-ID with global significance. The access or trunk interface is implicitly associated with VLANs based on VLAN membership. Access or trunk interfaces are typically used in enterprise-centric applications.



NOTE: The difference between access interfaces and trunk interfaces is that access interfaces can be part of one VLAN only and the interface is normally attached to an end-user device (packets are implicitly associated with the configured VLAN). In contrast, trunk interfaces multiplex traffic from multiple VLANs and usually interconnect switches.

To configure a Layer 2 virtual switch, include the following statements:

```
[edit]
routing-instances {
  routing-instance-name (
    instance-type virtual-switch;
    vlans vlan-name{
      vlan-id (all | none | number);
      [...configure optional VLAN parameters]
    }
  }
}
```

To enable a virtual switch, you must specify **virtual-switch** as the **instance-type**.

The VLANs that are specified with the **vlan-id** statement are included in the virtual switch.

You can configure other optional VLAN parameters in the virtual switch.

Related Documentation

- [Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port on page 1848](#)

Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port

You can associate one or more Layer 2 trunk interfaces with a virtual switch.

A virtual switch configured with a Layer 2 trunk port also supports IRB within a VLAN. IRB provides simultaneous support for Layer 2 bridging and Layer 3 IP routing on the same interface. Only an interface configured with the **interface-mode (access | trunk)** statement can be associated with a virtual switch. An access interface enables you to accept packets with no VLAN identifier.

In addition, you can configure Layer 2 learning and forwarding properties for the virtual switch.

To configure a virtual switch with a Layer 2 trunk interface, include the following statements:

```
[edit]
routing-instances {
  routing-instance-name {
    instance-type virtual-switch;
    interface interface-name;
    vlans name {
      vlan-id (all | none | number);
      [...configure optional VLAN parameters]
    }
  }
}
```

Related Documentation

- [Configuring a Layer 2 Virtual Switch on page 1847](#)

Configuring VLAN Encapsulation

To configure encapsulation on an interface, enter the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
encapsulation type;
```

The following list contains important notes regarding encapsulation:

- Ethernet interfaces in VLAN mode can have multiple logical interfaces. In CCC and VPLS modes, VLAN IDs from 1 through 511 are reserved for normal VLANs, and VLAN IDs 512 through 4094 are reserved for CCC or VPLS VLANs. For 4-port Fast Ethernet interfaces, you can use VLAN IDs 512 through 1024 for CCC or VPLS VLANs.
- For encapsulation type **flexible-ethernet-services**, all VLAN IDs are valid.
- For some encapsulation types, including flexible Ethernet services, Ethernet VLAN CCC, and VLAN VPLS, you can also configure the encapsulation type that is used inside the VLAN circuit itself. To do this, include the **encapsulation** statement:

```
encapsulation (vlan-ccc | vlan-tcc | vlan-vpls);
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**
- You cannot configure a logical interface with VLAN CCC or VLAN VPLS encapsulation unless you also configure the physical device with the same encapsulation or with flexible Ethernet services encapsulation. In general, the logical interface must have a VLAN ID of 512 or higher; if the VLAN ID is 511 or lower, it will be subject to the normal destination filter lookups in addition to source address filtering. However if you configure flexible Ethernet services encapsulation, this VLAN ID restriction is removed.

In general, you configure an interface's encapsulation at the **[edit interfaces *interface-name*]** hierarchy level.

Example: Configuring VLAN Encapsulation on a Gigabit Ethernet Interface

Configure VLAN CCC encapsulation on a Gigabit Ethernet interface:

```
interfaces ge-2/1/0 {  
  vlan-tagging;  
  encapsulation vlan-ccc;  
  unit 0 {  
    encapsulation vlan-ccc;  
    vlan-id 600;  
  }  
}
```

Example: Configuring VLAN Encapsulation on an Aggregated Ethernet Interface

Configure VLAN CCC encapsulation on an aggregated Gigabit Ethernet interface:

```
interfaces ae0 {  
  vlan-tagging;  
  encapsulation vlan-vpls;  
  unit 0 {  
    vlan-id 100;  
  }  
}
```

**Related
Documentation**

- [802.1Q VLANs Overview](#)
- [Junos® OS Ethernet Interfaces](#)

Configuring Restricted and Unrestricted Proxy ARP

To configure restricted or unrestricted proxy ARP, include the **proxy-arp** statement:

proxy-arp (restricted |unrestricted);

You can include this statement at the following hierarchy levels:

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

To return to the default—that is, to disable restricted or unrestricted proxy ARP—delete the **proxy-arp** statement from the configuration:

```
[edit]  
user@host# delete interfaces interface-name unit logical-unit-number proxy-arp
```

You can track the number of restricted or unrestricted proxy ARP requests processed by the router or switch by issuing the **show system statistics arp** operational mode command.



NOTE: When proxy ARP is enabled as default or unrestricted, the router or switch responds to any ARP request as long as the device has an active route to the target address of the ARP request. This gratuitous ARP behavior can result in an error when the receiving interface and target response interface are the same and the end device (for example, a client) performs a duplicate address check. To prevent this error, configure the router or switch interface with the `no-gratuitous-arp-reply` statement. See *Configuring Gratuitous ARP* for information about how to disable responses to gratuitous ARP requests.

Related Documentation

- [proxy-arp on page 1918](#)
- [Restricted and Unrestricted Proxy ARP Overview on page 1827](#)
- [Configuring Gratuitous ARP](#)
- [Junos® OS Ethernet Interfaces](#)

Rewriting a VLAN Tag and Adding a New Tag

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, and on Gigabit Ethernet and 10-Gigabit Ethernet interfaces on EX Series switches, to replace the outer VLAN tag of the incoming frame with a user-specified VLAN tag value, include the **swap-push** statement in the input VLAN map or output VLAN map:

swap-push

A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* **input-vlan-map**]
- [edit interfaces *interface-name* unit *logical-unit-number* **output-vlan-map**]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* **input-vlan-map**]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* **output-vlan-map**]

See *Configuring Inner and Outer TPIDs and VLAN IDs* and *Configuring Inner and Outer TPIDs and VLAN IDs* for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Related Documentation

- [input-vlan-map on page 1893](#)
- [output-vlan-map on page 1911](#)
- [swap-push on page 1926](#)

- *unit*
- *Junos® OS Ethernet Interfaces*

Configuring VLAN Translation with a VLAN ID List

In many cases, the VLAN identifiers on the frames of an interface's packets are not correct. VLAN translation, or VLAN rewrite, allows you to configure bidirectional VLAN identifier translation with a list on frames arriving on and leaving from a logical interface. This lets you use unique VLAN identifiers internally and maintain legacy VLAN identifiers on logical interfaces.

To perform VLAN translation on the packets on a trunk interface, insert the **vlan-rewrite** statement at the **[edit interfaces *interface-name* unit *unit-number*]** hierarchy level. You must also include the **interface-mode trunk** statement within the **[edit interfaces *interface-name* unit *unit-number* family ethernet-switching]** hierarchy because VLAN translation is only supported on trunk interfaces. The reverse translation takes place on traffic exiting the interface. In other words, if VLAN 200 is translated to 500 on traffic entering the interface, VLAN 500 is translated to VLAN 200 on traffic leaving the interface.

The following example translates incoming trunk packets from VLAN identifier 200 to 500 and 201 to 501 (other valid VLAN identifiers are not affected):

```
[edit interfaces ge-1/0/1]
unit 0 {
  ... # Other logical interface statements
  family ethernet-switching {
    interface-mode trunk # Translation is only for trunks
    inner-vlan-id-list [ 100 500–600 ];
    vlan-rewrite {
      translate 200 500;
      translate 201 501;
    }
    ... # Other ethernet-switching statements
  }
}
```



NOTE: This example also translates frame VLANs from 500 to 200 and 501 to 201 on egress.

Related Documentation

- *Rewriting a VLAN Tag and Adding a New Tag*

Configuring a Logical Interface for Access Mode

Enterprise network administrators can configure a single logical interface to accept untagged packets and forward the packets within a specified VLAN. A logical interface configured to accept untagged packets is called an *access interface* or *access port*.

`interface-mode access;`

You can include this statement at the following hierarchy levels:

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

When an untagged or tagged packet is received on an access interface, the packet is accepted, the VLAN ID is added to the packet, and the packet is forwarded within the VLAN that is configured with the matching VLAN ID.

The following example configures a logical interface as an access port with a VLAN ID of 20 on routers and switches that support the enhanced Layer 2 software:

```
[edit interfaces ge-1/2/0]
unit 1 {
  family ethernet-switching {
    interface-mode access;
    vlan members 20;
  }
}
```

- Related Documentation**
- [802.1Q VLANs Overview](#)
 - [Junos® OS Ethernet Interfaces](#)

Configuration Statements

- [\[edit dynamic-profiles\] Hierarchy Level on page 1853](#)
- [\[edit interfaces\] Hierarchy Level on page 1854](#)
- [\[edit logical-systems\] Hierarchy Level on page 1865](#)
- [\[edit protocols l2-learning\] Hierarchy Level on page 1870](#)
- [Layer 2 Routing Instances Configuration Hierarchy on page 1870](#)
- [\[edit switch-options\] Hierarchy Level on page 1873](#)
- [\[edit vlans\] Hierarchy Level on page 1873](#)

[\[edit dynamic-profiles\] Hierarchy Level](#)

```
dynamic-profiles {
  profile-name {
    class-of-service
    ... statements from those in [edit class-of-service] Hierarchy Level.
    firewall
    ... statements from those in [edit firewall] Hierarchy Level.
```

```
interfaces
... statements from those in [edit interfaces] Hierarchy Level.
policy-options
... statements from those in [edit policy-options] Hierarchy Level.
predefined-variable-defaults variable-name default-value
profile-variable-set variable-set-name dynamic-variable-name substitute-variable-name
protocols
... statements from those in [edit protocols] Hierarchy Level.
routing-instances
... statements from those in [edit routing-instances] Hierarchy Level.
routing-options
... statements from those in [edit routing-options] Hierarchy Level.
services
... statements from those in [edit services] Hierarchy Level.
variables {
  variable-name {
    default-value default-value;
    equals expression;
    mandatory;
  }
  uid;
  uid-reference;
}
}
```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit dynamic-profiles routing-instances] Hierarchy Level*
- *[edit dynamic-profiles routing-options] Hierarchy Level*
- *[edit dynamic-profiles variables] Hierarchy Level*

[edit interfaces] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```
interfaces {
  interface-name {
    ... the "interface-name" subhierarchy appears after the main [edit interfaces] hierarchy
    level ...
  }
  interface-set interface-set-name {
    interface interface-name {
      (unit unit-number | vlan-tags-outer vlan-tag);
    }
  }
}
irb (Interfaces) {
  accounting-profile name;
  description text;
  disable;

  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  hold-time up milliseconds down milliseconds;
```

```

mtu bytes;
no-gratuitous-arp-request;

traceoptions {
    flag flag;
}
(traps | no-traps);
unit logical-unit-number {
    accounting-profile name;
    bandwidth rate;
    description text;
    disable;
    encapsulation type;
    family inet {
        accounting {
            destination-class-usage;
            source-class-usage {
                input;
                output;
            }
        }
    }
    address ipv4-address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        preferred;
        primary;
        vrrp-group group-id {
            (accept-data | no-accept-data);
            advertise-interval seconds;
            advertisements-threshold number;
            authentication-key key;
            authentication-type authentication;
            fast-interval milliseconds;
            (preempt | no-preempt) {
                hold-time seconds;
            }
            priority number;
            track {
                interface interface-name {
                    bandwidth-threshold bits-per-second priority-cost priority;
                    priority-cost priority;
                }
                priority-hold-time seconds;
                route prefix/prefix-length routing-instance instance-name priority-cost priority;
            }
            virtual-address [ addresses ];
            vrrp-inherit-from vrrp-group;
        }
    }
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;

```

```
no-redirects;
primary;
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bandwidth priority-cost number;
                priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-inet6-address [addresses];
        virtual-link-local-address ipv6-address;
        vrrp-inherit-from {
            active-group group-number;
            active-interface interface-name;
        }
    }
}
}
(dad-disable | no-dad-disable);
filter {
    input filter-name;
```

```

        output filter-name;
    }
    mtu bytes;
    nd6-stale-time seconds;
    no-neighbor-learn;
    no-redirects;
    policer {
        input policer-name;
        output policer-name;
    }
    rpf-check {
        fail-filter filter-name;
        mode {
            loose;
        }
    }
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    filter {
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    policer {
        input policer-name;
        output policer-name;
    }
}
native-inner-vlan-id vlan-id;
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id-list [vlan-id's];
vlan-id-range [vlan-id-range];
}
}
traceoptions {
    file <filename> <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag <disable>;
    no-remote-trace;
}
}

interfaces {
    interface-name {
        disable;
        accounting-profile name;
        aggregated-ether-options {
            ethernet-switch-profile {
                tag-protocol-id [ hexadecimal-identifiers ];
            }
        }
        (flow-control | no-flow-control);
    }
}

```

```
lACP {
  (active | passive);
  admin-key key;
  fast-failover;
  link-protection {
    disable;
    (revertive | non-revertive);
  }
  periodic (fast | slow);
  system-id mac-address;
  system-priority priority;
}
(link-protection | no-link-protection);
link-speed (100m | 1g | 8g | 10g | 40g | 50g | 80g | 100g | oc192);
logical-interface-fpc-redundancy;
(loopback | no-loopback);
mc-ae {
  chassis-id chassis-id;
  events {
    iccp-peer-down {
      force-icl-down;
      prefer-status-control-active;
    }
  }
  mc-ae-id mc-ae-id;
  mode (active-active | active-standby);
  redundancy-group group-id;
  status-control (active | standby);
}
minimum-links number;
rebalance-periodic {
  start-time time;
  interval number;
}
source-address-filter {
  mac-address;
}
(source-filtering | no-source-filtering);
}
auto-configure {
  remove-when-no-subscribers;
  stacked-vlan-ranges {
    access-profile profile-name;
    authentication {
      password password-string;
      username-include {
        circuit-type;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        mac-address;
        option-82 ( circuit-id | remote-id);
        radius-realm radius-realm-string;
        user-prefix user-prefix-string;
      }
    }
  }
}
```



```

dynamic-profile profile-name {
    accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
    ranges (any | low-tag-high-tag), (any | low-tag-high-tag);
}
}
vlan-ranges {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-name;
            mac-address;
            option-82;
            radius-realm radius-realm-string;
            user-prefix user-prefix-string;
        }
    }
    dynamic-profile profile-name {
        accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
        ranges (any | low-tag)—(any | high-tag);
    }
}
override tag vlan-tag dynamic-profile profile name;
}
encapsulation (ethernet-bridge | ethernet-vpls | extended-vlan-bridge |
    extended-vlan-vpls | flexible-ethernet-services | vlan-vpls);
ether-options {
    802.3ad {
        aex;
        (backup | primary);
        lacp {
            force-up;
            port-priority
        }
    }
}
asynchronous-notification;
(auto-negotiation | no-auto-negotiation);
ethernet-switch-profile {
    ethernet-policer-profile {
        input-priority-map {
            ieee802.1p premium [ values ];
        }
        output-priority-map {
            classifier {
                premium {
                    forwarding-class class-name {
                        loss-priority (high | low);
                    }
                }
            }
        }
    }
    policer cos-policer-name {
        aggregate {

```

```

        bandwidth-limit bps;
        burst-size-limit bytes;
    }
    premium {
        bandwidth-limit bps;
        burst-size-limit bytes;
    }
}
tag-protocol-id;
}
(mac-learn-enable | no-mac-learn-enable);
}
(flow-control | no-flow-control);
ignore-l3-incompletes;
link-mode (automatic | full-duplex | half-duplex);
(lloopback | no-loopback);
keepalives <interval seconds> <down-count number> <up-count number>;
speed (1g | 10m | 100m | 10m-100m | auto-negotiation);
source-address-filter {
    mac-address;
}
source-filtering | no-source-filtering;
}
flexible-vlan-tagging;
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time (up milliseconds | down milliseconds);
interface-transmit-statistics;
(keepalives <down-count number> <interval seconds> <up-count number> |
no-keepalives);
layer2-policer {
    apply-groups [ group-names ];
    apply-groups-except [ group-names ];
}
link-mode (automatic | full-duplex);
mac mac-address;
mtu bytes;
multi-chassis-protection peer-ip-address {
    interface interface-name;
}
native-vlan-id number;
no-gratuitous-arp-request;
optics-options {
    alarm low-light-alarm {
        (link-down | syslog);
    }
    warning low-light-warning {
        (link-down | syslog);
    }
    wavelength nm;
}
passive-monitor-mode;
per-unit-scheduler;
speed (10m | 100m | 1g | auto | oc3 | oc12 | oc48);
stacked-vlan-tagging;
traceoptions {
    flag flag;

```

```

    }
    transmit-bucket {
        overflow discard;
        rate percentage;
        threshold bytes;
    }
    (traps | no-traps);
    unidirectional;
    vlan-tagging;
}

interface-name {
    unit logical-unit-number {
        disable;
        accept-source-mac {
            mac-address mac-address {
                policer {
                    input policer-name;
                    output policer-name;
                }
            }
        }
        account-layer2-overhead (Interface Level) {
            value;
            egress bytes;
            ingress bytes;
        }
        accounting-profile name;
        advisory-options {
            downstream-rate rate;
            upstream-rate rate;
        }
        arp-resp (restricted|unrestricted);
        bandwidth rate;
        clear-dont-fragment-bit;
        copy-tos-to-outer-ip-header;
        demux-destination family;
        encapsulation (vlan-bridge | vlan-vpls);
        epd-threshold cells plp1 cells;
        filter filter-name;
        inner-vlan-id-range start start-id end end-id;
        input-vlan-map {
            (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
            inner-tag-protocol-id tpid;
            inner-vlan-id number;
            tag-protocol-id tpid;
            vlan-id number;
        }
        interface-shared-with psdnumerical-index;
        layer2-policer {
            input-hierarchical-policer policer-name;
            input-policer policer-name;
            input-three-color policer-name;
            output-policer policer-name;
            output-three-color policer-name;
        }
    }
}

```

```

}
multi-chassis-protection peer-ip-address {
  interface interface-name;
}
native-inner-vlan-id number;
output-vlan-map {
  (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
  inner-tag-protocol-id tpid;
  inner-vlan-id number;
  tag-protocol-id tpid;
  vlan-id number;
}
peer-interface interface-name;
peer-unit unit-number;
plp-to-clp;
proxy-arp <restricted | unrestricted>;
rpm {
  (client | server);
  twamp-server;
}
swap-by-poppush;
vlan-id number;
vlan-id-list [ vlan-id vlan-id-vlan-id ];
vlan-id-range number-number;
vlan-tags (inner <tpid.>vlan-id | inner-list [ vlan-id vlan-id-vlan-id ] |
  inner-range <tpid.>vlan-id-vlan-id) outer <tpid.>vlan-id;
}

unit logical-unit-number {
  family ethernet-switching {
    filter {
      group filter-group-number;
      (input filter-name | input-list [ filter-names ]);
      (output filter-name | output-list [ filter-names ]);
      (inner-vlan-id-list [ vlan-ids ] | vlan-id number | vlan-id-list [ number
        number-number ]);
    }
    interface-mode (access | trunk);
    policer {
      input policer-name;
      output policer-name;
    }
    vlan-rewrite {
      translate old-vlan-id new-vlan-id;
    }
    vlan {
      members [ all vlan-identifiers ];
    }
  }
  family inet {
    filter {
      group filter-group-number;
      (input filter-name | input-list [ filter-names ]);
      (output filter-name | output-list [ filter-names ]);
    }
    input-hierarchical-policer policer-name;
    mac-validate (loose | strict);
  }
}

```

```

mtu bytes;
no-neighbor-learn;
no-redirects;
policer {
    arp policer-template-name;
    input policer-name;
    output policer-name;
}
primary;
receive-options-packets;
receive-ttl-exceeded;
rpf-check {
    fail-filter filter-name;
    mode loose;
}
sampling {
    (input | output | input output);
}
simple-filter {
    input filter-name;
}
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
unnumbered-address interface-name <destination address>
    <destination-profile profile-name> <preferred-source-address address>;
}

family inet6 {
    address ipv6-address {
        destination destination-address;
        eui-64;
        ndp ipv6-address <l2-interface interface-name> <(mac mac-address |
            multicast-mac multicast-mac-address) <publish>>;
        preferred;
        primary;
        vrrp-inet6-group group-number {
            (accept-data | no-accept-data);
            fast-interval milliseconds;
            inet6-advertise-interval seconds;
            (no-preempt; | ... the following preempt statement ...)
            preempt {
                hold-time seconds;
            }
            priority number;
            track {
                interface interface-name {
                    bandwidth-threshold bits-per-second priority-cost priority;
                    priority-cost priority;
                }
                priority-hold-time seconds;
                route ip-address-prefix/prefix-length routing-instance instance-name
                    priority-cost priority;
            }
        }
    }
}

```

```
    virtual-inet6-address [ addresses ];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
    }
}
(dad-disable | no-dad-disable);
filter {
    group filter-group-number;
    (input filter-name | input-list [ filter-names ]);
    (output filter-name | output-list [ filter-names ]);
}
input-hierarchical-policer policer-name;
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode loose;
}
sampling {
    (input | output | input output);
}
unnumbered-address interface-name preferred-source-address address;
}

family iso {
    address iso-address;
    mtu bytes;
}

family mlfr-end-to-end {
    bundle logical-interface-name;
}

family mpls {
    filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
    input-hierarchical-policer policer-name;
    maximum-labels maximum-labels;
    mtu bytes;
    policer {
        input policer-name;
        output policer-name;
    }
}
```

```

    }
  }

  family vpls {
    core-facing;
    filter {
      group filter-group-number;
      (input filter-name | input-list [ filter-names ]);
      (output filter-name | output-list [ filter-names ]);
    }
    policer {
      input policer-name;
      output policer-name;
    }
  }
}
}
}
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

[\[edit logical-systems\] Hierarchy Level](#)

The following lists the statements that can be configured at the **[edit logical-systems]** hierarchy level that are also documented in this manual. For more information about logical systems, see the *Junos OS Routing Protocols Configuration Guide*.

```

logical-systems logical-system-name {
  interfaces interface-name {
    unit logical-unit-number {
      accept-source-mac {
        mac-address mac-address {
          policer {
            input cos-policer-name;
            output cos-policer-name;
          }
        }
      }
    }
  }
  allow-any-vci;
  atm-scheduler-map (map-name | default);
  bandwidth rate;
  backup-options {
    interface interface-name;
  }
  cell-bundle-size cells;
  clear-dont-fragment-bit;
  compression {
    rtp {
      f-max-period number;
      port {
        minimum port-number;
        maximum port-number;
      }
    }
  }
}

```

```

    }
    queues [ queue-numbers ];
  }
}
compression-device interface-name;
description text;
interface {
  l2tp-interface-id name;
  (dedicated | shared);
}
dialer-options {
  activation-delay seconds;
  deactivation-delay seconds;
  dial-string [ dial-string-numbers ];
  idle-timeout seconds;
  initial-route-check seconds;
  load-threshold number;
  pool pool;
  remote-name remote-callers;
  watch-list {
    [ routes ];
  }
}
disable;
dlci dlci-identifier;
drop-timeout milliseconds;
dynamic-call-admission-control {
  activation-priority priority;
  bearer-bandwidth-limit kilobits-per-second;
}
encapsulation type;
epd-threshold cells plp1 cells;
fragment-threshold bytes;
input-vlan-map {
  inner-tag-protocol-id;
  inner-vlan-id;
  (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
  tag-protocol-id tpid;
  vlan-id number;
}
interleave-fragments;
inverse-arp;
layer2-policer {
  input-policer policer-name;
  input-three-color policer-name;
  output-policer policer-name;
  output-three-color policer-name;
}
link-layer-overhead percent;
minimum-links number;
mrru bytes;
multicast-dlci dlci-identifier;
multicast-vci vpi-identifier.vci-identifier;
multilink-max-classes number;
multipoint;
oam-liveness {

```



```

    up-count cells;
    down-count cells;
}
oam-period (seconds | disable);
output-vlan-map {
    inner-tag-protocol-id;
    inner-vlan-id;
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-swap);
    tag-protocol-id tpid;
    vlan-id number;
}
passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;
ppp-options {
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
    }
}
dynamic-profile profile-name;
pap {
    default-pap-password password;
    local-name name;
    local-password password;
    passive;
}
}
proxy-arp;
service-domain (inside | outside);
shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst length);
    queue-length number;
}
short-sequence;
transmit-weight number;
(traps | no-traps);
trunk-bandwidth rate;
trunk-id number;
tunnel {
    backup-destination address;
    destination address;
    key number;
    routing-instance {
        destination routing-instance-name;
    }
    source source-address;
    ttl number;

```

```
}
vci vpi-identifier.vci-identifier;
vlan-id number;
vlan-id-list [vlan-id vlan-id-vlan-id]
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
vlan-tags outer tpid.vlan-id inner-list [vlan-id vlan-id-vlan-id]
vpi vpi-identifier;
family family {
    accounting {
        destination-class-usage;
        source-class-usage {
            direction;
        }
    }
}
bundle interface-name;
filter {
    group filter-group-number;
    input filter-name;
    input-list {
        [filter-names ];
    }
    output filter-name;
    output-list {
        [filter-names ];
    }
}
ipsec-sa sa-name;
keep-address-and-control;
mtu bytes;
multicast-only;
no-redirects;
policer {
    arp policer-template-name;
    input policer-template-name;
    output policer-template-name;
}
primary;
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check <fail-filter filter-name > {
    <mode loose>;
}
sampling {
    direction;
}
service {
    input {
        service-set service-set-name <service-filter filter-name>;
        post-service-filter filter-name;
    }
    output {
        service-set service-set-name <service-filter filter-name>;
    }
}
```

```

(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
unnumbered-address interface-name destination address destination-profile
    profile-name;
address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    destination address;
    destination-profile name;
    eui-64;
    multipoint-destination address (dlci dlci-identifier | vci vci-identifier);
    multipoint-destination address {
        epd-threshold cells plp1 cells;
        inverse-arp;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
        oam-period (seconds | disable);
        shaping {
            (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained
                rate burst length);
            queue-length number;
        }
        vci vpi-identifier.vci-identifier;
    }
    preferred;
    primary;
    (vrp-group | vrp-inet6-group) group-number {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        authentication-type authentication;
        authentication-key key;
        fast-interval milliseconds;
        (preempt | no-preempt) {
            hold-time seconds;
        }
        priority-number number;
        track {
            priority-cost seconds;
            priority-hold-time interface-name {
                interface priority;
                bandwidth-threshold bits-per-second {
                    priority;
                }
            }
        }
        route ip-address/mask routing-instance instance-name priority-cost cost;
    }
}
virtual-address [ addresses ];
}
}
}
}
}

```

- Related Documentation**
- *Junos OS Hierarchy and RFC Reference*
 - *Junos® OS Ethernet Interfaces*
 - *Junos® OS Network Interfaces*

[\[edit protocols l2-learning\] Hierarchy Level](#)

```
protocols {
  l2-learning {
    global-mac-limit {
      limit;
      packet-action drop;
    }
    global-mac-statistics;
    global-mac-table-aging-time seconds;
    global-no-mac-learning;
  }
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

[Layer 2 Routing Instances Configuration Hierarchy](#)

Use the **vpls** routing instance type for point-to-multipoint LAN implementations between a set of sites in a VPN.

To configure routing instances for Layer 2 networks, include the following statements:

```
routing-instances {
  routing-instance-name {
    access {
      address-assignment {
        ... same statements as in the address-assignment subhierarchy in [edit access]
        Hierarchy Level ...
      }
      address-protection;
      description text;
      egress-protection {
        context-identifier context-id;
      }
      forwarding-options {
        ...forwarding-options...
      }
      instance-role role;
      instance-type type;
      interface interface-name;
      l2-domain-id-for-l3 id;
      l2vpn-id community;
      layer3-domain-identifier identifier;
      multicast-snooping-options {
        ... same statements as in [edit multicast-snooping-options] Hierarchy Level EXCEPT
        FOR ...
      }
    }
  }
}
```

```

    traceoptions {...} # NOT valid at this level
}
no-irb-layer-2-copy;
no-local-switching;
no-vrf-advertise;
no-vrf-propagate-ttl;
pbb-options {
    default-bvlan bvlan;
    peer-instance instance;
    vlan-id vlan-id isid-list [ isid-numbers ]
}
protocols {
    ... the protocols subhierarchy appears after the main [edit routing-instances
        routing-instance-name] hierarchy ...
}
provider-tunnel {
    ... the provider-tunnel subhierarchy appears after the main [edit routing-instances
        routing-instance-name] hierarchy ...
}
route-distinguisher (as-number:number | ip-address:number);
routing-interface interface;
routing-options {
    ... the routing-options subhierarchy appears after the main [edit routing-instances
        routing-instance-name] hierarchy ...
}
service-groups {
    service-group-name {
        pbb-service-options {
            default-isid isid-number;
            isid isid-number vlan-id-list [ vlan-ids ];
            mac-address mac-address;
        }
        service-type type;
    }
}
services {
    mobile-ip {
        ... same statements as in [edit services mobile-ip] Hierarchy Level ...
    }
}
switch-options {
    ... same statements as in [edit switch-options] Hierarchy Level ...
}
vlan-id (id | all | none);
vlan-model one-to-one;
vlan-tags outer <tpid>vlan-id inner <tpid>vlan-id;
[edit vlans] Hierarchy Level on page 445 {
    ... same statements as in [edit vlans] Hierarchy Level ...
}
vrf-advertise-selective {
    family {
        inet-mvpn;
        inet6-mvpn;
    }
}
vrf-export [ policy-names ];

```

```
vrf-import [ policy-names ];
vrf-propagate-ttl;
vrf-table-label;
vrf-target {
    export community-name;
    import community-name;
}
protocols {
    ... protocols-configuration ...
}
routing-options {
    ... routing-options-configuration ...
}
bridge-domains {
    bridge-domain-name {
        domain-type bridge;
        interface interface-name;
        routing-interface routing-interface-name;
        vlan-id (Bridge Domain or VLAN) (none | all | number);
        vlan-tags outer number inner number;
        bridge-options {
            interface-mac-limit limit {
                packet-action drop;
            }
            interface interface-name {
                interface-mac-limit limit {
                    packet-action drop;
                }
            }
            mac-statistics;
            mac-table-size limit {
                packet-action drop;
            }
            no-mac-learning;
            static-mac mac-address;
        }
    }
}
}
```

With the exception of the **instance-type virtual-switch** statement (which configures a virtual-switch routing instance), you can include the statements at the following hierarchy levels:

- **[edit]**
- **[edit logical-systems *logical-system-name*]**

The **instance-type virtual-switch** statement is not supported at the **[edit logical-systems *logical-system-name*]** hierarchy level.

- Related Documentation**
- *Routing Instances Overview*
 - *Layer 2 Routing Instance Types*

- [Configuring a Layer 2 Virtual Switch on page 1847](#)
- [Configuring a Layer 2 Control Protocol Routing Instance](#)

[\[edit switch-options\] Hierarchy Level](#)

```
switch-options {
  interface interface-name {
    interface-mac-limit {
      number-of-addresses;
      packet-action drop;
    }
    no-mac-learning;
  }
  interface-mac-limit {
    number-of-addresses;
    packet-action drop;
  }
  mac-statistics;
  mac-table-size {
    number-of-addresses;
    packet-action drop;
  }
  no-mac-learning;
  service-id number;
}
```

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

```
vlans {
  vlan-name {
    description text-description;
    domain-type bridge;
    forwarding-options {
      filter {
        input filter-name;
      }
      flood {
        input filter-name;
      }
    }
    interface interface-name;
    l3-interface interface-name;
    multicast-snooping-options {
      ... same statements as in multicast-snooping-options ...
    }
    no-irb-layer-2-copy;
    service-id number;
    switch-options {
      ... the switch-options subhierarchy appears after the main [edit vlans vlan-name]
        hierarchy ...
    }
    vlan-id (all | none | number);
    vlan-id-list [ vlan-id-numbers ];
    vlan-tags outer <tpid.>vlan-id <inner <tpid.>vlan-id>;
```

```
}  
  
vlan-name {  
  switch-options {  
    interface interface-name {  
      interface-mac-limit {  
        limit;  
        packet-action drop;  
      }  
      no-mac-learning;  
      static-mac mac-address {  
        vlan-id number;  
      }  
    }  
  }  
  interface-mac-limit {  
    limit;  
    packet-action drop;  
  }  
  mac-statistics;  
  mac-table-size {  
    number-of-addresses;  
    packet-action drop;  
  }  
  no-mac-learning;  
}  
}  
}
```


bpdu-destination-mac-address

Syntax	bpdu-destination-mac-address provider-bridge-group;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mvrp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvrp] (for virtual switch instance type), [edit protocols mvrp], [edit routing-instances <i>routing-instance-name</i> protocols mvrp] (for virtual switch instance type)
Release Information	Statement introduced in Junos OS Release 10.1 for MX Series routers.
Description	For Multiple VLAN Registration Protocol (MVRP) configurations, specifies the multicast address for MVRP. If configured, the provider MVRP multicast MAC address is used; otherwise, the Junos OS uses the customer MVRP multicast MAC address.
Default	By default, the provider MVRP multicast MAC address is used (if configured). Otherwise, the customer MVRP MAC address is used.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Automatic VLAN Administration Using MVRP • Configuring Multiple VLAN Registration Protocol (MVRP) on page 1845 • Verifying That MVRP Is Working Correctly on page 1937 • Understanding Multiple VLAN Registration Protocol (MVRP) on page 1824

bridge-priority

Syntax	<code>bridge-priority <i>priority</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (<i>mstp</i> <i>rstp</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <i>mstp msti msti-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <i>vstp vlan vlan-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<i>mstp</i> <i>rstp</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>mstp msti msti-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>vstp vlan vlan-id</i>],</p> <p>[edit protocols (<i>mstp</i> <i>rstp</i>)],</p> <p>[edit protocols <i>mstp msti msti-id</i>],</p> <p>[edit protocols <i>vstp vlan vlan-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<i>mstp</i> <i>rstp</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <i>mstp msti msti-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <i>vstp vlan vlan-id</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p>
Description	Determine which bridge is elected as the root bridge. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.
Options	<p><i>priority</i>—The bridge priority can be set only in increments of 4096.</p> <p>Range: 0 through 61,440</p> <p>Default: 32,768</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Bridge Priority for Election of Root Bridge and Designated Bridge on page 504

domain-type

Syntax	domain-type bridge;
Hierarchy Level	[edit bridge-domains <i>bridge-domain-name</i>], [edit vlans on page 445 <i>vlan-name</i>] [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>]
Release Information	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Define the type of domain for a Layer 2 bridge domain or VLAN.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring a Bridge Domain</i>• Configuring a VLAN on page 1829• <i>Configuring a Layer 2 Virtual Switch</i>• Configuring a Layer 2 Virtual Switch on page 1847

encapsulation (Physical Interface)

Syntax	encapsulation (atm-ccc-cell-relay atm-pvc cisco-hdlc cisco-hdlc-ccc cisco-hdlc-tcc ethernet-bridge ethernet-ccc ethernet-over-atm ethernet-tcc ethernet-vpls ethernet-vpls-fr ether-vpls-over-atm-llc ethernet-vpls-ppp extended-frame-relay-ccc extended-frame-relay-ether-type-tcc extended-frame-relay-tcc extended-vlan-bridge extended-vlan-ccc extended-vlan-tcc extended-vlan-vpls flexible-ethernet-services flexible-frame-relay frame-relay frame-relay-ccc frame-relay-ether-type frame-relay-ether-type-tcc frame-relay-port-ccc frame-relay-tcc generic-services multilink-frame-relay-uni-nni ppp ppp-ccc ppp-tcc vlan-ccc vlan-vci-ccc vlan-vpls);
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces rlsq <i>number:number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches (flexible-ethernet-services, ethernet-ccc, and ethernet-tcc options only).
Description	Specify the physical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.
Default	ppp—Use serial PPP encapsulation.
Options	<p>atm-ccc-cell-relay—Use ATM cell-relay encapsulation.</p> <p>atm-pvc—Use ATM PVC encapsulation.</p> <p>cisco-hdlc—Use Cisco-compatible High-Level Data Link Control (HDLC) framing.</p> <p>cisco-hdlc-ccc—Use Cisco-compatible HDLC framing on CCC circuits.</p> <p>cisco-hdlc-tcc—Use Cisco-compatible HDLC framing on TCC circuits for connecting different media.</p> <p>ethernet-bridge—Use Ethernet bridge encapsulation on Ethernet interfaces that have bridging enabled and that must accept all packets.</p> <p>ethernet-ccc—Use Ethernet CCC encapsulation on Ethernet interfaces that must accept packets carrying standard Tag Protocol ID (TPID) values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, CCC is not supported.</p> <p>ethernet-over-atm—For interfaces that carry IPv4 traffic, use Ethernet over ATM encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces. As defined in RFC 2684, <i>Multiprotocol Encapsulation over ATM Adaptation Layer 5</i>, this encapsulation type allows ATM interfaces to connect to devices that support only bridge protocol data units (BPDUs). Junos OS does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet LLC/SNAP frames with IP or ARP in the payload, and drops the rest. For packets destined to the Ethernet LAN, a route lookup is done using the destination</p>

IP address. If the route lookup yields a full address match, the packet is encapsulated with an LLC/SNAP and MAC header, and the packet is forwarded to the ATM interface.

ethernet-tcc—For interfaces that carry IPv4 traffic, use Ethernet TCC encapsulation on interfaces that must accept packets carrying standard TPID values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC is not supported.

ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard TPID values. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

ethernet-vpls-fr—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

ethernet-vpls-ppp—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer layer 2 PPP connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use it to forward the packet into a given VPLS instance.

ether-vpls-over-atm-llc—For ATM intelligent queuing (IQ) interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

extended-frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC.

extended-frame-relay-ether-type-tcc—Use extended Frame Relay ether type TCC for Cisco-compatible Frame Relay for DLCIs 1 through 1022. This encapsulation type is used for circuits with different media on either side of the connection.

extended-frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect different media. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC.

extended-vlan-bridge—Use extended VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q VLAN tagging and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

extended-vlan-ccc—Use extended VLAN encapsulation on CCC circuits with Gigabit Ethernet and 4-port Fast Ethernet interfaces that must accept packets carrying 802.1Q values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, extended VLAN CCC is not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC is not supported.

extended-vlan-tcc—For interfaces that carry IPv4 traffic, use extended VLAN encapsulation on TCC circuits with Gigabit Ethernet interfaces on which you want to use 802.1Q tagging. For 4-port Gigabit Ethernet PICs, extended VLAN TCC is not supported.

extended-vlan-vpls—Use extended VLAN VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.



NOTE: The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

flexible-ethernet-services—For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), use flexible Ethernet services encapsulation when you want to configure multiple per-unit Ethernet encapsulations. Aggregated Ethernet bundles can use this encapsulation type. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs), and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

flexible-frame-relay—For IQ interfaces only, use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value from 1 through 1022.

frame-relay—Use Frame Relay encapsulation.

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with the Cisco Frame Relay.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media.

frame-relay-port-ccc—Use Frame Relay port CCC encapsulation to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect different media.

generic-services—Use generic services encapsulation for services with a hierarchical scheduler.

multilink-frame-relay-uni-nni—Use MLFR UNI NNI encapsulation. This encapsulation is used on link services, voice services interfaces functioning as FRF.16 bundles, and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

ppp—Use serial PPP encapsulation.

ppp-ccc—Use serial PPP encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

ppp-tcc—Use serial PPP encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.

vlan-ccc—Use Ethernet VLAN encapsulation on CCC circuits.

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only. All logical interfaces configured on the Ethernet interface must also have the encapsulation type set to **vlan-vci-ccc**.

vlan-vpls—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN tagging and VPLS enabled. Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.



NOTE: Label-switched interfaces (LSIs) do not support VLAN VPLS encapsulation. Therefore, you can only use VLAN VPLS encapsulation on a PE-router-to-CE-router interface and not a core-facing interface.

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

**Related
Documentation**

- *Configuring Interface Encapsulation on Physical Interfaces*
- *Configuring CCC Encapsulation for Layer 2 VPNs*
- *Configuring Layer 2 Switching Cross-Connects Using CCC*
- *Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits*
- *Configuring ATM Interface Encapsulation*
- *Configuring ATM-to-Ethernet Interworking*
- *Configuring VLAN Encapsulation*
- *Configuring Extended VLAN Encapsulation*
- *Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces*
- *Configuring Interfaces for Layer 2 Circuits*
- *Configuring Interface Encapsulation on PTX Series Packet Transport Switches*
- *Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)*
- *Configuring MPLS LSP Tunnel Cross-Connects Using CCC*
- *Configuring TCC*
- [Configuring VPLS Interface Encapsulation on page 5337](#)
- [Configuring Interfaces for VPLS Routing on page 5336](#)
- *Defining the Encapsulation for Switching Cross-Connects*
- *Understanding Encapsulation on an Interface*

family

```

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

```

```

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

```

```

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

```

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



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

Options *family*—Protocol family:

- **any**—Protocol-independent family used for Layer 2 packet filtering



NOTE: This option is not supported on T4000 Type 5 FPCs.

- **ethernet-switching**—(M Series and T Series routers only) Configure only when the physical interface is configured with **ethernet-bridge** type encapsulation or when the logical interface is configured with **vlan-bridge** type encapsulation
- **ccc**—Circuit cross-connect protocol suite
- **inet**—Internet Protocol version 4 suite
- **inet6**—Internet Protocol version 6 suite
- **iso**—International Organization for Standardization Open Systems Interconnection (ISO OSI) protocol suite
- **mlfr-end-to-end**—Multilink Frame Relay FRF.15
- **mlfr-uni-nni**—Multilink Frame Relay FRF.16
- **multilink-ppp**—Multilink Point-to-Point Protocol
- **mpls**—Multiprotocol Label Switching (MPLS)
- **pppoe**—Point-to-Point Protocol over Ethernet
- **tcc**—Translational cross-connect protocol suite
- **tnp**—Trivial Network Protocol
- **vpls**—(M Series and T Series routers only) Virtual private LAN service


The remaining statements are explained separately.

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

Related Documentation

- *Configuring the Protocol Family*
- *Example: Configuring E-LINE and E-LAN Services for a PBB Network*
- *Junos Services Interfaces Configuration Release 11.2*

fast-aps-switch

Syntax	fast-aps-switch;
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options aps]
Release Information	Statement introduced in Junos OS Release 12.1.
Description	(M320 routers with Channelized OC3/STM1 Circuit Emulation PIC with SFP only and EX Series switches) Reduce the Automatic Protection Switching (APS) switchover time in Layer 2 circuits.
	<div>  <p>NOTE:</p> <ul style="list-style-type: none"> Configuring this statement reduces the APS switchover time only when the Layer 2 circuit encapsulation type for the interface receiving traffic from a Layer 2 circuit neighbor is SAToP. When the fast-aps-switch statement is configured in revertive APS mode, you must configure an appropriate value for revert time to achieve reduction in APS switchover time. To prevent the logical interfaces in the data path from being shut down, configure appropriate hold-time values on all the interfaces in the data path that support TDM. The fast-aps-switch statement cannot be configured when the APS annex-b option is configured. The interfaces that have the fast-aps-switch statement configured cannot be used in virtual private LAN service (VPLS) environments. </div>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> <i>Reducing APS Switchover Time in Layer 2 Circuits</i>

global-mac-limit

Syntax	<code>global-mac-limit <i>limit</i> { <code>packet-action</code> drop; }</code>
Hierarchy Level	[edit protocols l2-learning]
Release Information	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6.
Description	(MX Series routers and EX Series switches only) Limit the number of media access control (MAC) addresses learned from the logical interfaces on the router.
Default	393,215 MAC addresses
Options	<i>limit</i> —Number of MAC addresses that can be learned systemwide. Range: 20 through 1,048,575 The remaining statement is explained separately in the “Summary of Bridge Domain Configuration Statements” chapter.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Limiting the Number of MAC Addresses Learned from Each Logical Interface</i>

global-mac-move

Syntax	global-mac-move { notification-time <i>seconds</i> ; threshold-count <i>count</i> ; threshold-time <i>seconds</i> ; }
Hierarchy Level	[edit protocols l2-learning]
Release Information	Statement introduced in Junos OS Release 9.4. Support for logical systems added in Junos OS Release 9.6.
Description	(MX Series routers and EX Series switches only) Set parameters for media access control (MAC) address move reporting.
Default	By default, MAC moves notify every second, with a threshold time of 1 second and a threshold count of 50.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring MAC Move Parameters</i>

global-mac-statistics

Syntax	global-mac-statistics;
Hierarchy Level	[edit protocols l2-learning]
Release Information	Statement introduced in Junos OS Release 9.2. Support for logical systems added in Junos OS Release 9.6.
Description	(MX Series routers and EX Series switches only) Enable MAC accounting for the entire router or switch.
Default	disabled
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Enabling MAC Accounting</i>

global-mac-table-aging-time

Syntax	global-mac-table-aging-time <i>seconds</i> ;
Hierarchy Level	[edit protocols l2-learning]
Release Information	Statement introduced in Junos OS Release 9.2. Support for logical systems added in Junos OS Release 9.6.
Description	(MX Series routers and EX Series switches only) Configure the timeout interval for entries in the MAC table.
Default	300 seconds
Options	<i>seconds</i> —Time elapsed before MAC table entries are timed out and entries are deleted from the table. Range: 10 through 1 million
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the MAC Table Timeout Interval

global-no-mac-learning

Syntax	global-no-mac-learning;
Hierarchy Level	[edit protocols l2-learning]
Release Information	Statement introduced in Junos OS Release 9.2. Support for logical systems added in Junos OS Release 9.6.
Description	(MX Series routers and EX Series switches only) Disable MAC learning for the entire router or switch.
Default	MAC learning is enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Disabling Layer 2 Learning and Forwarding on page 1845

inner-tag-protocol-id

Syntax	<code>inner-tag-protocol-id <i>tpid</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure the IEEE 802.1Q TPID value to rewrite for the inner tag.</p> <p>All TPIDs you include in input and output VLAN maps must be among those you specify at the [edit interfaces <i>interface-name</i> gether-options ethernet-switch-profile tag-protocol-id [<i>tpids</i>]] hierarchy level.</p> <p>On MX Series routers, you can use this statement for Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs.</p>
Default	If the inner-tag-protocol-id statement is not configured, the TPID value is 0x8100.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Configuring Inner and Outer TPIDs and VLAN IDs</i>

inner-vlan-id

Syntax	<code>inner-vlan-id <i>number</i>;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code>input-vlan-map],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code>output-vlan-map]</code>
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>Specify the VLAN ID to rewrite for the inner tag of the final packet.</p> <p>On MX Series routers, you can use this statement for Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs.</p> <p>You cannot include the inner-vlan-id statement with the swap statement, swap-push statement, push-push statement, or push-swap statement and the inner-vlan-id statement at the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>output-vlan-map]</code> hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the inner-vlan-id statement you include at the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</code> hierarchy level.</p>
Options	<p><i>number</i>—VLAN ID number.</p> <p>Range: 0 through 4094</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Inner and Outer TPIDs and VLAN IDs</i>

input-vlan-map (Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP, and 10-Gigabit Ethernet SFP)

Syntax	<pre>input-vlan-map { (pop pop-pop pop-swap push push-push swap swap-push swap-swap); inner-tag-protocol-id <i>tpid</i>; inner-vlan-id <i>number</i>; tag-protocol-id <i>tpid</i>; vlan-id <i>number</i>; }</pre>
Hierarchy Level	<pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>pop-pop, pop-swap, push-push, swap-push, and swap-swap statements introduced in Junos OS Release 8.1.</p>
Description	<p>For Gigabit Ethernet IQ and 10-Gigabit Ethernet SFPP interfaces only, define the rewrite profile to be applied to incoming frames on this logical interface.</p> <p>The statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Stacking a VLAN Tag</i> • output-vlan-map (Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP, and 10-Gigabit Ethernet SFP) on page 1911

interface

Syntax	<code>interface <i>interface-name</i>;</code>
Hierarchy Level	[edit bridge-domains <i>bridge-domain-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>], [edit vlans on page 445 <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 8.4. Support for top-level configuration for the virtual-switch type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch. Support for logical systems added in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	(MX Series routers and EX Series switches only) Specify the logical interfaces to include in the bridge domain, VLAN, VPLS instance, or virtual switch.
Options	<i>interface-name</i> —Name of a logical interface. For more information about how to configure logical interfaces, see the <i>Junos® OS Network Interfaces</i> .
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring a Bridge Domain</i>• Configuring a VLAN on page 1829• <i>Configuring a Layer 2 Virtual Switch</i>• Configuring a Layer 2 Virtual Switch on page 1847

interface (MVRP)

Syntax	<pre>interface (all <i>interface-name</i>) { join-timer <i>milliseconds</i>; leave-timer <i>milliseconds</i>; leaveall-timer <i>milliseconds</i>; point-to-point; registration (forbidden normal restricted); }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mvrp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvrp] (for virtual switch instance type),</p> <p>[edit protocols mvrp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols mvrp] (for virtual switch instance type)</p>
Release Information	Statement introduced in Junos OS Release 10.1 for MX Series routers.
Description	Specify interfaces on which to configure Multiple VLAN Registration Protocol (MVRP).
Default	By default, MVRP is disabled.
Options	<p>all—All interfaces on the router or switch.</p> <p><i>interface-name</i>—Names of interface to be configured for MVRP.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Automatic VLAN Administration Using MVRP • Configuring Multiple VLAN Registration Protocol (MVRP) on page 1845 • Verifying That MVRP Is Working Correctly on page 1937 • Understanding Multiple VLAN Registration Protocol (MVRP) on page 1824

interface (Spanning Tree)

Syntax	<pre>interface <i>interface-name</i> { bpdv-timeout-action { alarm; block; } cost <i>cost</i>; edge; mode (p2p shared); no-root-port; priority <i>interface-priority</i>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols (<i>mstp</i> <i>rstp</i> <i>vstp</i>)], [edit logical-systems <i>logical-system-name</i> protocols <i>vstp</i> vlan <i>vlan-id</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<i>mstp</i> <i>rstp</i> <i>vstp</i>)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>vstp</i> vlan <i>vlan-id</i>], [edit protocols (<i>mstp</i> <i>rstp</i> <i>vstp</i>)], [edit protocols <i>vstp</i> vlan <i>vlan-id</i>], [edit routing-instances <i>routing-instance-name</i> protocols (<i>mstp</i> <i>rstp</i> <i>vstp</i>)], [edit routing-instances <i>routing-instance-name</i> protocols <i>vstp</i> vlan <i>vlan-id</i>]</pre>
Release Information	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6.
Description	Configure the interface to participate in the RSTP or MSTP instance.
Options	<i>interface-name</i> —Name of a Gigabit Ethernet or 10-Gigabit Ethernet interface. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Spanning-Tree Instance Interface on page 5042

interface-mac-limit

Syntax	<pre>interface-mac-limit <i>limit</i> { packet-action drop; }</pre>
Hierarchy Level	<pre>[edit bridge-domains <i>bridge-domain-name</i> bridge-options], [edit bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options], [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> switch-options], [edit logical-systems <i>logical-system-name</i> switch-options interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> switch-options], [edit routing-instances <i>routing-instance-name</i> switch-options interface <i>interface-name</i>], [edit switch-options], [edit switch-options on page 430], [edit switch-options interface <i>interface-name</i>], [edit switch-options on page 430 interface <i>interface-name</i>], [edit vlans on page 445 <i>vlan-name</i> switch-options], [edit vlans on page 445 <i>vlan-name</i> switch-options interface <i>interface-name</i>]</pre>
Release Information	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the switch-options statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the virtual-switch type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit switch-options], [edit switch-options interface <i>interface-name</i>], [edit vlans <i>vlan-name</i> switch-options], and [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>(MX Series routers or EX Series switches only) Configure a limit to the number of MAC addresses that can be learned from a bridge domain, VLAN, virtual switch, or set of bridge domains or VLANs.</p>
Default	1024 MAC addresses for each logical interface.

Options *limit*—Maximum number of MAC addresses learned from an interface.

Range: 1 through 131,071 MAC addresses per interface

The remaining statement is explained separately.


Required Privilege routing—To view this statement in the configuration.

Level routing-control—To add this statement to the configuration.

**Related
Documentation**

- *Layer 2 Learning and Forwarding for Bridge Domains Overview*
- [Layer 2 Learning and Forwarding for VLANs Overview on page 1822](#)
- *Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
- [Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port on page 1823](#)

interface-mode

Syntax	interface-mode (access trunk);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching] [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge]
Release Information	Statement introduced in Junos OS Release 9.2. Statement introduced in Junos Release 12.3R2 for EX Series switches.
Description	Determines whether the logical interface accepts or discards packets based on VLAN tags. Specify the trunk option to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the vlan-id-list statement, then forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the access option to accept packets with no VLAN ID, then forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the vlan-id statement.
<div style="display: flex; align-items: center;">  <div> <p>NOTE: On MX Series routers, if you want IGMP snooping to be functional for a bridge domain, then you should not configure interface-mode and irb for that bridge. Such a configuration commit succeeds, but IGMP snooping is not functional, and a message informing the same is displayed. For more information, see <i>Configuring a Trunk Interface on a Bridge Network</i>.</p> </div> </div>	
Options	<p>access—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the vlan-id statement.</p> <p>trunk—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the vlan-id-list statement.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring a Logical Interface for Access Mode</i> • <i>Configuring a Logical Interface for Trunk Mode</i>

interfaces

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

isid

Syntax	<code>isid isid-number vlan-id-list [vlan-ids];</code>
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> service-groups <i>service-group-name</i> pbb-service-options]
Release Information	Statement introduced in JUNOS Release 10.0.
Description	For IEEE 802.1ah provider backbone bridge (PBB) configurations, configure the service identifier (I-SID) for the customer routing instance (I-component) service group.
Options	<i>isid</i> —Service identifier. Enter an I-SID in the range from 256 through 16777214 . <i>vlan-id-list [vlan-ids]</i> —List of service VLANs (S-VLANs).
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring E-LINE and E-LAN Services for a PBB Network</i>

isis-list

Syntax	isis-list all-service-groups;
Hierarchy Level	[edit interfaces <i>pseudo-logical-interface-name</i> unit <i>logical-unit-number</i> family bridge]
Release Information	Statement introduced in JUNOS Release 10.0.
Description	For IEEE 802.1ah provider backbone bridge (PBB) configurations, map all service identifiers (I-SIDs) specified for the service groups.
Options	all-service-groups —Map all service identifiers (I-SIDs) for the specified service groups.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring E-LINE and E-LAN Services for a PBB Network</i>

join-timer (MVRP)

Syntax	<code>join-timer <i>milliseconds</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols <i>mvrp</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> <i>mvrp</i> interface (all <i>interface-name</i>)] (for virtual switch instance type),</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> <i>mvrp</i>] (for virtual switch instance type),</code> <code>[edit logical-systems <i>logical-system-name</i> protocols <i>mvrp</i> interface (all <i>interface-name</i>)],</code> <code>[edit protocols <i>mvrp</i> interface (all <i>interface-name</i>)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols <i>mvrp</i>] (for virtual switch instance</code> <code> type),</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols <i>mvrp</i> interface (all <i>interface-name</i>)]</code> <code> (for virtual switch instance type)</code>
Release Information	Statement introduced in Junos OS Release 10.1 for MX Series routers.
Description	For Multiple VLAN Registration Protocol (MVRP), configure the maximum interval interfaces must wait before sending MVRP protocol data units (PDUs).
Default	200 milliseconds
Options	<i>milliseconds</i> —Interval that the interface must wait before sending MVRP PDUs. Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Automatic VLAN Administration Using MVRP• Configuring Multiple VLAN Registration Protocol (MVRP) on page 1845• Verifying That MVRP Is Working Correctly on page 1937• Understanding Multiple VLAN Registration Protocol (MVRP) on page 1824• leaveall-timer on page 1904• leave-timer on page 1905• point-to-point on page 1914• registration on page 1921

l2-learning

Syntax	l2-learning { global-mac-limit <i>limit</i> ; global-mac-statistics; global-mac-table-aging-time <i>seconds</i> ; global-no-mac-learning; }
Hierarchy Level	[edit protocols]
Release Information	Statement introduced in Junos OS Release 8.4.
Description	(MX Series routers and EX Series switches only) Configure Layer 2 address learning and forwarding properties globally. The statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Layer 2 Learning and Forwarding Overview</i>

l3-interface

Syntax	l3-interface <i>interface-name-logical-unit-number</i> ;
Hierarchy Level	[edit interfaces <i>ge-chassis/slot/port</i> unit <i>logical-unit-number</i> family ethernet-switching] [edit vlans <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. [edit vlans <i>vlan-name</i>] hierarchy level introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers.
Description	Associate a Layer 3 interface with the VLAN. Configure Layer 3 interfaces on trunk ports to allow the interface to transfer traffic between multiple VLANs. Within a VLAN, traffic is bridged, while across VLANs, traffic is routed.
Default	No Layer 3 (routing) interface is associated with the VLAN.
Options	<i>interface-name-logical-unit-number</i> —Name of a logical interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

leaveall-timer (MVRP)

Syntax	<code>leaveall-timer <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mvrp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvrp interface (all <i>interface-name</i>)] (for virtual switch instance type),</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvrp] (for virtual switch instance type),</p> <p>[edit logical-systems <i>logical-system-name</i> protocols mvrp interface (all <i>interface-name</i>)],</p> <p>[edit protocols mvrp interface (all <i>interface-name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols mvrp] (for virtual switch instance type),</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols mvrp interface (all <i>interface-name</i>)] (for virtual switch instance type)</p>
Release Information	Statement introduced in Junos OS Release 10.1 for MX Series routers.
Description	For Multiple VLAN Registration Protocol (MVRP), configure the interval at which the LeaveAll state operates on the interface.
Default	10000 milliseconds
Options	<i>milliseconds</i> —Interval between the sending of Leave All messages. Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Automatic VLAN Administration Using MVRP • Configuring Multiple VLAN Registration Protocol (MVRP) on page 1845 • Verifying That MVRP Is Working Correctly on page 1937 • Understanding Multiple VLAN Registration Protocol (MVRP) on page 1824 • join-timer on page 1902 • leave-timer on page 1905 • point-to-point on page 1914 • registration on page 1921

leave-timer (MVRP)

Syntax	<code>leave-timer <i>milliseconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mvrp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvrp interface (all <i>interface-name</i>)] (for virtual switch instance type),</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvrp] (for virtual switch instance type),</p> <p>[edit logical-systems <i>logical-system-name</i> protocols mvrp interface (all <i>interface-name</i>)],</p> <p>[edit protocols mvrp interface (all <i>interface-name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols mvrp] (for virtual switch instance type),</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols mvrp interface (all <i>interface-name</i>)] (for virtual switch instance type)</p>
Release Information	Statement introduced in Junos OS Release 10.1 for MX Series routers.
Description	For Multiple VLAN Registration Protocol (MVRP), configure the number of milliseconds the switch retains a VLAN in the Leave state before the VLAN is unregistered. If the interface receives a join message before this timer expires, the VLAN remains registered.
Default	1000 milliseconds
Options	<i>milliseconds</i> —Interval that the switch retains a VLAN in the Leave state before the VLAN is unregistered. At a minimum, set the leave-timer interval at twice the join-timer interval. Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Automatic VLAN Administration Using MVRP • Configuring Multiple VLAN Registration Protocol (MVRP) on page 1845 • Verifying That MVRP Is Working Correctly on page 1937 • Understanding Multiple VLAN Registration Protocol (MVRP) on page 1824 • join-timer on page 1902 • leaveall-timer on page 1904 • point-to-point on page 1914 • registration on page 1921

mac-statistics

Syntax	mac-statistics;
Hierarchy Level	<p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options],</p> <p>[edit logical-systems <i>logical-system-name</i> switch-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> switch-options],</p> <p>[edit switch-options],</p> <p>[edit switch-options on page 430],</p> <p>[edit vlans on page 445 <i>vlan-name</i> switch-options]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the switch-options statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the virtual-switch type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit switch-options] and [edit vlans <i>vlan-name</i> switch-options] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	(MX Series routers and EX Series switches only) Enable MAC accounting either for a specific bridge domain or VLAN, or for a set of bridge domains or VLANs associated with a Layer 2 trunk port.
Default	disabled
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Layer 2 Learning and Forwarding for Bridge Domains Overview</i> • <i>Layer 2 Learning and Forwarding for VLANs Overview on page 1822</i> • <i>Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports</i> • <i>Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port on page 1823</i>

mac-table-size

Syntax	<pre>mac-table-size <i>limit</i> { packet-action drop; }</pre>
Hierarchy Level	<p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options],</p> <p>[edit logical-systems <i>logical-system-name</i> switch-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> switch-options],</p> <p>[edit switch-options],</p> <p>[edit switch-options on page 430],</p> <p>[edit vlans on page 445 <i>vlan-name</i> switch-options]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the switch-options statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the virtual-switch type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit switch-options] and [edit vlans <i>vlan-name</i> switch-options] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Modify the size of the MAC address table for the bridge domain or VLAN, a set of bridge domains or VLANs associated with a trunk port, or a virtual switch. The default is 5120 MAC addresses.</p>
Options	<p>limit—Specify the maximum number of addresses in the MAC address table.</p> <p>Range: 16 through 1,048,575 MAC addresses</p> <p>Default: 5120 MAC addresses</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Layer 2 Learning and Forwarding for Bridge Domains Overview</i> • Layer 2 Learning and Forwarding for VLANs Overview on page 1822 • <i>Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports</i>

- [Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port on page 1823](#)

mvrp

Syntax mvrp {
 bpdu-destination-mac-address provider-bridge-group;
 join-timer *milliseconds*;
 leave-timer *milliseconds*;
 leaveall-timer *milliseconds*;
 interface (all | *interface-name*) {
 join-timer *milliseconds*;
 leave-timer *milliseconds*;
 leaveall-timer *milliseconds*;
 point-to-point;
 registration (forbidden | normal | restricted);
 }
 no-dynamic-vlan;
 traceoptions (Spanning Trees) {
 file *filename* <files *number* > <size *size* > <no-stamp | world-readable |
 no-world-readable>;
 flag *flag*;
 }
 }
 }

Hierarchy Level [edit logical-systems *logical-system-name* protocols],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]
 (for virtual switch instance type)
 [edit protocols],
 [edit routing-instances *routing-instance-name* protocols] (for virtual switch instance type),

Release Information Statement introduced in Junos OS Release 10.1 for MX Series routers.

Description For Layer 2 networks, configure Multiple VLAN Registration Protocol (MVRP) to dynamically share VLAN information and dynamically configure needed VLANs. Maintaining VLAN configurations based on active VLANs reduces the amount of traffic traveling in the network, saving network resources. MVRP is configured on trunk interfaces.

The remaining statements are explained separately.


Default MVRP is disabled by default.

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


Related Documentation

- [Example: Configuring Automatic VLAN Administration Using MVRP](#)
- [Configuring Multiple VLAN Registration Protocol \(MVRP\) on page 1845](#)
- [Verifying That MVRP Is Working Correctly on page 1937](#)
- [Understanding Multiple VLAN Registration Protocol \(MVRP\) on page 1824](#)

native-vlan-id

Syntax	<code>native-vlan-id <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>ge-fpc/pic/port</i>] [edit interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>For 1-, 4-, and 8-port Gigabit Ethernet IQ2 and IQ2-E PICs, for 1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs configured for 802.1Q flexible VLAN tagging, for all Ethernet interfaces on MX Series routers, and for aggregated Ethernet interfaces on IQ2 and IQ2-E PICs or MX Series DPCs, configure mixed tagging support for untagged packets on a port. When the native-vlan-id statement is included with the <i>flexible-vlan-tagging</i> statement, untagged packets are accepted on the same mixed VLAN-tagged port.</p> <p>The logical interface on which untagged packets are received must be configured with the same native VLAN ID as that configured on the physical interface. To configure the logical interface, include the vlan-id statement (matching the native-vlan-id statement on the physical interface) at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.</p> <p>When the native-vlan-id statement is included with the interface-mode the statement, untagged packets are accepted and forwarded within the bridge domain that is configured with the matching VLAN ID.</p>
	<div>  <p>NOTE: As per 12.3 Ethernet Interfaces Configuration Guide, native-vlan-id enables the tagged interface to accept untagged frames. However, when native-vlan-id is configured, VLAN ID is popped in egress as long as the VID in the outgoing frame matches the native-vlan-id.</p> </div>
Options	<p>number—VLAN ID number.</p> <p>Range: (ACX Series routers and EX Series switches) 0 through 4094.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring Mixed Tagging Support for Untagged Packets Configuring a Logical Interface for Access Mode <i>flexible-vlan-tagging</i>

no-mac-learning

Syntax	no-mac-learning;
Hierarchy Level	<p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options],</p> <p>[edit logical-systems <i>logical-system-name</i> switch-options],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> switch-options],</p> <p>[edit switch-options],</p> <p>[edit switch-options on page 430],</p> <p>[edit switch-options on page 430 interface <i>interface-name</i>],</p> <p>[edit vlans on page 445 <i>vlan-name</i> switch-options],</p> <p>[edit vlans on page 445 <i>vlan-name</i> switch-options interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the switch-options statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the virtual-switch type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit switch-options], [edit switch-options interface <i>interface-name</i>], [edit vlans <i>vlan-name</i> switch-options], and [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	(MX Series routers and EX Series switches only) Disable MAC learning for a virtual switch, for a bridge domain or VLAN, for a specific logical interface in a bridge domain or VLAN, or for a set of bridge domains or VLANs associated with a Layer 2 trunk port.
	<div>  <p>NOTE: When MAC learning is disabled for a VPLS routing instance, traffic is not load balanced and only one of the equal-cost next hops is used.</p> </div>
Default	MAC learning is enabled. Use no-mac-learning to disable MAC learning.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Layer 2 Learning and Forwarding for Bridge Domains Overview • Layer 2 Learning and Forwarding for VLANs Overview on page 1822 • Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports • Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port on page 1823

output-vlan-map (Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP, and 10-Gigabit Ethernet SFP)

Syntax	<pre>output-vlan-map { (pop pop-pop pop-swap push push-push swap swap-push swap-swap); inner-tag-protocol-id <i>tpid</i>; inner-vlan-id <i>number</i>; tag-protocol-id <i>tpid</i>; vlan-id <i>number</i>; }</pre>
Hierarchy Level	<pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</pre>
Release Information	Statement introduced before Junos OS Release 7.4. pop-pop , pop-swap , push-push , swap-push , and swap-swap statements added in Junos OS Release 8.1.
Description	<p>For Gigabit Ethernet IQ and 10-Port 10-Gigabit Ethernet SFPP interfaces only, define the rewrite operation to be applied to outgoing frames on this logical interface.</p> <p>The statements are explained separately.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Stacking and Rewriting Gigabit Ethernet VLAN Tags • input-vlan-map (Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP, and 10-Gigabit Ethernet SFP) on page 1893

packet-action

Syntax packet-action drop;

Hierarchy Level [edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],
 [edit bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],
 [edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],
 [edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options **interface-mac-limit** *limit*],
 [edit logical-systems *logical-system-name* switch-options **interface-mac-limit** *limit*],
 [edit protocols **l2-learning** **global-mac-limit** *limit*],
 [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],
 [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],
 [edit routing-instances *routing-instance-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],
 [edit routing-instances *routing-instance-name* switch-options **interface-mac-limit** *limit*],
 [edit switch-options on page 430 interface *interface-name* **interface-mac-limit** *limit*],
 [edit switch-options on page 430 **interface-mac-limit** *limit*],
 [edit switch-options on page 430 **mac-table-size** *limit*],
 [edit vlans on page 445 *vlan-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],
 [edit vlans on page 445 *vlan-name* switch-options **interface-mac-limit** *limit*],
 [edit vlans on page 445 *vlan-name* switch-options **mac-table-size** *limit*]

Release Information Statement introduced in Junos OS Release 8.4.
 Support for the **switch-options** statement added in Junos OS Release 9.2.
 Support for top-level configuration for the **virtual-switch** type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.
 Support for logical systems added in Junos OS Release 9.6.
 [edit switch-options interface *interface-name* **interface-mac-limit** *limit*], [edit switch-options **interface-mac-limit** *limit*], [edit switch-options **mac-table-size** *limit*], [edit vlans *vlan-name* switch-options interface *interface-name* **interface-mac-limit** *limit*], [edit vlans *vlan-name* switch-options **interface-mac-limit** *limit*], and [edit vlans *vlan-name* switch-options **mac-table-size** *limit*] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.

Description	(MX Series routers and EX Series switches only) Specify that packets for new source MAC addresses be dropped after the MAC address limit is reached. If this statement is not configured, then packets for new source MAC addresses are forwarded by default.
Default	Disabled. The default is for packets for new source MAC addresses to be forwarded after the MAC address limit is reached.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Layer 2 Learning and Forwarding for Bridge Domains Overview</i>• Layer 2 Learning and Forwarding for VLANs Overview on page 1822• <i>Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports</i>• Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port on page 1823

point-to-point (MVRP)

Syntax	point-to-point;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mvrp interface (all <i>interface-name</i>)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvrp interface (all <i>interface-name</i>)] (for virtual switch instance type), [edit protocols mvrp interface (all <i>interface-name</i>)], [edit routing-instances <i>routing-instance-name</i> protocols mvrp interface (all <i>interface-name</i>)] (for virtual switch instance type)
Release Information	Statement introduced in Junos OS Release 10.1 for MX Series routers. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	(Optional) For Multiple VLAN Registration Protocol (MVRP) configurations, configure an interface to be recognized as a point-to-point connection. If specified, a point-to-point subset of the MRP state machine is used to provide a simpler and more efficient method to accelerate convergence on the network. Point-to-point must be enabled after enabling MVRP for the interface to be recognized as a point-to-point connection.
Default	MVRP is disabled by default. point-to-point is disabled by default.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Automatic VLAN Administration Using MVRP• Configuring Multiple VLAN Registration Protocol (MVRP) on page 1845• Verifying That MVRP Is Working Correctly on page 1937• Understanding Multiple VLAN Registration Protocol (MVRP) on page 1824• join-timer on page 1902• leaveall-timer on page 1904• leave-timer on page 1905• registration on page 1921

pop

Syntax	pop;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify the VLAN rewrite operation to remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed. On MX series routers, this statement can be used on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Removing a VLAN Tag</i>

pop-pop

Syntax	pop-pop;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>Specify the VLAN rewrite operation to remove both the outer and inner VLAN tags of the frame.</p> <p>On MX Series routers, you can use this statement on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs.,</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Removing the Outer and Inner VLAN Tags</i>

pop-swap

Syntax	pop-swap;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify the VLAN rewrite operation to remove the outer VLAN tag of the frame, and replace the inner VLAN tag of the frame with a user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. On MX Series routers, you can use this statement for Gigabit Ethernet IQ, IQ2, and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag</i>

proxy-arp

Syntax	proxy-arp (restricted unrestricted);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.6 for EX Series switches. restricted added in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 12.2 for the QFX Series..
Description	For Ethernet interfaces only, configure the router or switch to respond to any ARP request, as long as the router or switch has an active route to the ARP request's target address.
Default	Proxy ARP is not enabled. The router or switch responds to an ARP request only if the destination IP address is its own.
Options	<ul style="list-style-type: none">• none—The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address.• restricted—(Optional) The router or switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are in the same subnet. The router or switch must also have a route to the target IP address.• unrestricted—(Optional) The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address. <p>Default: unrestricted</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Restricted and Unrestricted Proxy ARP on page 1850• <i>Configuring Proxy ARP (CLI Procedure)</i>• <i>Example: Configuring Proxy ARP on an EX Series Switch</i>• <i>Configuring Gratuitous ARP</i>

push

Syntax	<code>push;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code> input-vlan-map],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code> output-vlan-map]</code>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Specify the VLAN rewrite operation to add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag.</p> <p>On MX Series routers, You can use this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces.</p> <p>If you include the push statement in the configuration, you must also include the pop statement at the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]</code> hierarchy level.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Stacking a VLAN Tag</i>

push-push

Syntax	push-push;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>Specify the VLAN rewrite operation to push two VLAN tags in front of the frame.</p> <p>On MX Series routers, You can use this command for Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs. .</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Stacking Two VLAN Tags</i>

registration

Syntax	registration (forbidden normal restricted);
Hierarchy Level	[edit protocols mvrp interface (all <i>interface-name</i>)]; [edit routing-instances <i>routing-instance-name</i> protocols mvrp interface (all <i>interface-name</i>)] (for virtual switch instance type); [edit logical-systems <i>logical-system-name</i> protocols mvrp interface (all <i>interface-name</i>)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvrp interface (all <i>interface-name</i>)] (for virtual switch instance type);
Release Information	Statement introduced in Junos OS Release 10.1 for MX Series routers.
Description	For Multiple VLAN Registration Protocol (MVRP) configurations, configure the registration mode for the interface.
Default	normal
Options	<p>forbidden—The interface or interfaces do not register and do not participate in MVRP.</p> <p>normal—The interface or interfaces accept MVRP messages and participate in MVRP.</p> <p>restricted—The interface or interfaces ignore all MVRP JOIN messages received for VLANs that are not statically configured for MVRP on the interface.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Automatic VLAN Administration Using MVRP • Configuring Multiple VLAN Registration Protocol (MVRP) on page 1845 • Verifying That MVRP Is Working Correctly on page 1937 • Understanding Multiple VLAN Registration Protocol (MVRP) on page 1824 • join-timer on page 1902 • leaveall-timer on page 1904 • leave-timer on page 1905 • point-to-point on page 1914

rstp

Syntax	<pre> rstp { bpd-block-on-edge; bpd-destination-mac-address provider-bridge-group; bridge-priority priority; extended-system-id; force-version stp; forward-delay seconds; hello-time seconds; max-age seconds; interface interface-name { bpd-timeout-action { alarm; block; } cost cost; edge; mode (p2p shared); no-root-port; priority interface-priority; } priority-hold-time seconds; traceoptions { file filename <files number> <size size> <world-readable no-world-readable>; flag flag <flag-modifier> <disable>; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],</p> <p>[edit protocols],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.4.</p> <p>bpd-block-on-edge statement added in Junos OS Release 9.4.</p> <p>bpd-timeout-action statement added in Junos OS Release 9.4.</p> <p>Support for logic systems added in Junos OS Release 9.6.</p>
Description	Configure RSTP parameters.
Options	The statements are explained separately.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Rapid Spanning-Tree Protocol on page 5073

service-id

Syntax	<code>service-id <i>number</i>;</code>
Hierarchy Level	[edit switch-options on page 430] [edit vlans on page 445 <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers.
Description	Specify a service identifier for each multi-chassis aggregated Ethernet interface that belongs to a link aggregation group (LAG).
Options	number —A number that identifies a particular service. Range: 1 through 65535
Required Privilege Level	system —To view this statement in the configuration. system control —To add this statement to the configuration.

static-mac

Syntax	<code>static-mac mac-address { vlan-id number; }</code>
Hierarchy Level	[edit bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>], [edit vlans on page 445 <i>vlan-name</i> switch-options interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6. [edit vlans <i>vlan-name</i> switch-options interface <i>interface name</i>] hierarchy level introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	(MX Series routers and EX Series switches only) Configure a static MAC address for a logical interface in a bridge domain or VLAN. The vlan-id option can be specified for static-macs only if vlan-id all is configured for the bridging domain or VLAN.
Options	mac-address —MAC address vlan-id number —(Optional) VLAN identifier to associate with static MAC address.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Layer 2 Learning and Forwarding for Bridge Domains Overview</i>• <i>Layer 2 Learning and Forwarding for VLANs Overview on page 1822</i>

swap

Syntax	swap;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify the VLAN rewrite operation to replace a VLAN tag. The outer VLAN tag of the frame is overwritten with the user-specified VLAN tag information. On MX Series routers, you can enter this statement for Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and aggregated Ethernet using Gigabit Ethernet IQ interfaces.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Rewriting the VLAN Tag on Tagged Frames on page 2290

swap-push

Syntax	swap-push;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>Specify the VLAN rewrite operation to replace the outer VLAN tag of the frame with a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.</p> <p>On MX Series routers, this command can be used on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Rewriting a VLAN Tag and Adding a New Tag on page 1851

swap-swap

Syntax	swap-swap;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify the VLAN rewrite operation to replace both the inner and the outer VLAN tags of the frame with a user-specified VLAN tag value. On MX Series routers, you can use this statement for Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Rewriting the Inner and Outer VLAN Tags on page 2289

switch-options

Syntax	<pre>switch-options { interface <i>interface-name</i> { interface-mac-limit <i>limit</i> { packet-action drop; } no-mac-learning; static-mac <i>static-mac-address</i> { vlan-id <i>number</i>; } } interface-mac-limit <i>limit</i> { packet-action drop; } mac-statistics; mac-table-size <i>limit</i> { packet-action drop; } no-mac-learning; }</pre>
Hierarchy Level	[edit vlans on page 445 <i>vlan--name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> vlans on page 445 <i>vlan-name</i>], [edit routing-instances <i>routing-instance-name</i> vlans on page 445 <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers.
Description	Configure Layer 2 learning and forwarding properties for a VLAN or a virtual switch. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

tag-protocol-id (TPID to Rewrite)

Syntax	<code>tag-protocol-id <i>tpid</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces only, configure the outer TPID value. All TPIDs you include in input and output VLAN maps must be among those you specify at the [edit interfaces <i>interface-name</i> gigether-options ethernet-switch-profile tag-protocol-id [<i>tpids</i>]] hierarchy level.</p> <p>For 10-Gigabit Ethernet LAN/WAN PIC interfaces on T Series routers, value the default TPID value (0x8100) is supported.</p>
Default	If the tag-protocol-id statement is not configured, the TPID value is 0x8100.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring Inner and Outer TPIDs and VLAN IDs

traceoptions (MVRP)

Syntax	<pre>traceoptions { file <i>name</i> <size <i>size</i>> <files <i>number</i>> <(world-readable no-world-readable)>; flag <i>flag</i> <flag-modifier> <disable>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mvrp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvrp] (for virtual switch instance type), [edit protocols mvrp], [edit routing-instances <i>routing-instance-name</i> protocols mvrp] (for virtual switch instance type)
Release Information	Statement introduced in Junos OS Release 10.1 for MX Series routers.
Description	For Multiple VLAN Registration Protocol (MVRP), configure tracing options.
Default	Traceoptions is disabled.
Options	<p>disable —(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. To include the file statement, you must specify a filename. Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place MVRP tracing output in the file <code>/var/log/mvrp-log</code>.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files, in the range from 2 through 1000. The default is 1 trace file. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>flag <i>flag</i>—Specify which tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none">• all—Enable all trace options flags.• error—Trace all failure conditions.• events—Trace process state change and cleanup events.• pdu—Trace RAPS PDU reception and transmission.• socket—Trace socket activity.• state-machine—Trace information about the state machine.• timers—Trace protocol timers.

no-world-readable—(Optional) Prevent any user from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the ***trace-file*** again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten. The file size range is from 10240 through 4294967295. The default file size is 1MB.

world-readable—(Optional) Allow any user to read the log file.

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

Related Documentation	• <i>Example: Configuring Automatic VLAN Administration Using MVRP</i>
	• Configuring Multiple VLAN Registration Protocol (MVRP) on page 1845
	• Verifying That MVRP Is Working Correctly on page 1937
	• Understanding Multiple VLAN Registration Protocol (MVRP) on page 1824

vlan-id (Bridge Domain or VLAN)

Syntax	<code>vlan-id (all none <i>number</i>);</code>
Hierarchy Level	<code>[edit bridge-domains <i>bridge-domain-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i></code> <code>bridge-domains <i>bridge-domain-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>]</code> <code>[edit vlans <i>vlan-name</i>]</code>
Release Information	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for Layer 2 trunk ports added in Junos OS Release 9.2.</p> <p>Support for SRX 3400, SRX 3600, SRX 5600, and SRX 5800 devices added in Junos OS Release 9.6.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Specify a VLAN identifier (VID) to include in the packets sent to and from the bridge domain or VLAN, or a VPLS routing instance.



NOTE: When configuring a VLAN identifier for provider backbone bridge (PBB) routing instances, dual-tagged VIDs and the none option are not permitted.

Options *number*—A valid VLAN identifier. If you configure multiple bridge domains or VLANs with a valid VLAN identifier, you must specify a unique VLAN identifier for each. However, you can use the same VLAN identifier for bridge domains or VLANs that belong to different virtual switches. Use this option to send single tagged frames with the specified VLAN identifier over VPLS VT interfaces.



NOTE: If you specify a VLAN identifier, you cannot also use the all option. They are mutually exclusive.

all—Specify that the bridge domain or VLAN spans all the VLAN identifiers configured on the member logical interfaces.



NOTE: You cannot specify the all option if you include a routing interface in the bridge domain or VLAN.

none—Specify to enable shared VLAN learning or to send untagged frames over VPLS VT interfaces.



NOTE: Multichassis link aggregation (MC-LAG) does not support the none option with the vlan-id statement with bridge domains or VLANs.

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


Related Documentation

- *Configuring a Bridge Domain*
- [Configuring a VLAN on page 1829](#)
- *Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances*
- [Configuring VLAN Identifiers for VLANs and VPLS Routing Instances on page 1830](#)
- *Configuring Bridge Domains as Switches for Layer 2 Trunk Ports*
- *Configuring a Layer 2 Virtual Switch*
- [Configuring a Layer 2 Virtual Switch on page 1847](#)
- *Example: Configuring E-LINE and E-LAN Services for a PBB Network*
- *Example: Configuring Interfaces and Routing Instances for a User Logical System*
- *bridge-domains*
- [vlans on page 445](#)

vlan-id (VLAN ID to Rewrite)

Syntax	<code>vlan-id number;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code>input-vlan-map],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code>output-vlan-map]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2, 10-Gigabit Ethernet LAN/WAN PIC, and IQ2-E interfaces and aggregated Ethernet using Gigabit Ethernet IQ interfaces, specify the line VLAN identifiers to be rewritten at the input or output interface.</p> <p>You cannot include the vlan-id statement with the swap statement, swap-push statement, push-push statement, or push-swap statement at the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]</code> hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the vlan-id statement that you include at the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</code> hierarchy level.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Rewriting the VLAN Tag on Tagged Frames on page 2290• Binding VLAN IDs to Logical Interfaces on page 2291

vlan-id-list

Syntax	<code>vlan-id-list [<i>vlan-id-numbers</i>];</code>
Hierarchy Level	<p>[edit bridge-domains <i>bridge-domain-name</i>], [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>] [edit interfaces <i>interface-name</i> unit 0], [edit vlans <i>vlan-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.4. Support for logical systems added in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Specify a VLAN identifier list to use for a bridge domain or VLAN in trunk mode.</p> <p>Specify the trunk option in the interface-mode statement to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the vlan-id-list statement to forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the access option to accept packets with no VLAN ID to forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the vlan-id statement.</p>
Options	<p><i>vlan-id-numbers</i>—Valid VLAN identifiers. You can combine individual numbers with range lists including a hyphen.</p> <p>Range: 0 through 4095</p>
<div>  <p>NOTE: On EX Series switches, the range is 0 through 4094.</p> </div>	
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Bridge Domain • Configuring a VLAN on page 1829 • Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances • Configuring VLAN Identifiers for VLANs and VPLS Routing Instances on page 1830

vlan-rewrite

Syntax	vlan-rewrite translate (200 500 201 501)
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>number</i> family bridge interface-mode trunk] [edit interfaces <i>interface-name</i> unit <i>number</i> family ethernet-switching interface-mode trunk]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Translates an incoming VLAN to a bridge-domain VLAN, corresponding counter translation at egress. Supports translation of VLAN 200 to VLAN 500 and VLAN 201 to VLAN 501. Other valid VLANs pass through without translation.
Options	translate 200 500 —Translates incoming packets with VLAN 200 to 500. translate 201 501 —Translates incoming packets with VLAN 201 to 501. translate 202 502 —Translates incoming packets with VLAN 202 to 502.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Rewriting a VLAN Tag and Adding a New Tag on page 1851

vlan-tags

Syntax	<code>vlan-tags outer <i>number</i> inner <i>number</i>;</code>
Hierarchy Level	[edit bridge-domains <i>bridge-domain-name</i>], [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>] [edit vlans on page 445 <i>vlan-name</i>]
Release Information	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify dual VLAN identifier tags for a bridge domain, VLAN, or VPLS routing instance.
Options	outer <i>number</i> —A valid VLAN identifier. inner <i>number</i> —A valid VLAN identifier.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring a Bridge Domain • Configuring a VLAN on page 1829 • Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances • Configuring VLAN Identifiers for VLANs and VPLS Routing Instances on page 1830 • Configuring a Layer 2 Virtual Switch. • Configuring a Layer 2 Virtual Switch on page 1847

Administration

- [Routine Monitoring on page 1937](#)
- [Operational Commands on page 1939](#)

Routine Monitoring

- [Verifying That MVRP Is Working Correctly on page 1937](#)

Verifying That MVRP Is Working Correctly

Purpose	After configuring your MX Series router or EX Series switch to participate in Multiple VLAN Registration Protocol (MVRP), verify that the configuration is properly set and that MVRP messages are being sent and received on your switch.
Action	1. Confirm that the router is declaring Virtual LANs (VLANs).

Show that MVRP is enabled:

```
user@host> show mvrp
MVRP configuration for routing instance 'default-switch'
MVRP dynamic VLAN creation : Enabled
MVRP BPDU MAC address      : Customer bridge group (01-80-C2-00-00-21)
MVRP timers (ms)
  Interface   Join   Leave  LeaveAll
  ge-11/3/0   200   800    10000
```

Show the MVRP applicant state:

```
user@host> show mvrp applicant-state
MVRP applicant state for routing instance 'default-switch'
(V0) Very anxious observer, (VP) Very anxious passive, (VA) Very anxious new,
(AN) Anxious new, (AA) Anxious active, (QA) Quiet active, (LA) Leaving active,
(AO) Anxious observer, (QO) Quiet observer, (LO) Leaving observer,
(AP) Anxious passive, (QP) Quiet passive
```

VLAN Id	Interface	State
100	ge-11/3/0	Declaring (QA)
200	ge-11/3/0	Declaring (QA)
300	ge-11/3/0	Declaring (QA)

2. Confirm that VLANs are registered on interfaces.

List VLANs in the registered state:

```
user@host> show mvrp registration-state
MVRP registration state for routing instance 'default-switch'
```

VLAN Id	Interface	Registrar State	Forced State	Managed State	STP State
100	ge-11/3/0	Registered	Registered	Normal	Forwarding
200	ge-11/3/0	Registered	Registered	Normal	Forwarding
300	ge-11/3/0	Empty	Empty	Normal	Forwarding

3. Display a list of VLANs created dynamically.

List dynamic VLAN membership:

```
user@host> show mvrp dynamic-vlan-memberships
MVRP dynamic vlans for routing instance 'default-switch'
(s) static vlan, (f) fixed registration
```

VLAN Id	Interfaces
100	ge-3/3/0 ge-3/0/5
200	ge-3/3/0 ge-3/0/5

Meaning The output of `show mvrp applicant-state` shows that trunk interface `ge-11/3/0` is declaring (sending out) interest in the VLAN IDs `100`, `200`, and `300` and MVRP is operating properly.

The output of `show mvrp registrant-state` shows the registrar state for VLANs `100` and `200` as **Registered**, indicating that these VLANs are receiving traffic from a customer site. VLAN `300` is in an **Empty** state and is not receiving traffic from a customer site.

The output of the **show mvrp dynamic-vlan-membership** shows that VLANs 100 and 200 are created dynamically (here, on an MX Series router operating as an aggregation switch between MX Series routers operating as edge switches). VLANs created statically are marked with an (s) (which is not indicated in this output).

**Related
Documentation**

- *Example: Configuring Automatic VLAN Administration Using MVRP on EX Series Switches*
- *Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)*
- *Controlling the Management State of a VLAN in MVRP Configurations (CLI Procedure)*

Operational Commands

show ethernet-switching flood

Syntax	<code>show ethernet-switching flood</code> <code><brief detail extensive></code> <code><event-queue></code> <code><instance <i>instance-name</i>></code> <code><logical-system <i>logical-system-name</i>></code> <code><route (all-ce-flood all-ve-flood alt-root-flood bd-flood mlp-flood re-flood)></code> <code><vlan-name <i>vlan-name</i>></code>
Release Information	Command introduced in Junos OS Release 12.3R2. Command introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	(MX Series routers and EX Series switches only) Display Ethernet-switching flooding information.
Options	<p>none—Display all Ethernet-switching flooding information for all VLANs.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>event-queue—(Optional) Display the queue of pending Ethernet-switching flood events.</p> <p>instance <i>instance-name</i>—(Optional) Display Ethernet-switching flooding information for the specified routing instance.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Display Ethernet-switching flooding information for the specified logical system.</p> <p>route (all-ce-flood all-ve-flood alt-root-flood bd-flood mlp-flood re-flood)—(Optional) Display the following:</p> <ul style="list-style-type: none">• all-ce-flood—Display the route for flooding traffic to all customer edge routers or switches if no-local-switching is enabled.• all-ve-flood—Display the route for flooding traffic to all VPLS edge routers or switches if no-local-switching is enabled.• alt-root-flood—Display the Spanning Tree Protocol (STP) alt-root flooding route used for the interface.• bd-flood—Display the route for flooding traffic of a VLAN if no-local-switching is not enabled.• mlp-flood—Display the route for flooding traffic to MAC learning chips.• re-flood—Display the route for Routing Engine flooding to all interfaces. <p>vlan-name <i>vlan-name</i>—(Optional) Display Ethernet-switching flooding information for the specified VLAN.</p>
Required Privilege Level	view
List of Sample Output	show ethernet-switching flood on page 1941

[show ethernet-switching flood brief on page 1941](#)
[show ethernet-switching flood detail on page 1941](#)
[show ethernet-switching flood extensive on page 1942](#)

Sample Output

show ethernet-switching flood

```

user@host> show ethernet-switching flood
Name: __juniper_private1__
CEs: 0
VEs: 0
Name: default-switch
CEs: 9
VEs: 0
VLAN Name: VLAN101
Flood Routes:
  Prefix    Type          Owner                NhType    NhIndex
  0x3057b/51 FLOOD_GRP_COMP_NH __all_ces__         comp      12866
  0x30004/51 FLOOD_GRP_COMP_NH __re_flood__        comp      12863
VLAN Name: VLAN102
Flood Routes:
  Prefix    Type          Owner                NhType    NhIndex
  0x3057c/51 FLOOD_GRP_COMP_NH __all_ces__         comp      12875
  0x30005/51 FLOOD_GRP_COMP_NH __re_flood__        comp      12872
VLAN Name: VLAN103
Flood Routes:
  Prefix    Type          Owner                NhType    NhIndex
  0x3057d/51 FLOOD_GRP_COMP_NH __all_ces__         comp      12884
  0x30006/51 FLOOD_GRP_COMP_NH __re_flood__        comp      12881

```

show ethernet-switching flood brief

```

user@host> show ethernet-switching flood brief
Name                Active CEs    Active VEs
__juniper_private1__ 0              0
default-switch       9              0

```

show ethernet-switching flood detail

```

user@host> show ethernet-switching flood detail
Name: __juniper_private1__
CEs: 0
VEs: 0
Name: default-switch
CEs: 9
VEs: 0
VLAN Name: VLAN101
Flood Routes:
  Prefix    Type          Owner                NhType    NhIndex
  0x3057b/51 FLOOD_GRP_COMP_NH __all_ces__         comp      12866
  0x30004/51 FLOOD_GRP_COMP_NH __re_flood__        comp      12863
VLAN Name: VLAN102
Flood Routes:
  Prefix    Type          Owner                NhType    NhIndex
  0x3057c/51 FLOOD_GRP_COMP_NH __all_ces__         comp      12875
  0x30005/51 FLOOD_GRP_COMP_NH __re_flood__        comp      12872
VLAN Name: VLAN103
Flood Routes:
  Prefix    Type          Owner                NhType    NhIndex

```

```

0x3057d/51 FLOOD_GRP_COMP_NH __all_ces__      comp      12884
0x30006/51 FLOOD_GRP_COMP_NH __re_flood__     comp      12881

```

show ethernet-switching flood extensive

```

user@host> show ethernet-switching flood extensive
Name: __juniper_private1__
CEs: 0
VEs: 0
Name: default-switch
CEs: 9
VEs: 0
VLAN Name: VLAN101
  Flood route prefix: 0x3057b/51
  Flood route type: FLOOD_GRP_COMP_NH
  Flood route owner: __all_ces__
  Flood group name: __all_ces__
  Flood group index: 1
  Nexthop type: comp
  Nexthop index: 12866
  Flooding to:
    Name      Type      NhType      Index
    __all_ces__ Group      comp        12860
    Composition: split-horizon
    Flooding to:
      Name      Type      NhType      Index
      ae20.0    CE        ucst        7605

  Flood route prefix: 0x30004/51
  Flood route type: FLOOD_GRP_COMP_NH
  Flood route owner: __re_flood__
  Flood group name: __re_flood__
  Flood group index: 65534
  Nexthop type: comp
  Nexthop index: 12863
  Flooding to:
    Name      Type      NhType      Index
    __all_ces__ Group      comp        12860
    Composition: split-horizon
    Flooding to:
      Name      Type      NhType      Index
      ae20.0    CE        ucst        7605

VLAN Name: VLAN102

  Flood route prefix: 0x3057c/51
  Flood route type: FLOOD_GRP_COMP_NH
  Flood route owner: __all_ces__
  Flood group name: __all_ces__
  Flood group index: 1
  Nexthop type: comp
  Nexthop index: 12875
  Flooding to:
    Name      Type      NhType      Index
    __all_ces__ Group      comp        12869
    Composition: split-horizon
    Flooding to:
      Name      Type      NhType      Index
      ae20.0    CE        ucst        7605

  Flood route prefix: 0x30005/51
  Flood route type: FLOOD_GRP_COMP_NH

```

```

Flood route owner: __re_flood__
Flood group name: __re_flood__
Flood group index: 65534
Nexthop type: comp
Nexthop index: 12872
Flooding to:
  Name      Type      NhType      Index
  __all_ces__  Group    comp        12869
  Composition: split-horizon
  Flooding to:
    Name      Type      NhType      Index
    ae20.0    CE        ucst        7605
VLAN Name: VLAN103

Flood route prefix: 0x3057d/51
Flood route type: FLOOD_GRP_COMP_NH
Flood route owner: __all_ces__
Flood group name: __all_ces__
Flood group index: 1
Nexthop type: comp
Nexthop index: 12884
Flooding to:
  Name      Type      NhType      Index
  __all_ces__  Group    comp        12878
  Composition: split-horizon
  Flooding to:
    Name      Type      NhType      Index
    ae20.0    CE        ucst        7605

Flood route prefix: 0x30006/51
Flood route type: FLOOD_GRP_COMP_NH
Flood route owner: __re_flood__
Flood group name: __re_flood__
Flood group index: 65534
Nexthop type: comp
Nexthop index: 12881
Flooding to:
  Name      Type      NhType      Index
  __all_ces__  Group    comp        12878
  Composition: split-horizon
  Flooding to:
    Name      Type      NhType      Index
    ae20.0    CE        ucst        7605
VLAN Name: VLAN104

```

show ethernet-switching interface

Syntax	show ethernet-switching interface <brief detail extensive> <interface-name>
Release Information	Command introduced in Junos OS Release 12.3R2. Command introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	(MX Series routers and EX Series switches only) Display Layer 2 learning information for all the interfaces.
Options	none —Display Ethernet-switching information for all interfaces. brief detail extensive —(Optional) Display the specified level of output. interface-name —(Optional) Display Ethernet-switching information for the specified interface.
Required Privilege Level	view
List of Sample Output	show ethernet-switching interface ae10.0 on page 1945
Output Fields	Table 133 on page 1944 describes the output fields for the show ethernet-switching interface command. Output fields are listed in the approximate order in which they appear.

Table 133: show ethernet-switching interface Output Fields

Field Name	Field Description
Logical interface	Name of the logical interface.
VLAN members	VLANs associated with this interface.
Tag	VLAN ID.
MAC limit	Number of MAC addresses that can be associated with the interface.
STP state	Spanning tree protocol (STP) state.
Logical interface flags	Status of Layer 2 learning properties for each interface: <ul style="list-style-type: none"> • DL—MAC learning is disabled. • LH—MAC interface limit has been reached.. • AD—Packets are dropped after the MAC interface limit is reached. • DN—The MAC interface is down.
Tagging	Tagging state of the VLAN.

Sample Output

show ethernet-switching interface ae10.0

```

user@host> show ethernet-switching interface ae10.0
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical  Vlan    TAG  MAC    STP    Logical    Tagging
interface members  limit state  interface flags
ae10.0
      VLAN70.. 701    8192
      VLAN70.. 702    1024    Forwarding
      VLAN70.. 703    1024    Forwarding
      VLAN70.. 704    1024    Forwarding
      VLAN70.. 705    1024    Forwarding
      VLAN70.. 706    1024    Forwarding
      VLAN70.. 707    1024    Forwarding
      VLAN70.. 708    1024    Forwarding
      VLAN70.. 709    1024    Forwarding
      VLAN71.. 710    1024    Forwarding
      VLAN71.. 711    1024    Forwarding
      VLAN71.. 712    1024    Forwarding
      VLAN71.. 713    1024    Forwarding
      VLAN71.. 714    1024    Forwarding
      VLAN71.. 715
[...output truncated...]

```

show ethernet-switching statistics

Syntax	show ethernet-switching statistics <instance <i>instance-name</i>> <logical-system <i>logical-system-name</i>> <vlan-name <i>vlan-name</i>>
Release Information	Command introduced in Junos OS Release 12.3R2. Command introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	(MX Series routers and EX Series switches only) Display Ethernet-switching statistics.
Options	none —Display Ethernet-switching statistics for all VLANs in all routing instances. instance <i>instance-name</i> —(Optional) Display statistics for the specified routing instance. logical-system <i>logical-system-name</i> —(Optional) Display Ethernet-switching statistics information for the specified logical system. vlan-name <i>vlan-name</i> —(Optional) Display statistics for the specified VLAN.
Required Privilege Level	view
List of Sample Output	show ethernet-switching statistics on page 1946

Sample Output

show ethernet-switching statistics

```
user@host> show ethernet-switching statistics
Local interface: ae1.0, Index: 1035
  Broadcast packets:          220
  Broadcast bytes   :        13720
  Multicast packets:          130
  Multicast bytes   :       11700
  Flooded packets   :           0
  Flooded bytes    :           0
  Unicast packets   :           0
  Unicast bytes    :           0
  Current MAC count:           0 (Limit 1024)
Local interface: vt-3/3/10.1048576, Index: 1280
  Broadcast packets:          0
  Broadcast bytes   :           0
  Multicast packets:          0
  Multicast bytes   :           0
  Flooded packets   :           2
  Flooded bytes    :          128
  Unicast packets   :          632
  Unicast bytes    :       39184
  Current MAC count:           2
Local interface: ge-3/1/2.0, Index: 1258
  Broadcast packets:          100
  Broadcast bytes   :          6800
  Multicast packets:          200
  Multicast bytes   :       18000
  Flooded packets   :           0
```



```

    Flooded bytes      :          0
    Unicast packets    :         632
    Unicast bytes      :       39184
    Current MAC count:         2 (Limit 1024)
Local interface: ae3.0, Index: 1043
    Broadcast packets:          0
    Broadcast bytes   :          0
    Multicast packets:          0
    Multicast bytes   :          0
    Flooded packets   :          0
    Flooded bytes     :          0
    Unicast packets   :          0
    Unicast bytes     :          0
    Current MAC count:         0 (Limit 1024)
Local interface: ge-3/3/8.0, Index: 1276
    Broadcast packets:          0
    Broadcast bytes   :          0
    Multicast packets:          0
    Multicast bytes   :          0
    Flooded packets   :          0
    Flooded bytes     :          0
    Unicast packets   :          0
    Unicast bytes     :          0
    Current MAC count:         0 (Limit 8192)
Local interface: ae5.0, Index: 1045
    Broadcast packets:          0
    Broadcast bytes   :          0
    Multicast packets:          0
    Multicast bytes   :          0
    Flooded packets   :          0
    Flooded bytes     :          0
    Unicast packets   :          0
    Unicast bytes     :          0
    Current MAC count:         0 (Limit 8192)
Local interface: ae4.0, Index: 1044
    Broadcast packets:         200
    Broadcast bytes   :      13600
    Multicast packets:          0
    Multicast bytes   :          0
    Flooded packets   :          0
    Flooded bytes     :          0
    Unicast packets   :          0
    Unicast bytes     :          0
    Current MAC count:         0 (Limit 8192)
Local interface: ae26.0, Index: 1042
    Broadcast packets:          0
    Broadcast bytes   :          0
    Multicast packets:          0
    Multicast bytes   :          0
    Flooded packets   :          0
    Flooded bytes     :          0
    Unicast packets   :          0
    Unicast bytes     :          0
    Current MAC count:         0 (Limit 8192)
Local interface: ae25.0, Index: 1041
    Broadcast packets:         133
    Broadcast bytes   :        7980
    Multicast packets:      369934
    Multicast bytes   :    59207572
    Flooded packets   :          0
    Flooded bytes     :          0

```

```
Unicast packets :          1433
Unicast bytes   :        119930
Current MAC count:           3 (Limit 8192)
Local interface: ae23.0, Index: 1040
Broadcast packets:          226
Broadcast bytes :        14464
Multicast packets:        585668
Multicast bytes :    153464476
Flooded packets :           0
Flooded bytes   :           0
Unicast packets :        26552
Unicast bytes   :    1947627
Current MAC count:           7 (Limit 8192)
Local interface: ae20.0, Index: 1039
Broadcast packets:          115
Broadcast bytes :        6900
Multicast packets:        395113
Multicast bytes :    61622869
Flooded packets :           0
Flooded bytes   :           0
Unicast packets :        1419
Unicast bytes   :    117924
Current MAC count:           4 (Limit 8192)
```

show ethernet-switching table

Syntax	<pre>show ethernet-switching table <brief count detail extensive> <address> <instance <i>instance-name</i>> <interface <i>interface-name</i>> <interface <i>interface-name</i>> isid <i>isid</i> <address> <vlan-id (all-vlan <i>vlan-id</i>)> <vlan-name (all <i>vlan-name</i>)></pre>
Release Information	<p>Command introduced in Junos OS Release 12.3R2.</p> <p>Command introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	(MX Series routers and EX Series switches only) Display Layer 2 MAC address information.
Options	<p>none—Display all learned Layer 2 MAC address information.</p> <p>brief count detail extensive—(Optional) Display the specified level of output.</p> <p>address—(Optional) Display the specified learned Layer 2 MAC address information.</p> <p>instance <i>instance-name</i>—(Optional) Display learned Layer 2 MAC addresses for the specified routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Display learned Layer 2 MAC addresses for the specified interface.</p> <p>isid <i>isid</i>—(Optional) Display learned Layer 2 MAC addresses for the specified ISID.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Display Ethernet-switching statistics information for the specified logical system.</p> <p>vlan-id (all-vlan <i>vlan-id</i>)—(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.</p> <p>vlan-name (all <i>vlan-name</i>)—(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.</p>
Additional Information	When Layer 2 protocol tunneling is enabled, the tunneling MAC address 01:00:0c:cd:cd:d0 is installed in the MAC table. When the Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP) is configured for Layer 2 protocol tunneling on an interface, the corresponding protocol MAC address is installed in the MAC table.
Required Privilege Level	view
List of Sample Output	<p>show ethernet-switching table on page 1950</p> <p>show ethernet-switching table brief on page 1952</p> <p>show ethernet-switching table count on page 1953</p>

[show ethernet-switching table extensive on page 1954](#)

Output Fields [Table 134 on page 1950](#) describes the output fields for the **show ethernet-switching table** command. Output fields are listed in the approximate order in which they appear.

Table 134: show ethernet-switching table Output fields

Field Name	Field Description
Routing instance	Name of the routing instance.
VLAN name	Name of the VLAN.
MAC address	MAC address or addresses learned on a logical interface.
MAC flags	Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> • S—Static MAC address is configured. • D—Dynamic MAC address is configured. • L—Locally learned MAC address is configured. • SE—MAC accounting is enabled. • NM—Non-configured MAC. • R—Locally learned MAC address is configured.
Logical interface	Name of the logical interface.
MAC count	Number of MAC addresses learned on the specific routing instance or interface.
Learning interface	Name of the logical interface on which the MAC address was learned.
Learning VLAN	VLAN ID of the routing instance or VLAN in which the MAC address was learned.
Layer 2 flags	Debugging flags signifying that the MAC address is present in various lists.
Epoch	Spanning-tree-protocol epoch number identifying when the MAC address was learned. Used for debugging.
Sequence number	Sequence number assigned to this MAC address. Used for debugging.
Learning mask	Mask of the Packet Forwarding Engines where this MAC address was learned. Used for debugging.
IPC generation	Creation time of the logical interface when this MAC address was learned. Used for debugging.

Sample Output**show ethernet-switching table**

```
user@host> show ethernet-switching table
```

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN101	88:e0:f3:bb:07:f0	D	-	ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN102	88:e0:f3:bb:07:f0	D	-	ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN103	88:e0:f3:bb:07:f0	D	-	ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN104	88:e0:f3:bb:07:f0	D	-	ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN1101	00:1f:12:32:f5:c1	D	-	ae0.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN1102	00:1f:12:32:f5:c1	D	-	ae0.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN1103	00:1f:12:32:f5:c1	D	-	ae0.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN1104	00:1f:12:32:f5:c1	D	-	ae0.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN1105	00:1f:12:32:f5:c1	D	-	ae0.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN1106	00:1f:12:32:f5:c1	D	-	ae0.0

[...output truncated...]

show ethernet-switching table brief

user@host> **show ethernet-switching table brief**

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN101	88:e0:f3:bb:07:f0	D	-	ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN102	88:e0:f3:bb:07:f0	D	-	ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN103	88:e0:f3:bb:07:f0	D	-	ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

```

Routing instance : default-switch
  Vlan      MAC      MAC      Age      Logical
  name      address   flags                   interface
  VLAN104   88:e0:f3:bb:07:f0 D        -        ae20.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
          SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
  Vlan      MAC      MAC      Age      Logical
  name      address   flags                   interface
  VLAN1101  00:1f:12:32:f5:c1 D        -        ae0.0
[...output truncated...]

```

show ethernet-switching table count

```

user@host> show ethernet-switching table count
0 MAC address learned in routing instance default-switch VLAN VLAN1000
ae26.0:1000

1 MAC address learned in routing instance default-switch VLAN VLAN101
ae20.0:101

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID   MAC count   Static MAC count
  101             1           0

1 MAC address learned in routing instance default-switch VLAN VLAN102
ae20.0:102

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID   MAC count   Static MAC count
  102             1           0

1 MAC address learned in routing instance default-switch VLAN VLAN103
ae20.0:103

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID   MAC count   Static MAC count
  103             1           0

1 MAC address learned in routing instance default-switch VLAN VLAN104
ae20.0:104

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID   MAC count   Static MAC count
  104             1           0

0 MAC address learned in routing instance default-switch VLAN VLAN105
ae20.0:105

0 MAC address learned in routing instance default-switch VLAN VLAN106
ae20.0:106

0 MAC address learned in routing instance default-switch VLAN VLAN107
ae20.0:107

0 MAC address learned in routing instance default-switch VLAN VLAN108
ae20.0:108

```

```
0 MAC address learned in routing instance default-switch VLAN VLAN109
ae20.0:109

0 MAC address learned in routing instance default-switch VLAN VLAN110
ae20.0:110

1 MAC address learned in routing instance default-switch VLAN VLAN1101
ae0.0:1101

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID      MAC count      Static MAC count
      1101              1              0

1 MAC address learned in routing instance default-switch VLAN VLAN1102
ae0.0:1102

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID      MAC count      Static MAC count
      1102              1              0
[...output truncated...]
```

show ethernet-switching table extensive

```
user@host> show ethernet-switching table extensive

MAC address: 88:e0:f3:bb:07:f0
Routing instance: default-switch
VLAN ID: 101
Learning interface: ae20.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                      Sequence number: 2
Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
Routing instance: default-switch
VLAN ID: 102
Learning interface: ae20.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                      Sequence number: 2
Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
Routing instance: default-switch
VLAN ID: 103
Learning interface: ae20.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                      Sequence number: 2
Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
Routing instance: default-switch
VLAN ID: 104
Learning interface: ae20.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                      Sequence number: 2
Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
Routing instance: default-switch
VLAN ID: 1101
```



```
Learning interface: ae0.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 2
Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
Routing instance: default-switch
VLAN ID: 1102
Learning interface: ae0.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 2
Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
Routing instance: default-switch
VLAN ID: 1103
Learning interface: ae0.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 2
Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
Routing instance: default-switch
VLAN ID: 1104
Learning interface: ae0.0
Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
Epoch: 0                               Sequence number: 2
Learning mask: 0x00000008
```

show interfaces (10-Gigabit Ethernet)

Syntax	<code>show interfaces <i>xe-fpc/pic/port</i></code> <code><brief detail extensive terse></code> <code><descriptions></code> <code><media></code> <code><snmp-index <i>snmp-index</i>></code> <code><statistics></code>
Release Information	Command introduced in Junos OS Release 8.0.
Description	(M320, M120, MX Series, and T Series routers and EX Series switches only) Display status information about the specified 10-Gigabit Ethernet interface.
Options	<p><code><i>xe-fpc/pic/port</i></code>—Display standard information about the specified 10-Gigabit Ethernet interface.</p> <p><code>brief detail extensive terse</code>—(Optional) Display the specified level of output.</p> <p><code>descriptions</code>—(Optional) Display interface description strings.</p> <p><code>media</code>—(Optional) Display media-specific information about network interfaces.</p> <p><code>snmp-index <i>snmp-index</i></code>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><code>statistics</code>—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
List of Sample Output	<p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2) on page 1971</p> <p>show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode) on page 1974</p> <p>show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC) on page 1976</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode) on page 1978</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only) on page 1978</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only) on page 1979</p>
Output Fields	See Table 135 on page 1957 for the output fields for the show interfaces (10-Gigabit Ethernet) command.

Table 135: show interfaces Gigabit Ethernet Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 2376 .	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
WAN-PHY mode	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: Enabled or Disabled for parent interface; Rx-only or Tx-only for child interfaces.	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under “Common Output Fields Description” on page 2376 .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 2376 .	All levels

Table 135: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Link flags	Information about the link. Possible values are described in the “Links Flags” section under “ Common Output Fields Description ” on page 2376.	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	detail extensive none
Schedulers	(Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified
Output Rate	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Egress accounting overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	detail extensive
Ingress accounting overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	detail extensive

Table 135: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Output bytes—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 135 on page 1957.</p>	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 135: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 135: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 135: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 135: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 136 on page 1971 • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—Number of frames that exceed 1518 octets. • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive
OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	extensive
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	extensive

Table 135: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 135: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 135: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 135: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 135: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under "Common Output Fields Description" on page 2376 .	All levels

Table 135: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • push-push—Two VLAN tags are pushed in from the incoming frame. • swap-push—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value. • pop-swap—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed. 	brief detail extensive none
Demux:	<p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> • Source Family Inet • Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under “ Common Output Fields Description ” on page 2376.	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the routing device.	extensive

Table 135: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the Output bytes and Output packets interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under “ Common Output Fields Description ” on page 2376.	detail extensive
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	detail extensive none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	detail extensive
Output Filters	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	detail extensive
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under “ Common Output Fields Description ” on page 2376.	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about address flag (possible values are described in the “Addresses Flags” section under “ Common Output Fields Description ” on page 2376.	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interlace.	detail extensive none

Table 135: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

For Gigabit Ethernet IQ PICs, traffic and MAC statistics output varies. [Table 136 on page 1971](#) describes the traffic and MAC statistics for two sample interfaces, each of which is sending traffic in packets of 500 bytes (including 478 bytes for the Layer 3 packet, 18 bytes for the Layer 2 VLAN traffic header, and 4 bytes for cyclic redundancy check [CRC] information). In [Table 136 on page 1971](#), the **ge-0/3/0** interface is the inbound physical interface, and the **ge-0/0/0** interface is the outbound physical interface. On both interfaces, traffic is carried on logical unit .50 (VLAN 50).

Table 136: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Inbound physical interface	show interfaces ge-0/3/0 extensive	Traffic statistics: Input bytes: 496 bytes per packet, representing the Layer 2 packet MAC statistics: Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes	The additional 4 bytes are for the CRC.
Inbound logical interface	show interfaces ge-0/3/0.50 extensive	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	
Outbound physical interface	show interfaces ge-0/0/0 extensive	Traffic statistics: Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes MAC statistics: Received octets: 478 bytes per packet, representing the Layer 3 packet	For input bytes, the additional 12 bytes includes 6 bytes for the destination MAC address + 4 bytes for VLAN + 2 bytes for the Ethernet type.
Outbound logical interface	show interfaces ge-0/0/0.50 extensive	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	

Sample Output

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2)

```

user@host> show interfaces xe-5/0/0 extensive
Physical interface: xe-5/0/0, Enabled, Physical link is Up
  Interface index: 177, SNMP ifIndex: 99, Generation: 178
  Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Speed: 10Gbps, Loopback:

```

```

None, Source filtering: Enabled,
Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags : None
CoS queues : 8 supported, 4 maximum usable queues
Schedulers : 1024
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:14:f6:b9:f1:f6, Hardware address: 00:14:f6:b9:f1:f6
Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes : 6970332384 0 bps
Output bytes : 0 0 bps
Input packets: 81050506 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Ingress traffic statistics at Packet Forwarding Engine:
Input bytes : 6970299398 0 bps
Input packets: 81049992 0 pps
Drop bytes : 0 0 bps
Drop packets: 0 0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0,
MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets

0 best-effort 81049992 81049992 0
1 expedited-fo 0 0 0
2 assured-forw 0 0 0
3 network-cont 0 0 0

Egress queues: 4 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets

0 best-effort 0 0 0
1 expedited-fo 0 0 0
2 assured-forw 0 0 0
3 network-cont 0 0 0

Active alarms : None
Active defects : None
PCS statistics Seconds
Bit errors 0
Errored blocks 0

```

```

MAC statistics:
Total octets          6970332384
Total packets        81050506
Unicast packets      81050000
Broadcast packets    506
Multicast packets    0
CRC/Align errors    0
FIFO errors          0
MAC control frames   0
MAC pause frames     0
Oversized frames     0
Jabber frames        0
Fragment frames      0
VLAN tagged frames   0
Code violations       0

Filter statistics:
Input packet count    81050506
Input packet rejects  506
Input DA rejects      0
Input SA rejects      0
Output packet count   0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0

Packet Forwarding Engine configuration:
Destination slot: 5

CoS information:
Direction : Output
CoS transmit queue   Bandwidth      Buffer Priority  Limit
                        %      bps      %      usec
0 best-effort        95    950000000    95      0      low  none
3 network-control    5     50000000     5      0      low  none

Direction : Input
CoS transmit queue   Bandwidth      Buffer Priority  Limit
                        %      bps      %      usec
0 best-effort        95    950000000    95      0      low  none
3 network-control    5     50000000     5      0      low  none

Logical interface xe-5/0/0.0 (Index 71) (SNMP ifIndex 95) (Generation 195)
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ] Encapsulation: ENET2
Egress accounting overhead: 100
Ingress accounting overhead: 90

Traffic statistics:
Input bytes : 0
Output bytes : 46
Input packets: 0
Output packets: 1

IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Local statistics:
Input bytes : 0
Output bytes : 46
Input packets: 0
Output packets: 1

Transit statistics:
Input bytes : 0
Output bytes : 0

```

```

Input packets:                0                0 pps
Output packets:               0                0 pps
IPv6 transit statistics:
  Input bytes :                0
  Output bytes :               0
  Input packets:              0
  Output packets:             0
Protocol inet, MTU: 1500, Generation: 253, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 192.1.1/24, Local: 192.1.1.1, Broadcast: 192.1.1.255,
Generation: 265
Protocol multiservice, MTU: Unlimited, Generation: 254, Route table: 0
  Flags: None
  Policer: Input: __default_arp_policer__

```

show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode)

```

user@host> show interfaces xe-1/0/0 extensive
Physical interface: xe-1/0/0, Enabled, Physical link is Up
  Interface index: 141, SNMP ifIndex: 34, Generation: 47
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Loopback: Disabled
  WAN-PHY mode
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Link flags     : None
  CoS queues     : 4 supported
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:a2:10:9d, Hardware address: 00:05:85:a2:10:9d
  Last flapped   : 2005-07-07 11:22:34 PDT (3d 12:28 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes :                0                0 bps
    Output bytes :               0                0 bps
    Input packets:              0                0 pps
    Output packets:             0                0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    HS Link CRC errors: 0, HS Link FIFO overflows: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0,
    Aged packets: 0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
    Resource errors: 0
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets
    0 best-effort   0                0                0
    1 expedited-fo  0                0                0
    2 assured-forw  0                0                0
    3 network-cont  0                0                0
  Active alarms : LOL, LOS, LBL
  Active defects: LOL, LOS, LBL, SEF, AIS-L, AIS-P
  PCS statistics
    Seconds  Count
    Bit errors  0        0
    Errored blocks  0        0
  MAC statistics:
    Receive  Transmit
    Total octets  0        0
    Total packets  0        0
    Unicast packets  0        0
    Broadcast packets  0        0
    Multicast packets  0        0

```

```

CRC/Align errors                0          0
FIFO errors                     0          0
MAC control frames              0          0
MAC pause frames                0          0
Oversized frames               0
Jabber frames                  0
Fragment frames                0
VLAN tagged frames             0
Code violations                 0
Filter statistics:
  Input packet count            0
  Input packet rejects          0
  Input DA rejects              0
  Input SA rejects              0
  Output packet count           0
  Output packet pad count       0
  Output packet error count     0
CAM destination filters: 0, CAM source filters: 0
PMA PHY:
  Seconds      Count  State
  PLL lock     0      0 OK
  PHY light    63159  1 Light Missing
WIS section:
  BIP-B1        0      0
  SEF           434430  434438 Defect Active
  LOS           434430  1 Defect Active
  LOF           434430  1 Defect Active
  ES-S          434430
  SES-S         434430
  SEFS-S        434430
WIS line:
  BIP-B2        0      0
  REI-L         0      0
  RDI-L         0      0 OK
  AIS-L         434430  1 Defect Active
  BERR-SF       0      0 OK
  BERR-SD       0      0 OK
  ES-L          434430
  SES-L         434430
  UAS-L         434420
  ES-LFE        0
  SES-LFE       0
  UAS-LFE       0
WIS path:
  BIP-B3        0      0
  REI-P         0      0
  LOP-P         0      0 OK
  AIS-P         434430  1 Defect Active
  RDI-P         0      0 OK
  UNEQ-P        0      0 OK
  PLM-P         0      0 OK
  ES-P          434430
  SES-P         434430
  UAS-P         434420
  ES-PFE        0
  SES-PFE       0
  UAS-PFE       0
Received path trace:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted path trace: orissa so-1/0/0
6f 72 69 73 73 61 20 73 6f 2d 31 2f 30 2f 30 00 orissa so-1/0/0.
Packet Forwarding Engine configuration:

```

```

Destination slot: 1
CoS information:
  CoS transmit queue      Bandwidth      Buffer      Priority  Limit
                           %      bps      %      bytes
  0 best-effort           95      950000000  95        0      low      none
  3 network-control       5       50000000   5         0      low      none

```

show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC)

```

user@host> show interfaces ge-7/0/0 extensive
Physical interface: ge-7/0/0, Enabled, Physical link is Down
Interface index: 143, SNMP ifIndex: 508, Generation: 208
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Enabled
Device flags   : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
Link flags     : None
Wavelength     : 1550.12 nm, Frequency: 193.40 THz
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:05:85:70:2b:72, Hardware address: 00:05:85:70:2b:72
Last flapped   : 2011-04-20 15:48:54 PDT (18:39:49 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 2, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort           0              0              0

  1 expedited-fo           0              0              0

  2 assured-forw           0              0              0

  3 network-cont
Queue number:      Mapped forwarding classes
  0                best-effort
  1                expedited-forwarding
  2                assured-forwarding
  3                network-control
Active alarms : LINK
Active defects : LINK
MAC statistics:
Total octets      Receive      Transmit
Total packets      0              0

```

```

Unicast packets          0          0
Broadcast packets        0          0
Multicast packets        0          0
CRC/Align errors         0          0
FIFO errors              0          0
MAC control frames       0          0
MAC pause frames         0          0
Oversized frames         0
Jabber frames            0
Fragment frames          0
VLAN tagged frames       0
Code violations           0
Total octets             0          0
Total packets            0          0
Unicast packets          0          0
Broadcast packets        0          0
Multicast packets        0          0
CRC/Align errors         0          0
FIFO errors              0          0
MAC control frames       0          0
MAC pause frames         0          0
Oversized frames         0
Jabber frames            0
Fragment frames          0
VLAN tagged frames       0
Code violations           0
OTN alarms                : None
OTN defects                : None
OTN FEC Mode              : GFEC
OTN Rate                  : Fixed Stuff Bytes 11.0957Gbps
OTN Line Loopback : Enabled
OTN FEC statistics :
  Corrected Errors          0
  Corrected Error Ratio (   0 sec average) 0e-0
OTN FEC alarms:      Seconds    Count  State
  FEC Degrade         0          0  OK
  FEC Excessive        0          0  OK
OTN OC:              Seconds    Count  State
  LOS                  2          1  OK
  LOF                  67164      2  Defect Active
  LOM                  67164      71 Defect Active
  Wavelength Lock      0          0  OK
OTN OTU:
  AIS                  0          0  OK
  BDI                  65919      4814 Defect Active
  IAE                  67158      1  Defect Active
  TTIM                 7          1  OK
  SF                   67164      2  Defect Active
  SD                   67164      3  Defect Active
  TCA-ES               0          0  OK
  TCA-SES              0          0  OK
  TCA-UAS              80         40  OK
  TCA-BBE              0          0  OK
  BIP                  0          0  OK
  BBE                  0          0  OK
  ES                   0          0  OK
  SES                  0          0  OK
  UAS                  587         0  OK
Received DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Received SAPI:

```

```

00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
OTN Received Overhead Bytes:
  APS/PCC0: 0x02, APS/PCC1: 0x42, APS/PCC2: 0xa2, APS/PCC3: 0x48
  Payload Type: 0x03
OTN Transmitted Overhead Bytes:
  APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
  Payload Type: 0x03
Filter statistics:
  Input packet count                0
  Input packet rejects              0
  Input DA rejects                  0
  Input SA rejects                  0
  Output packet count                0
  Output packet pad count            0
  Output packet error count          0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 7
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
      0 best-effort        95      9500000000    95      0      low
none
      3 network-control    5       500000000    5       0      low
none
...

```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode)

```

user@host> show interfaces xe-7/0/0 extensive
Physical interface: xe-7/0/0, Enabled, Physical link is Up
  Interface index: 173, SNMP ifIndex: 212, Generation: 174
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
  Unidirectional: Enabled,
  Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
...

```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only)

```

user@host> show interfaces xe-7/0/0-tx extensive
Physical interface: xe-7/0/0-tx, Enabled, Physical link is Up
  Interface index: 176, SNMP ifIndex: 137, Generation: 177
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
  Unidirectional: Tx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:73:e4:83, Hardware address: 00:05:85:73:e4:83
  Last flapped   : 2007-06-01 09:08:19 PDT (3d 02:31 ago)
  Statistics last cleared: Never
Traffic statistics:
  Input bytes   :                0                0 bps

```



```

Output bytes :      322891152287160      9627472888 bps
Input packets:              0              0 pps
Output packets:    328809727380      1225492 pps

...

Filter statistics:
  Output packet count      328810554250
  Output packet pad count      0
  Output packet error count    0
...

Logical interface xe-7/0/0-tx.0 (Index 73) (SNMP ifIndex 138) (Generation 139)

Flags: SNMP-Traps Encapsulation: ENET2
Egress accounting overhead: 100
Ingress accounting overhead: 90
Traffic statistics:
  Input bytes :              0
  Output bytes :    322891152287160
  Input packets:              0
  Output packets:    328809727380
IPv6 transit statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:            0
Local statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:            0
Transit statistics:
  Input bytes :              0              0 bps
  Output bytes :    322891152287160      9627472888 bps
  Input packets:              0              0 pps
  Output packets:    328809727380      1225492 pps
IPv6 transit statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:            0
Protocol inet, MTU: 1500, Generation: 147, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.11.12/24, Local: 10.11.12.13, Broadcast: 10.11.12.255,
Generation: 141
  Protocol multiservice, MTU: Unlimited, Generation: 148, Route table: 0
  Flags: None
  Policer: Input: __default_arp_policer__

```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only)

```

user@host> show interfaces xe-7/0/0-rx extensive
Physical interface: xe-7/0/0-rx, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 118, Generation: 175
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Rx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues

```

Hold-times : Up 0 ms, Down 0 ms
Current address: 00:05:85:73:e4:83, Hardware address: 00:05:85:73:e4:83
Last flapped : 2007-06-01 09:08:22 PDT (3d 02:31 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 322857456303482 9627496104 bps
Output bytes : 0 0 bps
Input packets: 328775413751 1225495 pps
Output packets: 0 0 pps

...

Filter statistics:
Input packet count 328775015056
Input packet rejects 1
Input DA rejects 0

...

Logical interface xe-7/0/0-rx.0 (Index 72) (SNMP ifIndex 120) (Generation 138)

Flags: SNMP-Traps Encapsulation: ENET2

Traffic statistics:

Input bytes : 322857456303482
Output bytes : 0
Input packets: 328775413751
Output packets: 0

IPv6 transit statistics:

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Local statistics:

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Transit statistics:

Input bytes : 322857456303482 9627496104 bps
Output bytes : 0 0 bps
Input packets: 328775413751 1225495 pps
Output packets: 0 0 pps

IPv6 transit statistics:

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Protocol inet, MTU: 1500, Generation: 145, Route table: 0

Addresses, Flags: Is-Preferred Is-Primary

Destination: 192.1.1/24, Local: 192.1.1.1, Broadcast: 192.1.1.255,

Generation: 139

Protocol multiservice, MTU: Unlimited, Generation: 146, Route table: 0

Flags: None

Policer: Input: __default_arp_policer__

show interfaces (Gigabit Ethernet)

Syntax	<pre>show interfaces <i>ge-fpc/pic/port</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre>
Release Information	Command introduced before Junos OS Release 7.4.
Description	(M Series, T Series, and MX Series routers and EX Series switches only) Display status information about the specified Gigabit Ethernet interface.
Options	<p><i>ge-fpc/pic/port</i>—Display standard information about the specified Gigabit Ethernet interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p>
Additional Information	In a logical system, this command displays information only about the logical interfaces and not about the physical interfaces.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> <i>Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration</i>
List of Sample Output	<p>show interfaces (Gigabit Ethernet) on page 1996</p> <p>show interfaces (Gigabit Ethernet on MX Series Routers) on page 1996</p> <p>show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration) on page 1997</p> <p>show interfaces brief (Gigabit Ethernet) on page 1997</p> <p>show interfaces detail (Gigabit Ethernet) on page 1998</p> <p>show interfaces extensive (Gigabit Ethernet IQ2) on page 1999</p> <p>show interfaces (Gigabit Ethernet Unnumbered Interface) on page 2002</p> <p>show interfaces (ACI Interface Set Configured) on page 2002</p>
Output Fields	<p>Table 137 on page 1982 describes the output fields for the show interfaces (Gigabit Ethernet) command. Output fields are listed in the approximate order in which they appear. For Gigabit Ethernet IQ and IQE PICs, the traffic and MAC statistics vary by interface type. For more information, see Table 138 on page 1995.</p>

Table 137: show interfaces Gigabit Ethernet Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 2376 .	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
WAN-PHY mode	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: Enabled or Disabled for parent interface; Rx-only or Tx-only for child interfaces.	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> Online—Autonegotiation is manually configured as online. Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under “Common Output Fields Description” on page 2376 .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 2376 .	All levels

Table 137: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Link flags	Information about the link. Possible values are described in the “Links Flags” section under “ Common Output Fields Description ” on page 2376.	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	detail extensive none
Schedulers	(Gigabit Ethernet intelligent queuing 2 [IQ2] interfaces only) Number of CoS schedulers configured.	extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds (ms).	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None
Output Rate	Output rate in bps and pps.	None
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Egress accounting overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	detail extensive
Ingress accounting overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	detail extensive

Table 137: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Output bytes—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the show interfaces (10-Gigabit Ethernet) command.</p>	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 137: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 137: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
Interface transmit statistics	<p>(On MX Series devices) Status of the interface-transmit-statistics configuration: Enabled or Disabled.</p> <ul style="list-style-type: none"> • Enabled—When the interface-transmit-statistics statement is included in the configuration. If this is configured, the interface statistics show the actual transmitted load on the interface. • Disabled—When the interface-transmit-statistics statement is not included in the configuration. If this is not configured, the interface statistics show the offered load on the interface. 	detail extensive
OTN FEC statistics	<p>The forward error correction (FEC) counters provide the following statistics:</p> <ul style="list-style-type: none"> • Corrected Errors—The count of corrected errors in the last second. • Corrected Error Ratio—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits. 	detail extensive
PCS statistics	<p>(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.</p> <ul style="list-style-type: none"> • Bit errors—High bit error rate. Indicates the number of bit errors when the PCS receiver is operating in normal mode. • Errored blocks—Loss of block lock. The number of errored blocks when the PCS receiver is operating in normal mode. 	detail extensive

Table 137: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the show interfaces (10-Gigabit Ethernet) command. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—There are two possible conditions regarding the number of oversized frames: <ul style="list-style-type: none"> • Packet length exceeds 1518 octets, or • Packet length exceeds MRU • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. <p>NOTE: The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the VLAN tagged frames field displays 0 when the show interfaces command is executed on a 20-port Gigabit Ethernet MIC. In other words, the number of VLAN tagged frames cannot be determined for the 20-port Gigabit Ethernet MIC.</p> <ul style="list-style-type: none"> • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive
OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	extensive
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	extensive

Table 137: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 137: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 137: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 137: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 137: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under " Common Output Fields Description " on page 2376.	All levels

Table 137: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • push-push—Two VLAN tags are pushed in from the incoming frame. • swap-push—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value. • pop-swap—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed. 	brief detail extensive none
Demux	<p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> • Source Family Inet • Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels
ACI VLAN: Dynamic Profile	Name of the dynamic profile that defines the agent circuit identifier (ACI) interface set. If configured, the ACI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information.	brief detail extensive none
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under “ Common Output Fields Description ” on page 2376.	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Dynamic Profile	(MX Series routers with Trio MPCs only) Name of the dynamic profile that was used to create this interface configured with a Point-to-Point Protocol over Ethernet (PPPoE) family.	detail extensive none
Service Name Table	(MX Series routers with Trio MPCs only) Name of the service name table for the interface configured with a PPPoE family.	detail extensive none
Max Sessions	(MX Series routers with Trio MPCs only) Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.	detail extensive none

Table 137: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Duplicate Protection	(MX Series routers with Trio MPCs only) State of PPPoE duplicate protection: On or Off . When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the router.	extensive
Transit statistics	<p>Number and rate of bytes and packets transiting the switch.</p> <p>NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the Output bytes and Output packets interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.</p>	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. Possible values are described in the "Family Flags" section under " Common Output Fields Description " on page 2376.	detail extensive
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	detail extensive none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	detail extensive

Table 137: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output Filters	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	detail extensive
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 2376 .	detail extensive none
protocol-family	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about the address flag. Possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 2376 .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Table 138: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Inbound physical interface	show interfaces ge-0/3/0 extensive	<p>Traffic statistics:</p> <p>Input bytes: 496 bytes per packet, representing the Layer 2 packet</p> <p>MAC statistics:</p> <p>Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes</p>	The additional 4 bytes are for the CRC.
Inbound logical interface	show interfaces ge-0/3/0.50 extensive	<p>Traffic statistics:</p> <p>Input bytes: 478 bytes per packet, representing the Layer 3 packet</p>	

Table 138: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type (*continued*)

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Outbound physical interface	show interfaces ge-0/0/0 extensive	Traffic statistics: Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes MAC statistics: Received octets: 478 bytes per packet, representing the Layer 3 packet	For input bytes, the additional 12 bytes include 6 bytes for the destination MAC address plus 4 bytes for VLAN plus 2 bytes for the Ethernet type.
Outbound logical interface	show interfaces ge-0/0/0.50 extensive	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	

Sample Output

show interfaces (Gigabit Ethernet)

```

user@host> show interfaces ge-3/0/2
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:05:85:4a:e9:7c, Hardware address: 00:05:85:4a:e9:7c
  Last flapped   : 2006-08-10 17:25:10 PDT (00:01:08 ago)
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms   : None
  Active defects  : None

Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69)
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
  0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  Egress accounting overhead: 100
  Ingress accounting overhead: 90
  Input packets : 0
  Output packets: 0
  Protocol ccc, MTU: 1522
  Flags: Is-Primary

```

show interfaces (Gigabit Ethernet on MX Series Routers)

```

user@host> show interfaces ge-2/2/2
Physical interface: ge-2/2/2, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 188
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, MAC-REWRITE Error: None,
  Loopback: Disabled,

```

```

Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags     : None
CoS queues    : 8 supported, 4 maximum usable queues
Schedulers    : 0
Current address: 00:1f:12:b7:d7:c0, Hardware address: 00:1f:12:b7:d6:76
Last flapped   : 2008-09-05 16:44:30 PDT (3d 01:04 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
Active alarms  : None
Active defects : None
Logical interface ge-2/2/2.0 (Index 82) (SNMP ifIndex 219)
  Flags: SNMP-Traps 0x20000000 Encapsulation: Ethernet-Bridge
  Egress accounting overhead: 100
  Ingress accounting overhead: 90
  Input packets : 0
  Output packets: 0
  Protocol aenet, AE bundle: ae0.0    Link Index: 4

```

show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration)

```

user@host> show interfaces ge-2/1/2 extensive | match "output|interface"
Physical interface: ge-2/1/2, Enabled, Physical link is Up
Interface index: 151, SNMP ifIndex: 530, Generation: 154
Interface flags: SNMP-Traps Internal: 0x4000
Output bytes   :      240614363944      772721536 bps
Output packets :      3538446506      1420444 pps
Direction : Output
Interface transmit statistics: Enabled

Logical interface ge-2/1/2.0 (Index 331) (SNMP ifIndex 955) (Generation 146)
Output bytes   :      195560312716      522726272 bps
Output packets :      4251311146      1420451 pps

```

show interfaces brief (Gigabit Ethernet)

```

user@host> show interfaces ge-3/0/2 brief
Physical interface: ge-3/0/2, Enabled, Physical link is Up
Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags     : None

Logical interface ge-3/0/2.0
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  ccc

Logical interface ge-3/0/2.32767
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2

```

show interfaces detail (Gigabit Ethernet)

```

user@host> show interfaces ge-3/0/2 detail
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35, Generation: 177
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:4a:e9:7c, Hardware address: 00:05:85:4a:e9:7c
  Last flapped   : 2006-08-09 17:17:00 PDT (01:31:33 ago)
  Statistics last cleared: Never
Traffic statistics:
  Input bytes :                0                0 bps
  Output bytes :                0                0 bps
  Input packets:                0                0 pps
  Output packets:                0                0 pps
Ingress traffic statistics at Packet Forwarding Engine:
  Input bytes :                0                0 bps
  Input packets:                0                0 pps
  Drop bytes :                0                0 bps
  Drop packets:                0                0 pps
Ingress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort                0                0                0
  1 expedited-fo                0                0                0
  2 assured-forw                0                0                0
  3 network-cont                0                0                0

Egress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort                0                0                0
  1 expedited-fo                0                0                0
  2 assured-forw                0                0                0
  3 network-cont                0                0                0

Active alarms : None
Active defects : None

Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69) (Generation 140)
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530)
Out(swap-push 0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  Egress accounting overhead: 100
  Ingress accounting overhead: 90
Traffic statistics:
  Input bytes :                0
  Output bytes :                0

```

```

Input packets:          0
Output packets:         0
Local statistics:
Input bytes :           0
Output bytes :          0
Input packets:          0
Output packets:         0
Transit statistics:
Input bytes :           0          0 bps
Output bytes :          0          0 bps
Input packets:          0          0 pps
Output packets:         0          0 pps
Protocol ccc, MTU: 1522, Generation: 149, Route table: 0
Flags: Is-Primary

```

```

Logical interface ge-3/0/2.32767 (Index 71) (SNMP ifIndex 70)
(Generation 139)
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2
Traffic statistics:
Input bytes :           0
Output bytes :          0
Input packets:          0
Output packets:         0
Local statistics:
Input bytes :           0
Output bytes :          0
Input packets:          0
Output packets:         0
Transit statistics:
Input bytes :           0          0 bps
Output bytes :          0          0 bps
Input packets:          0          0 pps
Output packets:         0          0 pps

```

show interfaces extensive (Gigabit Ethernet IQ2)

```

user@host> show interfaces ge-7/1/3 extensive
Physical interface: ge-7/1/3, Enabled, Physical link is Up
Interface index: 170, SNMP ifIndex: 70, Generation: 171
Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4004000
Link flags : None
CoS queues : 8 supported, 4 maximum usable queues
Schedulers : 256
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:14:f6:30:5e:74, Hardware address: 00:14:f6:30:5e:74
Last flapped : 2007-11-07 21:31:41 PST (02:03:33 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes :          38910844056          7952 bps
Output bytes :           7174605          8464 bps
Input packets:         418398473           11 pps
Output packets:          78903           12 pps
IPv6 transit statistics:
Input bytes :           0
Output bytes :           0
Input packets:          0
Output packets:          0

```

Ingress traffic statistics at Packet Forwarding Engine:

```

Input bytes :          38910799145          7952 bps
Input packets:         418397956           11 pps
Drop bytes :           0                0 bps
Drop packets:          0                0 pps

```

Input errors:

```

Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0

```

Output errors:

```

Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

```

```

FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0

```

```

Ingress queues: 4 supported, 4 in use

```

Queue counters:	Queued packets	Transmitted packets	Dropped packets
0 best-effort	418390823	418390823	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	7133	7133	0

```

Egress queues: 4 supported, 4 in use

```

Queue counters:	Queued packets	Transmitted packets	Dropped packets
0 best-effort	1031	1031	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	77872	77872	0

```

Active alarms : None

```

```

Active defects : None

```

MAC statistics:

	Receive	Transmit
Total octets	38910844056	7174605
Total packets	418398473	78903
Unicast packets	408021893366	1026
Broadcast packets	10	12
Multicast packets	418398217	77865
CRC/Align errors	0	0
FIFO errors	0	0
MAC control frames	0	0
MAC pause frames	0	0
Oversized frames	0	
Jabber frames	0	
Fragment frames	0	
VLAN tagged frames	0	
Code violations	0	OTN Received Overhead Bytes:
APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58		
Payload Type: 0x08		

OTN Transmitted Overhead Bytes:

```

APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
Payload Type: 0x08

```

Filter statistics:

```

Input packet count      418398473
Input packet rejects    479
Input DA rejects        479

```

```

Input SA rejects                                0
Output packet count                            78903
Output packet pad count                        0
Output packet error count                      0
CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link partner:
Link mode: Full-duplex, Flow control: Symmetric/Asymmetric,
Remote fault: OK
Local resolution:
Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
Destination slot: 7
CoS information:
Direction : Output
CoS transmit queue      Bandwidth      Buffer      Priority      Limit
                        %      bps      %      usec
0 best-effort           95      950000000  95      0
low none
3 network-control       5      500000000   5      0
low none
Direction : Input
CoS transmit queue      Bandwidth      Buffer      Priority      Limit
                        %      bps      %      usec
0 best-effort           95      950000000  95      0
low none
3 network-control       5      500000000   5      0
low none

Logical interface ge-7/1/3.0 (Index 70) (SNMP ifIndex 85) (Generation 150)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
Input bytes :      812400
Output bytes :    1349206
Input packets:     9429
Output packets:    9449
IPv6 transit statistics:
Input bytes :      0
Output bytes :      0
Input packets:      0
Output packets:     0
Local statistics:
Input bytes :      812400
Output bytes :    1349206
Input packets:     9429
Output packets:    9449
Transit statistics:
Input bytes :      0      7440 bps
Output bytes :      0      7888 bps
Input packets:      0      10 pps
Output packets:      0      11 pps
IPv6 transit statistics:
Input bytes :      0
Output bytes :      0
Input packets:      0
Output packets:     0
Protocol inet, MTU: 1500, Generation: 169, Route table: 0
Flags: Is-Primary, Mac-Validate-Strict
Mac-Validate Failures: Packets: 0, Bytes: 0
Addresses, Flags: Is-Preferred Is-Primary

```

```
Input Filters: F1-ge-3/0/1.0-in, F3-ge-3/0/1.0-in
Output Filters: F2-ge-3/0/1.0-out (53)
Destination: 10.74.2/24, Local: 10.74.2.2, Broadcast: 10.74.2.255,
Generation: 196
Protocol multiservice, MTU: Unlimited, Generation: 170, Route table: 0
Flags: Is-Primary
Policer: Input: __default_arp_policer__
```

NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics displayed in the **show interfaces** command output might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the interface counters. For detailed information, see the description of the logical interface **Transit statistics** fields in [Table 137 on page 1982](#).

show interfaces (Gigabit Ethernet Unnumbered Interface)

```
user@host> show interfaces ge-3/2/0
Physical interface: ge-3/2/0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 50
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Current address: 00:14:f6:11:26:f8, Hardware address: 00:14:f6:11:26:f8
  Last flapped   : 2006-10-27 04:42:23 PDT (08:01:52 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 624 bps (1 pps)
  Active alarms  : None
  Active defects : None

Logical interface ge-3/2/0.0 (Index 67) (SNMP ifIndex 85)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 0
  Output packets: 6
  Protocol inet, MTU: 1500
  Flags: Unnumbered
  Donor interface: lo0.0 (Index 64)
  Preferred source address: 22.22.22.22
```

show interfaces (ACI Interface Set Configured)

```
user@host> show interfaces ge-1/0/0.4001
Logical interface ge-1/0/0.4001 (Index 340) (SNMP ifIndex 548)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.4001 ] Encapsulation: PPP-over-

Ethernet
ACI VLAN:
  Dynamic Profile: aci-vlan-set-profile
  PPPoE:
    Dynamic Profile: aci-vlan-pppoe-profile,
    Service Name Table: None,
    Max Sessions: 32000, Max Sessions VSA Ignore: Off,
    Duplicate Protection: On, Short Cycle Protection: Off,
    AC Name: nbc
  Input packets : 9
```


Output packets: 8
Protocol multiservice, MTU: Unlimited

show interfaces irb

Syntax	<pre>show interfaces irb <brief detail extensive terse> <descriptions> <media> <routing-instance <i>instance-name</i>> <snmp-index <i>snmp-index</i>> <statistics></pre>
Release Information	<p>Command introduced in Junos OS Release 12.3R2.</p> <p>Command introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Display integrated routing and bridging interfaces information.
Options	<p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>routing-instance <i>instance-name</i>—(Optional) Display information for the interface with the specified SNMP index.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the interface with the specified SNMP index.</p> <p>statistics—(Optional) Display static interface statistics.</p>
Additional Information	Integrated routing and bridging (IRB) provides simultaneous support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route local packets to another routed interface or to another VLAN that has a Layer 3 protocol configured.
Required Privilege Level	view
List of Sample Output	<p>show interfaces irb extensive on page 2008</p> <p>show interfaces irb snmp-index on page 2009</p>
Output Fields	Table 139 on page 2004 lists the output fields for the show interfaces irb command. Output fields are listed in the approximate order in which they appear.

Table 139: show interfaces irb Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the physical interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 2376 .	All levels

Table 139: show interfaces irb Output Fields (*continued*)

Field Name	Field Description	Level of Output
Proto	Protocol configured on the interface.	terse
Interface index	Physical interface index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Type	Physical interface type.	detail extensive none
Link-level type	Encapsulation being used on the physical interface.	detail extensive brief none
MTU	MTU size on the physical interface.	detail extensive brief none
Clocking	Reference clock source: Internal or External . Always unspecified on IRB interfaces.	detail extensive brief
Speed	Speed at which the interface is running. Always unspecified on IRB interfaces.	detail extensive brief
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under " Common Output Fields Description " on page 2376.	detail extensive brief none
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under " Common Output Fields Description " on page 2376.	detail extensive brief none
Link type	Physical interface link type: full duplex or half duplex .	detail extensive none
Link flags	Information about the link. Possible values are described in the "Links Flags" section under " Common Output Fields Description " on page 2376.	detail extensive none
Physical Info	Physical interface information.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	MAC address of the hardware.	detail extensive none
Alternate link address	Backup address of the link.	detail extensive
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hours:minutes:seconds timezone (hours:minutes:seconds ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive

Table 139: show interfaces irb Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface • Output packets—Number of packets transmitted on the interface. 	detail extensive
IPv6 transit statistics	<p>Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Giants—Number of frames received that are larger than the giant threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • Resource errors—Sum of transmit drops. 	detail extensive
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the DPC is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	detail extensive

Logical Interface

Table 139: show interfaces irb Output Fields (*continued*)

Field Name	Field Description	Level of Output
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface (which reflects its initialization sequence).	detail extensive none
SNMP ifIndex	SNMP interface index number of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under "Common Output Fields Description" on page 2376 .	detail extensive
Encapsulation	Encapsulation on the logical interface.	detail extensive
Bandwidth	Speed at which the interface is running.	detail extensive
Routing Instance	Routing instance IRB is configured under.	detail extensive
Bridging Domain	Bridging domain IRB is participating in.	detail extensive
Traffic statistics	Number and rate of bytes and packets received and transmitted on the logical interface. <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface • Output packets—Number of packets transmitted on the interface. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine.	detail extensive
Transit statistics	Statistics for traffic transiting the router.	detail extensive
Protocol	Protocol family configured on the local interface. Possible values are described in the "Protocol Field" section under "Common Output Fields Description" on page 2376 .	detail extensive
MTU	Maximum transmission unit size on the logical interface.	detail extensive

Table 139: show interfaces irb Output Fields (*continued*)

Field Name	Field Description	Level of Output
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route table	Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive
Addresses, Flags	Information about address flags. Possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 2376 .	detail extensive
Policer	The policer that is to be evaluated when packets are received or transmitted on the interface.	detail extensive
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under “Common Output Fields Description” on page 2376 .	detail extensive

Sample Output

show interfaces irb extensive

```

user@host> show interfaces irb extensive
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 23, Generation: 130
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: Unspecified
  Device flags      : Present Running
  Interface flags:  SNMP-Traps
  Link type        : Full-Duplex
  Link flags       : None
  Physical info    : Unspecified
  Hold-times       : Up 0 ms, Down 0 ms
  Current address:  02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Alternate link address: Unspecified
  Last flapped     : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes :                0
    Output bytes:                0
    Input packets:               0
    Output packets:              0
  IPv6 transit statistics:
    Input bytes :                0
    Output bytes:                0
    Input packets:               0
    Output packets:              0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0

```

```

Logical interface irb.0 (Index 68) (SNMP ifIndex 70) (Generation 143)
  Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
  Bandwidth: 1000mbps
  Routing Instance: customer_0 Bridging Domain: bd0
  Traffic statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  IPv6 transit statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Local statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Transit statistics:
    Input bytes : 0 0 bps
    Output bytes : 0 0 bps
    Input packets: 0 0 pps
    Output packets: 0 0 pps
  IPv6 transit statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Protocol inet, MTU: 1500, Generation: 154, Route table: 0
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255,
      Generation: 155
  Protocol multiservice, MTU: 1500, Generation: 155, Route table: 0
    Flags: Is-Primary
    Policer: Input: __default_arp_policer

```

show interfaces irb snmp-index

```

user@host> show interfaces irb snmp-index 25
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 25
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514
  Device flags : Present Running
  Interface flags: SNMP-Traps
  Link type : Full-Duplex
  Link flags : None
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Last flapped : Never
    Input packets : 0
    Output packets: 0

Logical interface irb.0 (Index 68) (SNMP ifIndex 70)
  Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
  Bandwidth: 1000mbps
  Routing Instance: customer_0 Bridging Domain: bd0
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 1500
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary

```

Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255
Protocol multiservice, MTU: 1500
Flags: Is-Primary

show interfaces queue

Syntax	<pre>show interfaces queue <aggregate remaining-traffic> <both-ingress-egress> <egress> <forwarding-class forwarding-class> <ingress> <interface-name interface-name> <l2-statistics> <remaining-traffic></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>both-ingress-egress, egress, and ingress options introduced in Junos OS Release 7.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>l2-statistics option introduced in Junos OS Release 12.1.</p>
Description	Display class-of-service (CoS) queue information for physical interfaces.
Options	<p>none—Show detailed CoS queue statistics for all physical interfaces.</p> <p>aggregate—(Optional) Display the aggregated queuing statistics of all logical interfaces that have traffic-control profiles configured. (Not on the QFX Series.)</p> <p>both-ingress-egress—(Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics. (Not on the QFX Series.)</p> <p>egress—(Optional) Display egress queue statistics.</p> <p>forwarding-class forwarding-class—(Optional) Forwarding class name for this queue. Shows detailed CoS statistics for the queue associated with the specified forwarding class.</p> <p>ingress—(Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics. (Not on the QFX Series.)</p> <p>interface-name interface-name—(Optional) Show detailed CoS queue statistics for the specified interface.</p> <p>l2-statistics—(optional) Display layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles</p>

Overhead for Layer 2 Statistics

Transmitted packets and transmitted byte counts are displayed for the layer 2 level with the addition of encapsulation overheads applied for fragmentation, as shown in [Table 140 on page 2012](#). Others counters, such as packets and bytes queued (input) and drop counters, are displayed at the layer 3 level. In the case of link fragmentation and interleaving (LFI) for which not fragmentation is applied, corresponding layer 2 overheads are added, as shown in [Table 140 on page 2012](#).

Table 140: Layer 2 Overhead, Transmitted Packets/Bytes

Protocol	Fragmentation		LFI
	First fragmentation	Second to n fragmentations	
	Bytes	Bytes	
MLPPP (Long)	13	12	8
MLPPP (short)	11	10	8
MLFR (FRF15)	12	10	8
MFR (FRF16)	10	8	-
MCMLPPP(Long)	13	12	-
MCMLPPP(Short)	11	10	-

Layer 2 Statistics - Fragmentation Overhead Calculation

MLPPP/MC-MLPPP Overhead details:

=====

Fragment 1:

```

Outer PPP header           : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
Inner PPP header           : 1 byte
HDLC flag and FCS bytes    : 4 bytes

```

Fragments 2 .. n :

```

Outer PPP header           : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
HDLC flag and FCS bytes    : 4 bytes

```

MLFR (FRF15) Overhead details:

=====

Fragment 1:

```

Framereelay header        : 2 bytes
Control,NLPID             : 2 bytes
Fragmentaion header       : 2 bytes
Inner proto               : 2 bytes
HDLC flag and FCS         : 4 bytes

```

Fragments 2 ...n :

```

Framereelay header        : 2 bytes
Control,NLPID             : 2 bytes
Fragmentaion header       : 2 bytes
HDLC flag and FCS         : 4 bytes

```

MFR (FRF16) Overhead details:

=====

```

Fragment 1:
  Fragmentation header : 2 bytes
  Framereelay header   : 2 bytes
  Inner proto          : 2 bytes
  HDLC flag and FCS    : 4 bytes

Fragments 2 ...n :
  Fragmentation header : 2 bytes
  Framereelay header   : 2 bytes
  HDLC flag and FCS    : 4 bytes

```

Overhead with LFI

```

MLPPP(Long & short sequence):
=====
  Outer PPP header : 4 bytes
  HDLC flag and FCS : 4 bytes

MLFR (FRF15):
=====
  Framereelay header : 2 bytes
  Control,NLPID      : 2 bytes
  HDLC flag and FCS  : 4 bytes

```

The following examples show overhead for different cases:

- A 1000-byte packet is sent to a mlppp bundle without any fragmentation. At the layer 2 level, bytes transmitted is 1013 in 1 packet. This overhead is for MLPPP long sequence encap.
- A 1000-byte packet is sent to a mlppp bundle with a fragment threshold of 250byte. At the layer 2 level, bytes transmitted is 1061 bytes in 5 packets.
- A 1000-byte LFI packet is sent to an mlppp bundle. At the layer 2 level, bytes transmitted is 1008 in 1 packet.

remaining-traffic—(Optional) Display the queuing statistics of all logical interfaces that do not have traffic-control profiles configured. (Not on the QFX Series.)

Additional Information On M Series routers (except for the M320 and M120 routers), this command is valid only for a PIC installed on an enhanced Flexible PIC Concentrator (FPC).

Queue statistics for aggregated interfaces are supported on the M Series and T Series routers only. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can view the statistics for a child interface by using the **show interfaces statistics** command for that child interface.

When you configure tricolor marking on a 10-port 1-Gigabit Ethernet PIC, for queues 6 and 7 only, the output does not display the number of queued bytes and packets, or the number of bytes and packets dropped because of RED. If you do not configure tricolor marking on the interface, these statistics are available for all queues.

For the 4-port Channelized OC12 IQE PIC and 1-port Channelized OC48 IQE PIC, the **Packet Forwarding Engine Chassis Queues** field represents traffic bound for a particular

physical interface on the PIC. For all other PICs, the **Packet Forwarding Engine Chassis Queues** field represents the total traffic bound for the PIC.

For Gigabit Ethernet IQ2 PICs, the **show interfaces queue** command output does not display the number of tail-dropped packets. This limitation does not apply to Packet Forwarding Engine chassis queues.

When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (under the **Packet Forwarding Engine Chassis Queues** field) shows the prefragmentation values.

The behavior of the **egress** queues for the **Routing Engine-Generated Traffic** is not same as the configured queue for MLPPP and MFR configurations.

For information about how to configure CoS, see the *Junos® OS Network Interfaces*. For related CoS operational mode commands, see the *Junos OS Operational Mode Commands*.

Required Privilege Level view

List of Sample Output

- [show interfaces queue \(Aggregated Ethernet on a T320 Router\) on page 2018](#)
- [show interfaces queue \(Fast Ethernet on a J4300 Router\) on page 2020](#)
- [show interfaces queue \(Gigabit Ethernet on a T640 Router\) on page 2021](#)
- [show interfaces queue aggregate \(Gigabit Ethernet Enhanced DPC\) on page 2021](#)
- [show interfaces queue \(Gigabit Ethernet IQ2 PIC\) on page 2025](#)
- [show interfaces queue both-ingress-egress \(Gigabit Ethernet IQ2 PIC\) on page 2028](#)
- [show interfaces queue ingress \(Gigabit Ethernet IQ2 PIC\) on page 2030](#)
- [show interfaces queue egress \(Gigabit Ethernet IQ2 PIC\) on page 2031](#)
- [show interfaces queue remaining-traffic \(Gigabit Ethernet Enhanced DPC\) on page 2032](#)
- [show interfaces queue \(Channelized OC12 IQE Type 3 PIC in SONET Mode\) on page 2035](#)
- [show interfaces queue \(QFX Series\) on page 2045](#)
- [show interfaces queue l2-statistics \(lsq interface\) on page 2046](#)

Output Fields [Table 141 on page 2014](#) lists the output fields for the **show interfaces queue** command. Output fields are listed in the approximate order in which they appear.

Table 141: show interfaces queue Output Fields

Field Name	Field Description
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “ Common Output Fields Description ” on page 2376.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Forwarding classes supported	Total number of forwarding classes supported on the specified interface.

Table 141: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
Forwarding classes in use	Total number of forwarding classes in use on the specified interface.
Ingress queues supported	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues supported on the specified interface.
Ingress queues in use	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues in use on the specified interface.
Output queues supported	Total number of output queues supported on the specified interface.
Output queues in use	Total number of output queues in use on the specified interface.
Egress queues supported	Total number of egress queues supported on the specified interface.
Egress queues in use	Total number of egress queues in use on the specified interface.
Queue	Queue number.
Queue counters (Ingress)	<p>CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism.
Burst size	(Logical interfaces on IQ PICs only) Maximum number of bytes up to which the logical interface can burst. The burst size is based on the shaping rate applied to the interface.
Forwarding classes	Forwarding class name.
Queued Packets	<p>Number of packets queued to this queue.</p> <p>NOTE: For Gigabit Ethernet IQ2 interfaces, the Queued Packets count is calculated by the Junos OS interpreting one frame buffer as one packet. If the queued packets are very large or very small, the calculation might not be completely accurate for transit traffic. The count is completely accurate for traffic terminated on the router.</p>
Queued Bytes	<p>Number of bytes queued to this queue. The byte counts vary by PIC type. For more information, see Table 142 on page 2017.</p>
Transmitted Packets	<p>Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the Packet Forwarding Engine Chassis Queues field) shows the prefragmentation values.</p> <p>NOTE: For layer 2 statistics, see “Overhead for Layer 2 Statistics” on page 2011</p>

Table 141: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
Transmitted Bytes	<p>Number of bytes transmitted by this queue. The byte counts vary by PIC type. For more information, see Table 142 on page 2017.</p> <p>NOTE: On MX Series routers, this number can be inaccurate when you issue the command for a physical interface repeatedly and in quick succession, because the statistics for the child nodes are collected infrequently. Wait ten seconds between successive iterations to avoid this situation.</p> <p>NOTE: For layer 2 statistics, see “Overhead for Layer 2 Statistics” on page 2011</p>
Tail-dropped packets	Number of packets dropped because of tail drop.
RED-dropped packets	<p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> • (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> • Low, non-TCP—Number of low-loss priority non-TCP packets dropped because of RED. • Low, TCP—Number of low-loss priority TCP packets dropped because of RED. • High, non-TCP—Number of high-loss priority non-TCP packets dropped because of RED. • High, TCP—Number of high-loss priority TCP packets dropped because of RED. • (J Series routers and MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> • Low—Number of low-loss priority packets dropped because of RED. • Medium-low—Number of medium-low loss priority packets dropped because of RED. • Medium-high—Number of medium-high loss priority packets dropped because of RED. • High—Number of high-loss priority packets dropped because of RED.
RED-dropped bytes	<p>Number of bytes dropped because of RED. The byte counts vary by PIC type. For more information, see Table 142 on page 2017.</p> <ul style="list-style-type: none"> • (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> • Low, non-TCP—Number of low-loss priority non-TCP bytes dropped because of RED. • Low, TCP—Number of low-loss priority TCP bytes dropped because of RED. • High, non-TCP—Number of high-loss priority non-TCP bytes dropped because of RED. • High, TCP—Number of high-loss priority TCP bytes dropped because of RED. • (J Series routers only) The output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> • Low—Number of low-loss priority bytes dropped because of RED. • Medium-low—Number of medium-low loss priority bytes dropped because of RED. • Medium-high—Number of medium-high loss priority bytes dropped because of RED. • High—Number of high-loss priority bytes dropped because of RED.

Byte counts vary by PIC type. [Table 142 on page 2017](#) shows how the byte counts on the outbound interfaces vary depending on the PIC type. [Table 142 on page 2017](#) is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.

Table 142: Byte Count by PIC Type

PIC Type	Output Level	Byte Count Includes	Comments
Gigabit Ethernet IQ and IQE PICs	Interface	<p>Queued: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>Transmitted: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>RED dropped: 496 bytes per packet representing 478 bytes of Layer 3 packet + 18 bytes</p>	<p>The 12 additional bytes include 6 bytes for the destination MAC address + 4 bytes for the VLAN + 2 bytes for the Ethernet type.</p> <p>For RED dropped, 6 bytes are added for the source MAC address.</p>
	Packet forwarding component	<p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p>	—
Non-IQ PIC	Interface	<p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> • Queued: 478 bytes of Layer 3 packet. • Transmitted: 478 bytes of Layer 3 packet. <p>T4000 routers with Type 5 FPCs :</p> <ul style="list-style-type: none"> • Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Inter frame Gap. • Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Interframe Gap. <p>M Series routers:</p> <ul style="list-style-type: none"> • Queued: 478 bytes of Layer 3 packet. • Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead. <p>PTX Series Packet Transport Switches:</p> <ul style="list-style-type: none"> • Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes FCS + the full Layer 1 overhead of the MAC header DA + SA + EtherType (non-VLAN). • Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead of the MAC header DA + SA + EtherType (non-VLAN). • RED dropped: 478 bytes of Layer 3 packet + 22 bytes special header. To the TQ, this packet has 4 bytes more than queued or transmitted. 	<p>The Layer 2 overhead is 14 bytes for non-VLAN traffic and 18 bytes for VLAN traffic.</p>

Table 142: Byte Count by PIC Type (*continued*)

PIC Type	Output Level	Byte Count Includes	Comments
IQ and IQE PICs with a SONET/SDH interface	Interface	<p>Queued: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>Transmitted: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>RED dropped: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p>	The additional 4 bytes are for the Layer 2 Point-to-Point Protocol (PPP) header.
	Packet forwarding component	<p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 486 bytes per packet, representing 478 bytes of Layer 3 packet + 8 bytes</p>	For transmitted packets, the additional 8 bytes includes 4 bytes for the PPP header and 4 bytes for a cookie.
Non-IQ PIC with a SONET/SDH interface	Interface	<p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> Queued: 478 bytes of Layer 3 packet. Transmitted: 478 bytes of Layer 3 packet. <p>M Series routers:</p> <ul style="list-style-type: none"> Queued: 478 bytes of Layer 3 packet. Transmitted: 483 bytes per packet, representing 478 bytes of Layer 3 packet + 5 bytes RED dropped: 478 bytes per packet, representing 478 bytes of Layer 3 packet 	For transmitted packets, the additional 5 bytes includes 4 bytes for the PPP header and 1 byte for the packet loss priority (PLP).
Interfaces configured with Frame Relay Encapsulation	Interface	The default Frame Relay overhead is 7 bytes. If you configure the Frame Check Sequence (FCS) to 4 bytes, then the overhead increases to 10 bytes.	
1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs	Interface	<p>Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p> <p>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p>	The Layer 2 overhead is 18 bytes for non-VLAN traffic and 22 bytes for VLAN traffic.
4-port 1G IQ2 and IQ2-E PICs	Packet forwarding component	Queued: 478 bytes of Layer 3 packet.	—
8-port 1G IQ2 and IQ2-E PICs		Transmitted: 478 bytes of Layer 3 packet.	

Sample Output

show interfaces queue (Aggregated Ethernet on a T320 Router)

The following example shows that the aggregated Ethernet interface, **ae1**, has traffic on queues **af1** and **af12**:


```

user@host> show interfaces queue ae1
Physical interface: ae1, Enabled, Physical link is Up
Interface index: 158, SNMP ifIndex: 33 Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:
    Packets      :           5           0 pps
    Bytes        :          242           0 bps
  Transmitted:
    Packets      :           5           0 pps
    Bytes        :          242           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 1, Forwarding classes: af1
  Queued:
    Packets      :      42603765      595484 pps
    Bytes        :     5453281920     609776496 bps
  Transmitted:
    Packets      :      42603765      595484 pps
    Bytes        :     5453281920     609776496 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 2, Forwarding classes: ef1
  Queued:
    Packets      :           0           0 pps
    Bytes        :           0           0 bps
  Transmitted:
    Packets      :           0           0 pps
    Bytes        :           0           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 3, Forwarding classes: nc
  Queued:
    Packets      :           45           0 pps
    Bytes        :          3930           0 bps
  Transmitted:
    Packets      :           45           0 pps
    Bytes        :          3930           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 4, Forwarding classes: af11
  Queued:
    Packets      :           0           0 pps
    Bytes        :           0           0 bps
  Transmitted:
    Packets      :           0           0 pps
    Bytes        :           0           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 5, Forwarding classes: ef11
  Queued:
    Packets      :           0           0 pps
    Bytes        :           0           0 bps
  Transmitted:
    Packets      :           0           0 pps
    Bytes        :           0           0 bps

```

```

Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 6, Forwarding classes: af12
Queued:
Packets : 31296413 437436 pps
Bytes : 4005940864 447935200 bps
Transmitted:
Packets : 31296413 437436 pps
Bytes : 4005940864 447935200 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: nc2
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

show interfaces queue (Fast Ethernet on a J4300 Router)

```

user@host> show interfaces queue fe-4/0/0.0
Logical interface fe-4/0/0.0 (Index 71) (SNMP ifIndex 42)
Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
Queued:
Packets : 5240762 3404 pps
Bytes : 3020710354 15934544 bps
Transmitted:
Packets : 5240762 3404 pps
Bytes : 3020710354 15934544 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
Queue: 1, Forwarding classes: af1
Queued:
Packets : 2480391 1650 pps
Bytes : 1304685666 6945704 bps
Transmitted:
Packets : 2478740 1650 pps
Bytes : 1303817240 6945704 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 1651 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 1651 0 pps

```

```

RED-dropped bytes      :           868426          0 bps
Low                    :                   0          0 pps
Medium-low             :                   0          0 pps
Medium-high            :                   0          0 pps
High                   :           868426          0 pps

```

show interfaces queue (Gigabit Ethernet on a T640 Router)

```

user@host> show interfaces queue
Physical interface: ge-7/0/1, Enabled, Physical link is Up
Interface index: 150, SNMP ifIndex: 42
Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:
    Packets      :           13          0 pps
    Bytes        :          622          0 bps
  Transmitted:
    Packets      :           13          0 pps
    Bytes        :          622          0 bps
    Tail-dropped packets :           0          0 pps
    RED-dropped packets  :           0          0 pps
    RED-dropped bytes   :           0          0 bps
Queue: 1, Forwarding classes: af1
  Queued:
    Packets      :       1725947945       372178 pps
    Bytes        :    220921336960    381110432 bps
  Transmitted:
    Packets      :       1725947945       372178 pps
    Bytes        :    220921336960    381110432 bps
    Tail-dropped packets :           0          0 pps
    RED-dropped packets  :           0          0 pps
    RED-dropped bytes   :           0          0 bps
Queue: 2, Forwarding classes: ef1
  Queued:
    Packets      :           0          0 pps
    Bytes        :           0          0 bps
  Transmitted:
    Packets      :           0          0 pps
    Bytes        :           0          0 bps
    Tail-dropped packets :           0          0 pps
    RED-dropped packets  :           0          0 pps
    RED-dropped bytes   :           0          0 bps
Queue: 3, Forwarding classes: nc
  Queued:
    Packets      :           571          0 pps
    Bytes        :         49318        336 bps
  Transmitted:
    Packets      :           571          0 pps
    Bytes        :         49318        336 bps
    Tail-dropped packets :           0          0 pps
    RED-dropped packets  :           0          0 pps
    RED-dropped bytes   :           0          0 bps

```

show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 aggregate
Physical interface: ge-2/2/9, Enabled, Physical link is Up
Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use

```

```

Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :          148450735          947295 pps
    Bytes        :          8016344944        409228848 bps
  Transmitted:
    Packets      :          76397439          487512 pps
    Bytes        :         4125461868        210602376 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :          72053285          459783 pps
    Low          :          72053285          459783 pps
    Medium-low   :              0              0 pps
    Medium-high  :              0              0 pps
    High         :              0              0 pps
  RED-dropped bytes  :         3890877444        198626472 bps
    Low           :         3890877444        198626472 bps
    Medium-low    :              0              0 bps
    Medium-high   :              0              0 bps
    High         :              0              0 bps

Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :              0              0 pps
    Bytes        :              0              0 bps
  Transmitted:
    Packets      :              0              0 pps
    Bytes        :              0              0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :              0              0 pps
    Low          :              0              0 pps
    Medium-low   :              0              0 pps
    Medium-high  :              0              0 pps
    High         :              0              0 pps
  RED-dropped bytes  :              0              0 bps
    Low          :              0              0 bps
    Medium-low   :              0              0 bps
    Medium-high  :              0              0 bps
    High         :              0              0 bps

Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      :          410278257          473940 pps
    Bytes        :         22156199518        204742296 bps
  Transmitted:
    Packets      :          4850003           4033 pps
    Bytes        :         261900162        1742256 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :          405425693          469907 pps
    Low          :          405425693          469907 pps
    Medium-low   :              0              0 pps
    Medium-high  :              0              0 pps
    High         :              0              0 pps
  RED-dropped bytes  :         21892988124        203000040 bps
    Low           :         21892988124        203000040 bps
    Medium-low    :              0              0 bps
    Medium-high   :              0              0 bps
    High         :              0              0 bps

Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      :              0              0 pps
    Bytes        :              0              0 bps
  Transmitted:
    Packets      :              0              0 pps
    Bytes        :              0              0 bps

```

```

Tail-dropped packets : Not Available
RED-dropped packets :                0                0 pps
  Low                :                0                0 pps
  Medium-low         :                0                0 pps
  Medium-high        :                0                0 pps
  High               :                0                0 pps
RED-dropped bytes   :                0                0 bps
  Low                :                0                0 bps
  Medium-low         :                0                0 bps
  Medium-high        :                0                0 bps
  High               :                0                0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets          :                76605230          485376 pps
    Bytes            :                5209211400        264044560 bps
  Transmitted:
    Packets          :                76444631          484336 pps
    Bytes            :                5198235612        263478800 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :                160475           1040 pps
    Low              :                160475           1040 pps
    Medium-low       :                0                0 pps
    Medium-high      :                0                0 pps
    High             :                0                0 pps
  RED-dropped bytes   :                10912300          565760 bps
    Low              :                10912300          565760 bps
    Medium-low       :                0                0 bps
    Medium-high      :                0                0 bps
    High             :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :                0                0 pps
    Low              :                0                0 pps
    Medium-low       :                0                0 pps
    Medium-high      :                0                0 pps
    High             :                0                0 pps
  RED-dropped bytes   :                0                0 bps
    Low              :                0                0 bps
    Medium-low       :                0                0 bps
    Medium-high      :                0                0 bps
    High             :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets          :                4836136           3912 pps
    Bytes            :                333402032        2139056 bps
  Transmitted:
    Packets          :                3600866           1459 pps
    Bytes            :                244858888        793696 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :                1225034           2450 pps
    Low              :                1225034           2450 pps
    Medium-low       :                0                0 pps
    Medium-high      :                0                0 pps

```

```

      High : 0 0 pps
    RED-dropped bytes : 83302312 1333072 bps
      Low : 83302312 1333072 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
      High : 0 0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps

```

Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

```

  Queued:
    Packets : 77059796 486384 pps
    Bytes : 3544750624 178989576 bps
  Transmitted:
    Packets : 77059797 486381 pps
    Bytes : 3544750670 178988248 bps
  Tail-dropped packets : 0 0 pps
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps

```

Queue: 1, Forwarding classes: expedited-forwarding

```

  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Tail-dropped packets : 0 0 pps
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps

```

```

      Medium-high      :           0           0 bps
      High             :           0           0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets           :       4846580           3934 pps
    Bytes              :       222942680       1447768 bps
  Transmitted:
    Packets           :       4846580           3934 pps
    Bytes              :       222942680       1447768 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
      Low              :           0           0 pps
      Medium-low       :           0           0 pps
      Medium-high      :           0           0 pps
      High             :           0           0 pps
    RED-dropped bytes  :           0           0 bps
      Low              :           0           0 bps
      Medium-low       :           0           0 bps
      Medium-high      :           0           0 bps
      High             :           0           0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets           :           0           0 pps
    Bytes              :           0           0 bps
  Transmitted:
    Packets           :           0           0 pps
    Bytes              :           0           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
      Low              :           0           0 pps
      Medium-low       :           0           0 pps
      Medium-high      :           0           0 pps
      High             :           0           0 pps
    RED-dropped bytes  :           0           0 bps
      Low              :           0           0 bps
      Medium-low       :           0           0 bps
      Medium-high      :           0           0 bps
      High             :           0           0 bps

```

show interfaces queue (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-7/1/3
Physical interface: ge-7/1/3, Enabled, Physical link is Up
  Interface index: 170, SNMP ifIndex: 70 Forwarding classes: 16 supported, 4 in
  use Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets           :       418390039           10 pps
    Bytes              :       38910269752       7440 bps
  Transmitted:
    Packets           :       418390039           10 pps
    Bytes              :       38910269752       7440 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets           :           0           0 pps
    Bytes              :           0           0 bps
  Transmitted:
    Packets           :           0           0 pps

```

```

Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 7055 1 pps
Bytes : 451552 512 bps
Transmitted:
Packets : 7055 1 pps
Bytes : 451552 512 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Forwarding classes: 16 supported, 4 in use Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : 1031 0 pps
Bytes : 143292 0 bps
Transmitted:
Packets : 1031 0 pps
Bytes : 143292 0 bps
Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```


Queue: 3, Forwarding classes: network-control

Queued:

Packets	:	77009	11 pps
Bytes	:	6894286	7888 bps

Transmitted:

Packets	:	77009	11 pps
Bytes	:	6894286	7888 bps

Tail-dropped packets : Not Available

RL-dropped packets	:	0	0 pps
--------------------	---	---	-------

RL-dropped bytes	:	0	0 bps
------------------	---	---	-------

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

Queued:

Packets	:	1031	0 pps
Bytes	:	147328	0 bps

Transmitted:

Packets	:	1031	0 pps
Bytes	:	147328	0 bps

Tail-dropped packets : 0 0 pps

RED-dropped packets : 0 0 pps

Low, non-TCP	:	0	0 pps
--------------	---	---	-------

Low, TCP	:	0	0 pps
----------	---	---	-------

High, non-TCP	:	0	0 pps
---------------	---	---	-------

High, TCP	:	0	0 pps
-----------	---	---	-------

RED-dropped bytes : 0 0 bps

Low, non-TCP	:	0	0 bps
--------------	---	---	-------

Low, TCP	:	0	0 bps
----------	---	---	-------

High, non-TCP	:	0	0 bps
---------------	---	---	-------

High, TCP	:	0	0 bps
-----------	---	---	-------

Queue: 1, Forwarding classes: expedited-forwarding

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Tail-dropped packets : 0 0 pps

RED-dropped packets : 0 0 pps

Low, non-TCP	:	0	0 pps
--------------	---	---	-------

Low, TCP	:	0	0 pps
----------	---	---	-------

High, non-TCP	:	0	0 pps
---------------	---	---	-------

High, TCP	:	0	0 pps
-----------	---	---	-------

RED-dropped bytes : 0 0 bps

Low, non-TCP	:	0	0 bps
--------------	---	---	-------

Low, TCP	:	0	0 bps
----------	---	---	-------

High, non-TCP	:	0	0 bps
---------------	---	---	-------

High, TCP	:	0	0 bps
-----------	---	---	-------

Queue: 2, Forwarding classes: assured-forwarding

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Tail-dropped packets : 0 0 pps

RED-dropped packets : 0 0 pps

Low, non-TCP	:	0	0 pps
--------------	---	---	-------

```

        Low, TCP           :           0           0 pps
        High, non-TCP      :           0           0 pps
        High, TCP          :           0           0 pps
        RED-dropped bytes  :           0           0 bps
        Low, non-TCP       :           0           0 bps
        Low, TCP           :           0           0 bps
        High, non-TCP      :           0           0 bps
        High, TCP          :           0           0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets           :           94386           12 pps
  Bytes            :          13756799          9568 bps
Transmitted:
  Packets           :           94386           12 pps
  Bytes            :          13756799          9568 bps
  Tail-dropped packets :           0           0 pps
  RED-dropped packets :           0           0 pps
  Low, non-TCP      :           0           0 pps
  Low, TCP          :           0           0 pps
  High, non-TCP     :           0           0 pps
  High, TCP         :           0           0 pps
  RED-dropped bytes :           0           0 bps
  Low, non-TCP      :           0           0 bps
  Low, TCP          :           0           0 bps
  High, non-TCP     :           0           0 bps
  High, TCP         :           0           0 bps

```

show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 both-ingress-egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets           : Not Available
  Bytes            :           0           0 bps
Transmitted:
  Packets           :           254           0 pps
  Bytes            :          16274          0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :           0           0 pps
  RED-dropped bytes  :           0           0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets           : Not Available
  Bytes            :           0           0 bps
Transmitted:
  Packets           :           0           0 pps
  Bytes            :           0           0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :           0           0 pps
  RED-dropped bytes  :           0           0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets           : Not Available
  Bytes            :           0           0 bps
Transmitted:
  Packets           :           0           0 pps
  Bytes            :           0           0 bps

```

```

Tail-dropped packets : Not Available
RED-dropped packets  : 0 0 pps
RED-dropped bytes    : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets              : Not Available
Bytes                : 0 0 bps
Transmitted:
Packets              : 0 0 pps
Bytes                : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets  : 0 0 pps
RED-dropped bytes    : 0 0 bps
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets              : Not Available
Bytes                : 0 0 bps
Transmitted:
Packets              : 3 0 pps
Bytes                : 126 0 bps
Tail-dropped packets : Not Available
RED-dropped packets  : 0 0 pps
RED-dropped bytes    : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets              : Not Available
Bytes                : 0 0 bps
Transmitted:
Packets              : 0 0 pps
Bytes                : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets  : 0 0 pps
RED-dropped bytes    : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets              : Not Available
Bytes                : 0 0 bps
Transmitted:
Packets              : 0 0 pps
Bytes                : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets  : 0 0 pps
RED-dropped bytes    : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets              : Not Available
Bytes                : 0 0 bps
Transmitted:
Packets              : 0 0 pps
Bytes                : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets  : 0 0 pps
RED-dropped bytes    : 0 0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets              : 80564692 0 pps
Bytes                : 3383717100 0 bps

```

```

Transmitted:
  Packets      :      80564692      0 pps
  Bytes        :      3383717100    0 bps
  Tail-dropped packets :      0      0 pps
  RED-dropped packets :      0      0 pps
  RED-dropped bytes  :      0      0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets      :      80564685      0 pps
  Bytes        :      3383716770    0 bps
Transmitted:
  Packets      :      80564685      0 pps
  Bytes        :      3383716770    0 bps
  Tail-dropped packets :      0      0 pps
  RED-dropped packets :      0      0 pps
  RED-dropped bytes  :      0      0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets      :      0      0 pps
  Bytes        :      0      0 bps
Transmitted:
  Packets      :      0      0 pps
  Bytes        :      0      0 bps
  Tail-dropped packets :      0      0 pps
  RED-dropped packets :      0      0 pps
  RED-dropped bytes  :      0      0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets      :      9397      0 pps
  Bytes        :      3809052      232 bps
Transmitted:
  Packets      :      9397      0 pps
  Bytes        :      3809052      232 bps
  Tail-dropped packets :      0      0 pps
  RED-dropped packets :      0      0 pps
  RED-dropped bytes  :      0      0 bps

```

show interfaces queue ingress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 ingress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets      : Not Available
  Bytes        :      0      0 bps
Transmitted:
  Packets      :      288      0 pps
  Bytes        :      18450    0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :      0      0 pps
  RED-dropped bytes  :      0      0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets      : Not Available
  Bytes        :      0      0 bps
Transmitted:
  Packets      :      0      0 pps
  Bytes        :      0      0 bps

```

```

Tail-dropped packets : Not Available
RED-dropped packets  : 0 0 pps
RED-dropped bytes    : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets      : Not Available
  Bytes        : 0 0 bps
Transmitted:
  Packets      : 0 0 pps
  Bytes        : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets  : 0 0 pps
RED-dropped bytes    : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets      : Not Available
  Bytes        : 0 0 bps
Transmitted:
  Packets      : 0 0 pps
  Bytes        : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets  : 0 0 pps
RED-dropped bytes    : 0 0 bps

```

show interfaces queue egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets      : Not Available
  Bytes        : 0 0 bps
Transmitted:
  Packets      : 3 0 pps
  Bytes        : 126 0 bps
Tail-dropped packets : Not Available
RED-dropped packets  : 0 0 pps
RED-dropped bytes    : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets      : Not Available
  Bytes        : 0 0 bps
Transmitted:
  Packets      : 0 0 pps
  Bytes        : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets  : 0 0 pps
RED-dropped bytes    : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets      : Not Available
  Bytes        : 0 0 bps
Transmitted:
  Packets      : 0 0 pps
  Bytes        : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets  : 0 0 pps
RED-dropped bytes    : 0 0 bps

```

```

Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      : Not Available
    Bytes        :                      0          0 bps
  Transmitted:
    Packets      :                      0          0 pps
    Bytes        :                      0          0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                      0          0 pps
    RED-dropped bytes  :                      0          0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :          80564692          0 pps
    Bytes        :          3383717100        0 bps
  Transmitted:
    Packets      :          80564692          0 pps
    Bytes        :          3383717100        0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :          80564685          0 pps
    Bytes        :          3383716770        0 bps
  Transmitted:
    Packets      :          80564685          0 pps
    Bytes        :          3383716770        0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
  Transmitted:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      :          9538          0 pps
    Bytes        :          3819840          0 bps
  Transmitted:
    Packets      :          9538          0 pps
    Bytes        :          3819840          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps

```

show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 remaining-traffic
Physical interface: ge-2/2/9, Enabled, Physical link is Up
  Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort

```

```

Queued:
  Packets      :      110208969      472875 pps
  Bytes       :      5951284434    204282000 bps
Transmitted:
  Packets      :      110208969      472875 pps
  Bytes       :      5951284434    204282000 bps
Tail-dropped packets : Not Available
RED-dropped packets :      0      0 pps
  Low          :      0      0 pps
  Medium-low   :      0      0 pps
  Medium-high  :      0      0 pps
  High         :      0      0 pps
RED-dropped bytes  :      0      0 bps
  Low          :      0      0 bps
  Medium-low   :      0      0 bps
  Medium-high  :      0      0 bps
  High         :      0      0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets      :      0      0 pps
  Bytes       :      0      0 bps
Transmitted:
  Packets      :      0      0 pps
  Bytes       :      0      0 bps
Tail-dropped packets : Not Available
RED-dropped packets :      0      0 pps
  Low          :      0      0 pps
  Medium-low   :      0      0 pps
  Medium-high  :      0      0 pps
  High         :      0      0 pps
RED-dropped bytes  :      0      0 bps
  Low          :      0      0 bps
  Medium-low   :      0      0 bps
  Medium-high  :      0      0 bps
  High         :      0      0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets      :      0      0 pps
  Bytes       :      0      0 bps
Transmitted:
  Packets      :      0      0 pps
  Bytes       :      0      0 bps
Tail-dropped packets : Not Available
RED-dropped packets :      0      0 pps
  Low          :      0      0 pps
  Medium-low   :      0      0 pps
  Medium-high  :      0      0 pps
  High         :      0      0 pps
RED-dropped bytes  :      0      0 bps
  Low          :      0      0 bps
  Medium-low   :      0      0 bps
  Medium-high  :      0      0 bps
  High         :      0      0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets      :      0      0 pps
  Bytes       :      0      0 bps
Transmitted:
  Packets      :      0      0 pps
  Bytes       :      0      0 bps
Tail-dropped packets : Not Available

```

```

RED-dropped packets : 0 0 pps
  Low : 0 0 pps
  Medium-low : 0 0 pps
  Medium-high : 0 0 pps
  High : 0 0 pps
RED-dropped bytes : 0 0 bps
  Low : 0 0 bps
  Medium-low : 0 0 bps
  Medium-high : 0 0 bps
  High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets : 109355853 471736 pps
    Bytes : 7436199152 256627968 bps
  Transmitted:
    Packets : 109355852 471736 pps
    Bytes : 7436198640 256627968 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps

```



```

RED-dropped bytes      :                0          0 bps
  Low                  :                0          0 bps
  Medium-low           :                0          0 bps
  Medium-high          :                0          0 bps
  High                 :                0          0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets              :                0          0 pps
  Bytes                :                0          0 bps
Transmitted:
  Packets              :                0          0 pps
  Bytes                :                0          0 bps
Tail-dropped packets : Not Available
RED-dropped packets   :                0          0 pps
  Low                  :                0          0 pps
  Medium-low           :                0          0 pps
  Medium-high          :                0          0 pps
  High                 :                0          0 pps
RED-dropped bytes     :                0          0 bps
  Low                  :                0          0 bps
  Medium-low           :                0          0 bps
  Medium-high          :                0          0 bps
  High                 :                0          0 bps

```

show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode)

```

user@host> show interfaces queue t3-1/1/0:7
Physical interface: t3-1/1/0:7, Enabled, Physical link is Up

Interface index: 192, SNMP ifIndex: 1948

Description: full T3 interface connect to 6ce13 t3-3/1/0:7 for FR testing -
Lam

Forwarding classes: 16 supported, 9 in use

Egress queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

Queued:

  Packets              :            214886          13449 pps
  Bytes                :            9884756        5164536 bps

Transmitted:

  Packets              :            214886          13449 pps
  Bytes                :            9884756        5164536 bps
Tail-dropped packets :                0          0 pps
RED-dropped packets  :                0          0 pps
  Low                 :                0          0 pps
  Medium-low          :                0          0 pps
  Medium-high         :                0          0 pps

```

High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 1, Forwarding classes: REALTIME

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 2, Forwarding classes: PRIVATE

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 3, Forwarding classes: CONTROL

Queued:

Packets	:	60	0 pps
Bytes	:	4560	0 bps

Transmitted:

Packets	:	60	0 pps
Bytes	:	4560	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 4, Forwarding classes: CLASS_B_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 5, Forwarding classes: CLASS_C_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps

RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 6, Forwarding classes: CLASS_V_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 7, Forwarding classes: CLASS_S_OUTPUT, GETS

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps

RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Packet Forwarding Engine Chassis Queues:

Queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

Queued:

Packets	:	371365	23620 pps
Bytes	:	15597330	7936368 bps

Transmitted:

Packets	:	371365	23620 pps
Bytes	:	15597330	7936368 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps

High	:	0	0 bps
Queue: 1, Forwarding classes: REALTIME			
Queued:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Transmitted:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps
Queue: 2, Forwarding classes: PRIVATE			
Queued:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Transmitted:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps

High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 3, Forwarding classes: CONTROL

Queued:

Packets	:	32843	0 pps
Bytes	:	2641754	56 bps

Transmitted:

Packets	:	32843	0 pps
Bytes	:	2641754	56 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 4, Forwarding classes: CLASS_B_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 5, Forwarding classes: CLASS_C_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 6, Forwarding classes: CLASS_V_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 7, Forwarding classes: CLASS_S_OUTPUT, GETS

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps

RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

show interfaces queue (QFX Series)

```

user@switch> show interfaces queue xe-0/0/15
Physical interface: xe-0/0/15, Enabled, Physical link is Up
Interface index: 49165, SNMP ifIndex: 539
Forwarding classes: 12 supported, 8 in use
Egress queues: 12 supported, 8 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      : 0          0 pps
    Bytes        : 0          0 bps
  Transmitted:
    Packets      : 0          0 pps
    Bytes        : 0          0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0          0 pps
    Total-dropped bytes  : 0          0 bps
Queue: 3, Forwarding classes: fcoe
  Queued:
    Packets      : 0          0 pps
    Bytes        : 0          0 bps
  Transmitted:
    Packets      : 0          0 pps
    Bytes        : 0          0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0          0 pps
    Total-dropped bytes  : 0          0 bps
0 bps
Queue: 4, Forwarding classes: no-loss
  Queued:
    Packets      : 0          0 pps
    Bytes        : 0          0 bps
  Transmitted:
    Packets      : 0          0 pps
    Bytes        : 0          0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0          0 pps
    Total-dropped bytes  : 0          0 bps
Queue: 7, Forwarding classes: network-control
  Queued:
    Packets      : 0          0 pps
    Bytes        : 0          0 bps
  Transmitted:
    Packets      : 0          0 pps
    Bytes        : 0          0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0          0 pps
    Total-dropped bytes  : 0          0 bps
Queue: 8, Forwarding classes: mcast
  Queued:

```

Packets	:	0	0 pps
Bytes	:	0	0 bps
Transmitted:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets : Not Available			
Total-dropped packets:		0	0 pps
Total-dropped bytes :		0	0 bps

show interfaces queue l2-statistics (lsq interface)

```

user@switch> show interfaces queue lsq-2/2/0.2 l2-statistics
Logical interface lsq-2/2/0.2 (Index 69) (SNMP ifIndex 1598)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
  Queued:
    Packets      :           1      0 pps
    Bytes        :        1001      0 bps
  Transmitted:
    Packets      :           5      0 pps
    Bytes        :        1062      0 bps
    Tail-dropped packets :           0      0 pps
    RED-dropped packets :           0      0 pps
    RED-dropped bytes  :           0      0 bps
Queue: 1, Forwarding classes: ef
  Queued:
    Packets      :           1      0 pps
    Bytes        :        1500      0 bps
  Transmitted:
    Packets      :           6      0 pps
    Bytes        :        1573      0 bps
    Tail-dropped packets :           0      0 pps
    RED-dropped packets :           0      0 pps
    RED-dropped bytes  :           0      0 bps
Queue: 2, Forwarding classes: af
  Queued:
    Packets      :           1      0 pps
    Bytes        :         512      0 bps
  Transmitted:
    Packets      :           3      0 pps
    Bytes        :         549      0 bps
    Tail-dropped packets :           0      0 pps
    RED-dropped packets :           0      0 pps
    RED-dropped bytes  :           0      0 bps
Queue: 3, Forwarding classes: nc
  Queued:
    Packets      :           0      0 pps
    Bytes        :           0      0 bps
  Transmitted:
    Packets      :           0      0 pps
    Bytes        :           0      0 bps
    Tail-dropped packets :           0      0 pps
    RED-dropped packets :           0      0 pps
    RED-dropped bytes  :           0      0 bps
=====

```

show mvrp

Syntax	show mvrp
Release Information	Command introduced in Junos OS Release 10.1.
Description	For MX Series routers and EX Series switches, display Multiple VLAN Registration Protocol (MVRP) configuration information.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show mvrp applicant-state on page 2049 • show mvrp dynamic-vlan-memberships on page 2051 • show mvrp interface on page 2052 • show mvrp registration-state on page 2053 • show mvrp statistics on page 2055
List of Sample Output	show mvrp on page 2047
Output Fields	Table 143 on page 2047 lists the output fields for the show mvrp command. Output fields are listed in the approximate order in which they appear.

Table 143: show mvrp Output Fields

Field Name	Field Description
MVRP dynamic VLAN creation	Displays whether global MVRP dynamic Virtual LAN (VLAN) creation is Enabled or Disabled .
MVRP BPDU MAC address	Displays the multicast media access control (MAC) address for MVRP. If configured, the provider MVRP multicast MAC address is used; otherwise, the customer MVRP multicast MAC address is used.
MVRP timers (ms)	Displays MVRP timer information: <ul style="list-style-type: none"> • Interface—The interface on which MVRP is configured. • Join—The maximum number of milliseconds the interfaces must wait before sending VLAN advertisements. • Leave—The number of milliseconds an interface must wait after receiving a Leave message to remove the interface from the VLAN specified in the message. • LeaveAll—The interval at which LeaveAll messages are sent on interfaces. LeaveAll messages maintain current MVRP VLAN membership information in the network.

Sample Output

show mvrp

```
user@host> show mvrp
```

```
MVRP configuration for routing instance 'default-switch'
MVRP dynamic VLAN creation : Enabled
MVRP BPDU MAC address      : Customer bridge group (01-80-C2-00-00-21)
MVRP timers (ms)
  Interface      Join   Leave  LeaveAll
  ge-11/2/8      200    800    10000
  ge-11/0/9      200    800    10000
  ge-11/3/0      200    800    10000
```

show mvrp applicant-state

Syntax	show mvrp applicant-state
Release Information	Command introduced in Junos OS Release 10.1.
Description	For MX Series routers and EX Series switches, display Multiple VLAN Registration Protocol (MVRP) applicant state information.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show mvrp on page 2047 • show mvrp interface on page 2052 • show mvrp registration-state on page 2053 • show mvrp statistics on page 2055
List of Sample Output	show mvrp applicant-state on page 2050
Output Fields	Table 144 on page 2049 lists the output fields for the show mvrp applicant-state command. Output fields are listed in the approximate order in which they appear.

Table 144: show mvrp applicant-state Output Fields

Field Name	Field Description
VLAN Id	Displays the Virtual LAN (VLAN) ID number.
Interface	Displays the interface number associated with the VLAN ID.
State	Displays one of the following MVRP registrar states: <ul style="list-style-type: none"> • VO—Very anxious observer. • VP—Very anxious passive. • VA—Very anxious new. • AN—Anxious new. • AA—Anxious active. • QA—Quiet active. • LA—Leaving active. • AO—Anxious observer. • QO—Quiet observer. • LO—Leaving observer. • AP—Anxious passive. • QA—Quiet passive.

Sample Output

show mvrp applicant-state

```
user@host> show mvrp applicant-state
MVRP applicant state for routing instance 'default-switch'
(V0) Very anxious observer, (VP) Very anxious passive, (VA) Very anxious new,
(AN) Anxious new, (AA) Anxious active, (QA) Quiet active, (LA) Leaving active,
(A0) Anxious observer, (Q0) Quiet observer, (L0) Leaving observer,
(AP) Anxious passive, (QP) Quiet passive
```

VLAN Id	Interface	State
100	ge-11/3/0	Declaring (QA)
200	ge-11/3/0	Declaring (QA)
300	ge-11/3/0	Declaring (QA)

show mvrp dynamic-vlan-memberships

Syntax	show mvrp dynamic-vlan-memberships
Release Information	Command introduced in Junos OS Release 10.1.
Description	For MX Series routers and EX Series switches, display all Virtual LANs (VLANs) that have been created dynamically using Multiple VLAN Registration Protocol (MVRP) on the router or switch.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show mvrp on page 2047 • show mvrp applicant-state on page 2049 • show mvrp interface on page 2052 • show mvrp registration-state on page 2053 • show mvrp statistics on page 2055
List of Sample Output	show mvrp dynamic-vlan-memberships on page 2051
Output Fields	Table 145 on page 2051 lists the output fields for the show mvrp dynamic-vlan-memberships command. Output fields are listed in the approximate order in which they appear.

Table 145: show mvrp dynamic-vlan-memberships Output Fields

Field Name	Field Description
VLAN Id	The VLAN ID of the dynamically created VLAN.
Interfaces	The interface or interfaces that are bound to the dynamically created VLAN.

Sample Output

show mvrp dynamic-vlan-memberships

```

user@host> show mvrp dynamic-vlan-memberships
MVRP dynamic vlans for routing instance 'default-switch'
(s) static vlan, (f) fixed registration

VLAN Id      Interfaces
  100 (s)    ge-11/3/0
  200 (s)    ge-11/3/0
  300 (s)

```

show mvrp interface

Syntax	show mvrp interface
Release Information	Command introduced in Junos OS Release 10.1.
Description	For MX Series routers and EX Series switches, display Multiple VLAN Registration Protocol (MVRP) interface-specific information.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show mvrp on page 2047 • show mvrp applicant-state on page 2049 • show mvrp dynamic-vlan-memberships on page 2051 • show mvrp registration-state on page 2053 • show mvrp statistics on page 2055
List of Sample Output	show mvrp interface on page 2052
Output Fields	Table 146 on page 2052 lists the output fields for the show mvrp interface command. Output fields are listed in the approximate order in which they appear.

Table 146: show mvrp interface Output Fields

Field Name	Field Description
Interface	Interface on which MVRP is configured.
Status	Status of the MVRP: Enabled or Disabled .
Registration Mode	Registration for the interface: Fixed , Forbidden , or Normal .
Applicant Mode	Applicant mode.

Sample Output

show mvrp interface

```

user@host> show mvrp interface
MVRP interface information for routing instance 'default-switch'

Interface      Status      Registration Mode      Applicant
Mode
ge-11/2/8      Enabled     Normal                 Normal
ge-11/0/9      Enabled     Normal                 Normal
ge-11/3/0      Enabled     Normal                 Normal

```

show mvrp registration-state

Syntax	show mvrp registration-state
Release Information	Command introduced in Junos OS Release 10.1.
Description	For MX Series routers and EX Series switches, display Multiple VLAN Registration Protocol (MVRP) registration state information.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show mvrp on page 2047 • show mvrp dynamic-vlan-memberships on page 2051 • show mvrp interface on page 2052 • show mvrp statistics on page 2055
List of Sample Output	show mvrp registration-state on page 2053
Output Fields	Table 147 on page 2053 lists the output fields for the show mvrp registration-state command. Output fields are listed in the approximate order in which they appear.

Table 147: show mvrp registration-state Output Fields

Field Name	Field Description
VLAN Id	Displays the Virtual LAN (VLAN) ID number.
Interface	Displays the interface number associated with the VLAN ID.
Registrar State	Displays whether the registrar state is Registered or Empty .
Forced State	Displays whether the forced state is Registered or Empty .
Managed State	Displays one of the following states: <ul style="list-style-type: none"> • Fixed—VLANs always stay in a registered state and are declared as such on all other forwarding ports. • Normal—VLANs participate in the MVRP protocol and honor incoming join requests normally. • Forbidden—VLANs ignore the incoming join requests and always stay in an unregistered state.
STP State	Displays whether the Spanning Tree Protocol (STP) is Blocking or Forwarding .

Sample Output

show mvrp registration-state

```
user@host> show mvrp registration-state
MVRP registration state for routing instance 'default-switch'
```

VLAN Id	Interface	Registrar State	Forced State	Managed State	STP State
100	ge-11/2/8	Empty	Registered	Fixed	Forwarding
	ge-11/0/9	Empty	Empty	Normal	Forwarding
	ge-11/3/0	Registered	Registered	Normal	Forwarding
101	ge-11/2/8	Empty	Registered	Fixed	Forwarding
	ge-11/0/9	Empty	Empty	Normal	Forwarding
	ge-11/3/0	Registered	Registered	Normal	Forwarding

show mvrp statistics

Syntax	show mvrp statistics
Release Information	Command introduced in Junos OS Release 10.1.
Description	For MX Series routers and EX Series switches, display Multiple VLAN Registration Protocol (MVRP) statistics in the form of Multiple Registration Protocol data unit (MRPDU) messages.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show mvrp on page 2047 • show mvrp applicant-state on page 2049 • show mvrp dynamic-vlan-memberships on page 2051 • show mvrp interface on page 2052 • show mvrp registration-state on page 2053
List of Sample Output	show mvrp statistics on page 2055
Output Fields	Table 148 on page 2055 lists the output fields for the show mvrp statistics command. Output fields are listed in the approximate order in which they appear.

Table 148: show mvrp statistics Output Fields

Field Name	Field Description
interface name	Interface for which MVRP statistics are displayed.
VLAN IDs registered	Number of Virtual LAN (VLAN) IDs registered.
Sent MVRP PDUs	Number of MRPDU messages transmitted from the router.
Received MVRP PDUs without error	Number of MRPDU messages received on the router.
Received MVRP PDUs with error	Number of invalid MRPDU messages received on the router.

Sample Output

show mvrp statistics

```

user@host> show mvrp statistics
MVRP statistics for routing instance 'default-switch'

Interface name           : ge-11/2/8
VLAN IDs registered      : 0
Sent MVRP PDUs           : 1467

```

Received MVRP PDUs without error: 0
Received MVRP PDUs with error : 0

Interface name : ge-11/0/9
VLAN IDs registered : 0
Sent MVRP PDUs : 1418
Received MVRP PDUs without error: 702
Received MVRP PDUs with error : 0

Interface name : ge-11/3/0
VLAN IDs registered : 2
Sent MVRP PDUs : 1524
Received MVRP PDUs without error: 1366
Received MVRP PDUs with error : 0

show vlans

Syntax	<pre>show vlans <brief detail extensive> <instance <i>instance-name</i>> <logical-system <i>logical-system-name</i>> <operational> <<i>vlan-name</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 12.3R2.</p> <p>Command introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	(MX Series routers and EX Series switches only) Display VLAN information.
Options	<p>none—Display information for all VLANs.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display information for the specified routing instance.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Display Ethernet-switching statistics information for the specified logical system.</p> <p>operational—(Optional) Display information for the operational routing instances.</p> <p><i>vlan-name</i>— (Optional) Display information about the specified VLAN.</p>
Required Privilege Level	view
List of Sample Output	<p>[xref target has no title]</p> <p>show vlans detail on page 2058</p>

Sample Output

Routing instance	VLAN name	Tag	Interfaces
VPLS-1	__VPLS-1__	all	ae1.0
VPLS-2	__VPLS-2__	all	ae3.0 ge-3/1/2.0 vt-3/3/10.1048576
default-switch	VLAN1000	1000	ae26.0
default-switch	VLAN101	101	ae20.0
default-switch	VLAN102	102	ae20.0
default-switch	VLAN103	103	ae20.0
default-switch	VLAN104	104	ae20.0
default-switch	VLAN105	105	ae20.0
default-switch	VLAN106	106	ae20.0

```
default-switch      VLAN107      107      ae20.0
default-switch      VLAN108      108      ae20.0
[...output truncated...]
```

show vlans detail

```
user@host> show vlans detail
Routing instance: VPLS-1
  VLAN Name: __VPLS-1__      State: Active
  Tag: all
  Internal index: 2, Generation Index: , Origin: Dynamic
  Interfaces:
    ae1.0,tagged
  Number of interfaces: Tagged 1 , Untagged 0
  Total MAC count: 0

Routing instance: VPLS-2
  VLAN Name: __VPLS-2__      State: Active
  Tag: all
  Internal index: 3, Generation Index: , Origin: Dynamic
  Interfaces:
    ae3.0,tagged
    ge-3/1/2.0,tagged
    vt-3/3/10.1048576,tagged
  Number of interfaces: Tagged 3 , Untagged 0
  Total MAC count: 4

Routing instance: default-switch
  VLAN Name: VLAN1000      State: Active
  Tag: 1000
  Internal index: 4, Generation Index: 1, Origin: Static
  Layer 3 interface: irb.1000
  Interfaces:
    ae26.0,tagged,trunk
  Number of interfaces: Tagged 1 , Untagged 0
  Total MAC count: 0

Routing instance: default-switch
  VLAN Name: VLAN101      State: Active
  Tag: 101
  Internal index: 5, Generation Index: 2, Origin: Static
  Layer 3 interface: irb.101
  Interfaces:
    ae20.0,tagged,trunk
  Number of interfaces: Tagged 1 , Untagged 0
  Total MAC count: 1

Routing instance: default-switch
  VLAN Name: VLAN102      State: Active
  Tag: 102
  Internal index: 6, Generation Index: 3, Origin: Static
  Layer 3 interface: irb.102
  Interfaces:
    ae20.0,tagged,trunk
  Number of interfaces: Tagged 1 , Untagged 0
  Total MAC count: 1
[...output truncated...]
```


traceroute ethernet

Syntax	traceroute ethernet (<i>mac-address</i> <i>mep-id</i>) maintenance-association <i>ma-name</i> maintenance-domain <i>md-name</i> ttl <i>value</i> <wait seconds>
Release Information	Command introduced in Junos OS Release 9.0. mep-id option introduced in Junos OS Release 9.1.
Description	Triggers the linktrace protocol to trace the route between two maintenance points. The result of the traceroute protocol is stored in the path database. To display the path database, use the show oam ethernet connectivity-fault-management path-database command. Before using the traceroute command, you can verify the remote MEP's MAC address using the show oam ethernet connectivity-fault-management path-database command.
Options	mac-address —Destination unicast MAC address of the remote maintenance point. mep-id —MEP identifier of the remote maintenance point. The range of values is 1 through 8191. maintenance-association <i>ma-name</i> —Specifies an existing maintenance association from the set of configured maintenance associations. maintenance-domain <i>md-name</i> —Specifies an existing maintenance domain from the set of configured maintenance domains. ttl value —Number of hops to use in the linktrace request. The range is 1 to 255 hops. The default is 4. wait seconds —(Optional) Maximum time to wait for a response to the traceroute request. The range is 1 to 255 seconds. The default is 5.
Required Privilege Level	network
List of Sample Output	traceroute ethernet on page 2060
Output Fields	Table 149 on page 2059 lists the output fields for the traceroute ethernet command. Output fields are listed in the approximate order in which they appear.

Table 149: traceroute ethernet Output Fields

Field Name	Field Description
Linktrace to	MAC address of the destination maintenance point.
Interface	Local interface used to send the linktrace message (LTM).

Table 149: traceroute ethernet Output Fields (*continued*)

Field Name	Field Description
Maintenance Domain	Maintenance domain specified in the traceroute command.
Level	Maintenance domain level configured.
Maintenance Association	Maintenance association specified in the traceroute command.
Local Mep	The local maintenance end point identifier.
Transaction Identifier	4-byte identifier maintained by the MEP. Each LTM uses a transaction identifier. The transaction identifier is maintained globally across all Maintenance Domains. Use the transaction identifier to match an incoming linktrace response (LTR), with a previously sent LTM.
Hop	Sequential hop count of the linktrace path.
TTL	Number of hops remaining in the linktrace message. The time to live (TTL) is decremented at each hop.
Source MAC address	MAC address of the 802.1ag maintenance point that is sending the linktrace message.
Next-hop MAC address	MAC address of the 802.1ag node that is the next hop in the LTM path.

Sample Output

traceroute ethernet

```

user@host> traceroute ethernet maintenance-domain md1 maintenance-association ma1
00:90:69:7e:01:ff
Linktrace to 00:01:02:03:04:05, Interface : ge-5/0/0.0
  Maintenance Domain: MD1, Level: 7
  Maintenance Association: MA1, Local Mep: 1

Hop      TTL      Source MAC address      Next hop MAC address
Transaction Identifier:100001
1         63      00:00:aa:aa:aa:aa      00:00:bb:bb:bb:bb
2         62      00:00:bb:bb:bb:bb      00:00:cc:cc:cc:cc
3         61      00:00:cc:cc:cc:cc      00:01:02:03:04:05
4         60      00:01:02:03:04:05      00:00:00:00:00:00

```

CHAPTER 9

High Availability

- [Overview on page 2061](#)
- [Configuration on page 2116](#)
- [Administration on page 2270](#)
- [Troubleshooting on page 2278](#)

Overview

- [Graceful Routing Engine Switchover \(GRES\) on page 2061](#)
- [Nonstop Bridging \(NSB\) on page 2069](#)
- [Nonstop Active Routing \(NSR\) on page 2072](#)
- [Graceful Restart on page 2085](#)
- [Unified ISSU on page 2092](#)
- [VRRP on page 2109](#)

Graceful Routing Engine Switchover (GRES)

- [Understanding Graceful Routing Engine Switchover in the Junos OS on page 2061](#)
- [Graceful Routing Engine Switchover System Requirements on page 2065](#)
- [Requirements for Routers with a Backup Router Configuration on page 2068](#)

Understanding Graceful Routing Engine Switchover in the Junos OS

This topic contains the following sections:

- [Graceful Routing Engine Switchover Concepts on page 2061](#)
- [Effects of a Routing Engine Switchover on page 2064](#)

Graceful Routing Engine Switchover Concepts

Graceful Routing Engine switchover (GRES) feature in Junos OS enables a routing platform with redundant Routing Engines to continue forwarding packets, even if one Routing Engine fails. Graceful Routing Engine switchover preserves interface and kernel information. Traffic is not interrupted. However, graceful Routing Engine switchover does not preserve the control plane. Neighboring routers detect that the router has experienced a restart and react to the event in a manner prescribed by individual routing protocol specifications. To preserve routing during a switchover, graceful Routing Engine switchover

must be combined with either graceful restart protocol extensions or nonstop active routing. Any updates to the master Routing Engine are replicated to the backup Routing Engine as soon as they occur. If the kernel on the master Routing Engine stops operating, the master Routing Engine experiences a hardware failure, or the administrator initiates a manual switchover, mastership switches to the backup Routing Engine.



NOTE: To quickly restore or to preserve routing protocol state information during a switchover, graceful Routing Engine switchover must be combined with either graceful restart or nonstop active routing (NSR), respectively. For more information about graceful restart, see [“Graceful Restart Concepts” on page 2085](#). For more information about nonstop active routing, see [“Nonstop Active Routing Concepts” on page 2072](#).

If the backup Routing Engine does not receive a keepalive from the master Routing Engine after 2 seconds (4 seconds on M20 routers), it determines that the master Routing Engine has failed and takes mastership. The Packet Forwarding Engine seamlessly disconnects from the old master Routing Engine and reconnects to the new master Routing Engine. The Packet Forwarding Engine does not reboot, and traffic is not interrupted. The new master Routing Engine and the Packet Forwarding Engine then become synchronized. If the new master Routing Engine detects that the Packet Forwarding Engine state is not up to date, it resends state update messages.



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

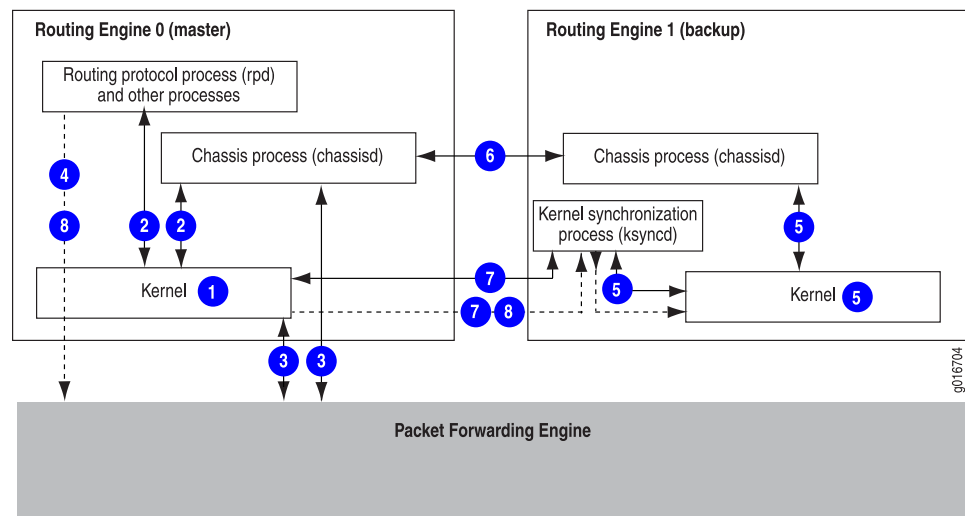
If the router displays a warning message similar to “Standby Routing Engine is not ready for graceful switchover. Packet Forwarding Engines that are not ready for graceful switchover might be reset,” do not attempt switchover. If you choose to proceed with switchover, only the Packet Forwarding Engines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



NOTE: We do not recommend performing a commit operation on the backup Routing Engine when graceful Routing Engine switchover is enabled on the router.

[Figure 19 on page 2063](#) shows the system architecture of graceful Routing Engine switchover and the process a routing platform follows to prepare for a switchover.

Figure 19: Preparing for a Graceful Routing Engine Switchover

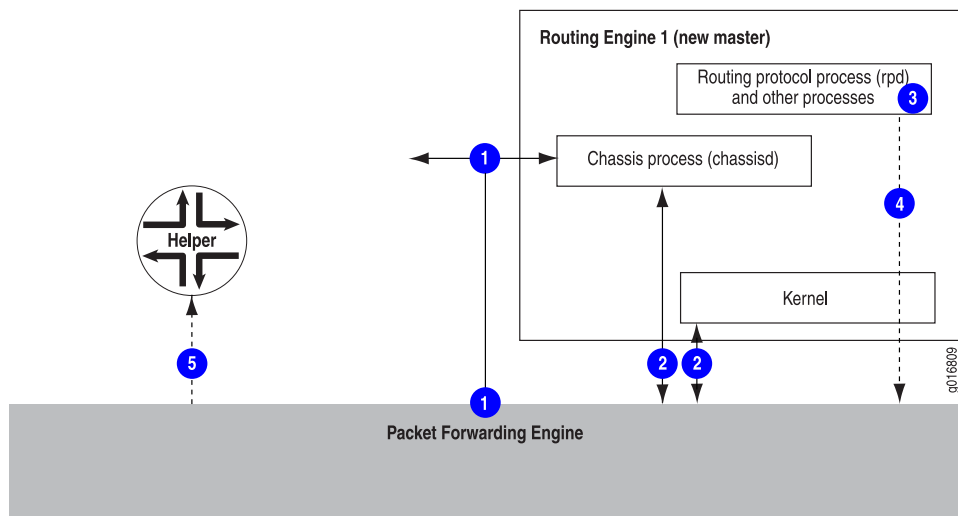


The switchover preparation process for graceful Routing Engine switchover follows these steps:

1. The master Routing Engine starts.
2. The routing platform processes (such as the chassis process [chassisd]) start.
3. The Packet Forwarding Engine starts and connects to the master Routing Engine.
4. All state information is updated in the system.
5. The backup Routing Engine starts.
6. The system determines whether graceful Routing Engine switchover has been enabled.
7. The kernel synchronization process (ksyncd) synchronizes the backup Routing Engine with the master Routing Engine.
8. After ksyncd completes the synchronization, all state information and the forwarding table are updated.

Figure 20 on page 2064 shows the effects of a switchover on the routing platform.

Figure 20: Graceful Routing Engine Switchover Process



When a switchover occurs, the switchover process follows these steps:

1. When keepalives from the master Routing Engine are lost, the system switches over gracefully to the backup Routing Engine.
2. The Packet Forwarding Engine connects to the backup Routing Engine, which becomes the new master.
3. Routing platform processes that are not part of graceful Routing Engine switchover (such as the routing protocol process [rpd]) restart.
4. State information learned from the point of the switchover is updated in the system.
5. If configured, graceful restart protocol extensions collect and restore routing information from neighboring peer *helper* routers.



NOTE: On T Series and M320 routers, the Switch Interface Boards (SIBs) are taken offline and restarted one by one during a graceful Routing Engine switchover. This is done to provide the SPMB that manages the SIB enough time to populate state information for its associated SIB. However, on a fully-populated chassis where all FPCs are sending traffic at full line rate, there might be momentary packet loss during the switchover.

Effects of a Routing Engine Switchover

Table 150 on page 2065 describes the effects of a Routing Engine switchover when no high availability features are enabled and when graceful Routing Engine switchover, graceful restart, and nonstop active routing features are enabled.

Table 150: Effects of a Routing Engine Switchover

Feature	Benefits	Considerations
Dual Routing Engines only (no features enabled)	When the switchover to the new master Routing Engine is complete, routing convergence takes place and traffic is resumed.	All physical interfaces are taken offline, Packet Forwarding Engines restart, the standby Routing Engine restarts the routing protocol process (rpd), and all hardware and interfaces are discovered by the new master Routing Engine. The switchover takes several minutes and all of the router's adjacencies are aware of the physical (interface alarms) and routing (topology) change.
Graceful Routing Engine switchover enabled	During the switchover, interface and kernel information is preserved. The switchover is faster because the Packet Forwarding Engines are not restarted.	The new master Routing Engine restarts the routing protocol process (rpd). All hardware and interfaces are acquired by a process that is similar to a warm restart. All adjacencies are aware of the router's change in state.
Graceful Routing Engine switchover and nonstop active routing enabled	Traffic is not interrupted during the switchover. Interface, kernel, and routing protocol information is preserved.	Unsupported protocols must be refreshed using the normal recovery mechanisms inherent in each protocol.
Graceful Routing Engine switchover and graceful restart enabled	Traffic is not interrupted during the switchover. Interface and kernel information is preserved. Graceful restart protocol extensions quickly collect and restore routing information from the neighboring routers.	Neighbors are required to support graceful restart and a wait interval is required. The routing protocol process (rpd) restarts. For certain protocols, a significant change in the network can cause graceful restart to stop.

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Graceful Routing Engine Switchover System Requirements on page 2065](#)
- [Configuring Graceful Routing Engine Switchover on page 2117](#)
- [Requirements for Routers with a Backup Router Configuration on page 2068](#)

Graceful Routing Engine Switchover System Requirements

Graceful Routing Engine switchover is supported on all routing platforms that contain dual Routing Engines. All Routing Engines configured for graceful Routing Engine switchover must run the same Junos OS Release. Hardware and software support for graceful Routing Engine switchover is described in the following sections:

- [Graceful Routing Engine Switchover Platform Support on page 2066](#)
- [Graceful Routing Engine Switchover Feature Support on page 2066](#)
- [Graceful Routing Engine Switchover DPC Support on page 2067](#)
- [Graceful Routing Engine Switchover and Subscriber Access on page 2068](#)
- [Graceful Routing Engine Switchover PIC Support on page 2068](#)

Graceful Routing Engine Switchover Platform Support

To enable graceful Routing Engine switchover, your system must meet these minimum requirements:

- M20 and M40e routers—Junos OS Release 5.7 or later
- M10i router—Junos OS Release 6.1 or later
- M320 router—Junos OS Release 6.2 or later
- T320 router, T640 router, and TX Matrix router—Junos OS Release 7.0 or later
- M120 router—Junos OS Release 8.2 or later
- MX960 router—Junos OS Release 8.3 or later
- MX480 router—Junos OS Release 8.4 or later (8.4R2 recommended)
- MX240 router—Junos OS Release 9.0 or later
- T1600 router—Junos OS Release 8.5 or later
- T4000 router—Junos OS Release 12.1R2 or later
- TX Matrix Plus router—Junos OS Release 9.6 or later

For more information about support for graceful Routing Engine switchover, see the sections that follow.

Graceful Routing Engine Switchover Feature Support

Graceful Routing Engine switchover supports most Junos OS features in Release 5.7 and later. Particular Junos OS features require specific versions of Junos OS. See [Table 151 on page 2066](#).

Table 151: Graceful Routing Engine Switchover Feature Support

Application	Junos OS Release
Aggregated Ethernet interfaces with Link Aggregation Control Protocol (LACP) and aggregated SONET interfaces	6.2
Asynchronous Transfer Mode (ATM) virtual circuits (VCs)	6.2
Logical systems	6.3
NOTE: In Junos OS Release 9.3 and later, the logical router feature is renamed to logical system.	
Multicast	6.4 (7.0 for TX Matrix router)
Multilink Point-to-Point Protocol (MLPPP) and Multilink Frame Relay (MLFR)	7.0

Table 151: Graceful Routing Engine Switchover Feature Support (*continued*)

Application	Junos OS Release
Automatic Protection Switching (APS)—The current active interface (either the designated working or the designated protect interface) remains the active interface during a Routing Engine switchover.	7.4
Point-to-multipoint Multiprotocol Label Switching MPLS LSPs (transit only)	7.4
Compressed Real-Time Transport Protocol (CRTP)	7.6
Virtual private LAN service (VPLS)	8.2
Ethernet Operation, Administration, and Management (OAM) as defined by IEEE 802.3ah	8.5
Extended DHCP relay agent	8.5
Ethernet OAM as defined by IEEE 802.1ag	9.0
Packet Gateway Control Protocol (PGCP) process (pgcpd) on Multiservices 500 PICs on T640 routers.	9.0
Subscriber access	9.4
Layer 2 Circuit and LDP-based VPLS pseudowire redundant configuration	9.6

The following constraints apply to graceful Routing Engine switchover feature support:

- When graceful Routing Engine switchover and aggregated Ethernet interfaces are configured in the same system, the aggregated Ethernet interfaces must not be configured for fast-polling LACP. When fast polling is configured, the LACP polls time out at the remote end during the Routing Engine mastership switchover. When LACP polling times out, the aggregated link and interface are disabled. The Routing Engine mastership change is fast enough that standard and slow LACP polling do not time out during the procedure. However, note that this restriction does not apply to MX Series Routers that are running Junos OS Release 9.4 or later and have distributed periodic packet management (PPM) enabled—which is the default configuration—on them. In such cases, you can configure graceful Routing Engine switchover and have aggregated Ethernet interfaces configured for fast-polling LACP on the same device.
- VRRP changes mastership when a Routing Engine switchover occurs, even when graceful Routing Engine switchover is configured.

Graceful Routing Engine Switchover DPC Support

Graceful Routing Engine switchover supports all Dense Port Concentrators (DPCs) on the MX Series 3D Universal Edge Routers running the appropriate version of Junos OS. For more information about DPCs, see the *MX Series DPC Guide*.

Graceful Routing Engine Switchover and Subscriber Access

Graceful Routing Engine switchover currently supports most of the features directly associated with dynamic DHCP and dynamic PPPoE subscriber access. Graceful Routing Engine switchover also supports the unified in-service software upgrade (ISSU) for the DHCP access model and the PPPoE access model used by subscriber access.

Graceful Routing Engine Switchover PIC Support

Graceful Routing Engine switchover is supported on most PICs, except for the services PICs listed in this section. The PIC must be on a supported routing platform running the appropriate version of Junos OS. For information about FPC types, FPC/PIC compatibility, and the initial Junos OS Release in which an FPC supported a particular PIC, see the PIC guide for your router platform.

The following constraints apply to graceful Routing Engine switchover support for services PICs:

- You can include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level on a router with Adaptive Services, Multiservices, and Tunnel Services PICs configured on it and successfully commit the configuration. However, all services on these PICs—except the Layer 2 service packages and extension-provider and SDK applications on Multiservices PICs—are reset during a switchover.
- Graceful Routing Engine switchover is not supported on any Monitoring Services PICs or Multilink Services PICs. If you include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level on a router with either of these PIC types configured on it and issue the **commit** command, the commit fails.
- Graceful Routing Engine switchover is not supported on Multiservices 400 PICs configured for monitoring services applications. If you include the **graceful-switchover** statement, the commit fails.



NOTE: When an unsupported PIC is online, you cannot enable graceful Routing Engine switchover. If graceful Routing Engine switchover is already enabled, an unsupported PIC cannot come online.

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Understanding Graceful Routing Engine Switchover in the Junos OS on page 2061](#)
- [Configuring Graceful Routing Engine Switchover on page 2117](#)
- [Requirements for Routers with a Backup Router Configuration on page 2068](#)

Requirements for Routers with a Backup Router Configuration

If your Routing Engine configuration includes a **backup-router** statement or an **inet6-backup-router** statement, you can also use the **destination** statement to specify a subnet address or multiple subnet addresses for the backup router. Include destination subnets for the backup Routing Engine at the **[edit system (backup-router]**

inet6-backup-router) address] hierarchy level. This requirement also applies to any T640 router connected to a TX Matrix router that includes a **backup-router** or **inet6-backup-router** statement.



NOTE: If you have a backup router configuration in which multiple static routes point to a gateway from the management Ethernet interface, you must configure prefixes that are more specific than the static routes or include the **retain** flag at the **[edit routing-options static route]** hierarchy level.

For example, if you configure the static route 172.16.0.0/12 from the management Ethernet interface for management purposes, you must specify the backup router configuration as follows:

```
backup-router 172.29.201.62 destination [172.16.0.0/13 172.16.128.0/13]
```

**Related
Documentation**

- [Understanding Graceful Routing Engine Switchover in the Junos OS on page 2061](#)
- [Graceful Routing Engine Switchover System Requirements on page 2065](#)

Nonstop Bridging (NSB)

- [Nonstop Bridging Concepts on page 2069](#)
- [Nonstop Bridging System Requirements on page 2071](#)

Nonstop Bridging Concepts


Nonstop bridging uses the same infrastructure as graceful Routing Engine switchover (GRES) to preserve interface and kernel information. However, nonstop bridging also saves Layer 2 Control Protocol (L2CP) information by running the Layer 2 Control Protocol process (l2cpd) on the backup Routing Engine.



NOTE: To use nonstop bridging, you must first enable graceful Routing Engine switchover on your routing platform. For more information about graceful Routing Engine switchover, see [“Understanding Graceful Routing Engine Switchover in the Junos OS” on page 2061](#).

Figure 21 on page 2070 shows the system architecture of nonstop bridging and the process a routing platform follows to prepare for a switchover.

Figure 21: Nonstop Bridging Switchover Preparation Process




The switchover preparation process for nonstop bridging follows these steps:

1. The master Routing Engine starts.
2. The routing platform processes on the master Routing Engine (such as the chassis process [chassisd] and the Layer 2 Control Protocol process [l2cpd]) start.
3. The Packet Forwarding Engine starts and connects to the master Routing Engine.
4. All state information is updated in the system.
5. The backup Routing Engine starts, including the chassis process (chassisd) and the Layer 2 Control Protocol process (l2cpd).
6. The system determines whether graceful Routing Engine switchover and nonstop bridging have been enabled.
7. The kernel synchronization process (ksyncd) synchronizes the backup Routing Engine with the master Routing Engine.
8. For supported protocols, state information is updated directly between the l2cpds on the master and backup Routing Engines.

Figure 22 on page 2071 shows the effects of a switchover on the routing platform.

Figure 22: Nonstop Bridging During a Switchover



The switchover process follows these steps:

1. When keepalives from the master Routing Engine are lost, the system switches over gracefully to the backup Routing Engine.
2. The Packet Forwarding Engine connects to the backup Routing Engine, which becomes the new master. Because the Layer 2 Control Protocol process (l2cpd) and chassis process (chassisd) are already running, these processes do not need to restart.
3. State information learned from the point of the switchover is updated in the system. Forwarding and bridging are continued during the switchover, resulting in minimal packet loss.

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Nonstop Bridging System Requirements on page 2071](#)
- [Configuring Nonstop Bridging on page 2180](#)
- [Configuring Nonstop Bridging on EX Series Switches \(CLI Procedure\)](#)

Nonstop Bridging System Requirements

This topic contains the following sections:

- [Platform Support on page 2071](#)
- [Protocol Support on page 2072](#)

Platform Support

Nonstop bridging is supported on MX Series 3D Universal Edge Routers. Your system must be running Junos OS Release 8.4 or later.

Nonstop bridging is supported on EX Series Ethernet Switch with redundant Routing Engines.

For a list of the EX Series switches and Layer 2 protocols that support nonstop bridging, see [“EX Series Switch Software Features Overview” on page 63](#).



NOTE: All Routing Engines configured for nonstop bridging must be running the same Junos OS release.

Protocol Support

Nonstop bridging is supported for the following Layer 2 control protocols:

- Spanning Tree Protocol (STP)
- Rapid Spanning Tree Protocol (RSTP)
- Multiple Spanning Tree Protocol (MSTP)

Related Documentation

- [Nonstop Bridging Concepts on page 2069](#)
- [Configuring Nonstop Bridging on page 2180](#)
- [Configuring Nonstop Bridging on EX Series Switches \(CLI Procedure\)](#)

Nonstop Active Routing (NSR)

- [Nonstop Active Routing Concepts on page 2072](#)
- [Nonstop Active Routing System Requirements on page 2075](#)

Nonstop Active Routing Concepts


Nonstop active routing (NSR) uses the same infrastructure as graceful Routing Engine switchover (GRES) to preserve interface and kernel information. However, nonstop active routing also saves routing protocol information by running the routing protocol process (rpd) on the backup Routing Engine. By saving this additional information, nonstop active routing is self-contained and does not rely on helper routers to assist the routing platform in restoring routing protocol information. Nonstop active routing is advantageous in networks where neighbor routers do not support graceful restart protocol extensions. As a result of this enhanced functionality, nonstop active routing is a natural replacement for graceful restart.



NOTE: To use nonstop active routing, you must first enable graceful Routing Engine switchover on your routing platform. For more information about graceful Routing Engine switchover, see [“Understanding Graceful Routing Engine Switchover in the Junos OS” on page 2061](#).

Figure 23 on page 2073 shows the system architecture of nonstop active routing and the process a routing platform follows to prepare for a switchover.

Figure 23: Nonstop Active Routing Switchover Preparation Process

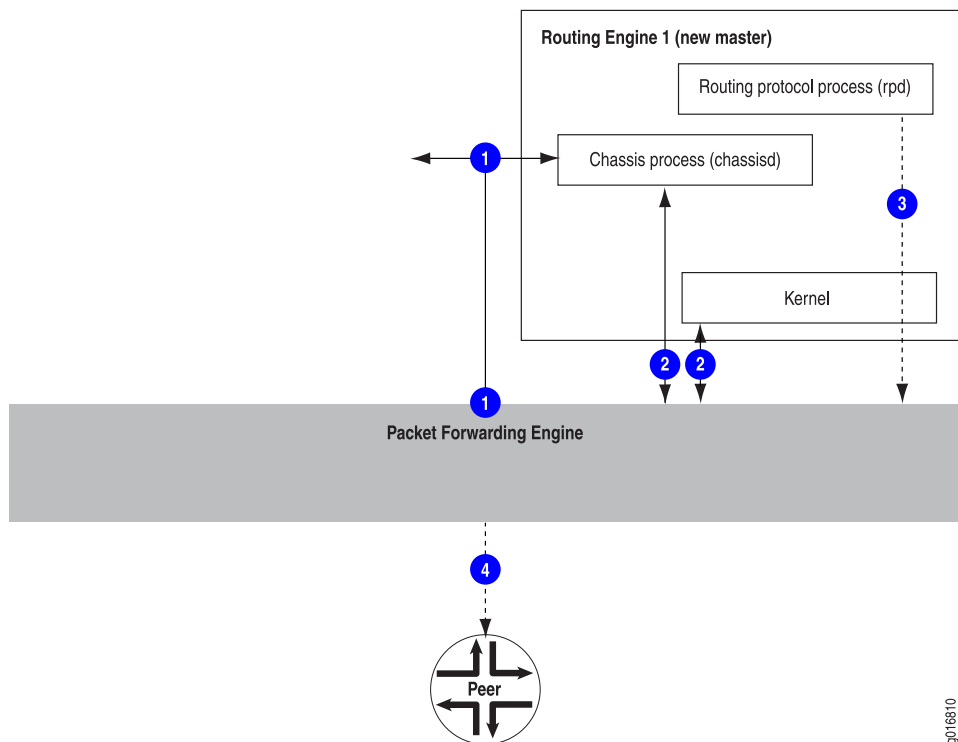


The switchover preparation process for nonstop active routing follows these steps:

1. The master Routing Engine starts.
2. The routing platform processes on the master Routing Engine (such as the chassis process [chassisd] and the routing protocol process [rpd]) start.
3. The Packet Forwarding Engine starts and connects to the master Routing Engine.
4. All state information is updated in the system.
5. The backup Routing Engine starts, including the chassis process (chassisd) and the routing protocol process (rpd).
6. The system determines whether graceful Routing Engine switchover and nonstop active routing have been enabled.
7. The kernel synchronization process (ksyncd) synchronizes the backup Routing Engine with the master Routing Engine.
8. For supported protocols, state information is updated directly between the routing protocol processes on the master and backup Routing Engines.

Figure 24 on page 2074 shows the effects of a switchover on the routing platform.

Figure 24: Nonstop Active Routing During a Switchover



The switchover process follows these steps:

1. When keepalives from the master Routing Engine are lost, the system switches over gracefully to the backup Routing Engine.
2. The Packet Forwarding Engine connects to the backup Routing Engine, which becomes the new master. Because the routing protocol process (rpd) and chassis process (chassisd) are already running, these processes do not need to restart.
3. State information learned from the point of the switchover is updated in the system. Forwarding and routing are continued during the switchover, resulting in minimal packet loss.
4. Peer routers continue to interact with the routing platform as if no change had occurred. Routing adjacencies and session state relying on underlying routing information are preserved and not reset.

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Nonstop Active Routing System Requirements on page 2075](#)
- [Configuring Nonstop Active Routing on page 2183](#)

Nonstop Active Routing System Requirements

This section contains the following topics:

- [Nonstop Active Routing Platform Support on page 2075](#)
- [Nonstop Active Routing Protocol and Feature Support on page 2076](#)
- [Nonstop Active Routing BFD Support on page 2078](#)
- [Nonstop Active Routing BGP Support on page 2079](#)
- [Nonstop Active Routing Layer 2 Circuit and VPLS Support on page 2080](#)
- [Nonstop Active Routing PIM Support on page 2080](#)
- [Nonstop Active Routing MSDP Support on page 2083](#)
- [Nonstop Active Routing Support for RSVP-TE LSPs on page 2083](#)

Nonstop Active Routing Platform Support

[Table 152 on page 2075](#) lists the platforms that support nonstop active routing (NSR).

Table 152: Nonstop Active Routing Platform Support

Platform	Junos OS Release
M10i router	8.4 or later
M20 router	8.4 or later
M40e router	8.4 or later
M120 router	9.0 or later
M320 router	8.4 or later
MX Series routers	9.0 or later
PTX Series Packet Transport switches	12.1R4 or later
<p>NOTE:</p> <p>Nonstop active routing (NSR) switchover on PTX series is supported only for the following MPLS and VPN protocols and applications using chained composite next hops:</p> <ul style="list-style-type: none"> • Labeled BGP • Layer 2 VPNs excluding Layer 2 interworking (Layer 2 stitching) • Layer 3 VPNs • LDP • RSVP 	
T320 router, T640 router, and TX Matrix router	8.4 or later
T1600 router	8.5 or later

Table 152: Nonstop Active Routing Platform Support (*continued*)

Platform	Junos OS Release
TX Plus Matrix router	10.0 or later



NOTE: All Routing Engines configured for nonstop active routing must be running the same Junos OS release.

Nonstop Active Routing Protocol and Feature Support

Table 153 on page 2076 lists the protocols that are supported by nonstop active routing.

Table 153: Nonstop Active Routing Protocol and Feature Support

Protocol	Junos OS Release
Aggregated Ethernet interfaces with Link Aggregation Control Protocol (LACP)	9.4 or later
Bidirectional Forwarding Detection (BFD)	8.5 or later
For more information, see “Nonstop Active Routing BFD Support” on page 2078.	
BGP	8.4 or later
For more information, see “Nonstop Active Routing BGP Support” on page 2079.	
Labeled BGP (PTX Series Packet Transport Switches only)	12.1R4 or later
IS-IS	8.4 or later
LDP	8.4 or later
LDP-based virtual private LAN service (VPLS)	9.3 or later
LDP OAM (operation, administration, and management) features	9.6 or later
LDP (PTX Series Packet Transport Switches only)	12.1R4 or later

Nonstop active routing support for LDP includes:

- LDP unicast transit LSPs
- LDP egress LSPs for labeled internal BGP (IBGP) and external BGP (EBGP)
- LDP over RSVP transit LSPs
- LDP transit LSPs with indexed next hops
- LDP transit LSPs with unequal cost load balancing

NOTE: Nonstop active routing is not supported for LDP Point-to-Multipoint LSPs and LDP ingress LSPs.

Table 153: Nonstop Active Routing Protocol and Feature Support (*continued*)

Protocol	Junos OS Release
Layer 2 circuits	(on LDP-based VPLS) 9.2 or later (on RSVP-TE LSP) 11.1 or later
Layer 2 VPNs	9.1 or later
Layer 2 VPNs (PTX Series Packet Transport Switches only)	12.1R4 or later
NOTE: Nonstop active routing is not supported for Layer 2 interworking (Layer 2 stitching).	
Layer 3 VPNs (see the first Note after this table for restrictions)	9.2 or later
Layer 3 VPNs (PTX Series Packet Transport Switches only)	12.1R4 or later
Multicast Source Discovery Protocol (MSDP)	12.1 or later
For more information, see “Nonstop Active Routing MSDP Support” on page 2083 .	
OSPF/OSPFv3	8.4 or later
Protocol Independent Multicast (PIM)	(for IPv4) 9.3 or later
For more information, see “Nonstop Active Routing PIM Support” on page 2080 .	(for IPv6) 10.4 or later
RIP and RIP next generation (RIPng)	9.0 or later
RSVP (PTX Series Packet Transport Switches only)	12.1R4 or later
Nonstop active routing support for RSVP includes:	
<ul style="list-style-type: none"> Point-to-Multipoint LSPs <ul style="list-style-type: none"> RSVP Point-to-Multipoint ingress, transit, and egress LSPs using existing non-chained next hop. RSVP Point-to-Multipoint transit LSPs using composite next hops for Point-to-Multipoint label routes. Point-to-Point LSPs <ul style="list-style-type: none"> RSVP Point-to-Point ingress, transit, and egress LSPs using non-chained next hops. RSVP Point-to-Point transit LSPs using chained composite next hops. 	
RSVP-TE LSP	9.5 or later
For more information, see “Nonstop Active Routing Support for RSVP-TE LSPs” on page 2083 .	

Table 153: Nonstop Active Routing Protocol and Feature Support (*continued*)

Protocol	Junos OS Release
VPLS	(LDP-based) 9.1 or later (RSVP-TE-based) 11.2 or later



NOTE: Layer 3 VPN support does not include dynamic GRE tunnels, multicast VPNs, or BGP flow routes.



NOTE: If you configure a protocol that is not supported by nonstop active routing, the protocol operates as usual. When a switchover occurs, the state information for the unsupported protocol is not preserved and must be refreshed using the normal recovery mechanisms inherent in the protocol.



NOTE: On routers that have logical systems configured on them, only the master logical system supports nonstop active routing.

Nonstop Active Routing BFD Support

Nonstop active routing supports the Bidirectional Forwarding Detection (BFD) protocol, which uses the topology discovered by routing protocols to monitor neighbors. The BFD protocol is a simple hello mechanism that detects failures in a network. Because BFD is streamlined to be efficient at fast liveness detection, when it is used in conjunction with routing protocols, routing recovery times are improved. With nonstop active routing enabled, BFD session states are not restarted when a Routing Engine switchover occurs.



NOTE: BFD session states are saved only for clients using aggregate or static routes or for BGP, IS-IS, OSPF/OSPFv3, or PIM.

When a BFD session is distributed to the Packet Forwarding Engine, BFD packets continue to be sent during a Routing Engine switchover. If nondistributed BFD sessions are to be kept alive during a switchover, you must ensure that the session failure detection time is greater than the Routing Engine switchover time. The following BFD sessions are not distributed to the Packet Forwarding Engine: multihop sessions, tunnel-encapsulated sessions, and sessions over integrated routing and bridging (IRB) interfaces.



NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping. The minimum-interval configuration statement is a BFD liveness detection parameter.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions, and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with nonstop active routing configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

Nonstop Active Routing BGP Support

Nonstop active routing BGP support is subject to the following conditions:

- You must include the **path-selection external-router-ID** statement at the **[edit protocols bgp]** hierarchy level to ensure consistent path selection between the master and backup Routing Engines during and after the nonstop active routing switchover.
- If the BGP peer in the master Routing Engine has negotiated address-family capabilities that are not supported for nonstop active routing, then the corresponding BGP neighbor state on the backup Routing Engine shows as idle. On switchover, the BGP session is reestablished from the new master Routing Engine.

Only the following address families are supported for nonstop active routing.



NOTE: Address families are supported only on the main instance of BGP. Only unicast is supported on VRF instances.

- inet unicast
- inet labeled-unicast
- inet multicast
- inet6 labeled-unicast
- inet6 multicast
- inet6 unicast

- route-target
 - l2vpn signaling
 - inet6-vpn unicast
 - inet-vpn unicast
 - inet-mdt
 - iso-vpn
- BGP route dampening does not work on the backup Routing Engine when nonstop active routing is enabled.

Nonstop Active Routing Layer 2 Circuit and VPLS Support

Nonstop active routing supports Layer 2 circuit and VPLS on both LDP-based and RSVP-TE-based networks. Nonstop active routing support enables the backup Routing Engine to track the label advertised by Layer 2 circuit and VPLS on the primary Routing Engine, and to use the same label after the Routing Engine switchover.

in Junos OS Release 9.6 and later, nonstop active routing support is extended to the Layer 2 circuit and LDP-based VPLS pseudowire redundant configurations.

Nonstop Active Routing PIM Support

Nonstop active routing supports Protocol Independent Multicast (PIM) with stateful replication on backup Routing Engines. State information replicated on the backup Routing Engine includes information about neighbor relationships, join and prune events, rendezvous point (RP) sets, synchronization between routes and next hops, multicast session states, and the forwarding state between the two Routing Engines.



NOTE: Nonstop active routing for PIM is supported for IPv4 on Junos OS Release 9.3 and later, and for IPv6 on Junos OS Release 10.4 and later. Starting with Release 11.1, Junos OS also supports nonstop active routing for PIM on devices that have both IPv4 and IPv6 configured on them.

To configure nonstop active routing for PIM, include the same statements in the configuration as for other protocols: the **nonstop-routing** statement at the **[edit routing-options]** hierarchy level and the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level. To trace PIM nonstop active routing events, include the **flag nsr-synchronization** statement at the **[edit protocols pim traceoptions]** hierarchy level.



NOTE: The **clear pim join**, **clear pim register**, and **clear pim statistics** operational mode commands are not supported on the backup Routing Engine when nonstop active routing is enabled.

Nonstop active routing support varies for different PIM features. The features fall into the following three categories: supported features, unsupported features, and incompatible features.

Supported features:

- Auto-RP



NOTE: Nonstop active routing PIM support on IPv6 does not support auto-RP because IPv6 does not support auto-RP.

- Bootstrap router (BSR)
- Static RPs
- Embedded RP on non-RP IPv6 routers
- Local RP



NOTE: RP set information synchronization is supported for local RP and BSR (on IPv4 and IPv6), autoRP (on IPv4), and embedded RP (on IPv6).

- BFD
- Dense mode
- Sparse mode
- Source-specific multicast (SSM)
- Draft Rosen multicast VPNs (MVPNs)
- Anycast RP (anycast RP set information synchronization and anycast RP register state synchronization on IPv4 and IPv6 configurations)
- Flow maps
- Unified ISSU
- Policy features such as neighbor policy, bootstrap router export and import policies, scope policy, flow maps, and reverse path forwarding (RPF) check policies
- Upstream assert synchronization
- PIM join load balancing

Starting with Release 12.2, Junos OS extends the nonstop active routing PIM support to draft Rosen MVPNs. Nonstop active routing PIM support for draft Rosen MVPNs enables nonstop active routing-enabled devices to preserve draft Rosen MPVN-related information—such as default and data multicast distribution tree (MDT) states—across switchovers. In releases earlier than 12.2, nonstop active routing PIM configuration was incompatible with draft Rosen MPVN configuration.

The backup Routing Engine sets up the default MDT based on the configuration and the information it receives from the master Routing Engine, and keeps updating the default MDT state information.

However, for data MDTs, the backup Routing Engine relies on the master Routing Engine to provide updates when data MDTs are created, updated, or deleted. The backup Routing

Engine neither monitors data MDT flow rates nor triggers a data MDT switchover based on variations in flow rates. Similarly, the backup Routing Engine does not maintain the data MDT delay timer or timeout timer. It does not send MDT join TLV packets for the data MDTs until it takes over as the master Routing Engine. After the switchover, the new master Routing Engine starts sending MDT join TLV packets for each data MDT, and also resets the data MDT timers. Note that the expiration time for the timers might vary from the original values on the previous master Routing Engine.

Starting with Release 12.3, Junos OS extends the Protocol Independent Multicast (PIM) nonstop active routing support to IGMP-only interfaces.

In Junos OS releases earlier than 12.3, the PIM joins created on IGMP-only interfaces were not replicated on the backup Routing Engine. Thus, the corresponding multicast routes were marked as pruned (meaning discarded) on the backup Routing Engine. Because of this limitation, after a switchover, the new master Routing Engine had to wait for the IGMP module to come up and start receiving reports to create PIM joins and to install multicast routes. This caused traffic loss until the multicast joins and routes were reinstated.

However, in Junos OS Release 12.3 and later, the multicast joins on the IGMP-only interfaces are mapped to PIM states, and these states are replicated on the backup Routing Engine. If the corresponding PIM states are available on the backup, the multicast routes are marked as forwarding on the backup Routing Engine. This enables uninterrupted traffic flow after a switchover. This enhancement covers IGMPv2, IGMPv3, MLDv1, and MLDv2 reports and leaves.

Unsupported features: You can configure the following PIM features on a router along with nonstop active routing, but they function as if nonstop active routing is not enabled. In other words, during Routing Engine switchover and other outages, their state information is not preserved, and traffic loss is to be expected.

- Internet Group Management Protocol (IGMP) exclude mode
- IGMP snooping

Incompatible features: Nonstop active routing does not support the following features, and you cannot configure them on a router enabled for PIM nonstop active routing. The commit operation fails if the configuration includes both nonstop active routing and one or more of these features:

- Next-generation MVPNs with PIM provider tunnels

Junos OS provides a configuration statement that disables nonstop active routing for PIM only, so that you can activate incompatible PIM features and continue to use nonstop active routing for the other protocols on the router. Before activating an incompatible PIM feature, include the **nonstop-routing disable** statement at the **[edit protocols pim]** hierarchy level. Note that in this case, nonstop active routing is disabled for all PIM features, not just incompatible features.

Nonstop Active Routing MSDP Support

Starting with Release 12.1, Junos OS extends nonstop active routing support to the Multicast Source Discovery Protocol (MSDP).

Nonstop active routing support for MSDP preserves the following MSDP-related information across the switchover:

- MSDP configuration and peer information
- MSDP peer socket information
- Source-active and related information

However, note that the following restrictions or limitations apply to nonstop active routing MSDP support:

- Because the backup Routing Engine learns the active source information by processing the source-active messages from the network, synchronizing of source active information between the master and backup Routing Engines might take up to 60 seconds. So, no planned switchover is allowed within 60 seconds of the initial replication of the sockets.
- Similarly, Junos OS does not support two planned switchovers within 240 seconds of each other.

Junos OS enables you to trace MSDP nonstop active routing events by including the **flag nsr-synchronization** statement at the **[edit protocols msdp traceoptions]** hierarchy level.

Nonstop Active Routing Support for RSVP-TE LSPs

Junos OS extends nonstop active routing support to label-switching routers (LSRs) and Layer 2 Circuits that are part of an RSVP-TE LSP. Nonstop active routing support on LSRs ensures that the master to backup Routing Engine switchover on an LSR remains transparent to the network neighbors and that the LSP information remains unaltered during and after the switchover.

You can use the **show rsvp version** command to view the nonstop active routing mode and state on an LSR. Similarly, you can use the **show mpls lsp** and **show rsvp session** commands on the standby Routing Engine to view the state recreated on the standby Routing Engine.

The Junos OS nonstop active routing feature is also supported on RSVP point-to-multipoint LSPs. Nonstop active routing support for RSVP point-to-multipoint egress and transit LSPs was added in Junos OS Release 11.4, and for ingress LSPs in Release 12.1. During the switchover, the LSP comes up on the backup Routing Engine that shares and synchronizes the state information with the master Routing Engine before and after the switchover. Nonstop active routing support for point-to-multipoint transit and egress LSPs ensures that the switchover remains transparent to the network neighbors, and preserves the LSP information across the switchover.

However, Junos OS nonstop active routing support for RSVP point-to-multipoint LSPs does not include support for dynamically created point-to-multipoint LSPs, such as VPLS and next-generation MVPNs.

The **show rsvp session detail** command enables you to check the point-to-multipoint LSP remerge state information (**P2MP LSP re-merge**; possible values are **head**, **member**, and **none**).

However, Junos OS does not support nonstop active routing for the following features:

- Generalized Multiprotocol Label Switching (GMPLS) and LSP hierarchy
- Interdomain or loose-hop expansion LSPs
- BFD liveness detection

Nonstop active routing support for RSVP-TE LSPs is subject to the following limitations and restrictions:

- Detour LSPs are not maintained across a switchover and so, detour LSPs might fail to come back online after the switchover.
- Control plane statistics corresponding to the **show rsvp statistics** and **show rsvp interface detail | extensive** commands are not maintained across Routing Engine switchovers.
- Statistics from the backup Routing Engine are not reported for **show mpls lsp statistics** and **monitor mpls label-switched-path** commands. However, if a switchover occurs, the backup Routing Engine, after taking over as the master, starts reporting statistics. Note that the **clear statistics** command issued on the old master Routing Engine does not have any effect on the new master Routing Engine, which reports statistics, including any uncleared statistics.
- State timeouts might take additional time during nonstop active routing switchover. For example, if a switchover occurs after a neighbor has missed sending two hello messages to the master, the new master Routing Engine waits for another three hello periods before timing out the neighbor.
- On the RSVP ingress router, if you configure auto-bandwidth functionality, the bandwidth adjustment timers are set in the new master after the switchover. This causes a one-time increase in the length of time required for the bandwidth adjustment after the switchover occurs.
- RSVP ingress LSPs that have BFD liveness detection enabled on them do not come up on the backup Routing Engine during the switchover. Such BFD-enabled LSPs have to be reestablished after the switchover.
- Backup LSPs —LSPs that are established between the point of local repair (PLR) and the merge point after a node or link failure—are not preserved during a Routing Engine switchover.
- When nonstop active routing is enabled, graceful restart is not supported. However, graceful restart helper mode is supported.

- Related Documentation**
- [Nonstop Active Routing Concepts on page 2072](#)
 - [Configuring Nonstop Active Routing on page 2183](#)

Graceful Restart

- [Graceful Restart Concepts on page 2085](#)
- [Graceful Restart System Requirements on page 2086](#)
- [Aggregate and Static Routes on page 2086](#)
- [Graceful Restart and Routing Protocols on page 2087](#)
- [Graceful Restart and MPLS-Related Protocols on page 2089](#)
- [Graceful Restart and Layer 2 and Layer 3 VPNs on page 2091](#)
- [Graceful Restart on Logical Systems on page 2092](#)

Graceful Restart Concepts

With routing protocols, any service interruption requires that an affected router recalculate adjacencies with neighboring routers, restore routing table entries, and update other protocol-specific information. An unprotected restart of a router can result in forwarding delays, route flapping, wait times stemming from protocol reconvergence, and even dropped packets. The main benefits of graceful restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful restart enables a router to pass through intermediate convergence states that are hidden from the rest of the network.

Three main types of graceful restart are available on Juniper Networks routing platforms:

- Graceful restart for aggregate and static routes and for routing protocols—Provides protection for aggregate and static routes and for Border Gateway Protocol (BGP), End System-to-Intermediate System (ES-IS), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), Routing Information Protocol (RIP), next-generation RIP (RIPng), and Protocol Independent Multicast (PIM) sparse mode routing protocols.
- Graceful restart for MPLS-related protocols—Provides protection for Label Distribution Protocol (LDP), Resource Reservation Protocol (RSVP), circuit cross-connect (CCC), and translational cross-connect (TCC).
- Graceful restart for virtual private networks (VPNs)—Provides protection for Layer 2 and Layer 3 VPNs.

Graceful restart works similarly for routing protocols and MPLS protocols and combines components of these protocol types to enable graceful restart in VPNs. The main benefits of graceful restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful restart thus enables a router to pass through intermediate convergence states that are hidden from the rest of the network.

Most graceful restart implementations define two types of routers—the restarting router and the helper router. The restarting router requires rapid restoration of forwarding state information so it can resume the forwarding of network traffic. The helper router assists

the restarting router in this process. Graceful restart configuration statements typically affect either the restarting router or the helper router.

**Related
Documentation**

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Graceful Restart System Requirements on page 2086](#)
- [Aggregate and Static Routes on page 2086](#)
- [Graceful Restart and Routing Protocols on page 2087](#)
- [Graceful Restart and MPLS-Related Protocols on page 2089](#)
- [Graceful Restart and Layer 2 and Layer 3 VPNs on page 2091](#)
- [Graceful Restart on Logical Systems on page 2092](#)
- [Example: Configuring Graceful Restart on page 2139](#)
- [Configuring Graceful Restart for QFabric Systems](#)

Graceful Restart System Requirements

Graceful restart is supported on all routing platforms. To implement graceful restart for particular features, your system must meet these minimum requirements:

- Junos OS Release 5.3 or later for aggregate route, BGP, IS-IS, OSPF, RIP, RIPng, or static route graceful restart.
- Junos OS Release 5.5 or later for RSVP on egress provider edge (PE) routers.
- Junos OS Release 5.5 or later for LDP graceful restart.
- Junos OS Release 5.6 or later for the CCC, TCC, Layer 2 VPN, or Layer 3 VPN implementations of graceful restart.
- Junos OS Release 6.1 or later for RSVP graceful restart on ingress PE routers.
- Junos OS Release 6.4 or later for PIM sparse mode graceful restart.
- Junos OS Release 7.4 or later for ES-IS graceful restart (J Series Services Routers).
- Junos OS Release 8.5 or later for BFD session (helper mode only)—If a node is undergoing a graceful restart and its BFD sessions are distributed to the Packet Forwarding Engine, the peer node can help the peer with the graceful restart.
- Junos OS Release 9.2 or later for BGP to support helper mode without requiring that graceful restart be configured.

**Related
Documentation**

- [Graceful Restart Concepts on page 2085](#)

Aggregate and Static Routes

When you include the **graceful-restart** statement at the **[edit routing-options]** hierarchy level, any static routes or aggregated routes that have been configured are protected. Because no helper router assists in the restart, these routes are retained in the forwarding table while the router restarts (rather than being discarded or refreshed).

Related Documentation

- [Graceful Restart Concepts on page 2085](#)
- [Graceful Restart System Requirements on page 2086](#)
- [Enabling Graceful Restart on page 2127](#)
- [Verifying Graceful Restart Operation on page 2271](#)
- [Example: Configuring Graceful Restart on page 2139](#)

Graceful Restart and Routing Protocols

This section covers the following topics:

- [BGP on page 2087](#)
- [ES-IS on page 2087](#)
- [IS-IS on page 2087](#)
- [OSPF and OSPFv3 on page 2088](#)
- [PIM Sparse Mode on page 2089](#)
- [RIP and RIPng on page 2089](#)

BGP

When a router enabled for BGP graceful restart restarts, it retains BGP peer routes in its forwarding table and marks them as stale. However, it continues to forward traffic to other peers (or receiving peers) during the restart. To reestablish sessions, the restarting router sets the “restart state” bit in the BGP OPEN message and sends it to all participating peers. The receiving peers reply to the restarting router with messages containing end-of-routing-table markers. When the restarting router or switch receives all replies from the receiving peers, the restarting router performs route selection, the forwarding table is updated, and the routes previously marked as stale are discarded. At this point, all BGP sessions are reestablished and the restarting peer can receive and process BGP messages as usual.

While the restarting router does its processing, the receiving peers also temporarily retain routing information. When a receiving peer detects a TCP transport reset, it retains the routes received and marks the routes as stale. After the session is reestablished with the restarting router or switch, the stale routes are replaced with updated route information.

ES-IS

When graceful restart for ES-IS is enabled, the routes to end systems or intermediate systems are not removed from the forwarding table. The adjacencies are reestablished after restart is complete.



NOTE: ES-IS is supported only on the J Series Services Router.

IS-IS

Normally, IS-IS routers move neighbor adjacencies to the down state when changes occur. However, a router enabled for IS-IS graceful restart sends out Hello messages

with the Restart Request (RR) bit set in a restart type length value (TLV) message. This indicates to neighboring routers that a graceful restart is in progress and to leave the IS-IS adjacency intact. The neighboring routers must interpret and implement restart signaling themselves. Besides maintaining the adjacency, the neighbors send complete sequence number PDUs (CSNPs) to the restarting router and flood their entire database.

The restarting router never floods any of its own link-state PDUs (LSPs), including pseudonode LSPs, to IS-IS neighbors while undergoing graceful restart. This enables neighbors to reestablish their adjacencies without transitioning to the down state and enables the restarting router to reinitiate a smooth database synchronization.

OSPF and OSPFv3

When a router enabled for OSPF graceful restart restarts, it retains routes learned before the restart in its forwarding table. The router does not allow new OSPF link-state advertisements (LSAs) to update the routing table. This router continues to forward traffic to other OSPF neighbors (or helper routers), and sends only a limited number of LSAs during the restart period. To reestablish OSPF adjacencies with neighbors, the restarting router must send a grace LSA to all neighbors. In response, the helper routers enter helper mode and send an acknowledgement back to the restarting router. If there are no topology changes, the helper routers continue to advertise LSAs as if the restarting router had remained in continuous OSPF operation.

When the restarting router receives replies from all the helper routers, the restarting router selects routes, updates the forwarding table, and discards the old routes. At this point, full OSPF adjacencies are reestablished and the restarting router receives and processes OSPF LSAs as usual. When the helper routers no longer receive grace LSAs from the restarting router or the topology of the network changes, the helper routers also resume normal operation.



NOTE: For more information about the standard helper mode implementation, see RFC 3623, *Graceful OSPF Restart*.

Starting with Release 11.3, Junos OS supports the restart signaling-based helper mode for OSPF graceful restart configurations. The helper modes, both standard and restart signaling-based, are enabled by default. In restart signaling-based helper mode implementations, the restarting router relays the restart status to its neighbors only after the restart is complete. When the restart is complete, the restarting router sends hello messages to its helper routers with the **restart signal (RS)** bit set in the hello packet header. When a helper router receives a hello packet with the **RS** bit set in the header, the helper router returns a hello message to the restarting router. The reply hello message from the helper router contains the **ResyncState** flag and the **ResyncTimeout** timer that enable the restarting router to keep track of the helper routers that are syncing up with it. When all helpers complete the synchronization, the restarting router exits the restart mode.

**NOTE:**

For more information about restart signaling-based graceful restart helper mode implementation, see RFC 4811, *OSPF Out-of-Band Link State Database (LSDB) Resynchronization*, RFC 4812, *OSPF Restart Signaling*, and RFC 4813, *OSPF Link-Local Signaling*.

Restart signaling-based graceful restart helper mode is not supported for OSPFv3 configurations.

PIM Sparse Mode

PIM sparse mode uses a mechanism called a *generation identifier* to indicate the need for graceful restart. Generation identifiers are included by default in PIM hello messages. An initial generation identifier is created by each PIM neighbor to establish device capabilities. When one of the PIM neighbors restarts, it sends a new generation identifier to its neighbors. All neighbors that support graceful restart and are connected by point-to-point links assist by sending multicast updates to the restarting neighbor.

The restart phase completes when either the PIM state becomes stable or when the restart interval timer expires. If the neighbors do not support graceful restart or connect to each other using multipoint interfaces, the restarting router uses the restart interval timer to define the restart period.

RIP and RIPng

When a router enabled for RIP graceful restart restarts, routes that have been configured are protected. Because no helper router assists in the restart, these routes are retained in the forwarding table while the router restarts (rather than being discarded or refreshed).

Related Documentation

- [Graceful Restart Concepts on page 2085](#)
- [Graceful Restart System Requirements on page 2086](#)
- [Configuring Routing Protocols Graceful Restart on page 2128](#)
- [Verifying Graceful Restart Operation on page 2271](#)
- [Example: Configuring Graceful Restart on page 2139](#)

Graceful Restart and MPLS-Related Protocols

This section contains the following topics:

- [LDP on page 2089](#)
- [RSVP on page 2090](#)
- [CCC and TCC on page 2090](#)

LDP

LDP graceful restart enables a router whose LDP control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. It also

enables a router on which helper mode is enabled to assist a neighboring router that is attempting to restart LDP.

During session initialization, a router advertises its ability to perform LDP graceful restart or to take advantage of a neighbor performing LDP graceful restart by sending the graceful restart TLV. This TLV contains two fields relevant to LDP graceful restart: the reconnect time and the recovery time. The values of the reconnect and recovery times indicate the graceful restart capabilities supported by the router.

The reconnect time is configured in Junos OS as 60 seconds and is not user-configurable. The reconnect time is how long the helper router waits for the restarting router to establish a connection. If the connection is not established within the reconnect interval, graceful restart for the LDP session is terminated. The maximum reconnect time is 120 seconds and is not user-configurable. The maximum reconnect time is the maximum value that a helper router accepts from its restarting neighbor.

When a router discovers that a neighboring router is restarting, it waits until the end of the recovery time before attempting to reconnect. The recovery time is the length of time a router waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This time period is also typically the length of time that a neighboring router maintains its information about the restarting router, so it can continue to forward traffic.

You can configure LDP graceful restart both in the master instance for the LDP protocol and for a specific routing instance. You can disable graceful restart at the global level for all protocols, at the protocol level for LDP only, and for a specific routing instance only.

RSVP

RSVP graceful restart enables a router undergoing a restart to inform its adjacent neighbors of its condition. The restarting router requests a grace period from the neighbor or peer, which can then cooperate with the restarting router. The restarting router can still forward MPLS traffic during the restart period; convergence in the network is not disrupted. The restart is not visible to the rest of the network, and the restarting router is not removed from the network topology. RSVP graceful restart can be enabled on both transit routers and ingress routers. It is available for both point-to-point LSPs and point-to-multipoint LSPs.

CCC and TCC

CCC and TCC graceful restart enables Layer 2 connections between customer edge (CE) routers to restart gracefully. These Layer 2 connections are configured with the **remote-interface-switch** or **lsp-switch** statements. Because these CCC and TCC connections have an implicit dependency on RSVP LSPs, graceful restart for CCC and TCC uses the RSVP graceful restart capabilities.

RSVP graceful restart must be enabled on the provider edge (PE) routers and provider (P) routers to enable graceful restart for CCC and TCC. Also, because RSVP is used as the signaling protocol for signaling label information, the neighboring router must use helper mode to assist with the RSVP restart procedures.

- Related Documentation**
- [Graceful Restart Concepts on page 2085](#)
 - [Graceful Restart System Requirements on page 2086](#)
 - [Configuring Graceful Restart for MPLS-Related Protocols on page 2134](#)
 - [Example: Configuring Graceful Restart on page 2139](#)

Graceful Restart and Layer 2 and Layer 3 VPNs

VPN graceful restart uses three types of restart functionality:

1. BGP graceful restart functionality is used on all PE-to-PE BGP sessions. This affects sessions carrying any service signaling data for network layer reachability information (NLRI), for example, an IPv4 VPN or Layer 2 VPN NLRI.
2. OSPF, IS-IS, LDP, or RSVP graceful restart functionality is used in all core routers. Routes added by these protocols are used to resolve Layer 2 and Layer 3 VPN NLRI.
3. Protocol restart functionality is used for any Layer 3 protocol (RIP, OSPF, LDP, and so on) used between the PE and customer edge (CE) routers. This does not apply to Layer 2 VPNs because Layer 2 protocols used between the CE and PE routers do not have graceful restart capabilities.

Before VPN graceful restart can work properly, all of the components must restart gracefully. In other words, the routers must preserve their forwarding states and request neighbors to continue forwarding to the router in case of a restart. If all of the conditions are satisfied, VPN graceful restart imposes the following rules on a restarting router:

- The router must wait to receive all BGP NLRI information from other PE routers before advertising routes to the CE routers.
- The router must wait for all protocols in all routing instances to converge (or complete the restart process) before it sends CE router information to other PE routers. In other words, the router must wait for all instance information (whether derived from local configuration or advertisements received from a remote peer) to be processed before it sends this information to other PE routers.
- The router must preserve all forwarding state in the `instance.mpls.0` tables until the new labels and transit routes are allocated and announced to other PE routers (and CE routers in a carrier-of-carriers scenario).

If any condition is not met, VPN graceful restart does not succeed in providing uninterrupted forwarding between CE routers across the VPN infrastructure.

- Related Documentation**
- [Graceful Restart Concepts on page 2085](#)
 - [Graceful Restart System Requirements on page 2086](#)
 - [Configuring Logical System Graceful Restart on page 2137](#)
 - [Verifying Graceful Restart Operation on page 2271](#)
 - [Example: Configuring Graceful Restart on page 2139](#)

Graceful Restart on Logical Systems

Graceful restart for a logical system functions much as graceful restart does in the main router. The only difference is the location of the **graceful-restart** statement:

- For a logical system, include the **graceful-restart** statement at the **[edit logical-systems *logical-system-name* routing-options]** hierarchy level.
- For a routing instance inside a logical system, include the **graceful-restart** statement at both the **[edit logical-systems *logical-system-name* routing-options]** and **[edit logical-systems *logical-system-name* routing-instances *instance-name* routing-options]** hierarchy levels.

Related Documentation

- [Graceful Restart Concepts on page 2085](#)
- [Graceful Restart System Requirements on page 2086](#)
- [Configuring Logical System Graceful Restart on page 2137](#)
- [Verifying Graceful Restart Operation on page 2271](#)
- [Example: Configuring Graceful Restart on page 2139](#)

Unified ISSU

- [Upgrading Routers Using ISSU on page 2092](#)
- [Unified ISSU Concepts on page 2092](#)
- [Unified ISSU System Requirements on page 2097](#)

Upgrading Routers Using ISSU

Unified in-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. ISSU is only supported by dual Routing Engine platforms. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled.

For additional information about using ISSU, see the [Junos OS High Availability Guide](#).

Unified ISSU Concepts

A unified in-service software upgrade (unified ISSU) enables you to upgrade between two different Junos OS Releases with no disruption on the control plane and with minimal disruption of traffic. Unified ISSU is only supported on dual Routing Engine platforms. In addition, the graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled.

A unified ISSU provides the following benefits:

- Eliminates network downtime during software image upgrades
- Reduces operating costs, while delivering higher service levels
- Allows fast implementation of new features



NOTE: The master Routing Engine and backup Routing Engine must be running the same software version before you can perform a unified ISSU.

You cannot take any PICs online or offline during a unified ISSU.



NOTE: You can verify the unified ISSU-compatibility of the software, hardware, and the configuration on a device by issuing the **request system software validate in-service-upgrade** command. This command runs the validation checks, and shows whether the operating system, device components, and configurations are ISSU compatible or not. For more information about the **request system software validate in-service-upgrade** command, see *Junos OS Operational Mode Commands*.



NOTE: Unicast RPF-related statistics are not saved across a unified ISSU, and the unicast RPF counters are reset to zero during a unified ISSU.

To perform a unified ISSU, complete the following steps:

1. Enable graceful Routing Engine switchover and nonstop active routing. Verify that the Routing Engines and protocols are synchronized.
2. Download the new software package from the Juniper Networks Support website and then copy the package to the router.
3. Issue the **request system software in-service-upgrade** command on the master Routing Engine.

A Junos OS Release package comprises three distinct systems:

- Juniper Networks Operating System, which provides system control and all the features and functions of the Juniper Networks router that executes in the Routing Engines
- Juniper Networks Packet Forwarding Engine, which supports the high-performance traffic forwarding and packet handling capabilities
- Interface control

After the **request system software in-service-upgrade** command is issued, the following process occurs.




NOTE: In the illustrations, a solid line indicates the high-speed internal link between a Routing Engine and a Packet Forwarding Engine. A dotted line indicates the chassis process (chassisd), another method of communication between a Routing Engine and a Packet Forwarding Engine. RE0m and RE1s indicate master and backup (or standby) Routing Engines, respectively.



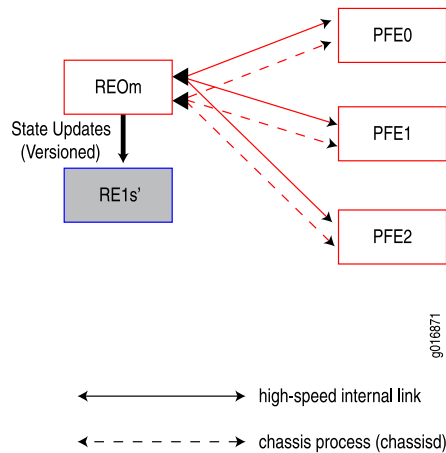
NOTE: The following process pertains to all supported routing platforms except the TX Matrix router. For information about the unified ISSU process on the TX Matrix router, see *Unified ISSU Process on the TX Matrix Router*. On M320 and T320 routers and on T640 and T1600 routers, the Packet Forwarding Engine resides on an FPC. However, on an M120 router, the Forwarding Engine Board (FEB) replaces the functions of a Packet Forwarding Engine. In the illustrations and steps, when considering an M120 router, you can regard the PFE as an FPC. As an additional step on an M120 router, after the FPCs and PICs have been upgraded, the FEBs are upgraded.

1. The master Routing Engine validates the router configuration to ensure that it can be committed using the new software version. Checks are made for disk space available for the `/var` file system on both Routing Engines, unsupported configurations, and for unsupported Physical Interface Cards (PICs). If there is not sufficient disk space available on either of the Routing Engines, the unified ISSU process fails and returns an error message saying that the Routing Engine does not have enough disk space available. However, unsupported PICs do not prevent a unified ISSU. The software issues a warning to indicate that these PICs will restart during the upgrade. Similarly, an unsupported protocol configuration does not prevent a unified ISSU. The software issues a warning that packet loss may occur for the protocol during the upgrade.

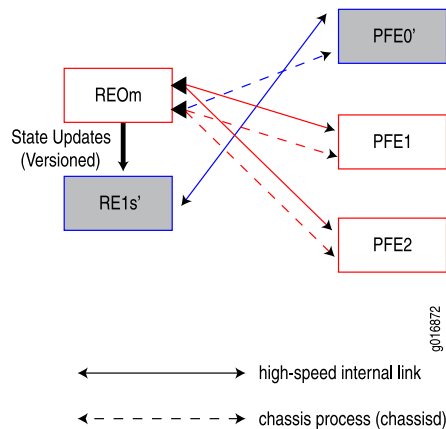


2. When the validation succeeds, the kernel state synchronization daemon (ksyncd) synchronizes the kernel on the backup Routing Engine with the master Routing Engine.
3. The backup Routing Engine is upgraded with the new software image. Before being upgraded, the backup Routing Engine gets the configuration file from the master Routing Engine and validates the configuration to ensure that it can be committed using the new software version. After being upgraded, it is resynchronized with the

master Routing Engine. In the illustration, an apostrophe (') indicates the device is running the new version of software.




4. The chassis process (chassisd) on the master Routing Engine prepares other software processes for the unified ISSU. When all the processes are ready, chassisd sends an ISSU_PREPARE message to the Flexible PIC Concentrators (FPCs) installed in the router.
5. The Packet Forwarding Engine on each FPC saves its state and downloads the new software image from the backup Routing Engine. Next, each Packet Forwarding Engine sends an ISSU_READY message to the chassis process (chassisd).




6. After receiving an ISSU_READY message from a Packet Forwarding Engine, the chassis process (chassisd) sends an ISSU_REBOOT message to the FPC on which the Packet Forwarding Engine resides. The FPC reboots with the new software image. After the FPC is rebooted, the Packet Forwarding Engine restores the FPC state and a high-speed internal link is established with the backup Routing Engine running the new software. The chassis process (chassisd) is also reestablished with the master Routing Engine.
7. After all Packet Forwarding Engines have sent a READY message using the chassis process (chassisd) on the master Routing Engine, other software processes are

prepared for a Routing Engine switchover. The system is ready for a switchover at this point.




NOTE: In the case of an M120 router, the FEBs are upgraded at this point. When all FEBs have been upgraded, the system is ready for a switchover.

- The Routing Engine switchover occurs and the backup Routing Engine becomes the new master Routing Engine.



9. The new backup Routing Engine is now upgraded to the new software image. (This step is skipped if the **no-old-master-upgrade** option is specified.)



10. When the backup Routing Engine has been successfully upgraded, the unified ISSU is complete.

Related Documentation

- [Unified ISSU Process on the TX Matrix Router](#)
- [Unified ISSU System Requirements on page 2097](#)
- [Best Practices on page 2197](#)
- [Before You Begin on page 2197](#)
- [Performing a Unified ISSU on page 2200](#)

Unified ISSU System Requirements

This section contains the following topics:

- [Unified ISSU Junos OS Release Support on page 2097](#)
- [Unified ISSU Platform Support on page 2098](#)
- [Unified ISSU Protocol Support on page 2098](#)
- [Unified ISSU Feature Support on page 2100](#)
- [Unified ISSU PIC Support on page 2100](#)
- [Unified ISSU Support on MX Series 3D Universal Edge Routers on page 2108](#)

Unified ISSU Junos OS Release Support

In order to perform a unified ISSU, your device must be running a Junos OS Release that supports unified ISSU for the specific platform. See [“Unified ISSU Platform Support” on page 2098](#). You can use unified ISSU to upgrade from an ISSU-capable software release to a newer software release. However, note that:

- The unified ISSU process is aborted if the Junos OS version specified for installation is a version earlier than the one currently running on the device. To downgrade from an ISSU-capable release to a previous software release (ISSU-capable or not), use the

request system add command. Unlike an upgrade using the unified ISSU process, a downgrade using the **request system add** command can cause network disruptions and loss of data. For more information about the use of the **request system add** command, see the *Installation and Upgrade Guide*.

- The unified ISSU process is aborted if the specified upgrade has conflicts with the current configuration, components supported, and so forth.
- Unified ISSU does not support extension application packages developed using the Junos SDK.

Unified ISSU Platform Support

Table 154 on page 2098 lists the platforms on which a unified ISSU is supported.

Table 154: Unified ISSU Platform Support

Platform	Junos OS Release
EX9200 switch	12.3R3 or later
M120 router	9.2 or later
M320 router	9.0 or later
M10i router with Enhanced Compact Forwarding Engine Board (CFEB-E)	9.5 or later
MX Series devices	9.3 or later
NOTE: Unified ISSU for MX Series routers does not support IEEE 802.1ag OAM and IEEE 802.3ah protocols.	11.2 or later on MX Series 3D Universal Edge Routers (with Trio Modular Port Concentrator/Modular Interface Card (MPC/MIC) interfaces).
T320 router	9.0 or later
T640 router	9.0 or later
T1600 router	9.1 or later
TX Matrix router	9.3 or later

Unified ISSU Protocol Support

Unified ISSU is dependent on nonstop active routing. Table 155 on page 2098 lists the protocols that are supported during a unified ISSU.

Table 155: Unified ISSU Protocol Support

Protocol	Junos OS Release
BGP	9.0 or later

Table 155: Unified ISSU Protocol Support (*continued*)

Protocol	Junos OS Release
DHCP access model (subscriber access)	11.2 or later
IS-IS	9.0 or later
LDP	9.0 or later
LDP-based virtual private LAN service (VPLS)	9.3 or later
Layer 2 circuits	9.2 or later
Layer 3 VPNs using LDP	9.2 or later
Link Aggregation Control Protocol (LACP) on MX Series routers	9.4 or later
OSPF/OSPFv3	9.0 or later
PPPoE access model (subscriber access)	11.4 or later
Protocol Independent Multicast (PIM)	9.3 or later
Routing Information Protocol (RIP)	9.1 or later

Unified ISSU Support for the Layer 2 Control Protocol Process

Unified ISSU supports the Layer 2 Control Protocol process (l2cpd) on MX Series 3D Universal Edge Routers. In a Layer 2 bridge environment, spanning tree protocols share information about port roles, bridge IDs, and root path costs between bridges using special data frames called Bridge Protocol Data Units (BPDUs). The transmission of BPDUs is controlled by the l2cpd process. Transmission of hello BPDUs is important in maintaining adjacencies on the peer systems.

The transmission of periodic packets on behalf of the l2cpd process is carried out by periodic packet management (PPM), which, by default, is configured to run on the Packet Forwarding Engine. The ppmd process on the Packet Forwarding Engine ensures that the BPDUs are transmitted even when the l2cpd process control plane is unavailable, and keeps the remote adjacencies alive during unified ISSU. However, if you want the distributed PPM (ppmd) process to run on the Routing Engine instead of the Packet Forwarding Engine, you can disable the ppmd process on the Packet Forwarding Engine, by including the **no-delegate-processing** statement at the [edit routing-options ppm] hierarchy level.



NOTE: The `delegate-processing` statement at the `[edit routing-options ppm]` hierarchy level, which was used to enable the `ppmd` process on the Packet Forwarding Engine in Junos OS Release 9.3 and earlier, has been deprecated as the `ppmd` process is enabled on the Packet Forwarding Engine by default in Junos OS Release 9.4 and later.

Unified ISSU enhancements and nonstop active bridging support for the `l2cpd` process ensure that the new master Routing Engine is able to take control during unified ISSU without any disruptions in the control plane and minimize the disruptions in the Layer 2 data plane during unified ISSU.

Unified ISSU Feature Support

Unified ISSU supports most Junos OS features in Junos OS Release 9.0. However, the following constraints apply:

- Link Aggregation Control Protocol (LACP)—Link changes are not processed until after the unified ISSU is complete.
- Automatic Protection Switching (APS)—Network changes are not processed until after the unified ISSU is complete.
- Ethernet Operation, Administration, and Management (OAM) as defined by IEEE 802.3ah and by IEEE 802.1ag—When a Routing Engine switchover occurs, the OAM hello times out, triggering protocol convergence.
- Ethernet circuit cross-connect (CCC) encapsulation—Circuit changes are not processed until after the unified ISSU is complete.
- Logical systems—On devices that have logical systems configured on them, only the master logical system supports unified ISSU.

Unified ISSU PIC Support

The following sections list the Physical Interface Cards (PICs) that are supported during a unified ISSU.

- [PIC Considerations on page 2101](#)
- [SONET/SDH PICs on page 2101](#)
- [Fast Ethernet and Gigabit Ethernet PICs on page 2103](#)
- [Channelized PICs on page 2104](#)
- [Tunnel Services PICs on page 2105](#)
- [ATM PICs on page 2105](#)
- [Serial PICs on page 2106](#)
- [DS3, E1, E3, and T1 PICs on page 2106](#)
- [Enhanced IQ PICs on page 2107](#)
- [Enhanced IQ2 Ethernet Services Engine \(ESE\) PIC on page 2107](#)



NOTE: For information about FPC types, FPC/PIC compatibility, and the initial Junos OS Release in which an FPC supported a particular PIC, see the PIC guide for your platform.

PIC Considerations

Take the following PIC restrictions into consideration before performing a unified ISSU:

- **Unsupported PICs**—If a PIC is not supported by unified ISSU, at the beginning of the upgrade the software issues a warning that the PIC will be brought offline. After the PIC is brought offline and the ISSU is complete, the PIC is brought back online with the new firmware.
- **PIC combinations**—For some PICs, newer Junos OS services can require significant Internet Processor ASIC memory, and some configuration rules might limit certain combinations of PICs on particular platforms. With a unified ISSU:
 - If a PIC combination is not supported by the software version that the device is being upgraded from, the upgrade will be aborted.
 - If a PIC combination is not supported by the software version to which the device is being upgraded, the in-service software upgrade will abort, even if the PIC combination is supported by the software version from which the device is being upgraded.
- **Interface statistics**—Interface statistics might be incorrect because:
 - During bootup of the new microkernel on the Packet Forwarding Engine (PFE), host-bound traffic is not handled and might be dropped, causing packet loss.
 - During the hardware update of the Packet Forwarding Engine and its interfaces, traffic is halted and discarded. (The duration of the hardware update depends on the number and type of interfaces and on the device configuration.)
 - During a unified ISSU, periodic statistics collection is halted. If hardware counters saturate or wrap around, the software does not display accurate interface statistics.
- **CIR oversubscription**—If oversubscription of committed rate information (CIR) is configured on logical interfaces:
 - And the sum of the CIR exceeds the physical interface's bandwidth, after a unified in-service software upgrade is performed, each logical interface might not be given its original CIR.
 - And the sum of the delay buffer rate configured on logical interfaces exceeds the physical interface's bandwidth, after a unified in-service software upgrade is performed, each logical interface might not receive its original delay-buffer-rate calculation.

SONET/SDH PICs

[Table 156 on page 2102](#) lists the SONET/SDH PICs that are supported during a unified ISSU.

Table 156: Unified ISSU PIC Support: SONET/SDH

PIC Type	Number of Ports	Model Number	Device
OC3c/STM1	4-port	PB-4OC3-SON-MM—(EOL)	M120 M320, T320, T640, T1600
		PB-4OC3-SON-SMIR—(EOL)	
		PE-4OC3-SON-MM—(EOL)	M10i
		PE-4OC3-SON-SMIR—(EOL)	
	2-port	PE-2OC3-SON-MM—(EOL)	
		PE-2OC3-SON-SMIR—(EOL)	
OC3c/STM1 with SFP	2-port	PE-2OC3-SON-SFP	M10i
OC3c/STM1, SFP (Multi-Rate)	4 OC3 ports, 4 OC12 ports	PB-4OC3-4OC12-SON-SFP	M120 M320, MX Series, T320, T640, T1600
	4 OC3 ports, 1 OC12 port	PB-4OC3-1OC12-SON-SFP	
		PB-4OC3-1OC12-SON2-SFP	M10i
		PE-4OC3-1OC12-SON-SFP	
OC12c/STM4	1-port	PE-1OC12-SON-SFP	M10i
		PE-1OC12-SON-MM—(EOL)	
		PE-1OC12-SON-SMIR—(EOL)	
		PB-1OC12-SON-MM—(EOL)	
		PB-1OC12-SON-SMIR—(EOL)	M120, M320, T320, T640, T1600, TX Matrix
	4-port	PB-4OC12-SON-MM	
		PB-4OC12-SON-SMIR	
OC12c/STM4, SFP	1-port	PB-1OC12-SON-SFP	M120, M320, T320, T640, T1600, TX Matrix
OC48c/STM16, SFP	1-port	PB-1OC48-SON-SFP	M120, M320, MX Series, T320, T640, T1600, TX Matrix
		PB-1OC48-SON-B-SFP	
	4-port	PC-4OC48-SON-SFP	
OC192/STM64	1-port	PC-1OC192-SON-VSR	MX Series routers

Table 156: Unified ISSU PIC Support: SONET/SDH (*continued*)

PIC Type	Number of Ports	Model Number	Device
OC192/STM64, XFP	1-port	PC-1OC192-SON-LR	M320, T320, T640, T1600
		PC-1OC192-SON-SR2	
		PC-1OC192-VSR	
OC192/STM64, XFP	4-port	PD-4OC192-SON-XFP	M120, T640, T1600
	1-port	PC-1OC192-SON-XFP	MX Series routers
OC768/STM256	1-port	PD-1OC768-SON-SR	T640, T1600

Fast Ethernet and Gigabit Ethernet PICs

Table 157 on page 2103 lists the Fast Ethernet and Gigabit Ethernet PICs that are supported during a unified ISSU.



NOTE: Starting with Junos OS Release 9.2, new Ethernet IQ2 PIC features might cause the software to reboot the PIC when a unified ISSU is performed. For information about applicable new Ethernet IQ2 PIC features, refer to the release notes for the specific Junos OS Release.

Table 157: Unified ISSU PIC Support: Fast Ethernet and Gigabit Ethernet

PIC Type	Number of Ports	Model Number	device
Fast Ethernet	4	PB-4FE-TX	M120, M320, T320, T640, T1600, TX Matrix
		PE-4FE-TX	M10i
	8	PB-8FE-FX	M120, M320
		PE-8FE-FX	M10i
	12	PB-12FE-TX-MDI	M120, M320, T320
		PB-12FE-TX-MDIX	
		PE-12FE-TX-MDI	M10i
		PE-12FE-TX-MDIX	
	48	PB-48FE-TX-MDI	M120, M320, T320
		PB-48FE-TX-MDIX	

Table 157: Unified ISSU PIC Support: Fast Ethernet and Gigabit Ethernet (*continued*)

PIC Type	Number of Ports	Model Number	device
Gigabit Ethernet, RJ-45	40	EX9200-40T	EX9200
Gigabit Ethernet, SFP	1	PE-1GE-SFP	M10i
		PB-1GE-SFP	M120, M320, T320, T640, T1600, TX Matrix
	2	PB-2GE-SFP	
	4	PB-4GE-SFP	
	10	PC-10GE-SFP	
	40	EX9200-40F	EX9200
Gigabit Ethernet IQ, SFP	1	PE-1GE-SFP-QPP	M10i
		PB-1GE-SFP-QPP	M120, M320, T320, T640, T1600, TX Matrix
	2	PB-2GE-SFP-QPP	
Gigabit Ethernet IQ2, SFP	4	PB-4GE-TYPE1-SFP-IQ2	M120, M320, T320, T640, T1600, TX Matrix
	8	PB-8GE-TYPE2-SFP-IQ2	
		PC-8GE-TYPE3-SFP-IQ2	
Gigabit Ethernet IQ2, XFP	1	PC-1XGE-TYPE3-XFP-IQ2	M120, M320, T320, T640, T1600, TX Matrix
10-Gigabit Ethernet, XENPAK	1	PC-1XGE-XENPAK	M120, M320, T320, T640, T1600, TX Matrix
10-Gigabit Ethernet, DWDM	1	PC-1XGE-DWDM-CBAND	M120, M320, T320, T640, T1600, TX Matrix
10-Gigabit Ethernet	4	PD-4XGE-XFP <i>NOTE:</i> This PIC goes offline during a unified ISSU if the PIC is inserted on T-1600-FPC4-ES with part number 710-013037 revision 12 or below.	T640, T1600, TX Matrix, TX Matrix Plus
	10	PD-5-10XGE-SFPP	T640, T1600

Channelized PICs

[Table 158 on page 2105](#) lists the channelized PICs that are supported during a unified ISSU.

Table 158: Unified ISSU PIC Support: Channelized

PIC Type	Number of Ports	Model Number	Platform
Channelized E1 IQ	10	PB-10CHE1-RJ48-QPP	M120, M320, T320, T640, T1600, TX Matrix
		PE-10CHE1-RJ48-QPP-N	M10i
Channelized T1 IQ	10	PB-10CHT1-RJ48-QPP	M320, T320, T640, T1600
		PE-10CHT1-RJ48-QPP	M10i
Channelized OC IQ	1	PB-1CHOC12SMIR-QPP	M120, M320, T320, T640, T1600, TX Matrix
		PB-1CHSTM1-SMIR-QPP	
		PB-1CHOC3-SMIR-QPP	
		PE-1CHOC12SMIR-QPP	M10i
		PE-1CHOC3-SMIR-QPP	
Channelized DS3 to DS0 IQ	4	PB-4CHDS3-QPP	M120, M320, T320, T640, T1600, TX Matrix
		PE-4CHDS3-QPP	M10i
Channelized STM 1	1	PE-1CHSTM1-SMIR-QPP	M10i

Tunnel Services PICs

[Table 159 on page 2105](#) lists the Tunnel Services PICs that are supported during a unified ISSU.

Table 159: Unified ISSU PIC Support: Tunnel Services

PIC Type	Model Number	Platform
1-Gbps Tunnel	PE-TUNNEL	M10i
	PB-TUNNEL-1	M120, M320, T320, T640, T1600, TX Matrix
4-Gbps Tunnel	PB-TUNNEL	
10-Gbps Tunnel	PC-TUNNEL	

ATM PICs

[Table 160 on page 2106](#) lists the ATM PICs that are supported during a unified ISSU. This includes support on Enhanced III FPCs.

Table 160: Unified ISSU PIC Support: ATM

PIC Type	Number of Ports	Model Number	Platform
DS3	4	PB-4DS3-ATM2	M120, M320, T320, T640, T1600, TX Matrix
		PE-4DS3-ATM2	M10i
E3	4	PB-4E3-ATM2	M120, M320, T320, T640, T1600, TX Matrix
	2	PE-2E3-ATM2	M10i
OC3/STM1	2	PB-2OC3-ATM2-MM	M120, M320, T320, T640, T1600, TX Matrix
		PB-2OC3-ATM2-SMIR	
		PE-2OC3-ATM2-MM	M10i
		PE-2OC3-ATM2-SMIR	
OC12/STM4	1	PB-1OC12-ATM2-MM	M120, M320, T320, T640, T1600, TX Matrix
		PB-1OC12-ATM2-SMIR	
	2	PB-2OC12-ATM2-MM	M120, M320, T320, T640, T1600, TX Matrix
		PB-2OC12-ATM2-SMIR	
	1	PE-1OC12-ATM2-MM	M10i
		PE-1OC12-ATM2-SMIR	
OC48/STM16	1	PB-1OC48-ATM2-SFP	M120, M320, T320, T640, T1600, TX Matrix

Serial PICs

Unified ISSU supports the following 2-port EIA-530 serial PICs:

- PB-2EIA530 on M320 routers with Enhanced III FPCs, and on M120 routers
- PE-2EIA530 on M10i routers

DS3, E1, E3, and T1 PICs

Unified ISSU supports the following PICs on M120, M320, and T320 routers; T640 and T1600 routers; and the TX Matrix router:

- 4-Port DS3 PIC (PB-4DS3)
- 4-Port E1 Coaxial PIC (PB-4E1-COAX)
- 4-Port E1 RJ48 PIC (PB-4E1-RJ48)

- 4-port E3 IQ PIC (PB-4E3-QPP)
- 4-Port T1 PIC (PB-4T1-RJ48)

Unified ISSU supports the following PICs on M10i routers:

- 2-Port DS3 PIC (PE-2DS3)
- 4-Port DS3 PIC (PE-4DS3)
- 4-Port E1 PICs (PE-4E1-COAX and PE-4E1-RJ48)
- 2-Port E3 PIC (PE-2E3)
- 4-Port T1 PIC (PE-4T1-RJ48)
- 4-Port E3 IQ PIC (PE-4E3-QPP)

Enhanced IQ PICs

Unified ISSU supports the following PICs on M120, M320, and T320 routers; T640 and T1600 routers; and the TX Matrix router:

- 1-port Channelized OC12/STM4 enhanced IQ PIC (PB-1CHOC12-STM4-IQE-SFP)
- 1-port nonchannelized OC12/STM4 enhanced IQ PIC (PB-1OC12-STM4-IQE-SFP)
- 4-port Channelized DS3/E3 enhanced IQ PIC (PB-4CHDS3-E3-IQE-BNC)
- 4-port nonchannelized DS3/E3 enhanced IQ PIC (PB-4DS3-E3-IQE-BNC)
- 4-port nonchannelized SONET/SDH OC48/STM16 Enhanced IQ (IQE) PIC with SFP (PC-4OC48-STM16-IQE-SFP)

Unified ISSU supports 1-port Channelized OC48/STM16 Enhanced IQ (IQE) PIC with SFP (PB-1CHOC48-STM16-IQE-SFP) on MX Series routers.

Enhanced IQ2 Ethernet Services Engine (ESE) PIC

Unified ISSU supports the enhanced IQ2 ESE PICs listed in [Table 161 on page 2107](#).

Table 161: Unified ISSU Support: Enhanced IQ2 Ethernet Services Engine (ESE) PIC

Model Number	Number of Ports	Platform
PC-8GE-TYPE3-SFP-IQ2E	8	M120, M320, T320, T640, and TX Matrix.
PB-8GE-TYPE2-SFP-IQ2E	8	M120, M320, T320, T640, and TX Matrix.
PB-4GE-TYPE1-SFP-IQ2E	4	M120, M320, T320, and T640.
PC-1XGE-TYPE3-XFP-IQ2E	1	M120, M320, T320, T640, and TX Matrix.
PB-1CHOC48-STM16-IQE	1	M120, M320, T320, T640, and TX Matrix.
PE-4GE-TYPE1-SFP-IQ2E	4	M10i.
PE-4GE-TYPE1-SFP-IQ2	4	M10i.

Unified ISSU Support on MX Series 3D Universal Edge Routers

The following sections list the Dense Port Concentrators (DPCs), Flexible PIC Concentrators (FPCs), Modular Port Concentrators (MPCs), and Modular Interface Cards (MICs) that are supported during a unified ISSU on MX Series 3D Universal Edge Routers.

- [Unified ISSU DPC and FPC Support on MX Series 3D Universal Edge Routers on page 2108](#)
- [Unified ISSU MIC and MPC Support on MX Series 3D Universal Edge Routers on page 2108](#)
- [Unified ISSU Limitation on MX Series 3D Universal Edge Routers on page 2109](#)

Unified ISSU DPC and FPC Support on MX Series 3D Universal Edge Routers

Unified ISSU supports all Dense Port Concentrators (DPCs) except the Multiservices DPC on MX Series routers. Unified ISSU also supports Type 2 FPC (MX-FPC2) and Type 3 FPC (MX-FPC3) on MX Series routers. For more information about PICs supported on MX-FPC2 and MX-FPC3, see [Table 156 on page 2102](#). For more information about DPCs and FPCs on MX Series routers, go to http://www.juniper.net/techpubs/en_US/release-independent/junos/information-products/pathway-pages/mx-series/.

Unified ISSU MIC and MPC Support on MX Series 3D Universal Edge Routers

Unified ISSU supports all the Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) listed in [Table 162 on page 2108](#) and [Table 163 on page 2109](#). Unified ISSU is not supported on MX80 routers nor in an MX Series Virtual Chassis.

In the MPCs on MX Series routers, interface-specific and firewall filter statistics are preserved across a unified ISSU. During the unified ISSU, counter and policer operations are disabled.

To preserve statistics across a unified ISSU on MX Series routers with MPC/MIC interfaces, the router stores the statistics data as binary large objects. The router collects the statistics before the unified ISSU is initialized, and restores the statistics after the unified ISSU completes. No statistics are collected during the unified ISSU process.

To verify that statistics are preserved across the unified ISSU, you can issue CLI operational commands such as **show interfaces statistics** after the unified ISSU completes.

Table 162: Unified ISSU Support: MX Series 3D Universal Edge Routers

MPC Type	Number of Ports	Model Number	Platform
30-Gigabit Ethernet MPC	—	MX-MPC1-3D	MX Series 3D Universal Edge Routers
30-Gigabit Ethernet Queuing MPC	—	MX-MPC1-3D-Q	MX Series 3D Universal Edge Routers
60-Gigabit Ethernet MPC	—	MX-MPC2-3D	MX Series 3D Universal Edge Routers
60-Gigabit Ethernet Queuing MPC	—	MX-MPC2-3D-Q	MX Series 3D Universal Edge Routers

Table 162: Unified ISSU Support: MX Series 3D Universal Edge Routers (*continued*)

MPC Type	Number of Ports	Model Number	Platform
60-Gigabit Ethernet Enhanced Queuing MPC	—	MX-MPC2-3D-EQ	MX Series 3D Universal Edge Routers
10-Gigabit Ethernet MPC with SFP+	16	MPC-3D-16XGE-SFPP	MX Series 3D Universal Edge Routers

Table 163: Unified ISSU Support: MX Series 3D Universal Edge Routers

MIC Type	Number of Ports	Model Number	Platform
Gigabit Ethernet MIC with SFP	20	MIC-3D-20GE-SFP	MX Series 3D Universal Edge Routers
10-Gigabit Ethernet MICs with XFP	2	MIC-3D-2XGE-XFP	MX Series 3D Universal Edge Routers
10-Gigabit Ethernet MICs with XFP	4	MIC-3D-4XGE-XFP	MX Series 3D Universal Edge Routers
Tri-Rate Copper Ethernet MIC	40	MIC-3D-40GE-TX	MX Series 3D Universal Edge Routers



NOTE: Note that unified ISSU is supported only by the MICs listed in Table 163 on page 2109.

Unified ISSU Limitation on MX Series 3D Universal Edge Routers

Unified in-service software upgrade (unified ISSU) is currently not supported when clock synchronization is configured for Synchronous Ethernet, Precision Time Protocol (PTP), and hybrid mode on MX80 3D Universal Edge Routers and on the MICs and MPCEs on MX240, MX480, and MX960 routers.

Related Documentation

- [Unified ISSU Concepts on page 2092](#)
- [Unified ISSU Process on the TX Matrix Router](#)
- [Before You Begin on page 2197](#)
- [Performing a Unified ISSU on page 2200](#)

VRRP

- [Understanding VRRP on page 2110](#)
- [Junos OS Support for VRRPv3 on page 2111](#)
- [Improving the Convergence Time for VRRP on page 2114](#)

Understanding VRRP

For Ethernet, Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, and logical interfaces, you can configure the Virtual Router Redundancy Protocol (VRRP) or VRRP for IPv6. VRRP enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts. The VRRP routing platforms share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routing platforms is the master (active) and the others are backups. If the master fails, one of the backup routers or switches becomes the new master router, providing a virtual default routing platform and enabling traffic on the LAN to be routed without relying on a single routing platform. Using VRRP, a backup router can take over a failed default router within a few seconds. This is done with minimum VRRP traffic and without any interaction with the hosts.

Routers running VRRP dynamically elect master and backup routers. You can also force assignment of master and backup routers using priorities from 1 through 255, with 255 being the highest priority. In VRRP operation, the default master router sends advertisements to backup routers at regular intervals. The default interval is 1 second. If a backup router does not receive an advertisement for a set period, the backup router with the next highest priority takes over as master and begins forwarding packets.




NOTE: To minimize network traffic, VRRP is designed in such a way that only the router that is acting as the master sends out VRRP advertisements at any given point in time. The backup routers do not send any advertisement until and unless they take over mastership.

VRRP for IPv6 provides a much faster switchover to an alternate default router than IPv6 Neighbor Discovery (ND) procedures. Typical deployments use only one backup router.

Figure 25 on page 2110 illustrates a basic VRRP topology. In this example, Routers A, B, and C are running VRRP and together they make up a virtual router. The IP address of this virtual router is 10.10.0.1 (the same address as the physical interface of Router A).

Figure 25: Basic VRRP



Because the virtual router uses the IP address of the physical interface of Router A, Router A is the master VRRP router, while Routers B and C function as backup VRRP routers.

Clients 1 through 3 are configured with the default gateway IP address of 10.10.0.1. As the master router, Router A forwards packets sent to its IP address. If the master virtual router fails, the router configured with the higher priority becomes the master virtual router and provides uninterrupted service for the LAN hosts. When Router A recovers, it becomes the master virtual router again.

VRRP is defined in RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*. VRRP for IPv6 is defined in Internet draft draft-ietf-vrrp-ipv6-spec-08.txt, *Virtual Router Redundancy Protocol for IPv6*. See also Internet draft draft-ietf-vrrp-unified-mib-06.txt, *Definitions of Managed Objects for the VRRP over IPv4 and IPv6*.



NOTE: Even though VRRP, as defined in RFC 3768, does not support authentication, the Junos OS implementation of VRRP supports authentication as defined in RFC 2338. This support is achieved through the backward compatibility options in RFC 3768.

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Configuring Basic VRRP Support on page 2220](#)

Junos OS Support for VRRPv3

Prior to Junos OS Release 12.2, Junos OS supported RFC 3768, *Virtual Router Redundancy Protocol (VRRP) for IPv4* and Internet draft draft-ietf-vrrp-ipv6-spec-08, *Virtual Router Redundancy Protocol for IPv6*.

Starting with Junos OS Release 12.2, Junos OS supports RFC 3768, *Virtual Router Redundancy Protocol (VRRP) for IPv4*. The support for VRRPv3 is implemented in compliance with RFC 5798, *Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv6*. Junos OS Release 12.2 also supports VRRP MIB for VRRPv3. The support for VRRP MIB for VRRPv3 is implemented in compliance with RFC 6527, *Definitions of Managed Objects for the Virtual Router Redundancy Protocol Version 3 (VRRPv3)*.

When you configure VRRP for IPv4 or IPv6 networks, you can enable VRRPv3 by configuring the **version-3** statement at the **[edit protocols vrrp]** hierarchy level.

Understanding VRRPv3 Behavioral Differences

You must consider the following aspects when enabling VRRPv3 for your IPv4 or IPv6 networks:

- Prior to Junos OS Release 12.2, when VRRP for IPv6 is configured without enabling VRRPv3, the VRRP checksum is calculated according to section 5.3.8 of RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*. However, when VRRPv3 is enabled, the VRRP checksum is calculated according to section 5.3.7 of draft-ietf-vrrp-ipv6-spec-08.txt, *Virtual Router Redundancy Protocol for IPv6*.
- Starting with Junos OS Release 12.2, when VRRP for IPv6 is configured without enabling VRRPv3, the VRRP checksum is calculated according to section 5.2.8 of RFC 5798, *Virtual Router Redundancy Protocol (VRRP)*. However, when VRRPv3 is enabled, the

VRRP checksum is calculated according to section 5.3.7 of draft-ietf-vrrp-ipv6-spec-08.txt, *Virtual Router Redundancy Protocol for IPv6*.



NOTE: When enabling VRRPv3, you must ensure that VRRPv3 is enabled on all the VRRP routers in the network. This is because VRRPv3 does not interoperate with the previous versions of VRRP.

- The **tcpdump** utility calculates the VRRP checksum according to draft-ietf-vrrp-ipv6-spec-08.txt. Therefore, when **tcpdump** parses IPv6 VRRP packets that are received from older Junos OS releases (prior to Junos OS Release 12.2), the **bad vrrp cksum** message is displayed:

```
23:20:32.657328 Out
...
-----original packet-----
00:00:5e:00:02:03 > 33:33:00:00:00:12, ethertype IPv6 (0x86dd), length
94: (class 0xc0, hlim 255, next-header: VRRP (112), length: 40)
fe80::224:dcff:fe47:57f > ff02::12: VRRPv3-advertisement 40: vrid=3 prio=100
intvl=100(centisec) (bad vrrp cksum b4e2!) addrs(2):
fe80::200:5eff:fe00:3,2001:4818:f000:14::1
3333 0000 0012 0000 5e00 0203 86dd 6c00
0000 0028 70ff fe80 0000 0000 0000 0224
dcff fe47 057f ff02 0000 0000 0000 0000
0000 0000 0012 3103 6402 0064 b4e2 fe80
0000 0000 0000 0200 5eff fe00 0003 2001
4818 f000 0014 0000 0000 0000 0001
```

You can ignore this message because it does not indicate VRRP failure.

- When VRRPv3 is enabled, the **authentication-type** and **authentication-key** statements (for IPv4 VRRP) cannot be configured for any VRRP groups. Therefore, if authentication is required, you need to configure alternative non-VRRP authentication mechanisms.
- When VRRPv3 is enabled, the **advertise-interval** statement (for IPv4 VRRP) and the **inet6-advertise-interval** statement (for IPv6 VRRP) cannot be used to configure advertisement intervals. Instead, use the **fast-interval** statement to configure advertisement intervals.

Understanding VRRP Interoperability

Prior to 12.2, VRRP (IPv6) followed Internet draft draft-ietf-vrrp-ipv6-spec-08, but checksum was calculated based on RFC 3768 Section 5.3.8. Starting with 12.2, VRRP (IPv6) follows RFC 5798 and checksum is calculated based on RFC 5798 Section 5.2.8.

Moreover, when VRRPv3 is enabled, pseudo-header is included only when calculating IPv6 VRRP checksum. Pseudo-header is not included when calculating IPv4 VRRP checksum. Therefore, care must be taken to correctly calculate the IPv4 and IPv6 checksum values when VRRPv3 is enabled.



NOTE: Because of the differences in VRRP checksum calculations, IPv6 VRRP configured on routers that use Junos OS Release 12.2 and later releases does not interoperate with IPv6 VRRP configured in releases before Junos OS Release 12.2.

- If you have configured VRRPv3 (IPv4 or IPv6) on routers that use Junos OS Release 12.2 or later releases, it will not operate with the routers that use Junos OS Release 12.1 or earlier releases. This is because only Junos OS Release 12.2 and later releases support VRRPv3.
- VRRP (IPv4 or IPv6) configured on routers that use Junos OS Release 12.2 and later releases, interoperate with VRRP (IPv4 or IPv6) configured on routers that use releases prior to Junos OS Release 12.2.
- VRRPv3 for IPv4 does not interoperate with the previous versions of VRRP. If VRRPv2 IPv4 advertisement packets are received by a router on which VRRPv3 is enabled, the router transitions itself to the backup state to avoid creating multiple masters in the network. Due to this behavior, you must be cautious when enabling VRRPv3 on your existing VRRPv2 networks. See [“Understanding VRRPv2 to VRRPv3 Transition” on page 2113](#) for more information.



NOTE: VRRPv3 advertisement packets are ignored by the routers on which previous versions of VRRP are configured.

Understanding VRRPv2 to VRRPv3 Transition

You must enable VRRPv3 in your network only if VRRPv3 can be enabled on all the VRRP routers in your network. Even if VRRPv3 can be enabled on all the VRRP routers in your network, care must be taken to avoid traffic loss when you transition your network to VRRPv3. This is because it is practically not possible to configure VRRPv3 on all routers simultaneously. There is a small time frame in the transition period during which VRRPv2 and VRRPv3 coexist in the network. During this period, to avoid having multiple masters in the network, the VRRPv3 IPv4 routers switch to the backup state when they receive a VRRPv2 IPv4 advertisement packet. VRRPv2 IPv4 packets are always given the highest priority. Additionally, to avoid having multiple masters in your IPv6 network due to checksum differences, you need to disable VRRP for IPv6 on the backup routers.



NOTE: Configuration change from VRRPv2 to VRRPv3 (or VRRPv3 to VRRPv2) restarts the VRRP state machine on all the configured VRRP groups.

The following example illustrates the steps and events that take place during a VRRPv2 to VRRPv3 transition:

Consider a scenario where two VRRPv2 routers, R1 and R2, are configured in two groups, G1 and G2. The R1 router acts as the master for G1 and the R2 router acts as the master for G2. [Table 164 on page 2114](#) lists the transition steps and events for this setup:

Table 164: Example: VRRPv2 to VRRPv3 Transition Steps and Events

1. Upgrade the R1 router with Junos OS Release 12.2 or later.
 - R2 becomes master for both G1 and G2.
 - After the upgrade of the R1 router is completed, R1 becomes the master for G1. R2 remains as the master for G2.
2. Upgrade the R2 router with Junos OS Release 12.2 or later.
 - R1 becomes master for both G1 and G2.
 - After the upgrade of R2 router is completed, R2 becomes the master for G2. R1 remains as the master for G1.

For IPv4	For IPv6
3. Enable VRRPv3 on the R1 router. <ul style="list-style-type: none"> • Because VRRPv2 IPv4 advertisement packets are given higher priority, R1 becomes the backup for both G1 and G2. 	3. Deactivate the G1 and G2 groups on the R2 router. <ul style="list-style-type: none"> • G1 and G2 groups on the R1 router become master.
4. Enable VRRPv3 on the R2 router. <ul style="list-style-type: none"> • R1 becomes the master for G1 and R2 becomes the master for G2. 	4. Enable VRRPv3 on the R1 router. <ul style="list-style-type: none"> • R1 becomes master for both G1 and G2.
	5. Enable VRRPv3 on the R2 router.
	6. Activate G1 and G2 groups on the R2 router. <ul style="list-style-type: none"> • R2 becomes master for G2. • R1 remains as the master for G1.

Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Configuring Basic VRRP Support on page 2220](#)
- [VRRP Configuration Hierarchy](#)
- [VRRP for IPv6 Configuration Hierarchy](#)
- [authentication-type on page 2248](#)
- [authentication-key on page 2247](#)
- [fast-interval on page 2251](#)
- [inet6-advertise-interval on page 2254](#)
- [version-3 on page 2265](#)
- [virtual-link-local-address on page 2267](#)

Improving the Convergence Time for VRRP

You can enable faster convergence time for the configured Virtual Router Redundancy Protocol (VRRP), thereby reducing the traffic restoration time to less than 1 second. To improve the convergence time for the VRRP, perform the following tasks:

- **Configure the distributed periodic packet management process**—When the VRRP process is busy and does not send VRRP advertisements, the backup VRRP routers might assume that the master router is down and take over as the master router, causing unnecessary flaps. To address this problem and to reduce the load on the VRRP process, Junos OS uses the distributed periodic packet management (PPM) process to send VRRP advertisements on behalf of the VRRP process.

To configure the distributed PPM process, include the **delegate-processing** statement at the **[edit protocols vrrp]** hierarchy level.

- **Disable the skew timer**—The skew timer in VRRP is used to ensure that two backup routers do not switch to the master state at the same time in case of a failover situation. When there is only one master router and one backup router in the network deployment, you can disable the skew timer, thereby reducing the time required to transition to the master state.

To disable the skew timer, include the **skew-timer-disable** statement at the **[edit protocols vrrp]** hierarchy level.

- **Configure the number of fast advertisements that can be missed by a backup router before it starts transitioning to the master state**—The backup router waits until a certain number of advertisement packets are lost after which it transitions to the master state. This waiting time can be fatal in scenarios such as router failure or link failure. To avoid such a situation and to enable faster convergence time, in Junos OS Release 12.2 and later, you can configure a fast advertisement interval value that specifies the number of fast advertisements that can be missed by a backup router before it starts transitioning to the master state.

To configure the fast advertisement interval, include the **global-advertisements-threshold** statement at the **[edit protocols vrrp]** hierarchy level.

- **Configure inheritance of VRRP groups**—Junos OS enables you to configure VRRP groups on the various subnets of a virtual LAN (VLAN) to inherit the state and configuration of one of the groups, which is known as the active VRRP group. When the **vrrp-inherit-from** statement is included in the configuration, only the active VRRP group, from which the other VRRP groups inherit the state, sends out frequent VRRP advertisements and processes incoming VRRP advertisements. Use inherit groups for scaled configurations. For example, if you have 1000 VRRP groups with an advertisement interval of 100 ms, then use inherit groups.

To configure inheritance for a VRRP group, include the **vrrp-inherit-from** statement at the **[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id]** hierarchy level.



NOTE:

- The reduction in convergence time is not applicable when VRRP is configured over integrated routing and bridging (IRB) interfaces, aggregated Ethernet interfaces, and multichassis link aggregation group (MC-LAG) interfaces.
 - Compared to other routers, the convergence time and the traffic restoration time are less for MX Series routers with MPCs.
 - Reduction in convergence time is applicable for all types of configurations at the physical interface but the convergence time might not be less than 1 second for all the configurations. The convergence time depends on the number of groups that are transitioning from the backup to the master state and the interval at which these groups are transitioning.
-

Related Documentation

- [Configuring Inheritance for a VRRP Group on page 2232](#)
- [Configuring VRRP to Improve Convergence Time on page 2236](#)
- [delegate-processing on page 2250](#)
- [global-advertisements-threshold on page 2252](#)
- [skew-timer-disable on page 2261](#)

Configuration

- [Configuration: GRES on page 2116](#)
- [Configuration Statements: GRES on page 2118](#)
- [Configuration: Graceful Restart on page 2127](#)
- [Configuration Statements: Graceful Restart on page 2167](#)
- [Configuration: NSB on page 2180](#)
- [Configuration Statements: NSB on page 2181](#)
- [Configuration: NSR on page 2182](#)
- [Configuration Statements: NSR on page 2191](#)
- [Configuration: Unified ISSU on page 2197](#)
- [Configuration Statements: Unified ISSU on page 2214](#)
- [Configuration: VRRP on page 2219](#)
- [Configuration Statements: VRRP on page 2242](#)

Configuration: GRES

- [Configuring Graceful Routing Engine Switchover on page 2117](#)
- [Resetting Local Statistics on page 2118](#)

Configuring Graceful Routing Engine Switchover

This section contains the following topics:

- [Enabling Graceful Routing Engine Switchover on page 2117](#)
- [Synchronizing the Routing Engine Configuration on page 2117](#)
- [Verifying Graceful Routing Engine Switchover Operation on page 2117](#)

Enabling Graceful Routing Engine Switchover

By default, graceful Routing Engine switchover is disabled. To configure graceful Routing Engine switchover, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level.

```
[edit chassis redundancy]
graceful-switchover;
```

When you enable graceful Routing Engine switchover, the command-line interface (CLI) indicates which Routing Engine you are using. For example:

```
{master} [edit]
user@host#
```

To disable graceful Routing Engine switchover, delete the **graceful-switchover** statement from the **[edit chassis redundancy]** hierarchy level.

Synchronizing the Routing Engine Configuration



NOTE: A newly inserted backup Routing Engine automatically synchronizes its configuration with the master Routing Engine configuration.

When you configure graceful Routing Engine switchover, you can bring the backup Routing Engine online after the master Routing Engine is already running. There is no requirement to start the two Routing Engines simultaneously.

Verifying Graceful Routing Engine Switchover Operation

To verify whether graceful Routing Engine switchover is enabled, on the backup Routing Engine, issue the **show system switchover** command. When the output of the command indicates that the **Graceful switchover** field is set to **on**, graceful Routing Engine switchover is operational. The status of the kernel database and configuration database synchronization between Routing Engines is also provided. For example:

```
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady state
```



NOTE: You must issue the **show system switchover** command on the backup Routing Engine. This command is not supported on the master Routing Engine.

For more information about the **show system switchover** command, see the *Junos OS Operational Mode Commands*.

**Related
Documentation**

- [Understanding Graceful Routing Engine Switchover in the Junos OS on page 2061](#)
- [Graceful Routing Engine Switchover System Requirements on page 2065](#)
- [Requirements for Routers with a Backup Router Configuration on page 2068](#)
- [Resetting Local Statistics on page 2118](#)
- [graceful-switchover on page 2127](#)

Resetting Local Statistics

When you enable graceful Routing Engine switchover, the master Routing Engine configuration is copied and loaded to the backup Routing Engine. User files, accounting information, and trace options information are not replicated to the backup Routing Engine.

When a graceful Routing Engine switchover occurs, local statistics such as process statistics and networking statistics are displayed as a cumulative value from the time the process first came online. Because processes on the master Routing Engine can start at different times from the processes on the backup Routing Engine, the statistics on the two Routing Engines for the same process might differ. After a graceful Routing Engine switchover, we recommend that you issue the **clear interface statistics (*interface-name* | all)** command to reset the cumulative values for local statistics. Forwarding statistics are not affected by graceful Routing Engine switchover.

For information about how to use the **clear** command to clear statistics and protocol database information, see the *Junos OS Operational Mode Commands*.



NOTE: The **clear firewall** command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for graceful Routing Engine switchover.

**Related
Documentation**

- [Understanding Graceful Routing Engine Switchover in the Junos OS on page 2061](#)
- [Configuring Graceful Routing Engine Switchover on page 2117](#)

Configuration Statements: GRES

- [\[edit chassis\] Hierarchy Level on page 2118](#)

[edit chassis] Hierarchy Level

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
    }
    lacp {
      link-protection {
```

```

        non-revertive;
    }
    system-priority;
}
}
sonet {
    device-count number;
}
maximum-links maximum-links-limit;
}
alarm {
    ds1 {
        ais (ignore | red | yellow);
        ylw (ignore | red | yellow);
    }
    ethernet {
        link-down (ignore | red | yellow);
    }
    integrated-services {
        failure (ignore | red | yellow);
    }
    management-ethernet {
        link-down (ignore | red | yellow);
    }
    relay
    input {
        port port-number {
            mode (close | open);
            trigger (ignore | red | yellow);
        }
    }
    output {
        port port-number {
            input-relay input-relay;
            mode (close | open);
            temperature;
        }
    }
    serial {
        cts-absent (ignore | red | yellow);
        dcd-absent (ignore | red | yellow);
        dsr-absent (ignore | red | yellow);
        loss-of-rx-clock (ignore | red | yellow);
        loss-of-tx-clock (ignore | red | yellow);
        tm-absent (ignore | red | yellow);
    }
    services {
        hw-down (ignore | red | yellow);
        linkdown (ignore | red | yellow);
        pic-hold-reset (ignore | red | yellow);
        pic-reset (ignore | red | yellow);
        rx-errors (ignore | red | yellow);
        sw-down (ignore | red | yellow);
        tx-errors (ignore | red | yellow);
    }
    sonet {

```

```

        (ais-l | ais-p | ber-sd | ber-sf | locd | lol | lop-p | los | pll | plm-p | rfi-l | rfl-p | uneq-p)
        (ignore | red | yellow);
    }
    t3 {
        (ais | exz | ferf | idle | lcv | lof | los | pll | ylw) (ignore | red | yellow);
    }
}
cluster {
    control-link-recovery;
    control-ports {
        fpc slot-number port port-number;
    }
    heartbeat-interval milliseconds;
    heartbeat-threshold number;
    redundancy-group {
        ... the redundancy-group subhierarchy appears at the end of the [edit chassis cluster]
            hierarchy ...
    }
    reth-count number;
    traceoptions {
        file <filename> <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
        flag flag;
        level severity;
        no-remote-trace;
    }
    redundancy-group group-number {
        gratuitous-arp-count number;
        hold-down-interval seconds;
        interface-monitor {
            interface-name weight number;
        }
        ip-monitoring {
            family {
                inet {
                    ipv4-address {
                        interface rethindex.logical-unit-number secondary-ip-address ipv4-address;
                        weight number;
                    }
                }
            }
            global-threshold number;
            global-weight number;
            retry-count count;
            retry-interval interval;
        }
        node node-number priority priority-number;
        preempt;
    }
    config-button {
        no-clear;
        no-rescue;
    }
    container-devices {
        device-count number;
    }
}

```

```

craft-lockout;
disable-power-management;
disk-partition partition-name (/config | /var) {
    level (full | high) {
        free-space threshold-value (mb | percent);
    }
}
enhanced-policer;
extended-statistics;
fabric {
    degraded {
        action-fpc-restart-disable;
        degraded-fabric-detection-enable
        degraded-fpc-bad-plane-threshold number-bad-planes;
    }
    redundancy-mode (increased-bandwidth | redundant);
}
filter;
fpc slot-number {
    ... the fpc subhierarchy appears after the main [edit chassis] hierarchy ...
}
fpc-feb-connectivity {
    fpc slot-number feb (slot-number | none);
}
fpc-resync;
fru-poweron-sequence sequence;
lcc index {
    ... the lcc subhierarchy appears after the main [edit chassis] hierarchy ...
}
maximum-ecmp value;
memory-enhanced {
    filter;
    route;
    vpn-label;
}
network-services (ethernet | enhanced-ethernet | ip | enhanced-ip);
(packet-scheduling | no-packet-scheduling);
pem {
    minimum number;
}
policer-drop-probability-low;
ppp-subscriber-services (disable | enable);
redundancy {
    cfeb slot (always | preferred);
    failover {
        on-disk-failure;
        on-loss-of-keepalives;
    }
    feb {
        redundancy-group group-name {
            description description;
            feb slot-number <backup | primary>;
            no-auto-failover;
        }
    }
}
graceful-switchover;

```

```
    keepalive-time seconds;
    routing-engine slot-number (backup | disabled | master);
    sfm slot-number (always | preferred);
    ssb slot-number (always | preferred);
}
route-memory-enhanced;
route-localization {
    inet (chassis);
    inet6;
}
routing-engine {
    bios {
        no-auto-upgrade;
    }
    on-disk-failure disk-failure-action (halt | reboot);
}
sfm slot-number {
    power off;
}
sib {
    minimum number;
}
(source-route | no-source-route);
state [
    cb-upgrade [on | off];
]
synchronization { # for M Series and T Series routers
    primary (external-a | external-b);
    secondary (external-a | external-b);
    signal-type (e1 | t1);
    switching-mode (non-revertive | revertive);
    transmitter-enable;
    validation-interval seconds;
    y-cable-line-termination;
}
synchronization { # for MX80 and MX240 routers
    clock-mode (auto-select | free-run);
    esmc-transmit {
        interfaces (all | interface-name);
    }
    hold-interval {
        configuration-change seconds;
        restart seconds;
        switchover seconds;
    }
    network-option (option-1 | option-2);
    quality-mode-enable;
    selection-mode (configured-quality|received-quality);
    source {
        (external-a | external-b) {
            priority number;
            quality-level (prc | prs | sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
            request (force-switch | lockout);
        }
        interfaces interface-name {
            priority number;
        }
    }
}
```



```

        quality-level (prc | prs | sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
        request (force-switch | lockout);
        wait-to-restore minutes;
    }
}
switchover-mode (revertive | non-revertive);
}
synchronization { # for ACX Series routers
    clock-mode (auto-select | free-run);
    esmc-transmit {
        interfaces (all | interface-name);
    }
    hold-interval {
        configuration-change seconds;
        restart seconds;
        switchover seconds;
    }
    network-option (option-1 | option-2);
    quality-mode-enable;
    selection-mode (configured-quality | received-quality);
    source {
        (bits | gps) {
            priority number;
            quality-level (prc | prs | sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
            request (force-switch | lockout);
        }
        interfaces interface-name {
            priority number;
            quality-level (prc | prs | sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
            request (force-switch | lockout);
            wait-to-restore minutes;
        }
    }
}
switchover-mode(non-revertive | revertive);
}
system-domains {
    protected-system-domains psdnumerical-index {
        control-plane-bandwidth-percent percent;
        control-slot-numbers [ slot-numbers ];
        control-system-id control-system-id;
        description description;
        fpcs [ slot-numbers ];
    }
    root-domain-id root-domain-id;
}
vrf-mtu-check;
}

chassis {
    fpc slot-number {
        number-of-ports active-ports;
        offline;
        pic slot-number {
            ... the pic subhierarchy appears after the main [edit chassis fpc slot-number] hierarchy
            ...
        }
    }
}

```

```
port-mirror-instance port-mirror-instance-name;  
power (off | on);  
sampling-instance instance-name;  
}  
  
fpc slot-number {  
  pic slot-number {  
    adaptive-services {  
      service-package (layer-2 | layer-3 | ...the following extension-provider subhierarchy  
        ...);  
    extension-provider {  
      control-cores number;  
      data-cores number;  
      data-flow-affinity {  
        hash-key (layer-3 | layer-4);  
      }  
      channelization;  
      forwarding-db-size megabytes;  
      object-cache-size megabytes;  
      package package-name;  
      policy-db-size megabytes;  
      syslog {  
        facility {  
          severity;  
          destination (pic-console | routing-engine);  
        }  
      }  
      wired-process-mem-size megabytes;  
    }  
  }  
  aggregated-devices {  
    ima {  
      device-count number;  
    }  
  }  
  aggregate-ports;  
  atm-cell-relay-accumulation;  
  atm-l2circuit-mode (aal5 | cell | trunk trunk);  
  cel {  
    el port-number {  
      channel-group group-number timeslots slot-number;  
    }  
  }  
  ct3 {  
    port port-number {  
      tl link-number {  
        channel-group group-number timeslots slot-number;  
      }  
    }  
  }  
  ethernet {  
    pic-mode (enhanced-switching | routing | switching);  
  }  
  fibre-channel {  
    port port-number;  
    port-range port-range-low port-range-high
```

```

}
egress-policer-overhead bytes;
forwarding-mode {
    sa-multicast;
    vlan-steering {
        vlan-rule (high-low | odd-even);
    }
}
framing (e1 | e3 | sdh | sonet | t1 | t3);
idle-cell-format {
    itu-t;
    payload-pattern payload-pattern-byte;
}
ingress-policer-overhead bytes;
inline-services {
    bandwidth (1g | 10g);
}
linerate-mode;
max-queues-per-interface (4 | 8);
mlfr-uni-nni-bundles number;
no-concatenate;
no-multi-rate;
port port-number {
    framing (e1 | e3 | sdh | sonet | t1 | t3);
    forwarding-mode {
        sa-multicast;
    }
    speed ( oc3-stm1 | oc12-stm4 | oc48-stm16);
}
port-mirror-instance port-mirror-instance-name;
q-pic-large-buffer {
    (large-scale | small-scale);
}
red-buffer-occupancy {
    weighted-averaged <instant-usage-weight-exponent weight-value>;
}
shdsl {
    pic-mode (1-port-atm | 2-port-atm);
}
sparse-dlcis;
traffic-manager {
    egress-shaping-overhead number;
    ingress-shaping-overhead number;
    mode {
        egress-only;
        ingress-and-egress;
        session-shaping;
    }
}
tunnel-queuing;
tunnel-services {
    bandwidth (1g | 10g | 20g | 40g);
    tunnel-only;
}
vtmapping (itu-t | klm);
}

```

```
}
}

chassis {
  lcc index {
    fpc slot-number {
      ... the fpc subhierarchy appears after the main [edit chassis lcc index] hierarchy ...
    }
    offline;
    online-expected;
  }
}

lcc index {
  fpc slot-number {
    pic slot-number {
      ... the pic subhierarchy appears after the main [edit chassis lcc index fpc slot-number]
      hierarchy ...
    }
    power (off | on);
    sampling-instance instance-name;
  }

  fpc slot-number {
    pic slot-number {
      aggregate-ports;
      atm-cell-relay-accumulation;
      atm-l2circuit-mode (aal5 | cell | trunk trunk);
      framing (e1 | e3 | sdh | sonet | t1 | t3);
      idle-cell-format {
        itu-t;
        payload-pattern payload-pattern-byte;
      }
      linerate-mode;
      max-queues-per-interface (4 | 8);
      no-concatenate;
      no-mcast-replication;
      no-pre-classifier;
      port port-number {
        framing (e1 | e3 | sdh | sonet | t1 | t3);
      }
      q-pic-large-buffer {
        (large-scale | small-scale);
      }
      red-buffer-occupancy {
        weighted-averaged <instant-usage-weight-exponent weight-value>;
      }
      shdsl {
        pic-mode (1-port-atm | 2-port-atm);
      }
      traffic-manager {
        egress-shaping-overhead bytes;
        ingress-shaping-overhead bytes;
        mode {
          egress-only;
          ingress-and-egress;
        }
      }
    }
  }
}
```

```

    }
  }
}

```

Related Documentation • [Notational Conventions Used in Junos OS Configuration Hierarchies](#)

graceful-switchover

Syntax	<code>graceful-switchover;</code>
Hierarchy Level	[edit chassis redundancy]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For routing platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine without interruption to packet forwarding.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	• Configuring Graceful Routing Engine Switchover on page 2117

Configuration: Graceful Restart

- [Enabling Graceful Restart on page 2127](#)
- [Configuring Routing Protocols Graceful Restart on page 2128](#)
- [Configuring Graceful Restart for MPLS-Related Protocols on page 2134](#)
- [Configuring VPN Graceful Restart on page 2136](#)
- [Configuring Logical System Graceful Restart on page 2137](#)
- [Example: Configuring Graceful Restart on page 2139](#)
- [Example: Managing Helper Modes for OSPF Graceful Restart on page 2164](#)

Enabling Graceful Restart

Graceful restart is disabled by default. You must configure graceful restart at the **[edit routing-options]** hierarchy level to enable the feature.

For graceful restart to function properly, graceful restart must be enabled on the **[edit routing-instance *instance-name* routing-options]** or **[edit routing-options]** hierarchy level.

For example:

```

routing-options {
  graceful-restart;
}

```



NOTE: If you configure graceful restart after a BGP or LDP session has been established, the BGP or LDP session restarts and the peers negotiate graceful restart capabilities.

To disable graceful restart, include the **disable** statement. To configure a time period for complete restart, include the **restart-duration** statement. You can specify a number between 120 and 900.

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

When you include the **graceful-restart** statement at the **[edit routing-options]** hierarchy level, graceful restart is also enabled for aggregate and static routes.

**Related
Documentation**

- [Graceful Restart Concepts on page 2085](#)
- [Graceful Restart System Requirements on page 2086](#)
- [Aggregate and Static Routes on page 2086](#)
- [Example: Configuring Graceful Restart on page 2139](#)

Configuring Routing Protocols Graceful Restart

This topic includes the following sections:

- [Enabling Graceful Restart on page 2128](#)
- [Configuring Graceful Restart Options for BGP on page 2129](#)
- [Configuring Graceful Restart Options for ES-IS on page 2130](#)
- [Configuring Graceful Restart Options for IS-IS on page 2130](#)
- [Configuring Graceful Restart Options for OSPF and OSPFv3 on page 2131](#)
- [Configuring Graceful Restart Options for RIP and RIPng on page 2132](#)
- [Configuring Graceful Restart Options for PIM Sparse Mode on page 2133](#)
- [Tracking Graceful Restart Events on page 2134](#)

Enabling Graceful Restart

By default, graceful restart is disabled. To enable graceful restart, include the **graceful-restart** statement at the **[edit routing-instance *instance-name* routing-options]** or **[edit routing-options]** hierarchy level.

For example:

```
routing-options {  
  graceful-restart;  
}
```

To configure the duration of the graceful restart period, include the **restart-duration** at the **[edit routing-options graceful-restart]** hierarchy level.



NOTE: Helper mode (the ability to assist a neighboring router attempting a graceful restart) is enabled by default when you start the routing platform, even if graceful restart is not enabled. You can disable helper mode on a per-protocol basis.

```
[edit]
routing-options {
  graceful-restart {
    disable;
    restart-duration seconds;
  }
}
```

To disable graceful restart globally, include the **disable** statement at the **[edit routing-options graceful-restart]** hierarchy level.

When graceful restart is enabled for all routing protocols at the **[edit routing-options graceful-restart]** hierarchy level, you can disable graceful restart on a per-protocol basis.



NOTE: If you configure graceful restart after a BGP or LDP session has been established, the BGP or LDP session restarts and the peers negotiate graceful restart capabilities. Also, the BGP peer routing statistics are reset to zero.

Configuring Graceful Restart Options for BGP

To configure the duration of the BGP graceful restart period, include the **restart-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the **stale-routes-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.

```
[edit]
protocols {
  bgp {
    graceful-restart {
      disable;
      restart-time seconds;
      stale-routes-time seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable BGP graceful restart capability for all BGP sessions, include the **disable** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.



NOTE: To set BGP graceful restart properties or disable them for a group, include the desired statements at the `[edit protocols bgp group group-name graceful-restart]` hierarchy level.

To set BGP graceful restart properties or disable them for a specific neighbor in a group, include the desired statements at the `[edit protocols bgp group group-name neighbor ip-address graceful-restart]` hierarchy level.



NOTE: Configuring graceful restart for BGP resets the BGP peer routing statistics to zero. Also, existing BGP sessions restart, and the peers negotiate graceful restart capabilities.

Configuring Graceful Restart Options for ES-IS

On J Series Services Routers, to configure the duration of the ES-IS graceful restart period, include the `restart-duration` statement at the `[edit protocols esis graceful-restart]` hierarchy level.

```
[edit]
protocols {
  esis {
    graceful-restart {
      disable;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable ES-IS graceful restart capability, include the `disable` statement at the `[edit protocols esis graceful-restart]` hierarchy level.

Configuring Graceful Restart Options for IS-IS

To configure the duration of the IS-IS graceful restart period, include the `restart-duration` statement at the `[edit protocols isis graceful-restart]` hierarchy level.

```
[edit]
protocols {
  isis {
    graceful-restart {
      disable;
      helper-disable;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```


To disable IS-IS graceful restart helper capability, include the **helper-disable** statement at the **[edit protocols isis graceful-restart]** hierarchy level. To disable IS-IS graceful restart capability, include the **disable** statement at the **[edit protocols isis graceful-restart]** hierarchy level.



NOTE: If you configure Bidirectional Forwarding Detection (BFD) and graceful restart for IS-IS, graceful restart might not work as expected.



NOTE: You can also track graceful restart events with the **traceoptions** statement at the **[edit protocols isis]** hierarchy level. For more information, see [“Tracking Graceful Restart Events” on page 2134](#).

Configuring Graceful Restart Options for OSPF and OSPFv3

To configure the duration of the OSPF/OSPFv3 graceful restart period, include the **restart-duration** statement at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level. To specify the length of time for which the router notifies helper routers that it has completed graceful restart, include the **notify-duration** at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level. Strict OSPF link-state advertisement (LSA) checking results in the termination of graceful restart by a helping router. To disable strict LSA checking, include the **no-strict-lsa-checking** statement at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level.

```
[edit]
protocols {
  ospf | ospfv3 {
    graceful-restart {
      disable;
      helper-disable
      no-strict-lsa-checking;
      notify-duration seconds;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable OSPF/OSPFv3 graceful restart, include the **disable** statement at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level.

Starting with Release 11.3, the Junos OS supports both the standard (based on RFC 3623, *Graceful OSPF Restart*) and the restart signaling-based (as specified in RFC 4811, RFC 4812, and RFC 4813) helper modes for OSPF version 2 graceful restart configurations. Both the standard and restart signaling-based helper modes are enabled by default. To disable the helper mode for OSPF version 2 graceful restart configurations, include the **helper-disable <both | restart-signaling | standard>** statement at the **[edit protocols ospf graceful-restart]** hierarchy level. Note that the last committed statement always takes precedence over the previous one.

```
[edit protocols ospf]
  graceful-restart {
    helper-disable <both | restart-signaling | standard>
  }
```

To reenble the helper mode, delete the **helper-disable** statement from the configuration by using the **delete protocols ospf graceful-restart helper-disable <restart-signaling | standard | both>** command. In this case also, the last executed command takes precedence over the previous ones.



NOTE:

Restart signaling-based helper mode is not supported for OSPFv3 configurations. To disable helper mode for OSPFv3 configurations, include the **helper-disable** statement at the **[edit protocols ospfv3 graceful-restart]** hierarchy level.



TIP: You can also track graceful restart events with the **traceoptions** statement at the **[edit protocols (ospf | ospf3)]** hierarchy level. For more information, see [“Tracking Graceful Restart Events” on page 2134](#).



NOTE: You cannot enable OSPFv3 graceful restart between a routing platform running Junos OS Release 7.5 and earlier and a routing platform running Junos OS Release 7.6 or later. As a workaround, make sure both routing platforms use the same Junos OS version.



NOTE: If you configure BFD and graceful restart for OSPF, graceful restart might not work as expected.

Configuring Graceful Restart Options for RIP and RIPng

To configure the duration of the RIP or RIPng graceful restart period, include the **restart-time** statement at the **[edit protocols (rip | ripng) graceful-restart]** hierarchy level.

```
[edit]
protocols {
  (rip | ripng) {
    graceful-restart {
      disable;
      restart-time seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable RIP or RIPng graceful restart capability, include the **disable** statement at the **[edit protocols (rip | ripng) graceful-restart]** hierarchy level.

Configuring Graceful Restart Options for PIM Sparse Mode

PIM sparse mode continues to forward existing multicast packet streams during a graceful restart, but does not forward new streams until after the restart is complete. After a restart, the routing platform updates the forwarding state with any updates that were received from neighbors and occurred during the restart period. For example, the routing platform relearns the join and prune states of neighbors during the restart, but does not apply the changes to the forwarding table until after the restart.

PIM sparse mode-enabled routing platforms generate a unique 32-bit random number called a generation identifier. Generation identifiers are included by default in PIM hello messages, as specified in the IETF Internet draft *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*. When a routing platform receives PIM hellos containing generation identifiers on a point-to-point interface, Junos OS activates an algorithm that optimizes graceful restart.

Before PIM sparse mode graceful restart occurs, each routing platform creates a generation identifier and sends it to its multicast neighbors. If a PIM sparse mode-enabled routing platform restarts, it creates a new generation identifier and sends it to its neighbors. When a neighbor receives the new identifier, it resends multicast updates to the restarting router to allow it to exit graceful restart efficiently. The restart phase completes when either the PIM state becomes stable or when the restart interval timer expires.

If a routing platform does not support generation identifiers or if PIM is enabled on multipoint interfaces, the PIM sparse mode graceful restart algorithm does not activate, and a default restart timer is used as the restart mechanism.

To configure the duration of the PIM graceful restart period, include the **restart-duration** statement at the **[edit protocols pim graceful-restart]** hierarchy level:

```
[edit]
protocols {
  pim {
    graceful-restart {
      disable;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable PIM sparse mode graceful restart capability, include the **disable** statement at the **[edit protocols pim graceful-restart]** hierarchy level.



NOTE: Multicast forwarding can be interrupted in two ways. First, if the underlying routing protocol is unstable, multicast reverse-path-forwarding (RPF) checks can fail and cause an interruption. Second, because the forwarding table is not updated during the graceful restart period, new multicast streams are not forwarded until graceful restart is complete.

Tracking Graceful Restart Events

To track the progress of a graceful restart event, you can configure graceful restart trace options flags for IS-IS and OSPF/OSPFv3. To configure graceful restart trace options, include the **graceful-restart** statement at the **[edit protocols protocol traceoptions flag]** hierarchy level:

```
[edit protocols]
isis {
  traceoptions {
    flag graceful-restart;
  }
}
(ospf | ospf3) {
  traceoptions {
    flag graceful-restart;
  }
}
```

Related Documentation

- [Graceful Restart Concepts on page 2085](#)
- [Graceful Restart System Requirements on page 2086](#)
- [Graceful Restart and Routing Protocols on page 2087](#)
- [Verifying Graceful Restart Operation on page 2271](#)
- [Example: Configuring Graceful Restart on page 2139](#)

Configuring Graceful Restart for MPLS-Related Protocols

This section contains the following topics:

- [Configuring Graceful Restart Globally on page 2134](#)
- [Configuring Graceful Restart Options for RSVP, CCC, and TCC on page 2135](#)
- [Configuring Graceful Restart Options for LDP on page 2135](#)

Configuring Graceful Restart Globally

To configure graceful restart globally for all MPLS-related protocols, include the **graceful-restart** statement at the **[edit routing-options]** hierarchy level. To configure the duration of the graceful restart period, include the **restart-duration** at the **[edit routing-options graceful-restart]** hierarchy level:

```
[edit]
routing-options {
  graceful-restart {
    disable;
```

```

    restart-duration seconds;
  }
}

```

To disable graceful restart globally, include the **disable** statement at the **[edit routing-options graceful-restart]** hierarchy level.

Configuring Graceful Restart Options for RSVP, CCC, and TCC

Because CCC and TCC rely on RSVP, you must modify these three protocols as a single group.

To configure how long the router retains the state of its RSVP neighbors while they undergo a graceful restart, include the **maximum-helper-recovery-time** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level. This value is applied to all neighboring routers, so it should be based on the time required by the slowest RSVP neighbor to recover.

To configure the delay between when the router discovers that a neighboring router has gone down and when it declares the neighbor down, include the **maximum-helper-restart-time** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level. This value is applied to all neighboring routers, so it should be based on the time required by the slowest RSVP neighbor to restart.

```

[edit]
protocols {
  rsvp {
    graceful-restart {
      disable;
      helper-disable;
      maximum-helper-recovery-time;
      maximum-helper-restart-time;
    }
  }
}
routing-options {
  graceful-restart;
}

```

To disable RSVP, CCC, and TCC graceful restart, include the **disable** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level. To disable RSVP, CCC, and TCC helper capability, include the **helper-disable** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level.

Configuring Graceful Restart Options for LDP

When configuring graceful restart for LDP, you can include the following optional statements at the **[edit protocols ldp graceful-restart]** hierarchy level:

```

[edit protocols ldp graceful-restart]
disable;
helper-disable;
maximum-neighbor-reconnect-time seconds;
maximum-neighbor-recovery-time seconds;
reconnect-time seconds;
recovery-time seconds;

```

```
[edit routing-options]
graceful-restart;
```

The statements have the following effects on the graceful restart process:

- To configure the length of time required to reestablish a session after a graceful restart, include the **reconnect-time** statement; the range is 30 through 300 seconds. To limit the maximum reconnect time allowed from a restarting neighbor router, include the **maximum-neighbor-reconnect-time** statement; the range is 30 through 300 seconds.
- To configure the length of time that helper routers are required to maintain the old forwarding state during a graceful restart, include the **recovery-time** statement; the range is 120 through 1800 seconds. On the helper router, you can configure a statement that overrides the request from the restarting router and sets the maximum length of time the helper router will maintain the old forwarding state. To configure this feature, include the **maximum-neighbor-recovery-time** statement; the range is 140 through 1900 seconds.



NOTE: The value for the **recovery-time** and **maximum-neighbor-recovery-time** statements at the **[edit protocols ldp graceful-restart]** hierarchy level should be approximately 80 seconds longer than the value for the **restart-duration** statement at the **[edit routing-options graceful-restart]** hierarchy level. Otherwise, a warning message appears when you try to commit the configuration.

- To disable LDP graceful restart capability, include the **disable** statement. To disable LDP graceful restart helper capability, include the **helper-disable** statement.

Configuring VPN Graceful Restart

Graceful restart allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. Without graceful restart, a control plane restart disrupts any VPN services provided by the router. Graceful restart is supported on Layer 2 VPNs, Layer 3 VPNs, virtual-router routing instances, and VPLS.

To implement graceful restart for a Layer 2 VPN or Layer 3 VPN, perform the configuration tasks described in the following sections:

- [Configuring Graceful Restart Globally on page 2136](#)
- [Configuring Graceful Restart for the Routing Instance on page 2137](#)

Configuring Graceful Restart Globally

To enable graceful restart, include the **graceful-restart** statement at the **[edit routing-options]** hierarchy level. To configure a global duration for the graceful restart period, include the **restart-duration** statement at the **[edit routing-options graceful-restart]** hierarchy level.

```
[edit]
routing-options {
  graceful-restart {
```

```

    disable;
    restart-duration seconds;
  }
}

```

To disable graceful restart globally, include the **disable** statement at the **[edit routing-options graceful-restart]** hierarchy level.

Configuring Graceful Restart for the Routing Instance

For Layer 3 VPNs only, you must also configure graceful restart for all routing and MPLS-related protocols within a routing instance by including the **graceful-restart** statement at the **[edit routing-instances instance-name routing-options]** hierarchy level. Because you can configure multi-instance BGP and multi-instance LDP, graceful restart for a carrier-of-carriers scenario is supported. To configure the duration of the graceful restart period for the routing instance, include the **restart-duration** statement at the **[edit routing-instances instance-name routing-options]**.

```

[edit]
routing-instances {
  instance-name {
    routing-options {
      graceful-restart {
        disable;
        restart-duration seconds;
      }
    }
  }
}

```

You can disable graceful restart for individual protocols with the **disable** statement at the **[edit routing-instances instance-name protocols protocol-name graceful-restart]** hierarchy level.

Related Documentation

- [Graceful Restart Concepts on page 2085](#)
- [Graceful Restart System Requirements on page 2086](#)
- [Graceful Restart and Layer 2 and Layer 3 VPNs on page 2091](#)
- [Verifying Graceful Restart Operation on page 2271](#)
- [Example: Configuring Graceful Restart on page 2139](#)

Configuring Logical System Graceful Restart

Graceful restart for a logical system functions much as graceful restart does in the main router. The only difference is the location of the **graceful-restart** statement.

The following topics describe what to configure to implement graceful restart in a logical system:

- [Enabling Graceful Restart Globally on page 2138](#)
- [Configuring Graceful Restart for a Routing Instance on page 2138](#)

Enabling Graceful Restart Globally

To enable graceful restart in a logical system, include the **graceful-restart** statement at the **[edit logical-systems *logical-system-name* routing-options]** hierarchy level. To configure a global duration of the graceful restart period, include the **restart-duration** statement at the **[edit logical-systems *logical-system-name* routing-options graceful-restart]** hierarchy level.

```
[edit]
logical-systems {
  logical-system-name {
    routing-options {
      graceful-restart {
        disable;
        restart-duration seconds;
      }
    }
  }
}
```

To disable graceful restart globally, include the **disable** statement at the **[edit logical-systems *logical-system-name* routing-options graceful-restart]** hierarchy level.

Configuring Graceful Restart for a Routing Instance

For Layer 3 VPNs only, you must also configure graceful restart globally for a routing instance inside a logical system. To configure, include the **graceful-restart** statement at the **[edit logical-systems *logical-system-name* routing-instances *instance-name* routing-options]** hierarchy level. Because you can configure multi-instance BGP and multi-instance LDP, graceful restart for a carrier-of-carriers scenario is supported. To configure the duration of the graceful restart period for the routing instance, include the **restart-duration** statement at the **[edit logical-systems *logical-system-name* routing-instances *instance-name* routing-options]**.

```
[edit]
logical-systems {
  logical-system-name {
    routing-instances {
      instance-name {
        routing-options {
          graceful-restart {
            disable;
            restart-duration seconds;
          }
        }
      }
    }
  }
}
```

To disable graceful restart for individual protocols with the **disable** statement at the **[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols *protocol-name* graceful-restart]** hierarchy level.

Related Documentation

- [Graceful Restart Concepts on page 2085](#)
- [Graceful Restart System Requirements on page 2086](#)

- [Graceful Restart on Logical Systems on page 2092](#)
- [Verifying Graceful Restart Operation on page 2271](#)
- [Example: Configuring Graceful Restart on page 2139](#)

Example: Configuring Graceful Restart

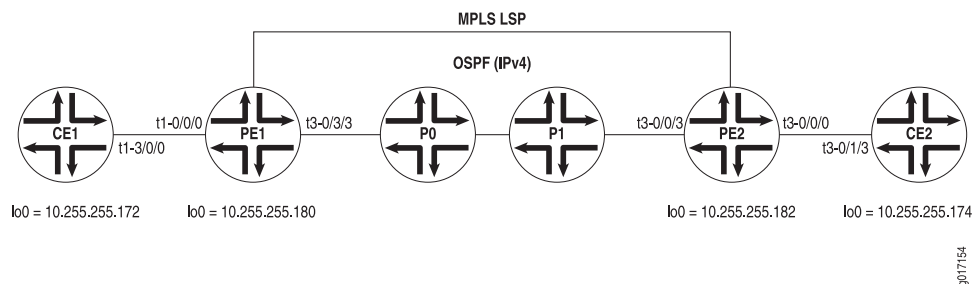
To enable graceful restart, include the **graceful-restart** statement at the **[edit routing-instance *instance-name* routing-options]** or **[edit routing-options]** hierarchy level as well as in the protocol level.

For example:

```
protocols {
  bgp {
    group ext {
      graceful-restart;
    }
  }
}
routing-options {
  graceful-restart;
}
```

Figure 26 on page 2139 shows a standard MPLS VPN network. Routers CE1 and CE2 are customer edge routers, PE1 and PE2 are provider edge routers, and P0 is a provider core router. Several Layer 3 VPNs are configured across this network, as well as one Layer 2 VPN. Interfaces are shown in the diagram and are not included in the configuration example that follows.

Figure 26: Layer 3 VPN Graceful Restart Topology



Router CE1 On Router CE1, configure the following protocols on the logical interfaces of **t3-3/1/0**: OSPF on unit 101, RIP on unit 102, BGP on unit 103, and IS-IS on unit 512. Also configure graceful restart, BGP, IS-IS, OSPF, and RIP on the main instance to be able to connect to the routing instances on Router PE1.

```
[edit]
interfaces {
  t3-3/1/0 {
    encapsulation frame-relay;
    unit 100 {
      dlci 100;
      family inet {
        address 10.96.100.2/30;
      }
    }
  }
}
```

```
    }
  }
  unit 101 {
    dlci 101;
    family inet {
      address 10.96.101.2/30;
    }
  }
  unit 102 {
    dlci 102;
    family inet {
      address 10.96.102.2/30;
    }
  }
  unit 103 {
    dlci 103;
    family inet {
      address 10.96.103.2/30;
    }
  }
  unit 512 {
    dlci 512;
    family inet {
      address 10.96.252.1/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.245.14.172/32;
      primary;
    }
    address 10.96.110.1/32;
    address 10.96.111.1/32;
    address 10.96.112.1/32;
    address 10.96.113.1/32;
    address 10.96.116.1/32;
  }
  family iso {
    address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4172.00;
  }
}
}
routing-options {
  graceful-restart;
  autonomous-system 65100;
}
protocols {
  bgp {
    group CE-PE-INET {
      type external;
      export BGP_INET_LB_DIRECT;
      neighbor 10.96.103.1 {
        local-address 10.96.103.2;
        family inet {
```

```

        unicast;
    }
    peer-as 65103;
}
}
isis {
    export ISIS_L2VPN_LB_DIRECT;
    interface t3-3/1/0.512;
}
ospf {
    export OSPF_LB_DIRECT;
    area 0.0.0.0 {
        interface t3-3/1/0.101;
    }
}
rip {
    group RIP {
        export RIP_LB_DIRECT;
        neighbor t3-3/1/0.102;
    }
}
}
policy-options {
    policy-statement OSPF_LB_DIRECT {
        term direct {
            from {
                protocol direct;
                route-filter 10.96.101.0/30 exact;
                route-filter 10.96.111.1/32 exact;
            }
            then accept;
        }
        term final {
            then reject;
        }
    }
    policy-statement RIP_LB_DIRECT {
        term direct {
            from {
                protocol direct;
                route-filter 10.96.102.0/30 exact;
                route-filter 10.96.112.1/32 exact;
            }
            then accept;
        }
        term final {
            then reject;
        }
    }
}
policy-statement BGP_INET_LB_DIRECT {
    term direct {
        from {
            protocol direct;
            route-filter 10.96.103.0/30 exact;
            route-filter 10.96.113.1/32 exact;

```

```
    }
    then accept;
  }
  term final {
    then reject;
  }
}
policy-statement ISIS_L2VPN_LB_DIRECT {
  term direct {
    from {
      protocol direct;
      route-filter 10.96.116.1/32 exact;
    }
    then accept;
  }
  term final {
    then reject;
  }
}
```

Router PE1 On Router PE1, configure graceful restart in the master instance, along with BGP, OSPF, MPLS, and LDP. Next, configure several protocol-specific instances of graceful restart. By including instances for BGP, OSPF, Layer 2 VPNs, RIP, and static routes, you can observe the wide range of options available when you implement graceful restart. Configure the following protocols in individual instances on the logical interfaces of **t3-0/0/0**: a static route on unit 100, OSPF on unit 101, RIP on unit 102, BGP on unit 103, and Frame Relay on unit 512 for the Layer 2 VPN instance.

```
[edit]
interfaces {
  t3-0/0/0 {
    dce;
    encapsulation frame-relay-ccc;
    unit 100 {
      dlci 100;
      family inet {
        address 10.96.100.1/30;
      }
      family mpls;
    }
    unit 101 {
      dlci 101;
      family inet {
        address 10.96.101.1/30;
      }
      family mpls;
    }
    unit 102 {
      dlci 102;
      family inet {
        address 10.96.102.1/30;
      }
      family mpls;
    }
    unit 103 {
```

```

        dlci 103;
        family inet {
            address 10.96.103.1/30;
        }
        family mpls;
    }
    unit 512 {
        encapsulation frame-relay-ccc;
        dlci 512;
    }
}
t1-0/1/0 {
    unit 0 {
        family inet {
            address 10.96.0.2/30;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.245.14.176/32;
        }
        family iso {
            address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4176.00;
        }
    }
}
}
routing-options {
    graceful-restart;
    router-id 10.245.14.176;
    autonomous-system 69;
}
protocols {
    mpls {
        interface all;
    }
    bgp {
        group PEPE {
            type internal;
            neighbor 10.245.14.182 {
                local-address 10.245.14.176;
                family inet-vpn {
                    unicast;
                }
                family l2vpn {
                    unicast;
                }
            }
        }
    }
}
ospf {
    area 0.0.0.0 {
        interface t1-0/1/0.0;
    }
}

```

```
        interface fxp0.0 {
            disable;
        }
        interface lo0.0 {
            passive;
        }
    }
    ldp {
        interface all;
    }
}
policy-options {
    policy-statement STATIC-import {
        from community STATIC;
        then accept;
    }
    policy-statement STATIC-export {
        then {
            community add STATIC;
            accept;
        }
    }
    policy-statement OSPF-import {
        from community OSPF;
        then accept;
    }
    policy-statement OSPF-export {
        then {
            community add OSPF;
            accept;
        }
    }
    policy-statement RIP-import {
        from community RIP;
        then accept;
    }
    policy-statement RIP-export {
        then {
            community add RIP;
            accept;
        }
    }
    policy-statement BGP-INET-import {
        from community BGP-INET;
        then accept;
    }
    policy-statement BGP-INET-export {
        then {
            community add BGP-INET;
            accept;
        }
    }
    policy-statement L2VPN-import {
        from community L2VPN;
        then accept;
    }
}
```

```

}
policy-statement L2VPN-export {
  then {
    community add L2VPN;
    accept;
  }
}
community BGP-INET members target:69:103;
community L2VPN members target:69:512;
community OSPF members target:69:101;
community RIP members target:69:102;
community STATIC members target:69:100;
}
routing-instances {
  BGP-INET {
    instance-type vrf;
    interface t3-0/0/0.103;
    route-distinguisher 10.245.14.176:103;
    vrf-import BGP-INET-import;
    vrf-export BGP-INET-export;
    routing-options {
      graceful-restart;
      autonomous-system 65103;
    }
    protocols {
      bgp {
        group BGP-INET {
          type external;
          export BGP-INET-import;
          neighbor 10.96.103.2 {
            local-address 10.96.103.1;
            family inet {
              unicast;
            }
            peer-as 65100;
          }
        }
      }
    }
  }
}
L2VPN {
  instance-type l2vpn;
  interface t3-0/0/0.512;
  route-distinguisher 10.245.14.176:512;
  vrf-import L2VPN-import;
  vrf-export L2VPN-export;
  protocols {# There is no graceful-restart statement for Layer 2 VPN instances.
    l2vpn {
      encapsulation-type frame-relay;
      site CE1-ISIS {
        site-identifier 512;
        interface t3-0/0/0.512 {
          remote-site-id 612;
        }
      }
    }
  }
}

```

```
    }  
  }  
  OSPF {  
    instance-type vrf;  
    interface t3-0/0/0.101;  
    route-distinguisher 10.245.14.176:101;  
    vrf-import OSPF-import;  
    vrf-export OSPF-export;  
    routing-options {  
      graceful-restart;  
    }  
    protocols {  
      ospf {  
        export OSPF-import;  
        area 0.0.0.0 {  
          interface all;  
        }  
      }  
    }  
  }  
  RIP {  
    instance-type vrf;  
    interface t3-0/0/0.102;  
    route-distinguisher 10.245.14.176:102;  
    vrf-import RIP-import;  
    vrf-export RIP-export;  
    routing-options {  
      graceful-restart;  
    }  
    protocols {  
      rip {  
        group RIP {  
          export RIP-import;  
          neighbor t3-0/0/0.102;  
        }  
      }  
    }  
  }  
  STATIC {  
    instance-type vrf;  
    interface t3-0/0/0.100;  
    route-distinguisher 10.245.14.176:100;  
    vrf-import STATIC-import;  
    vrf-export STATIC-export;  
    routing-options {  
      graceful-restart;  
      static {  
        route 10.96.110.1/32 next-hop t3-0/0/0.100;  
      }  
    }  
  }  
}
```


Router P0 On Router P0, configure graceful restart in the main instance, along with OSPF, MPLS, and LDP. This allows the protocols on the PE routers to reach one another.

```
[edit]
interfaces {
  t3-0/1/3 {
    unit 0 {
      family inet {
        address 10.96.0.5/30;
      }
      family mpls;
    }
  }
  t1-0/2/0 {
    unit 0 {
      family inet {
        address 10.96.0.1/30;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.245.14.174/32;
      }
      family iso {
        address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4174.00;
      }
    }
  }
}
routing-options {
  graceful-restart;
  router-id 10.245.14.174;
  autonomous-system 69;
}
protocols {
  mpls {
    interface all;
  }
  ospf {
    area 0.0.0.0 {
      interface t1-0/2/0.0;
      interface t3-0/1/3.0;
      interface fxp0.0 {
        disable;
      }
      interface lo0.0 {
        passive;
      }
    }
  }
  ldp {
    interface all;
  }
}
```

```
}
```

Router PE2 On Router PE2, configure BGP, OSPF, MPLS, LDP, and graceful restart in the master instance. Configure the following protocols in individual instances on the logical interfaces of **t1-0/1/3**: a static route on unit 200, OSPF on unit 201, RIP on unit 202, BGP on unit 203, and Frame Relay on unit 612 for the Layer 2 VPN instance. Also configure protocol-specific graceful restart in all routing instances, except the Layer 2 VPN instance.

```
[edit]
interfaces {
  t3-0/0/0 {
    unit 0 {
      family inet {
        address 10.96.0.6/30;
      }
      family mpls;
    }
  }
  t1-0/1/3 {
    dce;
    encapsulation frame-relay-ccc;
    unit 200 {
      dlci 200;
      family inet {
        address 10.96.200.1/30;
      }
      family mpls;
    }
    unit 201 {
      dlci 201;
      family inet {
        address 10.96.201.1/30;
      }
      family mpls;
    }
    unit 202 {
      dlci 202;
      family inet {
        address 10.96.202.1/30;
      }
      family mpls;
    }
    unit 203 {
      dlci 203;
      family inet {
        address 10.96.203.1/30;
      }
      family mpls;
    }
    unit 612 {
      encapsulation frame-relay-ccc;
      dlci 612;
    }
  }
  lo0 {
    unit 0 {
```

```

        family inet {
            address 10.245.14.182/32;
        }
        family iso {
            address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4182.00;
        }
    }
}
routing-options {
    graceful-restart;
    router-id 10.245.14.182;
    autonomous-system 69;
}
protocols {
    mpls {
        interface all;
    }
    bgp {
        group PEPE {
            type internal;
            neighbor 10.245.14.176 {
                local-address 10.245.14.182;
                family inet-vpn {
                    unicast;
                }
                family l2vpn {
                    unicast;
                }
            }
        }
    }
}
ospf {
    area 0.0.0.0 {
        interface t3-0/0/0.0;
        interface fxp0.0 {
            disable;
        }
        interface lo0.0 {
            passive;
        }
    }
}
ldp {
    interface all;
}
policy-options {
    policy-statement STATIC-import {
        from community STATIC;
        then accept;
    }
    policy-statement STATIC-export {
        then {
            community add STATIC;
            accept;
        }
    }
}

```

```
}
policy-statement OSPF-import {
  from community OSPF;
  then accept;
}
policy-statement OSPF-export {
  then {
    community add OSPF;
    accept;
  }
}
policy-statement RIP-import {
  from community RIP;
  then accept;
}
policy-statement RIP-export {
  then {
    community add RIP;
    accept;
  }
}
policy-statement BGP-INET-import {
  from community BGP-INET;
  then accept;
}
policy-statement BGP-INET-export {
  then {
    community add BGP-INET;
    accept;
  }
}
policy-statement L2VPN-import {
  from community L2VPN;
  then accept;
}
policy-statement L2VPN-export {
  then {
    community add L2VPN;
    accept;
  }
}
community BGP-INET members target:69:103;
community L2VPN members target:69:512;
community OSPF members target:69:101;
community RIP members target:69:102;
community STATIC members target:69:100;
}
routing-instances {
  BGP-INET {
    instance-type vrf;
    interface t1-0/1/3.203;
    route-distinguisher 10.245.14.182:203;
    vrf-import BGP-INET-import;
    vrf-export BGP-INET-export;
    routing-options {
      graceful-restart;
    }
  }
}
```

```

    autonomous-system 65203;
  }
  protocols {
    bgp {
      group BGP-INET {
        type external;
        export BGP-INET-import;
        neighbor 10.96.203.2 {
          local-address 10.96.203.1;
          family inet {
            unicast;
          }
        }
        peer-as 65200;
      }
    }
  }
}
L2VPN {
  instance-type l2vpn;
  interface t1-0/1/3.612;
  route-distinguisher 10.245.14.182:612;
  vrf-import L2VPN-import;
  vrf-export L2VPN-export;
  protocols {# There is no graceful-restart statement for Layer 2 VPN instances.
    l2vpn {
      encapsulation-type frame-relay;
      site CE2-ISIS {
        site-identifier 612;
        interface t1-0/1/3.612 {
          remote-site-id 512;
        }
      }
    }
  }
}
OSPF {
  instance-type vrf;
  interface t1-0/1/3.201;
  route-distinguisher 10.245.14.182:201;
  vrf-import OSPF-import;
  vrf-export OSPF-export;
  routing-options {
    graceful-restart;
  }
  protocols {
    ospf {
      export OSPF-import;
      area 0.0.0.0 {
        interface all;
      }
    }
  }
}
RIP {
  instance-type vrf;

```

```
interface t1-0/1/3.202;
route-distinguisher 10.245.14.182:202;
vrf-import RIP-import;
vrf-export RIP-export;
routing-options {
    graceful-restart;
}
protocols {
    rip {
        group RIP {
            export RIP-import;
            neighbor t1-0/1/3.202;
        }
    }
}
}
STATIC {
    instance-type vrf;
    interface t1-0/1/3.200;
    route-distinguisher 10.245.14.182:200;
    vrf-import STATIC-import;
    vrf-export STATIC-export;
    routing-options {
        graceful-restart;
        static {
            route 10.96.210.1/32 next-hop t1-0/1/3.200;
        }
    }
}
}
```

Router CE2 On Router CE2, complete the Layer 2 and Layer 3 VPN configuration by mirroring the protocols already set on Routers PE2 and CE1. Specifically, configure the following on the logical interfaces of **t1-0/0/3**: OSPF on unit 201, RIP on unit 202, BGP on unit 203, and IS-IS on unit 612. Finally, configure graceful restart, BGP, IS-IS, OSPF, and RIP on the main instance to be able to connect to the routing instances on Router PE2.

```
[edit]
interfaces {
    t1-0/0/3 {
        encapsulation frame-relay;
        unit 200 {
            dlci 200;
            family inet {
                address 10.96.200.2/30;
            }
        }
        unit 201 {
            dlci 201;
            family inet {
                address 10.96.201.2/30;
            }
        }
        unit 202 {
            dlci 202;
```

```

        family inet {
            address 10.96.202.2/30;
        }
    }
    unit 203 {
        dlci 203;
        family inet {
            address 10.96.203.2/30;
        }
    }
    unit 512 {
        dlci 512;
        family inet {
            address 10.96.252.2/30;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.245.14.180/32 {
                primary;
            }
            address 10.96.210.1/32;
            address 10.96.111.1/32;
            address 10.96.212.1/32;
            address 10.96.213.1/32;
            address 10.96.216.1/32;
        }
        family iso {
            address 47.0005.80ff.f800.0000.0108.0001.0102.4501.4180.00;
        }
    }
}
}
routing-options {
    graceful-restart;
    autonomous-system 65200;
}
protocols {
    bgp {
        group CE-PE-INET {
            type external;
            export BGP_INET_LB_DIRECT;
            neighbor 10.96.203.1 {
                local-address 10.96.203.2;
                family inet {
                    unicast;
                }
            }
            peer-as 65203;
        }
    }
}
isis {
    export ISIS_L2VPN_LB_DIRECT;
    interface t1-0/0/3.612;
}

```

```
}
ospf {
  export OSPF_LB_DIRECT;
  area 0.0.0.0 {
    interface t1-0/0/3.201;
  }
}
rip {
  group RIP {
    export RIP_LB_DIRECT;
    neighbor t1-0/0/3.202;
  }
}
}
policy-options {
  policy-statement OSPF_LB_DIRECT {
    term direct {
      from {
        protocol direct;
        route-filter 10.96.201.0/30 exact;
        route-filter 10.96.211.1/32 exact;
      }
      then accept;
    }
    term final {
      then reject;
    }
  }
  policy-statement RIP_LB_DIRECT {
    term direct {
      from {
        protocol direct;
        route-filter 10.96.202.0/30 exact;
        route-filter 10.96.212.1/32 exact;
      }
      then accept;
    }
    term final {
      then reject;
    }
  }
  policy-statement BGP_INET_LB_DIRECT {
    term direct {
      from {
        protocol direct;
        route-filter 10.96.203.0/30 exact;
        route-filter 10.96.213.1/32 exact;
      }
      then accept;
    }
    term final {
      then reject;
    }
  }
  policy-statement ISIS_L2VPN_LB_DIRECT {
    term direct {
```



```

        from {
            protocol direct;
            route-filter 10.96.216.1/32 exact;
        }
        then accept;
    }
    term final {
        then reject;
    }
}
}

```

Router PE1 Status Before a Restart The following example displays neighbor relationships on Router PE1 before a restart happens:

```

user@PE1> show bgp neighbor
Peer: 10.96.103.2+3785 AS 65100 Local: 10.96.103.1+179 AS 65103
  Type: External   State: Established   Flags: <>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ BGP-INET-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-unicast
  Local Address: 10.96.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.96.110.1      Local ID: 10.96.103.1      Active Holdtime: 90
  Keepalive Interval: 30
  Local Interface: t3-0/0/0.103
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-unicast
  NLRI peer can save forwarding state: inet-unicast
  NLRI that peer saved forwarding for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Table BGP-INET.inet.0 Bit: 30001
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Suppressed due to damping: 0
  Last traffic (seconds): Received 8    Sent 3    Checked 3
  Input messages: Total 15    Updates 0    Refreshes 0    Octets 321
  Output messages: Total 18    Updates 2    Refreshes 0    Octets 450
  Output Queue[2]: 0

Peer: 10.245.14.182+4701 AS 69   Local: 10.245.14.176+179 AS 69
  Type: Internal   State: Established   Flags: <>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
  Address families configured: inet-vpn-unicast 12vpn

```

```
Local Address: 10.245.14.176 Holdtime: 90 Preference: 170
Number of flaps: 1
Peer ID: 10.245.14.182    Local ID: 10.245.14.176    Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast l2vpn
NLRI advertised by peer: inet-vpn-unicast l2vpn
NLRI for this session: inet-vpn-unicast l2vpn
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast l2vpn
NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
NLRI that restart is negotiated for: inet-vpn-unicast l2vpn
NLRI of all end-of-rib markers sent: inet-vpn-unicast l2vpn
Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Suppressed due to damping: 0
Table bgp.l2vpn.0 Bit: 20000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Suppressed due to damping: 0
Table L2VPN.l2vpn.0 Bit: 90000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
```

```

Send state: in sync
Active prefixes:      1
Received prefixes:    1
Suppressed due to damping: 0
Last traffic (seconds): Received 28   Sent 28   Checked 28
Input messages: Total 2   Updates 0   Refreshes 0   Octets 86
Output messages: Total 13  Updates 10   Refreshes 0   Octets 1073
Output Queue[0]: 0
Output Queue[1]: 0
Output Queue[2]: 0
Output Queue[3]: 0
Output Queue[4]: 0
Output Queue[5]: 0
Output Queue[6]: 0
Output Queue[7]: 0
Output Queue[8]: 0

```

user@PE1> show route instance detail

master:

```

Router ID: 10.245.14.176
Type: forwarding      State: Active
Restart State: Complete Path selection timeout: 300
Tables:
  inet.0                : 17 routes (15 active, 0 holddown, 1 hidden)
  Restart Complete
  inet.3                 : 2 routes (2 active, 0 holddown, 0 hidden)
  Restart Complete
  iso.0                  : 1 routes (1 active, 0 holddown, 0 hidden)
  Restart Complete
  mpls.0                 : 19 routes (19 active, 0 holddown, 0 hidden)
  Restart Complete
  bgp.l3vpn.0            : 10 routes (10 active, 0 holddown, 0 hidden)
  Restart Complete
  inet6.0                 : 2 routes (2 active, 0 holddown, 0 hidden)
  Restart Complete
  bgp.l2vpn.0            : 1 routes (1 active, 0 holddown, 0 hidden)
  Restart Complete

```

BGP-INET:

```

Router ID: 10.96.103.1
Type: vrf              State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.103
Route-distinguisher: 10.245.14.176:103
Vrf-import: [ BGP-INET-import ]
Vrf-export: [ BGP-INET-export ]
Tables:
  BGP-INET.inet.0        : 4 routes (4 active, 0 holddown, 0 hidden)
  Restart Complete

```

L2VPN:

```

Router ID: 0.0.0.0
Type: l2vpn            State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.512
Route-distinguisher: 10.245.14.176:512
Vrf-import: [ L2VPN-import ]
Vrf-export: [ L2VPN-export ]
Tables:
  L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
  Restart Complete

```

```

OSPF:
  Router ID: 10.96.101.1
  Type: vrf                      State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.245.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0                  : 8 routes (7 active, 0 holddown, 0 hidden)
    Restart Complete
RIP:
  Router ID: 10.96.102.1
  Type: vrf                      State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.245.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0                   : 6 routes (6 active, 0 holddown, 0 hidden)
    Restart Complete
STATIC:
  Router ID: 10.96.100.1
  Type: vrf                      State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.245.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0                : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
__juniper_private1__:
  Router ID: 0.0.0.0
  Type: forwarding              State: Active

user@PE1> show route protocol l2vpn
inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
Restart Complete
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
BGP-INET.inet.0: 5 destinations, 6 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
OSPF.inet.0: 7 destinations, 8 routes (7 active, 0 holddown, 0 hidden)
Restart Complete
RIP.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Restart Complete
STATIC.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
mpls.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
800003                *[L2VPN/7] 00:06:00
                      > via t3-0/0/0.512, Pop      Offset: 4
t3-0/0/0.512          *[L2VPN/7] 00:06:00

```

```

> via t1-0/1/0.0, Push 800003, Push 100004(top) Offset: -4
bgp.l3vpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Restart Complete
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
L2VPN.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.245.14.176:512:512:611/96
    *[L2VPN/7] 00:06:01
    Discard

bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

```

Router PE1 Status During a Restart

Before you can verify that graceful restart is working, you must simulate a router restart. To cause the routing process to refresh and simulate a restart, use the **restart routing** operational mode command:

```

user@PE1> restart routing
Routing protocol daemon started, pid 3558

```

The following sample output is captured during the router restart:

```

user@PE1> show bgp neighbor
Peer: 10.96.103.2      AS 65100 Local: 10.96.103.1      AS 65103
  Type: External      State: Active      Flags: <ImportEval>
  Last State: Idle      Last Event: Start
  Last Error: None
  Export: [ BGP-INET-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-unicast
  Local Address: 10.96.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.245.14.182+179 AS 69      Local: 10.245.14.176+2131 AS 69
  Type: Internal      State: Established      Flags: <ImportEval>
  Last State: OpenConfirm      Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
  Address families configured: inet-vpn-unicast l2vpn
  Local Address: 10.245.14.176 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.245.14.182      Local ID: 10.245.14.176      Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-vpn-unicast l2vpn
  NLRI advertised by peer: inet-vpn-unicast l2vpn
  NLRI for this session: inet-vpn-unicast l2vpn
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-vpn-unicast l2vpn
  NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
  NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
  NLRI that restart is negotiated for: inet-vpn-unicast l2vpn
  NLRI of received end-of-rib markers: inet-vpn-unicast l2vpn
Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress

```

```

    Send state: in sync
    Active prefixes:      10
    Received prefixes:    10
    Suppressed due to damping: 0
Table bgp.l2vpn.0 Bit: 20000
    RIB State: BGP restart in progress
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:      1
    Received prefixes:    1
    Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
    RIB State: BGP restart in progress
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:      2
    Received prefixes:    2
    Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
    RIB State: BGP restart is complete
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:      2
    Received prefixes:    2
    Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
    RIB State: BGP restart is complete
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:      2
    Received prefixes:    2
    Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
    RIB State: BGP restart is complete
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:      1
    Received prefixes:    1
    Suppressed due to damping: 0
Table L2VPN.l2vpn.0 Bit: 90000
    RIB State: BGP restart is complete
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:      1
    Received prefixes:    1
    Suppressed due to damping: 0
Last traffic (seconds): Received 0    Sent 0    Checked 0
Input messages: Total 14    Updates 13    Refreshes 0    Octets 1053
Output messages: Total 3    Updates 0    Refreshes 0    Octets 105
Output Queue[0]: 0
Output Queue[1]: 0
Output Queue[2]: 0
Output Queue[3]: 0
Output Queue[4]: 0
Output Queue[5]: 0
Output Queue[6]: 0
Output Queue[7]: 0
Output Queue[8]: 0

```

```

user@PE1> show route instance detail
master:

```

```

Router ID: 10.245.14.176
Type: forwarding          State: Active
Restart State: Pending    Path selection timeout: 300
Tables:
  inet.0                  : 17 routes (15 active, 1 holddown, 1 hidden)
  Restart Pending: OSPF LDP
  inet.3                  : 2 routes (2 active, 0 holddown, 0 hidden)
  Restart Pending: OSPF LDP
  iso.0                   : 1 routes (1 active, 0 holddown, 0 hidden)
  Restart Complete
  mpls.0                  : 23 routes (23 active, 0 holddown, 0 hidden)
  Restart Pending: LDP VPN
  bgp.l3vpn.0             : 10 routes (10 active, 0 holddown, 0 hidden)
  Restart Pending: BGP VPN
  inet6.0                 : 2 routes (2 active, 0 holddown, 0 hidden)
  Restart Complete
  bgp.l2vpn.0             : 1 routes (1 active, 0 holddown, 0 hidden)
  Restart Pending: BGP VPN
BGP-INET:
  Router ID: 10.96.103.1
  Type: vrf                State: Active
  Restart State: Pending    Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.245.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0       : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn              State: Active
  Restart State: Pending    Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.245.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0         : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN L2VPN
OSPF:
  Router ID: 10.96.101.1
  Type: vrf                State: Active
  Restart State: Pending    Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.245.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0           : 8 routes (7 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF VPN
RIP:
  Router ID: 10.96.102.1
  Type: vrf                State: Active
  Restart State: Pending    Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.245.14.176:102

```

```

Vrf-import: [ RIP-import ]
Vrf-export: [ RIP-export ]
Tables:
  RIP.inet.0          : 8 routes (6 active, 2 holddown, 0 hidden)
  Restart Pending: RIP VPN
STATIC:
  Router ID: 10.96.100.1
  Type: vrf           State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.245.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0      : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
__juniper_private1__:
  Router ID: 0.0.0.0
  Type: forwarding     State: Active

```

user@PE1> show route instance summary

Instance	Type	Primary rib	Active/holddown/hidden
master	forwarding		
		inet.0	15/0/1
		iso.0	1/0/0
		mpls.0	35/0/0
		l3vpn.0	0/0/0
		inet6.0	2/0/0
		l2vpn.0	0/0/0
		l2circuit.0	0/0/0
BGP-INET	vrf		
		BGP-INET.inet.0	5/0/0
		BGP-INET.iso.0	0/0/0
		BGP-INET.inet6.0	0/0/0
L2VPN	l2vpn		
		L2VPN.inet.0	0/0/0
		L2VPN.iso.0	0/0/0
		L2VPN.inet6.0	0/0/0
		L2VPN.l2vpn.0	2/0/0
OSPF	vrf		
		OSPF.inet.0	7/0/0
		OSPF.iso.0	0/0/0
		OSPF.inet6.0	0/0/0
RIP	vrf		
		RIP.inet.0	6/0/0
		RIP.iso.0	0/0/0
		RIP.inet6.0	0/0/0
STATIC	vrf		
		STATIC.inet.0	4/0/0
		STATIC.iso.0	0/0/0
		STATIC.inet6.0	0/0/0
__juniper_private1__	forwarding		
		__juniper_priva.inet.0	0/0/0
		__juniper_privat.iso.0	0/0/0
		__juniper_priv.inet6.0	0/0/0

user@PE1> show route protocol l2vpn

```

inet.0: 16 destinations, 17 routes (15 active, 1 holddown, 1 hidden)
Restart Pending: OSPF LDP

```



```

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Pending: OSPF LDP

BGP-INET.inet.0: 5 destinations, 6 routes (5 active, 0 holddown, 0 hidden)
Restart Pending: VPN

OSPF.inet.0: 7 destinations, 8 routes (7 active, 1 holddown, 0 hidden)
Restart Pending: OSPF VPN

RIP.inet.0: 6 destinations, 8 routes (6 active, 2 holddown, 0 hidden)
Restart Pending: RIP VPN

STATIC.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Pending: VPN

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)
Restart Pending: LDP VPN
+ = Active Route, - = Last Active, * = Both

800001          *[L2VPN/7] 00:00:13
                 > via t3-0/0/0.512, Pop          Offset: 4
t3-0/0/0.512    *[L2VPN/7] 00:00:13
                 > via t1-0/1/0.0, Push 800003, Push 100004(top) Offset: -4

bgp.l3vpn.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Restart Pending: BGP VPN

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

L2VPN.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Pending: VPN L2VPN
+ = Active Route, - = Last Active, * = Both

10.245.14.176:512:512:611/96
                 *[L2VPN/7] 00:00:13
                 Discard
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Pending: BGP VPN

```

Related Documentation

- [Enabling Graceful Restart on page 2127](#)
- [Configuring Routing Protocols Graceful Restart on page 2128](#)
- [Configuring Graceful Restart for MPLS-Related Protocols on page 2134](#)
- [Configuring VPN Graceful Restart on page 2136](#)
- [Configuring Logical System Graceful Restart on page 2137](#)
- [Verifying Graceful Restart Operation on page 2271](#)

Example: Managing Helper Modes for OSPF Graceful Restart

- [Requirements on page 2164](#)
- [Overview on page 2164](#)
- [Configuration on page 2164](#)
- [Verification on page 2165](#)

Requirements

M Series or T Series routers running Junos OS Release 11.4 or later and EX Series switches.

Overview

Junos OS Release 11.4 extends OSPF graceful restart support to include restart signaling-based helper mode. Both standard (RFC 3623-based) and restart signaling-based helper modes are enabled by default, irrespective of the graceful-restart configuration status on the routing device.

Junos OS, however, enables you to choose between the helper modes with the **helper-disable <standard | restart-signaling | both>** statement.

Configuration

Both standard and restart signaling-based helper modes are enabled by default, irrespective of the graceful-restart configuration status on the routing device. Junos OS allows you to disable or enable the helper modes based on your requirements.

To configure the helper mode options for graceful restart:

1. To enable graceful restart, add the **graceful-restart** statement at the **[edit routing-options]** hierarchy level.

```
[edit routing-options]
user@host# set graceful-restart
```

The helper modes, both standard and restart signaling-based, are enabled by default.

2. To disable one or both of the helper modes, add the **helper-disable <both | restart-signaling | standard>** statement at the **[edit protocols ospf graceful-restart]** hierarchy level.

- To disable both standard and restart signaling-based helper modes:

```
[edit protocols ospf graceful-restart]
user@host# set helper-disable both
```

- To disable only the restart signaling-based helper mode:

```
[edit protocols ospf graceful-restart]
user@host# set helper-disable restart-signaling
```

- To disable only the standard helper mode:

```
[edit protocols ospf graceful-restart]
user@host# set helper-disable standard
```



NOTE: You must commit the configuration before the change takes effect.

The last committed statement always takes precedence over the previous one.

3. To enable one or both of the helper modes when the helper modes are disabled, delete the **helper-disable <both | restart-signaling | standard>** statement from the **[edit protocols ospf graceful-restart]** hierarchy level.

- To enable both standard and restart signaling-based helper modes:

```
[edit protocols ospf graceful-restart]
user@host# delete helper-disable
```

- To enable the restart signaling-based helper mode:

```
[edit protocols ospf graceful-restart]
user@host# delete helper-disable restart-signaling
```

- To enable the standard helper mode:

```
[edit protocols ospf graceful-restart]
user@host# delete helper-disable standard
```



NOTE: You must commit the configuration before the change takes effect.

The last committed statement always takes precedence over the previous one.

Verification

Confirm that the configuration is working properly.

- [Verifying OSPF Graceful Restart and Helper Mode Configuration on page 2165](#)

Verifying OSPF Graceful Restart and Helper Mode Configuration

Purpose Verify the OSPF graceful restart and helper mode configuration on a router.

- Action**
- Enter the **run show ospf overview** command from configuration mode.

```
user@host# run show ospf overview
```

```
~
~
~
Restart: Enabled
  Restart duration: 180 sec
  Restart grace period: 210 sec
  Graceful restart helper mode: Enabled
  Restart-signaling helper mode: Enabled
~
~
~
```

Meaning The output shows that graceful restart and both of the helper modes are enabled.

- Related Documentation**
- *Understanding Restart Signaling-Based Helper Mode Support for OSPF Graceful Restart*
 - [Tracing Restart Signaling-Based Helper Mode Events for OSPF Graceful Restart on page 2270](#)
 - *helper-disable (OSPF)*

Configuration Statements: Graceful Restart

disable

Syntax	disable;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (bgp isis ldp ospf ospf3 pim rip ripng rsvp) graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (bgp ldp ospf ospf3 pim) graceful-restart],</p> <p>[edit protocols (bgp esis isis ospf ospf3 ldp pim rip ripng rsvp) graceful-restart],</p> <p>[edit protocols bgp group <i>group-name</i> graceful-restart],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>ip-address</i> graceful-restart],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (bgp ldp ospf ospf3 pim) graceful-restart],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options graceful-restart],</p> <p>[edit routing-options graceful-restart]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Disable graceful restart.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Enabling Graceful Restart on page 2127 • Configuring Routing Protocols Graceful Restart on page 2128 • Configuring Graceful Restart for MPLS-Related Protocols on page 2134 • Configuring VPN Graceful Restart on page 2136 • Configuring Logical System Graceful Restart on page 2137 • Graceful Restart Configuration Statements • Configuring Graceful Restart for QFabric Systems

graceful-restart (Enabling Globally)

Syntax	<pre>graceful-restart { disable; helper-disable; maximum-helper-recovery-time <i>seconds</i>; maximum-helper-restart-time <i>seconds</i>; notify-duration <i>seconds</i>; recovery-time <i>seconds</i>; restart-duration <i>seconds</i>; stale-routes-time <i>seconds</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure graceful restart globally to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally.</p> <p>For VPNs, the graceful-restart statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers.</p> <p>For BGP, if you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.</p>
Default	Graceful restart is disabled by default.
Options	The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling Graceful Restart on page 2127• Configuring Routing Protocols Graceful Restart on page 2128• Configuring Graceful Restart for MPLS-Related Protocols on page 2134• Configuring VPN Graceful Restart on page 2136• Configuring Logical System Graceful Restart on page 2137• Graceful Restart Configuration Statements• Configuring Graceful Restart for QFabric Systems

helper-disable (Multiple Protocols)

Syntax	helper-disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols (isis ldp ospf ospf3 rsvp) graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ldp ospf ospf3) graceful-restart], [edit protocols (isis ldp ospf ospf3 rsvp) graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols (ldp ospf ospf3) graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Disable helper mode for graceful restart. When helper mode is disabled, a router or switch cannot help a neighboring router that is attempting to restart.
Default	Helper mode is enabled by default for these supported protocols: IS-IS, LDP, OSPF/OSPFv3, and RSVP.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Routing Protocols Graceful Restart on page 2128 • Configuring Graceful Restart for MPLS-Related Protocols on page 2134

maximum-helper-recovery-time

Syntax	maximum-helper-recovery-time <i>seconds</i> ;
Hierarchy Level	[edit protocols rsvp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols rsvp graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the length of time the router or switch retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart.
Options	<p><i>seconds</i>—Length of time that the router retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart.</p> <p>Range: 1 through 3600</p> <p>Default: 180</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for RSVP, CCC, and TCC on page 2135 • maximum-helper-restart-time (RSVP) on page 2170

maximum-helper-restart-time (RSVP)

Syntax	maximum-helper-restart-time <i>seconds</i> ;
Hierarchy Level	[edit protocols rsvp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols rsvp graceful-restart]
Release Information	Statement introduced in Junos OS Release 8.3.
Description	Specify the length of time the router or switch waits after it discovers that a neighboring router has gone down before it declares the neighbor down. This value is applied to all RSVP neighbor routers and should be based on the time that the slowest RSVP neighbor requires for restart.
Options	seconds —The time the router or switch waits after it discovers that a neighboring router has gone down before it declares the neighbor down. Range: 1 through 1800 Default: 60
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for RSVP, CCC, and TCC on page 2135• maximum-helper-recovery-time on page 2169

maximum-neighbor-reconnect-time

Syntax	maximum-neighbor-reconnect-time <i>seconds</i> ;
Hierarchy Level	[edit protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced in Junos OS Release 9.1.
Description	Specify the maximum length of time allowed to reestablish connection from a restarting neighbor.
Options	seconds —Maximum time allowed for reconnection. Range: 30 through 300
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for LDP on page 2135

maximum-neighbor-recovery-time

Syntax	<code>maximum-neighbor-recovery-time seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement changed from maximum-recovery-time to maximum-neighbor-recovery-time in Junos OS Release 9.1.
Description	Specify the maximum amount of time to wait before giving up an attempt to gracefully restart.
Options	seconds —Configure the maximum recovery time, in seconds. Range: 120 through 1800 seconds Default: 140 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Recovery Time and Maximum Recovery Time on page 3547 • Configuring Graceful Restart Options for LDP on page 2135 • no-strict-lsa-checking on page 2172 • recovery-time on page 2175

no-strict-lsa-checking

Syntax	no-strict-lsa-checking;
Hierarchy Level	[edit protocols (ospf ospf3) graceful-restart]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router or switch.
Default	By default, LSA checking is enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for OSPF and OSPFv3 on page 2131• <i>Configuring Graceful Restart for QFabric Systems</i>• maximum-neighbor-recovery-time on page 2171• recovery-time on page 2175

notify-duration

Syntax	<code>notify-duration seconds;</code>
Hierarchy Level	<p>[edit protocols (ospf ospf3) graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols (ospf ospf3) graceful-restart],</p> <p>[edit routing-instances <i>instance-name</i> protocols (ospf ospf3) graceful-restart]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Specify the length of time the router or switch notifies helper OSPF routers that it has completed graceful restart.
Options	<p>seconds—Length of time in the router notifies helper OSPF routers that it has completed graceful restart.</p> <p>Range: 1 through 3600</p> <p>Default: 30</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for OSPF and OSPFv3 on page 2131 • Configuring Graceful Restart for QFabric Systems • restart-duration on page 2176

reconnect-time

Syntax	<code>reconnect-time seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced in Junos OS Release 9.1.
Description	Specify the length of time required to reestablish a Label Distribution Protocol (LDP) session after graceful restart.
Options	seconds —Time required for reconnection. Range: 30 through 300 Default: 60 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring LDP Graceful Restart on page 3545 on LDP Configuration Guide• Configuring Graceful Restart Options for LDP on page 2135

recovery-time

Syntax	<code>recovery-time seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the length of time a router or switch waits for Label Distribution Protocol (LDP) neighbors to assist it with a graceful restart.
Options	seconds —Time the router waits for LDP to restart gracefully. Range: 120 through 1800 Default: 160
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for LDP on page 2135 • maximum-neighbor-recovery-time on page 2171 • no-strict-lsa-checking on page 2172

restart-duration

Syntax	<code>restart-duration seconds;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols (isis ospf ospf3 pim) graceful-restart],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3 pim) graceful-restart],</code> <code>[edit protocols (esis isis ospf ospf3 pim) graceful-restart],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3 pim) graceful-restart],</code> <code>[edit routing-options graceful-restart]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure the grace period for graceful restart globally.</p> <p>Additionally, you can individually configure the duration of the graceful restart period for the End System-to-Intermediate System (ES-IS), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), and OSPFv3 protocols and for Protocol Independent Multicast (PIM) sparse mode.</p>
Options	<p>seconds—Time for the graceful restart period.</p> <p>Range:</p> <p>The range of values varies according to whether the graceful restart period is being set globally or for a particular protocol:</p> <ul style="list-style-type: none">• [edit routing-options graceful-restart] (global setting)—120 through 900• ES-IS—30 through 300• IS-IS—30 through 300• OSPF/OSPFv3—1 through 3600• PIM—30 through 300 <p>Default:</p> <p>The default value varies according to whether the graceful restart period is being set globally or for a particular protocol:</p> <ul style="list-style-type: none">• [edit routing-options graceful-restart] (global setting)—300• ES-IS—180• IS-IS—210• OSPF/OSPFv3—180• PIM—60
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Enabling Graceful Restart on page 2127](#)
 - [Configuring Routing Protocols Graceful Restart on page 2128](#)
 - [Configuring Graceful Restart for MPLS-Related Protocols on page 2134](#)
 - [Configuring VPN Graceful Restart on page 2136](#)
 - *Configuring Graceful Restart for VPNs*
 - [Configuring Logical System Graceful Restart on page 2137](#)
 - *Graceful Restart Configuration Statements*
 - *Configuring Graceful Restart for QFabric Systems*

restart-time (BGP Graceful Restart)

Syntax	<code>restart-time seconds;</code>
Hierarchy Level	<p>[edit protocols (bgp rip ripng) graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols (bgp rip ripng) graceful-restart (Enabling Globally)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Configure the duration of the BGP, RIP, or next-generation RIP (RIPng) graceful restart period.
Options	<p>seconds—Length of time for the graceful restart period. Range: 1 through 600 seconds Default: Varies by protocol:</p> <ul style="list-style-type: none"> • BGP—120 seconds • RIP and RIPng—60 seconds
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for BGP on page 2129 • Configuring Graceful Restart Options for RIP and RIPng on page 2132 • <i>Configuring Graceful Restart for QFabric Systems</i> • stale-routes-time on page 2178

stale-routes-time

Syntax	<code>stale-routes-time seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-routing-name</i> protocols bgp graceful-restart], [edit logical-systems <i>logical-routing-name</i> routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart], [edit protocols bgp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols bgp graceful-restart]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify the maximum time that stale routes are kept during a restart. The stale-routes-time statement allows you to set the length of time the routing device waits to receive messages from restarting neighbors before declaring them down.
Options	seconds —Time the router device waits to receive messages from restarting neighbors before declaring them down. Range: 1 through 600 seconds Default: 300 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Graceful Restart Options for BGP on page 2129• Configuring Graceful Restart for QFabric Systems• restart-time (BGP Graceful Restart) on page 2177

traceoptions (Protocols)

Syntax	<pre> traceoptions { file <i>name</i> <size <i>size</i>> <files <i>number</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } </pre>
Hierarchy Level	[edit protocols isis], [edit protocols (ospf ospf3)]
Release Information	Statement introduced before Junos OS Release 7.4. graceful-restart flag for IS-IS and OSPF/OSPFv3 added in Junos OS Release 8.4.
Description	<p>Define tracing operations that graceful restart functionality in the router or switch.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	If you do not include this statement, no global tracing operations are performed.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place global routing protocol tracing output in the file routing-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000 files Default: 2 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. The nonstop active routing tracing option is:</p> <ul style="list-style-type: none"> • graceful-restart—Tracing operations for nonstop active routing <p>no-world-readable—Restrict users from reading the log file.</p> <p>size <i>size</i>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches this size, it is renamed trace-file.0. When the trace-file again reaches its maximum size, trace-file.0 is renamed trace-file.1 and trace-file is renamed trace-file.0. This renaming scheme continues</p>

until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

world-readable—Allow users to read the log file.

Required Privilege Level	routing and trace—To view this statement in the configuration.
	routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Tracking Graceful Restart Events on page 2134

Configuration: NSB

- [Configuring Nonstop Bridging on page 2180](#)
- [Resetting Local Statistics on page 2181](#)

Configuring Nonstop Bridging

This section includes the following topics:

- [Enabling Nonstop Bridging on page 2180](#)
- [Synchronizing the Routing Engine Configuration on page 2180](#)
- [Verifying Nonstop Bridging Operation on page 2181](#)

Enabling Nonstop Bridging

Nonstop bridging requires you to configure graceful Routing Engine switchover (GRES). To enable graceful Routing Engine switchover, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level:

```
[edit chassis redundancy]
graceful-switchover;
```

By default, nonstop bridging is disabled. To enable nonstop bridging, include the **nonstop-bridging** statement at the **[edit protocols layer2-control]** hierarchy level:

```
[edit protocols layer2-control]
nonstop-bridging;
```

To disable nonstop active routing, remove the **nonstop-bridging** statement from the **[edit protocols layer2-control]** hierarchy level.

Synchronizing the Routing Engine Configuration

When you configure nonstop bridging, you must also include the **commit synchronize** statement at the **[edit system]** hierarchy level so that, by default, when you issue the **commit** command, the configuration changes are synchronized on both Routing Engines.

If you issue the **commit synchronize** command at the **[edit]** hierarchy level on the backup Routing Engine, the Junos OS displays a warning and commits the candidate configuration.



NOTE: A newly inserted backup Routing Engine automatically synchronizes its configuration with the master Routing Engine configuration.

When you configure nonstop bridging, you can bring the backup Routing Engine online after the master Routing Engine is already running. There is no requirement to start the two Routing Engines simultaneously.

Verifying Nonstop Bridging Operation

When you enable nonstop bridging, you can issue Layer 2 Control Protocol-related operational mode commands on the backup Routing Engine. However, the output of the commands might not match the output of the same commands issued on the master Routing Engine.

Related Documentation

- [Nonstop Bridging Concepts on page 2069](#)
- [Nonstop Bridging System Requirements on page 2071](#)
- [nonstop-bridging on page 2182](#)
- [Configuring Nonstop Bridging on EX Series Switches \(CLI Procedure\)](#)

Resetting Local Statistics

After a graceful Routing Engine switchover, we recommend that you issue the **clear interface statistics** (*interface-name* | **all**) command to reset the cumulative values for local statistics on the new master Routing Engine.

Related Documentation

- [Configuring Nonstop Active Routing on page 2183](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 2185](#)

Configuration Statements: NSB

- [\[edit protocols layer2-control\] Hierarchy Level on page 2181](#)

[\[edit protocols layer2-control\] Hierarchy Level](#)

The following statement hierarchy can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```
protocols {
  layer2-control {
    bpdu-block {
      disable-timeout seconds;
      interface [ interface-names ];
    }
    mac-rewrite {
      interface interface-name {
        protocol {
```

```
        cdp;
        stp;
        vtp;
    }
}
nonstop-bridging;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
    flag flag <disable>;
}
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

nonstop-bridging

Syntax	nonstop-bridging;
Hierarchy Level	[edit protocols layer2-control]
Release Information	Statement introduced in Junos OS Release 8.4.
Description	For routing platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and preserve Layer 2 Control Protocol (L2CP) information.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Synchronizing the Routing Engine Configuration on page 2183• Configuring Nonstop Bridging on page 2180• <i>Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)</i>

Configuration: NSR

- [Configuring Nonstop Active Routing on page 2183](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 2185](#)
- [Example: Configuring Nonstop Active Routing on page 2189](#)

Configuring Nonstop Active Routing

This section includes the following topics:

- [Enabling Nonstop Active Routing on page 2183](#)
- [Synchronizing the Routing Engine Configuration on page 2183](#)
- [Verifying Nonstop Active Routing Operation on page 2184](#)

Enabling Nonstop Active Routing

Nonstop active routing (NSR) requires you to configure graceful Routing Engine switchover (GRES). To enable graceful Routing Engine switchover, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level:

```
[edit chassis redundancy]
graceful-switchover;
```

By default, nonstop active routing is disabled. To enable nonstop active routing, include the **nonstop-routing** statement at the **[edit routing-options]** hierarchy level:

```
[edit routing-options]
nonstop-routing;
```

To disable nonstop active routing, remove the **nonstop-routing** statement from the **[edit routing-options]** hierarchy level.



NOTE: When you enable nonstop active routing, you cannot enable automatic route distinguishers for multicast VPN routing instances. Automatic route distinguishers are enabled by configuring the **route-distinguisher-id** statement at the **[edit routing-instances instance-name]** hierarchy level; for more information, see the *Junos OS VPNs Configuration Guide*.

To enable the routing platform to switch over to the backup Routing Engine when the routing protocol process (rpd) fails rapidly three times in succession, include the **other-routing-engine** statement at the **[edit system processes routing failover]** hierarchy level.

For more information about the **other-routing-engine** statement, see the *Junos OS System Basics Configuration Guide*.

Synchronizing the Routing Engine Configuration

When you configure nonstop active routing, you must also include the **commit synchronize** statement at the **[edit system]** hierarchy level so that configuration changes are synchronized on both Routing Engines:

```
[edit system]
commit synchronize;
```

If you try to commit the nonstop active routing configuration without including the **commit synchronize** statement, the commit fails.

If you configure the **commit synchronize** statement at the **[edit system]** hierarchy level and issue a commit in the master Routing Engine, the master configuration is automatically synchronized with the backup.

However, if the backup Routing Engine is down when you issue the commit, the Junos OS displays a warning and commits the candidate configuration in the master Routing Engine. When the backup Routing Engine comes up, its configuration will automatically be synchronized with the master.



NOTE: A newly inserted backup Routing Engine automatically synchronizes its configuration with the master Routing Engine configuration.

When you configure nonstop active routing, you can bring the backup Routing Engine online after the master Routing Engine is already running. There is no requirement to start the two Routing Engines simultaneously.

Verifying Nonstop Active Routing Operation

To see whether or not nonstop active routing is enabled, issue the **show task replication** command. For BGP nonstop active routing, you can also issue the **show bgp replication** command.

For more information about these commands, see the *Junos OS Operational Mode Commands* and *Junos OS Operational Mode Commands*, respectively.

When you enable nonstop active routing or graceful Routing Engine switchover and issue routing-related operational mode commands on the backup Routing Engine (such as **show route**, **show bgp neighbor**, **show ospf database**, and so on), the output might not match the output of the same commands issued on the master Routing Engine. For example, it is normal for the routing table on the backup Routing Engine to contain persistent phantom routes that are not present in the routing table on the master Routing Engine.

To display BFD state replication status, issue the **show bfd session** command. The **replicated** flag appears in the output for this command when a BFD session has been replicated to the backup Routing Engine. For more information, see the *Junos OS Operational Mode Commands*.

Related Documentation

- [Nonstop Active Routing Concepts on page 2072](#)
- [Nonstop Active Routing System Requirements on page 2075](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 2185](#)
- [Resetting Local Statistics on page 2181](#)
- [Example: Configuring Nonstop Active Routing on page 2189](#)
- [nonstop-routing on page 2194](#)

Tracing Nonstop Active Routing Synchronization Events

To track the progress of nonstop active routing synchronization between Routing Engines, you can configure nonstop active routing trace options flags for each supported protocol and for BFD sessions and record these operations to a log file.

To configure nonstop active routing trace options for supported routing protocols, include the **nsr-synchronization** statement at the **[edit protocols *protocol-name* traceoptions flag]** hierarchy level and optionally specify one or more of the **detail**, **disable**, **receive**, and **send** options:

```
[edit protocols]
bgp {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
isis {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
ldp {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
mpls {
  traceoptions {
    flag nsr-synchronization;
    flag nsr-synchronization-detail;
  }
}
msdp {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
(ospf | ospf3) {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
rip {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
ripng {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
pim {
  traceoptions {
```

```
        flag nsr-synchronization <detail> <disable> <receive> <send>;
    }
}
```

To configure nonstop active routing trace options for BFD sessions, include the **nsr-synchronization** and **nsr-packet** statements at the **[edit protocols bfd traceoptions flag]** hierarchy level.

```
[edit protocols]
bfd {
  traceoptions {
    flag nsr-synchronization;
    flag nsr-packet;
  }
}
```

To trace the Layer 2 VPN signaling state replicated from routes advertised by BGP, include the **nsr-synchronization** statement at the **[edit routing-options traceoptions flag]** hierarchy level. This flag also traces the label and logical interface association that VPLS receives from the kernel replication state.

```
[edit routing-options]
traceoptions {
  flag nsr-synchronization;
}
```

Related Documentation

- [Configuring Nonstop Active Routing on page 2183](#)
- *Configuring Nonstop Active Routing on EX Series Switches (CLI Procedure)*
- [Example: Configuring Nonstop Active Routing on page 2189](#)
- *Example: Configuring Nonstop Active Routing on EX Series Switches*

traceoptions (Routing Options)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <disable>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options],</p> <p>[edit routing-options flow],</p> <p>[edit routing-options multicast]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>nsr-synchronization flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.</p> <p>nsr-synchronization and nsr-packet flags for BFD sessions added in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>nsr-synchronization flag for RIP and RIPng added in Junos OS Release 9.0.</p> <p>nsr-synchronization flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.</p> <p>nsr-synchronization flag for PIM added in Junos OS Release 9.3.</p> <p>nsr-synchronization flag for MPLS added in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>nsr-synchronization flag for MSDP added in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Define tracing operations that track all routing protocol functionality in the routing device.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	If you do not include this statement, no global tracing operations are performed.
Options	<p>Values:</p> <p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place global routing protocol tracing output in the file routing-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and</p>

so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the global routing protocol tracing options:

- **all**—All tracing operations
- **condition-manager**—Condition-manager events
- **config-internal**—Configuration internals
- **general**—All normal operations and routing table changes (a combination of the **normal** and **route** trace operations)
- **graceful-restart**—Graceful restart operations
- **normal**—All normal operations
- **nsr-packet**—Detailed trace information for BFD nonstop active routing only
- **nsr-synchronization**—Tracing operations for nonstop active routing
- **nsr-synchronization-detail**—(MPLS only) Tracing operations for nonstop active routing in detail
- **parse**—Configuration parsing
- **policy**—Routing policy operations and actions
- **regex-parse**—Regular-expression parsing
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

no-world-readable—(Optional) Prevent any user from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

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

Related Documentation

- *Example: Tracing Global Routing Protocol Operations*
- [Tracing Nonstop Active Routing Synchronization Events on page 2185](#)

Example: Configuring Nonstop Active Routing

The following example enables graceful Routing Engine switchover, nonstop active routing, and nonstop active routing trace options for BGP, IS-IS, and OSPF.

```
[edit]
system commit {
  synchronize;
}
chassis {
  redundancy {
    graceful-switchover; # This enables graceful Routing Engine switchover on
                        # the routing platform.
  }
}
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        address 10.0.1.1/30;
      }
      family iso;
    }
  }
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.1.1.1/30;
      }
      family iso;
    }
  }
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.2.1.1/30;
      }
      family iso;
    }
  }
  so-0/0/3 {
    unit 0 {
      family inet {
        address 10.3.1.1/30;
      }
      family iso;
    }
  }
}
```

```
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 192.168.2.1/32;
      }
      family iso {
        address 49.0004.1921.6800.2001.00;
      }
    }
  }
}
routing-options {
  nonstop-routing; # This enables nonstop active routing on the routing platform.
  router-id 192.168.2.1;
  autonomous-system 65432;
}
protocols {
  bgp {
    traceoptions {
      flag nsr-synchronization detail; # This logs nonstop active routing
      # events for BGP.
    }
    local-address 192.168.2.1;
    group external-group {
      type external;
      export BGP_export;
      neighbor 192.168.1.1 {
        family inet {
          unicast;
        }
        peer-as 65103;
      }
    }
    group internal-group {
      type internal;
      neighbor 192.168.10.1;
      neighbor 192.168.11.1;
      neighbor 192.168.12.1;
    }
  }
}
isis {
  traceoptions {
    flag nsr-synchronization detail; # This logs nonstop active routing events
    # for IS-IS.
  }
  interface all;
  interface fxp0.0 {
    disable;
  }
  interface lo0.0 {
    passive;
  }
}
ospf {
```

```

traceoptions {
    flag nsr-synchronization detail; # This logs nonstop active routing events
    # for OSPF.
}
area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
        disable;
    }
    interface lo0.0 {
        passive;
    }
}
}
}
policy-options {
    policy-statement BGP_export {
        term direct {
            from {
                protocol direct;
            }
            then accept;
        }
        term final {
            then reject;
        }
    }
}
}

```

- Related Documentation**
- [Configuring Nonstop Active Routing on page 2183](#)
 - [Tracing Nonstop Active Routing Synchronization Events on page 2185](#)

Configuration Statements: NSR

- [\[edit protocols layer2-control\] Hierarchy Level on page 2191](#)

[\[edit protocols layer2-control\] Hierarchy Level](#)

The following statement hierarchy can also be included at the [\[edit logical-systems *logical-system-name*\]](#) hierarchy level.

```

protocols {
    layer2-control {
        bpdu-block {
            disable-timeout seconds;
            interface [ interface-names ];
        }
        mac-rewrite {
            interface interface-name {
                protocol {
                    cdp;
                    stp;
                    vtp;
                }
            }
        }
    }
}

```

```
    }  
  }  
  nonstop-bridging;  
  traceoptions {  
    file filename <files number> <size maximum-file-size> <world-readable |  
      no-world-readable>;  
    flag flag <disable>;  
  }  
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

commit synchronize

Syntax	commit synchronize;
Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 10.4 for EX Series switches.
Description	For devices with multiple Routing Engines only. Configure the commit command to automatically result in a commit synchronize action between dual Routing Engines within the same chassis. The Routing Engine on which you execute the commit command (the requesting Routing Engine) copies and loads its candidate configuration to the other (the responding) Routing Engine. Each Routing Engine then performs a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines.



NOTE: When you configure nonstop active routing (NSR), you must include the **commit synchronize** statement. Otherwise, the commit operation fails.

On the TX Matrix router, synchronization only occurs between the Routing Engines within the same chassis and when synchronization is complete, the new configuration is then distributed to the Routing Engines on the T640 routers. That is, the master Routing Engine on the TX Matrix router distributes the configuration to the master Routing Engine on each T640 router. Likewise, the backup Routing Engine on the TX Matrix router distributes the configuration to the backup Routing Engine on each T640 router.

In EX Series Virtual Chassis configurations:

- On EX4200 switches in Virtual Chassis, synchronization occurs between the switch in the master role and the switch in the backup role.
- On EX8200 switches in a Virtual Chassis, synchronization occurs only between the master and backup XRE200 External Routing Engines.

Options	<p>and-quit—(Optional) (EX Series only) Quit configuration mode if the commit synchronization succeeds.</p> <p>comment—(Optional) (EX Series only) Write a message to the commit log.</p> <p>and-force—(Optional) (EX Series only) Force a commit synchronization on the other Routing Engine (ignore warnings).</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Synchronizing the Routing Engine Configuration on page 2183

- *Configuring Multiple Routing Engines to Synchronize Committed Configurations Automatically*

nonstop-routing

Syntax	nonstop-routing;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced in Junos OS Release 8.4. Statement introduced in Junos OS Release 10.4 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	For routing platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and to preserve routing protocol information.
Default	disabled
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Nonstop Active Routing on page 2183• <i>Configuring Nonstop Active Routing on EX Series Switches (CLI Procedure)</i>

traceoptions (Routing Options)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <disable>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options],</p> <p>[edit routing-options flow],</p> <p>[edit routing-options multicast]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>nsr-synchronization flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.</p> <p>nsr-synchronization and nsr-packet flags for BFD sessions added in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>nsr-synchronization flag for RIP and RIPng added in Junos OS Release 9.0.</p> <p>nsr-synchronization flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.</p> <p>nsr-synchronization flag for PIM added in Junos OS Release 9.3.</p> <p>nsr-synchronization flag for MPLS added in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>nsr-synchronization flag for MSDP added in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Define tracing operations that track all routing protocol functionality in the routing device.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	If you do not include this statement, no global tracing operations are performed.
Options	<p>Values:</p> <p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place global routing protocol tracing output in the file routing-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and</p>

so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the global routing protocol tracing options:

- **all**—All tracing operations
- **condition-manager**—Condition-manager events
- **config-internal**—Configuration internals
- **general**—All normal operations and routing table changes (a combination of the **normal** and **route** trace operations)
- **graceful-restart**—Graceful restart operations
- **normal**—All normal operations
- **nsr-packet**—Detailed trace information for BFD nonstop active routing only
- **nsr-synchronization**—Tracing operations for nonstop active routing
- **nsr-synchronization-detail**—(MPLS only) Tracing operations for nonstop active routing in detail
- **parse**—Configuration parsing
- **policy**—Routing policy operations and actions
- **regex-parse**—Regular-expression parsing
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

no-world-readable—(Optional) Prevent any user from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Tracing Global Routing Protocol Operations • Tracing Nonstop Active Routing Synchronization Events on page 2185

Configuration: Unified ISSU

- [Best Practices on page 2197](#)
- [Before You Begin on page 2197](#)
- [Performing a Unified ISSU on page 2200](#)
- [Verifying a Unified ISSU on page 2212](#)
- [Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 2212](#)

Best Practices

When you are planning to perform a unified in-service software upgrade (unified ISSU), choose a time when your network is as stable as possible. As with a normal upgrade, Telnet sessions, SNMP, and CLI access are briefly interrupted. In addition, the following restrictions apply:

- The master Routing Engine and backup Routing Engine must be running the same software version before you can perform a unified ISSU.
- During a unified ISSU, you cannot bring any PICs online or offline.
- Unicast RPF-related statistics are not saved across a unified ISSU, and the unicast RPF counters are reset to zero during a unified ISSU.

Related Documentation	<ul style="list-style-type: none"> • Before You Begin on page 2197 • Performing a Unified ISSU on page 2200 • Verifying a Unified ISSU on page 2212 • Troubleshooting Unified ISSU Problems on page 2278
------------------------------	--

Before You Begin

Before you begin a unified ISSU, complete the tasks in the following sections:

1. [Verify That the Master and Backup Routing Engines Are Running the Same Software Version on page 2198](#)
2. [Back Up the Router Software on page 2198](#)
3. [Verify That Graceful Routing Engine Switchover and Nonstop Active Routing Are Configured on page 2199](#)

Verify That the Master and Backup Routing Engines Are Running the Same Software Version

To verify that both Routing Engines are running the same version of software, issue the following command:

```
{master}
user@host> show version invoke-on all-routing-engines
re0:
-----
Hostname: host
Model: m320
JUNOS Base OS boot [9.0-20071210.0]
JUNOS Base OS Software Suite [9.0-20071210.0]
JUNOS Kernel Software Suite [9.0-20071210.0]
JUNOS Crypto Software Suite [9.0-20071210.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20071210.0]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20071210.0]
JUNOS Online Documentation [9.0-20071210.0]
JUNOS Routing Software Suite [9.0-20071210.0]
re1:
-----
Hostname: host
Model: m320
JUNOS Base OS boot [9.0-20071210.0]
JUNOS Base OS Software Suite [9.0-20071210.0]
JUNOS Kernel Software Suite [9.0-20071210.0]
JUNOS Crypto Software Suite [9.0-20071210.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20071210.0]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20071210.0]
JUNOS Online Documentation [9.0-20071210.0]
JUNOS Routing Software Suite [9.0-20071210.0]
```

If both Routing Engines are not running the same software version, issue the **request system software add** command on the desired Routing Engine so that the software version is the same. For more information, see the *Installation and Upgrade Guide*.

Back Up the Router Software

As a preventive measure in case any problems occur during an upgrade, issue the **request system snapshot** command on *each* Routing Engine to back up the system software to the router's hard disk. The following is an example of issuing the command on the master Routing Engine:

```
{master}
user@host> request system snapshot
Verifying compatibility of destination media partitions...
Running newfs (220MB) on hard-disk media / partition (ad1s1a)...
Running newfs (24MB) on hard-disk media /config partition (ad1s1e)...
Copying '/dev/ad0s1a' to '/dev/ad1s1a' .. (this may take a few minutes)
Copying '/dev/ad0s1e' to '/dev/ad1s1e' .. (this may take a few minutes)
The following filesystems were archived: / /config
```



NOTE: The root file system is backed up to `/altroot`, and `/config` is backed up to `/altconfig`. After you issue the `request system snapshot` command, the router's flash and hard disks are identical. You can return to the previous version of the software only by booting the router from removable media. For more information about the `request system snapshot` command, see the *Junos OS System Basics Configuration Guide*.

Verify That Graceful Routing Engine Switchover and Nonstop Active Routing Are Configured

Before you begin a unified ISSU, ensure that graceful Routing Engine switchover and nonstop active routing are configured on your router.

1. To verify graceful Routing Engine switchover is configured, on the backup Routing Engine (**re1**) issue the `show system switchover` command. The output should be similar to the following example. The **Graceful switchover** field state must be **On**.

```
{backup}
user@host> show system switchover
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
```

2. To verify nonstop active routing is configured, on the master Routing Engine (**re0**) issue the `show task replication` command. The output should be similar to the following example.

```
{master}
user@host> show task replication
Stateful Replication: Enabled
RE mode: Master

Protocol                Synchronization Status
OSPF                    Complete
IS-IS                   Complete
```

If graceful Routing Engine switchover and nonstop active routing are not configured, complete the following steps:

1. On the master Routing Engine (**re0**), enable graceful Routing Engine switchover. Include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level.
2. On the master Routing Engine, enable nonstop active routing. Include the **commit synchronize** statement at the **[edit system]** hierarchy level and the **nonstop-routing** statement at the **[edit routing-options]** hierarchy level.
3. On the master Router Engine, issue the **commit** command.

The system provides the following confirmation that the master and backup Routing Engines are synchronized:

```
re0:
configuration check succeeds
```

```
rel:
commit complete
re0:
commit complete
```

Related Documentation

- [Unified ISSU Concepts on page 2092](#)
- [Unified ISSU Process on the TX Matrix Router](#)
- [Unified ISSU System Requirements on page 2097](#)
- [Best Practices on page 2197](#)
- [Performing a Unified ISSU on page 2200](#)
- [Verifying a Unified ISSU on page 2212](#)
- [Troubleshooting Unified ISSU Problems on page 2278](#)

Performing a Unified ISSU

You can perform a unified ISSU in one of three ways:

1. [Upgrading and Rebooting Both Routing Engines Automatically on page 2200](#)
2. [Upgrading Both Routing Engines and Rebooting the New Backup Routing Engine Manually on page 2204](#)
3. [Upgrading and Rebooting Only One Routing Engine on page 2209](#)

Upgrading and Rebooting Both Routing Engines Automatically

When you issue the **request system software in-service-upgrade** command with the **reboot** option, the system automatically upgrades both Routing Engines to the newer software and reboots both Routing Engines. This option enables you to complete the unified ISSU with a single command.

To perform a unified ISSU using the **request system software in-service-upgrade package-name reboot** command, complete the following steps:

1. Download the software package from the Juniper Networks Support website, <http://www.juniper.net/support/>. Choose the Canada and U.S., Worldwide, or Junos-FIPS edition. Place the package on a local server. To download the package, you must have a service contract and an access account. If you do not have an access account, complete the registration form at the Juniper Networks website: <https://www.juniper.net/registration/Register.jsp>.
2. Copy the package to the router. We recommend that you copy it to the **/var/tmp** directory, which is a large file system on the hard disk.

user@host>file copy ftp://username:prompt@ftp.hostname.net/filename/var/tmp/filename
3. To verify the current software version running on both Routing Engines, on the master Routing Engine issue the **show version invoke-on all-routing-engines** command. The following example shows that both Routing Engines are running an image of Junos OS, Release 9.0, that was built on December 11, 2007:

```
{backup}
```

```
user@host> show version invoke-on all-routing-engines
re0:
```

```
-----
Hostname: host
Model: m320
JUNOS Base OS boot [9.0-20071211.2]
JUNOS Base OS Software Suite 9.0-20071211.2]
JUNOS Kernel Software Suite [9.0-20071211.2]
JUNOS Crypto Software Suite [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20071211.2]
JUNOS Online Documentation [9.0-20071211.2]
JUNOS Routing Software Suite [9.0-20071211.2]
```

```
re1:
```

```
-----
Hostname: host1
Model: m320
JUNOS Base OS boot [9.0-20071211.2]
JUNOS Base OS Software Suite [9.0-20071211.2]
JUNOS Kernel Software Suite [9.0-20071211.20]
JUNOS Crypto Software Suite [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20071211.2]
JUNOS Online Documentation [9.0-20071211.2]
JUNOS Routing Software Suite [9.0-20071211.2]
```

4. On the master Routing Engine, issue the **request system software in-service-upgrade package-name reboot** command. The following example upgrades the current version to an image of Junos OS, Release 9.0, that was built on January 14, 2008:

```
{master}
```

```
user@host> request system software in-service-upgrade
/var/tmp/jinstall-9.0-20080114.2-domestic-signed.tgz reboot
ISSU: Validating Image
PIC 0/3 will be offlined (In-Service-Upgrade not supported)
Do you want to continue with these actions being taken ? [yes,no] (no) yes

ISSU: Preparing Backup RE
Pushing bundle to re1
Checking compatibility with configuration
Initializing...
Using jbase-9.0-20080114.2
Verified manifest signed by PackageProduction_9_0_0
Using /var/tmp/jinstall-9.0-20080114.2-domestic-signed.tgz
Verified jinstall-9.0-20080114.2-domestic.tgz signed by PackageProduction_9_0_0
Using jinstall-9.0-20080114.2-domestic.tgz
Using jbundle-9.0-20080114.2-domestic.tgz
Checking jbundle requirements on /
Using jbase-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jkernel-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jcrypto-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jpfe-9.0-20080114.2.tgz
Using jdocs-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jroute-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Hardware Database regeneration succeeded
```

```
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-9.0-20080114.2-domestic-signed.tgz' ...
Verified jinstall-9.0-20080114.2-domestic.tgz signed by PackageProduction_9_0_0
Adding jinstall...
Verified manifest signed by PackageProduction_9_0_0

WARNING: This package will load JUNOS 9.0-20080114.2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-9.0-20080114.2-domestic-signed.tgz
...
Saving state for rollback ...
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU started
ISSU: Backup RE Prepare Done
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
  FPC 1         Online (ISSU)
  FPC 2         Online (ISSU)
  FPC 6         Online (ISSU)
  FPC 7         Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
Installing package '/var/tmp/paKEuy' ...
Verified jinstall-9.0-20080114.2-domestic.tgz signed by PackageProduction_9_0_0
Adding jinstall...
Verified manifest signed by PackageProduction_9_0_0
```



```

WARNING: This package will load JUNOS 9.0-20080114.2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

```

```

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

```

```

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

```

```

Saving package file in /var/sw/pkg/jinstall-9.0-20080114.2-domestic-signed.tgz
...
cp: /var/tmp/paKEuy is a directory (not copied).
Saving state for rollback ...
ISSU: Old Master Upgrade Done
ISSU: IDLE
Shutdown NOW!
Reboot consistency check bypassed - jinstall 9.0-20080114.2 will complete
installation upon reboot
[pid 30227]

```

```

*** FINAL System shutdown message from root@host ***

```

```

System going down IMMEDIATELY

```

Connection to host closed.

When the new backup (old master) Routing Engine is rebooted, you are logged out from the router.

5. After waiting a few minutes, log in to the router again. You are logged in to the new backup Routing Engine (**re0**). To verify that both Routing Engines have been upgraded, issue the following command:

```
{backup}
```

```

user@host> show version invoke-on all-routing-engines
re0:

```

```

-----
Hostname: host
Model: m320
JUNOS Base OS boot [9.0-20080114.2]
JUNOS Base OS Software Suite 9.0-20080114.2]
JUNOS Kernel Software Suite [9.0-20080114.2]
JUNOS Crypto Software Suite [9.0-20080114.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20080114.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20080114.2]
JUNOS Online Documentation [9.0-20080114.2]
JUNOS Routing Software Suite [9.0-20080114.2]

```

```

re1:
-----

```

```
Hostname: host1
Model: m320
JUNOS Base OS boot [9.0-20080114.2]
JUNOS Base OS Software Suite [9.0-20080114.2]
JUNOS Kernel Software Suite [9.0-20080114.2]
JUNOS Crypto Software Suite [9.0-20080114.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20080114.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20080114.2]
JUNOS Online Documentation [9.0-20080114.2]
JUNOS Routing Software Suite [9.0-20080114.2]
```

6. To make **re0** the master Routing Engine, issue the following command:

```
{backup}

user@host> request chassis routing-engine master acquire
Attempt to become the master routing engine ? [yes,no] (no) yes

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

{master}
user@host>
```

7. Issue the **request system snapshot** command on *each* Routing Engine to back up the system software to the router's hard disk.



NOTE: The root file system is backed up to **/altroot**, and **/config** is backed up to **/altconfig**. After you issue the **request system snapshot** command, the router's flash and hard disks are identical. You can return to the previous version of the software only by booting the router from removable media.

Upgrading Both Routing Engines and Rebooting the New Backup Routing Engine Manually

When you issue the **request system software in-service-upgrade** command without any options, the system upgrades and reboots the new master Routing Engine to the newer software. The new software is placed on the new backup (old master) Routing Engine; however, to complete the upgrade, you must issue the **request system reboot** command on the new backup Routing Engine.

To perform a unified ISSU using the **request system software in-service-upgrade package-name** command without any options, complete the following steps:

1. Download the software package from the Juniper Networks Support website, <http://www.juniper.net/support/>. Choose the Canada and U.S., Worldwide, or Junos-FIPS edition. Place the package on a local server. To download the package, you must have a service contract and an access account. If you do not have an access account, complete the registration form at the Juniper Networks website: <https://www.juniper.net/registration/Register.jsp>.
2. Copy the package to the router. We recommend that you copy it to the **/var/tmp** directory, which is a large file system on the hard disk.

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

3. To verify the current software version running on both Routing Engines, on the master Routing Engine, issue the **show version invoke-on all-routing-engines** command. The following example shows that both Routing Engines are running Junos OS Release 9.0R1:

```
{master}
```

```
user@host> show version invoke-on all-routing-engines
re0:
```

```
-----
Hostname: host
Model: m320
JUNOS Base OS boot [9.0R1]
JUNOS Base OS Software Suite [9.0R1]
JUNOS Kernel Software Suite [9.0R1]
JUNOS Crypto Software Suite [9.0R1]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0R1]
JUNOS Packet Forwarding Engine Support (M320) [9.0R1]
JUNOS Online Documentation [9.0R1]
JUNOS Routing Software Suite [9.0R1]
```

```
re1:
```

```
-----
Hostname: host1
Model: m320
JUNOS Base OS boot [9.0R1]
JUNOS Base OS Software Suite [9.0R1]
JUNOS Kernel Software Suite [9.0R1]
JUNOS Crypto Software Suite [9.0R1]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0R1]
JUNOS Packet Forwarding Engine Support (M320) [9.0R1]
JUNOS Online Documentation [9.0R1]
JUNOS Routing Software Suite [9.0R1]
```

4. On the master Routing Engine, issue the **request system software in-service-upgrade package-name** command. The following example upgrades the current version to Junos OS Release 9.0R1.2:

```
user@host> request system software in-service-upgrade
/var/tmp/jinstall-9.0R1.2-domestic-signed.tgz
ISSU: Validating Image
FPC 4 will be offlined (In-Service-Upgrade not supported)
Do you want to continue with these actions being taken ? [yes,no] (no) yes
```

```
ISSU: Preparing Backup RE
Pushing bundle to re1
Checking compatibility with configuration
Initializing...
Using jbase-9.0-20080117.0
Verified manifest signed by PackageProduction_9_0_0
Using /var/tmp/jinstall-9.0R1.2-domestic-signed.tgz
Verified jinstall-9.0R1.2-domestic.tgz signed by PackageProduction_9_0_0
Using jinstall-9.0R1.2-domestic.tgz
Using jbundle-9.0R1.2-domestic.tgz
Checking jbundle requirements on /
Using jbase-9.0R1.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jkernel-9.0R1.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jcrypto-9.0R1.2.tgz
```

```
Verified manifest signed by PackageProduction_9_0_0
Using jpf-9.0R1.2.tgz
Using jdocs-9.0R1.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jroute-9.0R1.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-9.0R1.2-domestic-signed.tgz' ...
Verified jinstall-9.0R1.2-domestic.tgz signed by PackageProduction_9_0_0
Adding jinstall...
Verified manifest signed by PackageProduction_9_0_0
```

```
WARNING: This package will load JUNOS 9.0R1.2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.
```

```
Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...
```

```
WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.
```

```
Saving package file in /var/sw/pkg/jinstall-9.0R1.2-domestic-signed.tgz ...
Saving state for rollback ...
Backup upgrade done
Rebooting Backup RE
```

```
Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU started
ISSU: Backup RE Prepare Done
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
```

Item	Status	Reason
FPC 0	Online (ISSU)	
FPC 1	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 3	Online (ISSU)	
FPC 4	Offline	Offlined by cli command
FPC 5	Online (ISSU)	

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

```

ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
Installing package '/var/tmp/paeBi5' ...
Verified jinstall-9.0R1.2-domestic.tgz signed by PackageProduction_9_0_0
Adding jinstall...
Verified manifest signed by PackageProduction_9_0_0

```

```

WARNING: This package will load JUNOS 9.0R1.2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

```

```

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

```

```

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

```

```

Saving package file in /var/sw/pkg/jinstall-9.0R1.2-domestic-signed.tgz ...
cp: /var/tmp/paeBi5 is a directory (not copied).
Saving state for rollback ...
ISSU: Old Master Upgrade Done
ISSU: IDLE

```

5. Issue the **show version invoke-on all-routing-engines** command to verify that the new backup (old master) Routing Engine (**re0**), is still running the previous software image, while the new master Routing Engine (**re1**) is running the new software image:

```
{backup}
```

```

user@host> show version
re0:

```

```

-----
Hostname: user
Model: m320
JUNOS Base OS boot [9.0R1]
JUNOS Base OS Software Suite [9.0R1]
JUNOS Kernel Software Suite [9.0R1]
JUNOS Crypto Software Suite [9.0R1]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0R1]
JUNOS Packet Forwarding Engine Support (M320) [9.0R1]
JUNOS Online Documentation [9.0R1]
JUNOS Routing Software Suite [9.0R1]
labpkg [7.0]
JUNOS Installation Software [9.0R1.2]

```

```
re1:
```

```

-----
Hostname: user1
Model: m320
JUNOS Base OS boot [9.0R1.2]
JUNOS Base OS Software Suite [9.0R1.2]
JUNOS Kernel Software Suite [9.0R1.2]
JUNOS Crypto Software Suite [9.0R1.2]

```

```
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0R1.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0R1.2]
JUNOS Online Documentation [9.0R1.2]
JUNOS Routing Software Suite [9.0R1.2]
```

6. At this point, if you choose not to install the newer software version on the new backup Routing Engine (**re1**), you can issue the **request system software delete jinstall** command on it. Otherwise, to complete the upgrade, go to the next step.
7. Reboot the new backup Routing Engine (**re0**) by issuing the **request system reboot** command:

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

Shutdown NOW!
Reboot consistency check bypassed - jinstall 9.0R1.2 will complete installation
upon reboot
[pid 6170]

{backup}
user@host>
System going down IMMEDIATELY
```

```
Connection to host closed by remote host.
Connection to host closed.
```

If you are not on the console port, you are disconnected from the router session.

8. After waiting a few minutes, log in to the router again. You are logged in to the new backup Routing Engine (**re0**). To verify that both Routing Engines have been upgraded, issue the following command:

```
{backup}
user@host> show version invoke-on all-routing-engines
re0:
-----
Hostname: host
Model: m320
JUNOS Base OS boot [9.0R1.2]
JUNOS Base OS Software Suite [9.0R1.2]
JUNOS Kernel Software Suite [9.0R1.2]
JUNOS Crypto Software Suite [9.0R1.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0R1.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0R1.2]
JUNOS Online Documentation [9.0R1.2]
JUNOS Routing Software Suite [9.0R1.2]

re1:
-----
Hostname: host1
Model: m320
JUNOS Base OS boot [9.0R1.2]
JUNOS Base OS Software Suite [9.0R1.2]
JUNOS Kernel Software Suite [9.0R1.2]
JUNOS Crypto Software Suite [9.0R1.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0R1.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0R1.2]
```

JUNOS Online Documentation [9.0R1.2]
 JUNOS Routing Software Suite [9.0R1.2]

9. To make **re0** the master Routing Engine, issue the following command:

```
{backup}

user@host> request chassis routing-engine master acquire
Attempt to become the master routing engine ? [yes,no] (no) yes

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

{master}
user@host>
```

10. Issue the **request system snapshot** command on *each* Routing Engine to back up the system software to the router's hard disk.



NOTE: The root file system is backed up to **/altroot**, and **/config** is backed up to **/altconfig**. After you issue the **request system snapshot** command, the router's flash and hard disks are identical. You can return to the previous version of the software only by booting the router from removable media.

Upgrading and Rebooting Only One Routing Engine

When you issue the **request system software in-service-upgrade** command with the **no-old-master-upgrade** option, the system upgrades and reboots only the new master Routing Engine. To upgrade the new backup (former master) Routing Engine, you must issue the **request system software add** command.

To perform a unified ISSU using the **request system software in-service-upgrade package-name no-old-master-upgrade** commands, complete the following steps:

1. Download the software package from the Juniper Networks Support website, <http://www.juniper.net/support/>. Choose the Canada and U.S., Worldwide, or Junos-FIPS edition. Place the package on a local server. To download the package, you must have a service contract and an access account. If you do not have an access account, complete the registration form at the Juniper Networks website: <https://www.juniper.net/registration/Register.jsp>.

2. Copy the package to the router. We recommend that you copy it to the **/var/tmp** directory, which is a large file system on the hard disk.

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

3. To verify the current software version running on both Routing Engines, on the master Routing Engine issue the **show version invoke-on all-routing-engines** command. The following example shows that both Routing Engines are running an image of Junos OS Release 9.0 that was built on December 11, 2007:

```
{backup}

user@host> show version invoke-on all-routing-engines
```

```
re0:
```

```
-----
Hostname: host
Model: m320
JUNOS Base OS boot [9.0-20071211.2]
JUNOS Base OS Software Suite 9.0-20071211.2]
JUNOS Kernel Software Suite [9.0-20071211.2]
JUNOS Crypto Software Suite [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20071211.2]
JUNOS Online Documentation [9.0-20071211.2]
JUNOS Routing Software Suite [9.0-20071211.2]
```

```
re1:
```

```
-----
Hostname: host1
Model: m320
JUNOS Base OS boot [9.0-20071211.2]
JUNOS Base OS Software Suite [9.0-20071211.2]
JUNOS Kernel Software Suite [9.0-20071211.20]
JUNOS Crypto Software Suite [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.0-20071211.2]
JUNOS Packet Forwarding Engine Support (M320) [9.0-20071211.2]
JUNOS Online Documentation [9.0-20071211.2]
JUNOS Routing Software Suite [9.0-20071211.2]
```

4. On the master Routing Engine, issue the **request system software in-service-upgrade package-name no-old-master-upgrade** command. The following example upgrades the current version to an image of Junos OS Release 9.0 that was built on January 16, 2008:

```
{master}
```

```
user@host> request system software in-service-upgrade
/var/tmp/jinstall-9.0-20080116.2-domestic-signed.tgz no-old-master-upgrade
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re1
Checking compatibility with configuration
Initializing...
Using jbase-9.0-20080116.2
Verified manifest signed by PackageProduction_9_0_0
Using /var/tmp/jinstall-9.0-20080116.2-domestic-signed.tgz
Verified jinstall-9.0-20080116.2-domestic.tgz signed by PackageProduction_9_0_0
Using jinstall-9.0-20080116.2-domestic.tgz
Using jbundle-9.0-20080116.2-domestic.tgz
Checking jbundle requirements on /
Using jbase-9.0-20080116.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jkernel-9.0-20080116.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jcrypto-9.0-20080116.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jpfe-9.0-20080116.2.tgz
Using jdocs-9.0-20080116.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jroute-9.0-20080116.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
```



```

Validation succeeded
Installing package '/var/tmp/jinstall-9.0-20080116.2-domestic-signed.tgz' ...
Verified jinstall-9.0-20080116.2-domestic.tgz signed by PackageProduction_9_0_0
Adding jinstall...
Verified manifest signed by PackageProduction_9_0_0

WARNING: This package will load JUNOS 9.0-20080116.2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-9.0-20080116.2-domestic-signed.tgz
...
Saving state for rollback ...
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU started
ISSU: Backup RE Prepare Done
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
  FPC 1         Online (ISSU)
  FPC 2         Online (ISSU)
  FPC 3         Online (ISSU)
  FPC 5         Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
Skipping Old Master Upgrade
ISSU: IDLE

```

```
{backup}  
user@host>
```

5. You are now logged in to the new backup (old master Routing Engine). If you want to install the new software version on the new backup Routing Engine, issue the **request system software add /var/tmp/jinstall-9.0-20080116.2-domestic-signed.tgz** command.

Related Documentation

- [Unified ISSU System Requirements on page 2097](#)
- [Best Practices on page 2197](#)
- [Before You Begin on page 2197](#)
- [Verifying a Unified ISSU on page 2212](#)
- [Troubleshooting Unified ISSU Problems on page 2278](#)
- [Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 2212](#)

Verifying a Unified ISSU

To verify the status of FPCs and their corresponding PICs after the most recent unified ISSU, issue the **show chassis in-service-upgrade** command on the master Routing Engine:

```
user@host> show chassis in-service-upgrade  
Item           Status      Reason  
FPC 0           Online  
FPC 1           Online  
FPC 2           Online  
  PIC 0         Online  
  PIC 1         Online  
FPC 3           Offline     Offlined by CLI command  
FPC 4           Online  
  PIC 1         Online  
FPC 5           Online  
  PIC 0         Online  
FPC 6           Online  
  PIC 3         Online  
FPC 7           Online
```

For more information about the **show chassis in-service-upgrade** command, see the *Junos OS Operational Mode Commands*.

Related Documentation

- [Performing a Unified ISSU on page 2200](#)
- [Troubleshooting Unified ISSU Problems on page 2278](#)
- [Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 2212](#)

Managing and Tracing BFD Sessions During Unified ISSU Procedures

Bidirectional Forwarding Detection (BFD) sessions temporarily increase their detection and transmission timers during unified ISSU procedures. After the upgrade, these timers revert to the values in use before the unified ISSU started. The BFD process replicates the unified ISSU state and timer values to the backup Routing Engine for each session.

No additional configuration is necessary to enable unified ISSU for BFD. However, you can disable the BFD timer negotiation during the unified ISSU by including the **no-issu-timer-negotiation** statement at the **[edit protocols bfd]** hierarchy level:

```
[edit protocols bfd]
no-issu-timer-negotiation;
```

If you configure this statement, the BFD timers maintain their original values during unified ISSU.



CAUTION: The sessions might flap during unified ISSU or Routing Engine switchover, depending on the detection intervals.

For more information about BFD, see the *Junos OS Routing Protocols Configuration Guide*.

To configure unified ISSU trace options for BFD sessions, include the **issu** statement at the **[edit protocols bfd traceoptions flag]** hierarchy level.

```
[edit protocols]
bfd {
  traceoptions {
    flag issu;
  }
}
```

Related Documentation

- [Unified ISSU Concepts on page 2092](#)
- [Unified ISSU Process on the TX Matrix Router](#)
- [Unified ISSU System Requirements on page 2097](#)
- [Best Practices on page 2197](#)
- [Before You Begin on page 2197](#)
- [Performing a Unified ISSU on page 2200](#)
- [Verifying a Unified ISSU on page 2212](#)
- [Troubleshooting Unified ISSU Problems on page 2278](#)

Configuration Statements: Unified ISSU

bfd

Syntax	<pre> bfd { traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],</p> <p>[edit protocols],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure trace options for Bidirectional Forwarding Protocol (BFD) traffic.
Default	If you do not include this statement, no BFD tracing operations are performed.
Options	<p>disable—(Optional) Disable the BFD tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the /var/log directory. We recommend that you place global routing protocol tracing output in the routing-log file.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. These are the BFD protocol tracing options:</p> <ul style="list-style-type: none"> • adjacency—Trace adjacency messages. • all—Trace all options for BFD. • error—Trace all errors. • event—Trace all events. • issu—Trace in-service software upgrade (ISSU) packet activity.

- **nsr-packet**—Trace non-stop-routing (NSR) packet activity.
- **nsr-synchronization**—Trace NSR synchronization events.
- **packet**—Trace all packets.
- **pipe**—Trace pipe messages.
- **pipe-detail**—Trace pipe messages in detail.
- **ppm-packet**—Trace packet activity by periodic packet management (PPM).
- **state**—Trace state transitions.
- **timer**—Trace timer processing.

match *regular-expression*—(Optional) Regular expression for lines to be logged.

no-world-readable—(Optional) Prevent any user from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the trace file again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level	routing and trace—To view this statement in the configuration.
	routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring BFD for Static Routes on page 3197

no-issu-timer-negotiation

Syntax	no-issu-timer-negotiation;
Hierarchy Level	[edit protocols bfd], [edit logical-systems <i>logical-system-name</i> protocols bfd], [edit routing-instances <i>routing-instance-name</i> protocols bfd]
Release Information	Statement introduced in Junos OS Release 9.1.
Description	Disable unified ISSU timer negotiation for Bidirectional Forwarding Detection (BFD) sessions.



CAUTION: The sessions might flap during unified ISSU or Routing Engine switchover, depending on the detection intervals.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 2212• <i>Junos OS Routing Protocols Configuration Guide</i>.

traceoptions (Protocols BFD)

Syntax	<pre>traceoptions { file <i>name</i> <size <i>size</i>> <files <i>number</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; }</pre>
Hierarchy Level	[edit protocols bfd]
Release Information	Statement introduced before Junos OS Release 7.4. issu flag for BFD added in Junos OS Release 9.1.
Description	<p>Define tracing operations that track unified in-service software upgrade (ISSU) functionality in the router or switch.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	If you do not include this statement, no global tracing operations are performed.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place global routing protocol tracing output in the file routing-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000 files Default: 2 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>flag <i>flag</i>—Tracing operation to perform. There is only one unified ISSU tracing option:</p> <ul style="list-style-type: none">• issu—Trace BFD unified ISSU operations. <p>no-world-readable—Restrict users from reading the log file.</p> <p>size <i>size</i>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches this size, it is renamed trace-file.0. When the trace-file again reaches its maximum size, trace-file.0 is renamed trace-file.1 and trace-file is renamed trace-file.0. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p>

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

world-readable—Allow users to read the log file.

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

Related Documentation	<ul style="list-style-type: none"> • Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 2212
------------------------------	---

Configuration: VRRP

- [Configuring the Startup Period for VRRP Operations on page 2219](#)
- [Configuring Basic VRRP Support on page 2220](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 2222](#)
- [Configuring the Advertisement Interval for the VRRP Master Router on page 2224](#)
- [Configuring a Backup Router to Preempt the Master Router on page 2226](#)
- [Modifying the Preemption Hold-Time Value on page 2227](#)
- [Configuring Asymmetric Hold Time for VRRP Routers on page 2227](#)
- [Configuring an Interface to Accept Packets Destined for the Virtual IP Address on page 2228](#)
- [Configuring a Logical Interface to Be Tracked on page 2229](#)
- [Configuring a Route to Be Tracked on page 2231](#)
- [Configuring Inheritance for a VRRP Group on page 2232](#)
- [Configuring the Silent Period on page 2233](#)
- [Configuring Passive ARP Learning for Backup VRRP Routers on page 2234](#)
- [Enabling the Distributed Periodic Packet Management Process for VRRP on page 2235](#)
- [Configuring VRRP to Improve Convergence Time on page 2236](#)
- [Example: Configuring VRRP on page 2237](#)
- [Example: Configuring VRRP for IPv6 on page 2239](#)
- [Example: Configuring VRRP Route Tracking on page 2240](#)
- [Tracing VRRP Operations on page 2241](#)

Configuring the Startup Period for VRRP Operations

To configure the startup period for VRRP operations, include the **startup-silent-period** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
  startup-silent-period seconds;
```



NOTE: During the silent startup period, the `show vrrp detail` command output shows a value of 0 for Master priority, and your own IP address for Master router. These values indicate that the Master selection is not completed yet, and these values can be ignored.

Related Documentation

- [Understanding VRRP on page 2110](#)
- [VRRP Configuration Hierarchy](#)
- [Configuring Basic VRRP Support on page 2220](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 2222](#)
- [Example: Configuring VRRP on page 2237](#)

Configuring Basic VRRP Support

An interface can be a member of one or more VRRP groups. To configure basic VRRP support, configure VRRP groups on interfaces by including the **vrrp-group** statement:

```
vrrp-group group-id {
  priority number;
  virtual-address [ addresses ];
}
```

You can include this statement at the following hierarchy levels:

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

To configure basic VRRP for IPv6 support, configure VRRP group support on interfaces by including the **vrrp-inet6-group** statement:

```
vrrp-inet6-group group-id {
  priority number;
  virtual-inet6-address [ addresses ];
  virtual-link-local-address ipv6-address;
}
```

You can include this statement at the following hierarchy levels:

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

Within a VRRP group, the master virtual router and the backup virtual router must be configured on two different routing platforms.

For each VRRP group, you must configure the following:

- Group identifier—Assign a value from 0 through 255.

- Address of one or more virtual routers that are members of the VRRP group—Normally, you configure only one virtual IP address per group. However, you can configure up to eight addresses. Do not include a prefix length in a virtual IP address. The following considerations apply to configuring a virtual IP address:
 - The virtual router IP address must be the same for all routing platforms in the VRRP group.
 - If you configure a virtual IP address to be the same as the physical interface's address, the interface becomes the master virtual router for the group. In this case, you must configure the priority to be 255, and you must configure preemption by including the **preempt** statement.
 - If the virtual IP address you choose is not the same as the physical interface's address, you must ensure that the virtual IP address does not appear anywhere else in the routing platform's configuration. Verify that you do not use this address for other interfaces, for the IP address of a tunnel, or for the IP address of static ARP entries.
 - You cannot configure a virtual IP address to be the same as the interface's address for an aggregated Ethernet interface. This configuration is not supported.
 - For VRRP for IPv6, the EUI-64 option cannot be used. In addition, the Duplicate Address Detection (DAD) process will not run for virtual IPv6 addresses.
 - You cannot configure the same virtual IP address on interfaces that belong to the same logical system and routing instance combination. However, you can configure the same virtual IP address on interfaces that belong to different logical system and routing instance combinations.
- Virtual link-local address—(VRRP for IPv6 only) You must explicitly define a virtual link-local address for each VRRP for IPv6 group. Otherwise, when you attempt to commit the configuration, the commit request fails. The virtual link-local address must be on the same subnet as the physical interface address.

- Priority for this routing platform to become the master virtual router—Configure the value used to elect the master virtual router in the VRRP group. It can be a number from 1 through 255. The default value for backup routers is 100. A larger value indicates a higher priority. The routing platform with the highest priority within the group becomes the master router.



NOTE: If there are two or more backup routers with the same priority, the router that has the highest primary address becomes the master.



NOTE: Mixed tagging (configuring two logical interfaces on the same Ethernet port, one with single-tag framing and one with dual-tag framing) is supported only for interfaces on Gigabit Ethernet IQ2 and IQ PICs. If you include the `flexible-vlan-tagging` statement at the `[edit interfaces interface-name]` hierarchy level for a VRRP-enabled interface on a PIC that does not support mixed tagging, VRRP on that interface is disabled. In the output of the `show vrrp summary` command, the interface status is listed as Down.



NOTE: If you enable MAC source address filtering on an interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the `source-address-filter` statement at the `[edit interfaces interface-name]` hierarchy level. (For more information, see the *Junos® OS Network Interfaces*.) MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 2378. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

Related Documentation

- [Understanding VRRP on page 2110](#)
- [Junos OS Support for VRRPv3 on page 2111](#)
- [VRRP Configuration Hierarchy](#)
- [Configuring the Startup Period for VRRP Operations on page 2219](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 2222](#)
- [Configuring the Advertisement Interval for the VRRP Master Router on page 2224](#)
- [Example: Configuring VRRP on page 2237](#)

Configuring VRRP Authentication (IPv4 Only)

VRRP (IPv4 only) protocol exchanges can be authenticated to guarantee that only trusted routing platforms participate in routing in an autonomous system (AS). By default, VRRP authentication is disabled. You can configure one of the following authentication methods. Each VRRP group must use the same method.

- Simple authentication—Uses a text password included in the transmitted packet. The receiving routing platform uses an authentication key (password) to verify the packet.
- Message Digest 5 (MD5) algorithm—Creates the authentication data field in the IP authentication header. This header is used to encapsulate the VRRP PDU. The receiving routing platform uses an authentication key (password) to verify the authenticity of the IP authentication header and VRRP PDU.

To enable authentication and specify an authentication method, include the **authentication-type** statement:

authentication-type *authentication*;

authentication can be **simple** or **md5**. The authentication type must be the same for all routing platforms in the VRRP group.

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]

If you include the **authentication-type** statement, you can configure a key (password) on each interface by including the **authentication-key** statement:

authentication-key *key*;

key (the password) is an ASCII string. For simple authentication, it can be from 1 through 8 characters long. For MD5 authentication, it can be from 1 through 16 characters long. If you include spaces, enclose all characters in quotation marks (" "). The key must be the same for all routing platforms in the VRRP group.

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]



NOTE: When VRRPv3 is enabled, the **authentication-type** and **authentication-key** statements cannot be configured for any VRRP groups. Therefore, if authentication is required, you need to configure alternative non-VRRP authentication mechanisms.

Related Documentation

- [Understanding VRRP on page 2110](#)
- [Junos OS Support for VRRPv3 on page 2111](#)
- [VRRP Configuration Hierarchy](#)

- [Configuring Basic VRRP Support on page 2220](#)
- [Example: Configuring VRRP on page 2237](#)

Configuring the Advertisement Interval for the VRRP Master Router

By default, the master router sends VRRP advertisement packets every second to all members of the VRRP group. These packets indicate that the master router is still operational. If the master router fails or becomes unreachable, the backup router with the highest priority value becomes the new master router.

You can modify the advertisement interval in seconds or in milliseconds. The interval must be the same for all routing platforms in the VRRP group.

For VRRP for IPv6, you must configure IPv6 router advertisements for the interface on which VRRP is configured to send IPv6 router advertisements for the VRRP group. To do so, include the **interface *interface-name*** statement at the **[edit protocols router-advertisement]** hierarchy level. (For information about this statement and guidelines, see the *Junos OS Routing Protocols Configuration Guide*.) When an interface receives an IPv6 router solicitation message, it sends an IPv6 router advertisement to all VRRP groups configured on it. In the case of logical systems, IPv6 router advertisements are not sent to VRRP groups.



NOTE: The master VRRP for an IPv6 router must respond to a router solicitation message with the virtual IP address of the router. However, when the **interface *interface-name*** statement is included at the **[edit protocols router-advertisement]** hierarchy level, the backup VRRP for an IPv6 router might send a response before the VRRP master responds, so that the default route of the client is not set to the master VRRP router's virtual IP address. To avoid this situation, include the **virtual-router-only** statement at the **[edit protocols router-advertisement interface *interface-name*]** hierarchy level. When this statement is included, router advertisements are sent only for VRRP IPv6 groups configured on the interface (if the groups are in the master state). You must include this statement on both the master and backup VRRP for IPv6 routers.

This topic contains the following sections:

- [Modifying the Advertisement Interval in Seconds on page 2224](#)
- [Modifying the Advertisement Interval in Milliseconds on page 2225](#)

Modifying the Advertisement Interval in Seconds

To modify the time, in seconds, between the sending of VRRP advertisement packets, include the **advertise-interval** statement:

advertise-interval *seconds*;

The interval can be from 1 through 255 seconds.

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]



NOTE: When VRRPv3 is enabled, the `advertise-interval` statement cannot be used to configure advertisement intervals. Instead, use the `fast-interval` statement to configure advertisement intervals.

Modifying the Advertisement Interval in Milliseconds

To modify the time, in milliseconds, between the sending of VRRP advertisement packets, include the `fast-interval` statement:

```
fast-interval milliseconds;
```

The interval can be from 10 through 40,950 milliseconds.

You can include this statement at the following hierarchy levels:

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



NOTE: In the VRRP PDU, Junos OS sets the advertisement interval to 0. When you configure VRRP with other vendors' routers, the `fast-interval` statement works correctly only when the other routers also have an advertisement interval set to 0 in the VRRP PDUs. Otherwise, Junos OS interprets other routers' settings as advertisement timer errors.

To modify the time, in milliseconds, between the sending of VRRP for IPv6 advertisement packets, include the `inet6-advertise-interval` statement:

```
inet6-advertise-interval ms;
```

The range of values is from 100 through 40,950 milliseconds (ms).

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* vrrp-inet6-group *group-id*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* vrrp-inet6-group *group-id*]



NOTE: When VRRPv3 is enabled, the `inet6-advertise-interval` statement cannot be used to configure advertisement intervals. Instead, use the `fast-interval` statement to configure advertisement intervals.

**Related
Documentation**

- [Understanding VRRP on page 2110](#)
- [Junos OS Support for VRRPv3 on page 2111](#)
- [VRRP Configuration Hierarchy](#)
- [Configuring Basic VRRP Support on page 2220](#)
- [Configuring a Backup Router to Preempt the Master Router on page 2226](#)
- [Modifying the Preemption Hold-Time Value on page 2227](#)
- [Configuring Asymmetric Hold Time for VRRP Routers on page 2227](#)
- [Configuring the Silent Period on page 2233](#)
- [Example: Configuring VRRP on page 2237](#)

Configuring a Backup Router to Preempt the Master Router

By default, a higher-priority backup router preempts a lower-priority master router. To explicitly enable the master router to be preempted, include the **preempt** statement:

```
preempt;
```

You can include this statement at the following hierarchy levels:

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

To prohibit a higher-priority backup router from preempting a lower-priority master router, include the **no-preempt** statement:

```
no-preempt;
```

**Related
Documentation**

- [Understanding VRRP on page 2110](#)
- [VRRP Configuration Hierarchy](#)
- [Configuring the Advertisement Interval for the VRRP Master Router on page 2224](#)
- [Modifying the Preemption Hold-Time Value on page 2227](#)
- [Configuring Asymmetric Hold Time for VRRP Routers on page 2227](#)
- [Example: Configuring VRRP on page 2237](#)

Modifying the Preemption Hold-Time Value

The hold time is the maximum number of seconds that can elapse before a higher-priority backup router preempts the master router. You might want to configure a hold time so that all Junos OS components converge before preemption.

By default, the hold-time value is 0 seconds. A value of 0 means that preemption can occur immediately after the backup router comes online. Note that the hold time is counted from the time the backup router comes online. The hold time is only valid when the VRRP router is just coming online.

To modify the preemption hold-time value, include the **hold-time** statement:

```
hold-time seconds;
```

The hold time can be from 0 through 3600 seconds.

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-id* preempt]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-id* preempt]**

Related Documentation

- [VRRP Configuration Hierarchy](#)
- [Configuring the Advertisement Interval for the VRRP Master Router on page 2224](#)
- [Configuring a Backup Router to Preempt the Master Router on page 2226](#)
- [Configuring Asymmetric Hold Time for VRRP Routers on page 2227](#)
- [Example: Configuring VRRP on page 2237](#)

Configuring Asymmetric Hold Time for VRRP Routers

In Junos OS Release 9.5 and later, the **asymmetric-hold-time** statement at the **[edit protocols vrrp]** hierarchy level enables you to configure a VRRP master router to switch over to the backup router immediately—that is, without waiting for the priority hold time to expire—when a tracked interface or route goes down or when the bandwidth of a tracked interface decreases. Such events can cause an immediate reduction in the priority based on the configured priority cost for the event, and trigger a mastership election.

However, when the tracked route or interface comes up again, or when the bandwidth for a tracked interface increases, the backup (original master) router waits for the hold time to expire before it updates the priority and initiates the switchover if the priority is higher than the priority for the VRRP master (original backup) router.

If the **asymmetric-hold-time** statement is not configured, the VRRP master waits for the hold time to expire before it initiates a switchover when a tracked route goes down.

Example: Configuring Asymmetric Hold Time

```
[edit]
user@host# set protocols vrrp asymmetric-hold-time
[edit]
user@host# show protocols vrrp
asymmetric-hold-time;
```

Related Documentation

- [VRRP Configuration Hierarchy](#)
- [Configuring the Advertisement Interval for the VRRP Master Router on page 2224](#)
- [Configuring a Backup Router to Preempt the Master Router on page 2226](#)
- [Modifying the Preemption Hold-Time Value on page 2227](#)
- [Example: Configuring VRRP on page 2237](#)

Configuring an Interface to Accept Packets Destined for the Virtual IP Address

In VRRP implementations where the router acting as the master router is not the IP address owner—the IP address owner is the router that has the interface whose actual IP address is used as the virtual router's IP address (virtual IP address)—the master router accepts only the ARP packets from the packets that are sent to the virtual IP address. Junos OS enables you to override this limitation with the help of the **accept-data** configuration. When the **accept-data** statement is included in the configuration, the master router accepts all packets sent to the virtual IP address even when the master router is not the IP address owner.



NOTE: If the master router is the IP address owner or has its priority set to 255, the master router, by default, accepts all packets addressed to the virtual IP address. In such cases, the **accept-data** configuration is not required.

To configure an interface to accept all packets sent to the virtual IP address, include the **accept-data** statement:

```
accept-data;
```

You can include this statement at the following hierarchy levels:

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

To prevent a master router that is the IP address owner or has its priority set to 255 from accepting packets other than the ARP packets addressed to the virtual IP address, include the **no-accept-data** statement:

```
no-accept-data;
```

**NOTE:**

- If you want to restrict the incoming IP packets to ICMP packets only, you must configure firewall filters to accept only ICMP packets.
- If you include the `accept-data` statement, your routing platform configuration does not comply with RFC 3768 (see section 6.4.3 of RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*).

Related Documentation

- [Understanding VRRP on page 2110](#)
- [VRRP Configuration Hierarchy](#)
- [Example: Configuring VRRP on page 2237](#)

Configuring a Logical Interface to Be Tracked

VRRP can track whether a logical interface is up, down, or not present, and can also dynamically change the priority of the VRRP group based on the state of the tracked logical interface, triggering a new master router election. VRRP can also track the operational speed of a logical interface and dynamically update the priority of the VRRP group when the speed crosses a configured threshold.

When interface tracking is enabled, you cannot configure a priority of 255 (a priority of 255 designates the master router). For each VRRP group, you can track up to 10 logical interfaces.

To configure a logical interface to be tracked, include the following statements:

```
track {
  interface interface-name {
    bandwidth-threshold bits-per-second priority-cost priority;
    priority-cost priority;
  }
  priority-hold-time seconds;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* *vrrp-group group-id*]
- [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* *vrrp-inet6-group group-id*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address* *vrrp-group group-id*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* *vrrp-inet6-group group-id*]

The interface specified is the interface to be tracked for the VRRP group. The priority hold time is the minimum length of time that must elapse between dynamic priority changes. A tracking event, such as an interface state change (up or down) or a change in bandwidth, triggers one of the following responses:

- The first tracking event initiates the priority hold timer, and also initializes the pending priority based on the current priority and the priority cost. However, the current priority remains unchanged.
- A tracking event or a manual configuration change that occurs while the priority hold timer is on triggers a pending priority update. However, the current priority remains unchanged.

This ensures that Junos OS does not initiate mastership elections every time a tracked interface flaps.

When the priority hold time expires, the current priority inherits the value from the pending priority, and the pending priority ceases.



NOTE: If you have configured **asymmetric-hold-time**, VRRP does not wait for the priority hold time to expire before initiating mastership elections if a tracked interface fails (state changes from up to down), or if the available bandwidth for a tracked interface decreases. For more information about **asymmetric-hold-time**, see [“Configuring Asymmetric Hold Time for VRRP Routers” on page 2227](#).

The bandwidth threshold specifies a threshold for the tracked interface. When the bandwidth of the tracked interface drops below the configured bandwidth threshold value, the VRRP group uses the bandwidth threshold priority cost. You can track up to five bandwidth threshold statements for each tracked interface.

The priority cost is the value to be subtracted from the configured VRRP priority when the tracked logical interface goes down, forcing a new master router election. The value can be from 1 through 254. The sum of the costs for all tracked logical interfaces or routes must be less than or equal to the configured priority of the VRRP group.

If you are tracking more than one interface, the router applies the sum of the priority costs for the tracked interfaces (at most, only one priority cost for each tracked interface) to the VRRP group priority. However, the interface priority cost and bandwidth threshold priority cost values for each VRRP group are not cumulative. The router uses only one priority cost to a tracked interface as indicated in [Table 165 on page 2230](#):

Table 165: Interface State and Priority Cost Usage

Tracked Interface State	Priority Cost Usage
Down	priority-cost priority
Not down; media speed below one or more bandwidth thresholds	Priority-cost of the lowest applicable bandwidth threshold

You must configure an interface priority cost only if you have configured no bandwidth thresholds. If you have not configured an interface priority cost value, and the interface is down, the interface uses the bandwidth threshold priority cost value of the lowest bandwidth threshold.

Related Documentation

- [Understanding VRRP on page 2110](#)
- [VRRP Configuration Hierarchy](#)
- [Configuring a Route to Be Tracked on page 2231](#)
- [Example: Configuring VRRP on page 2237](#)

Configuring a Route to Be Tracked

VRRP can track whether a route is reachable (that is, the route exists in the routing table of the routing instance included in the configuration) and dynamically change the priority of the VRRP group based on the reachability of the tracked route, triggering a new master router election.

To configure a route to be tracked, include the following statements:

```
track {
  priority-hold-time seconds;
  route prefix/prefix-length routing-instance instance-name priority-cost priority;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* *vrrp-group group-id*]
- [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* *vrrp-inet6-group group-id*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address* *vrrp-group group-id*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* *vrrp-inet6-group group-id*]

The route prefix specified is the route to be tracked for the VRRP group. The priority hold time is the minimum length of time that must elapse between dynamic priority changes. A route tracking event, such as adding a route to or removing a route from the routing table, might trigger one or more of the following:

- The first tracking event initiates the priority hold timer, and also initializes the pending priority based on the current priority and the priority cost. However, the current priority remains unchanged.
- A tracking event or a manual configuration change that occurs while the priority hold timer is on triggers a pending priority update. However, the current priority remains unchanged.

When the priority hold time expires, the current priority inherits the value from the pending priority, and the pending priority ceases.

This ensures that Junos OS does not initiate mastership elections every time a tracked route flaps.



NOTE: If you have configured `asymmetric-hold-time`, VRRP does not wait for the priority hold time to expire before initiating mastership elections if a tracked route is removed from the routing table. For more information about `asymmetric-hold-time`, see [“Configuring Asymmetric Hold Time for VRRP Routers” on page 2227](#).

The routing instance is the routing instance in which the route is to be tracked. If the route is in the default, or global, routing instance, specify the instance name as **default**.



NOTE: Tracking a route that belongs to a routing instance from a different logical system is not supported.

The priority cost is the value to be subtracted from the configured VRRP priority when the tracked route goes down, forcing a new master router election. The value can be from 1 through 254. The sum of the costs for all tracked logical interfaces or routes must be less than or equal to the configured priority of the VRRP group.

Related Documentation

- [Understanding VRRP on page 2110](#)
- [VRRP Configuration Hierarchy](#)
- [Configuring a Logical Interface to Be Tracked on page 2229](#)
- [Example: Configuring VRRP Route Tracking on page 2240](#)

Configuring Inheritance for a VRRP Group

Junos OS enables you to configure VRRP groups on the various subnets of a VLAN to inherit the state and configuration of one of the groups, which is known as the *active VRRP group*. When the `vrrp-inherit-from` configuration statement is included in the configuration, only the active VRRP group, from which the other VRRP groups are inheriting the state, sends out frequent VRRP advertisements, and processes incoming VRRP advertisements. The groups that are inheriting the state do not process any incoming VRRP advertisement because the state is always inherited from the active VRRP group. However, the groups that are inheriting the state do send out VRRP advertisements once every 2 to 3 minutes to facilitate MAC address learning on the switches placed between the VRRP routers.

If the `vrrp-inherit-from` statement is not configured, each of the VRRP master groups in the various subnets on the VLAN sends out separate VRRP advertisements and adds to the traffic on the VLAN.

To configure inheritance for a VRRP group, include the **vrrp-inherit-from** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* *vrrp-group group-id*]** hierarchy level.

```
[edit interfaces interface-name unit logical-unit-number family inet address address
  vrrp-group group-id]
vrrp-inherit-from vrrp-group;
```

When you configure a group to inherit a state from another group, the inheriting groups and the active group must be on the same physical interface and logical system. However, the groups do not need to necessarily be on the same routing instance (as was in Junos OS releases earlier than 9.6), VLAN, or logical interface.

When you include the **vrrp-inherit-from** statement for a VRRP group, the VRRP group inherits the following parameters from the active group:

- **advertise-interval**
- **authentication-key**
- **authentication-type**
- **fast-interval**
- **preempt | no-preempt**
- **priority**
- **track interfaces**
- **track route**

However, you can configure the **accept-data | no-accept-data** statement for the group to specify whether the interface should accept packets destined for the virtual IP address.

Related Documentation

- [Understanding VRRP on page 2110](#)
- [VRRP Configuration Hierarchy](#)

Configuring the Silent Period

The silent period starts when the interface state is changed from down to up. During this period, the Master Down Event is ignored. Configure the silent period interval to avoid alarms caused by the delay or interruption of the incoming VRRP advertisement packets during the interface startup phase.

To configure the silent period interval that the Master Down Event timer ignores, include the **startup-silent-period** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
startup-silent-period seconds;
```



NOTE: During the silent startup period, the `show vrrp detail` command output shows a value of 0 for Master priority and your IP address for Master router. These values indicate that the Master selection is not completed yet, and these values can be ignored.

When you have configured **startup-silent-period**, the Master Down Event is ignored until the **startup-silent-period** expires.

For example, configure a VRRP group, *vrrp-group1*, with an advertise interval of 1 second, startup silent period of 10 seconds, and an interface *interface1* with a priority less than 255.

When *interface1* transitions from down to up:

- The *vrrp-group1* group moves to the backup state, and starts the Master Down Event timer (3 seconds; three times the value of the advertise interval, which is 1 second in this case).
- If no VRRP PDU is received during the 3-second period, the **startup-silent-period** (10 seconds in this case) is checked, and if the startup silent period has not expired, the Master Down Event timer is restarted. This is repeated until the **startup-silent-period** expires. In this example, the Master Down Event timer runs four times (12 seconds) by the time the 10-second startup silent period expires.
- If no VRRP PDU is received by the end of the fourth 3-second cycle, *vrrp-group1* takes over mastership.

Related Documentation

- [Understanding VRRP on page 2110](#)
- [VRRP Configuration Hierarchy](#)
- [startup-silent-period on page 2261](#)

Configuring Passive ARP Learning for Backup VRRP Routers

By default, the backup VRRP router drops ARP requests for the VRRP-IP to VRRP-MAC address translation. This means that the backup router does not learn the ARP (IP-to-MAC address) mappings for the hosts sending the requests. When it detects a failure of the master router and transitions to become the new master router, the backup router must learn all the entries that were present in the ARP cache of the master router. In environments with many directly attached hosts, such as metro Ethernet environments, the number of ARP entries to learn can be high. This can cause a significant transition delay, during which the traffic transmitted to some of the hosts might be dropped.

Passive ARP learning enables the ARP cache in the backup router to hold approximately the same contents as the ARP cache in the master router, thus preventing the problem of learning ARP entries in a burst. To enable passive ARP learning, include the **passive-learning** statement at the **[edit system arp]** hierarchy level:

```
[edit system arp]
passive-learning;
```


We recommend setting passive learning on both the backup and master VRRP routers. Doing so prevents the need to manually intervene when the master router becomes the backup router. While a router is operating as the master router, the passive learning configuration has no operational impact. The configuration takes effect only when the router is operating as a backup router.

For information about configuring gratuitous ARP and the ARP aging timer, see the *Junos OS System Basics Configuration Guide*.

**Related
Documentation**

- [Understanding VRRP on page 2110](#)
- [VRRP Configuration Hierarchy](#)

Enabling the Distributed Periodic Packet Management Process for VRRP

Typically, VRRP advertisements are sent by the VRRP process (vrrpd) on the master VRRP router at regular intervals to let other members of the group know that the VRRP master router is operational.

When the vrrpd process is busy and does not send VRRP advertisements, the backup VRRP routers might assume that the master router is down and take over as the master router, causing unnecessary flaps. This takeover might occur even though the original master router is still active and available, and might resume sending advertisements after the traffic has decreased. To address this problem and to reduce the load on the vrrpd process, Junos OS uses the periodic packet management process (ppmd) to send VRRP advertisements on behalf of the vrrpd process. However, you can further delegate the job of sending VRRP advertisements to the distributed ppm process that resides on the Packet Forwarding Engine.

The ability to delegate the sending of VRRP advertisements to the distributed ppm process ensures that the VRRP advertisements are sent even when the ppm process—which is now responsible for sending VRRP advertisements—is busy. Such delegation prevents the possibility of false alarms when the ppm process is busy. The ability to delegate the sending of VRRP advertisements to distributed ppm also adds to scalability because the load is shared across multiple ppm instances and is not concentrated on any single unit.



NOTE: CPU-intensive VRRP advertisements, such as advertisements with MD5 authentication or those sent and received over logical interfaces, such as Aggregated Ethernet interfaces, continue to be processed by the VRRP process on the Routing Engine even when distributed ppm is enabled.

To configure the distributed ppm process to send VRRP advertisements, include the **delegate-processing** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
  delegate-processing;
```

**Related
Documentation**

- [Understanding VRRP on page 2110](#)

- *VRRP Configuration Hierarchy*

Configuring VRRP to Improve Convergence Time

You can enable faster convergence time for the configured Virtual Router Redundancy Protocol (VRRP), thereby reducing the traffic restoration time to less than 1 second. To improve the convergence time for VRRP, perform the following tasks.

Before you begin, configure VRRP. See [“Example: Configuring VRRP” on page 2237](#).

1. Configure the distributed periodic packet management (PPM) process to send VRRP advertisements when the VRRP process is busy.

[edit]

user@host# **set protocols vrrp delegate-processing**

2. Disable the skew timer to reduce the time required to transition to the master state.

[edit]

user@host# **set protocols vrrp skew-timer-disable**



NOTE: When there is only one master router and one backup router in the network deployment, you can disable the skew timer, thereby reducing the time required to transition to the master state.

3. Configure the number of fast advertisements that can be missed by a backup router before it starts transitioning to the master state.

[edit]

user@host# **set protocols vrrp global-advertisement-threshold *advertisement-value***

4. Configure VRRP groups on the various subnets of a VLAN to inherit the state and to configure one of the groups.

[edit]

user@host# **set interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id***

5. Verify the configuration in operational mode.

[edit]

user@host> **show protocols vrrp**

**NOTE:**

- The reduction in convergence time is not applicable when VRRP is configured over integrated routing and bridging (IRB) interfaces, aggregated Ethernet interfaces, and multichassis link aggregation group (MC-LAG) interfaces.
- Compared to other routers, the convergence time and the traffic restoration time are less for MX Series routers with MPCs.
- Reduction in convergence time is applicable for all types of configurations at the physical interface, but the convergence time might not be less than 1 second for all the configurations. The convergence time depends on the number of groups that are transitioning from the backup to the master state and the interval at which these groups are transitioning.

Related Documentation

- [Improving the Convergence Time for VRRP on page 2114](#)
- [Configuring Inheritance for a VRRP Group on page 2232](#)
- [delegate-processing on page 2250](#)
- [global-advertisements-threshold on page 2252](#)
- [skew-timer-disable on page 2261](#)

Example: Configuring VRRP

Configure one master (Router A) and one backup (Router B) routing platform. The address configured in the **virtual-address** statements differs from the addresses configured in the **address** statements. When you configure multiple VRRP groups on an interface, you configure one to be the master virtual router for that group.

On Router A

```
[edit interfaces]
ge-0/0/0 {
  unit 0 {
    family inet {
      address 192.168.1.20/24 {
        vrrp-group 27 {
          virtual-address 192.168.1.15;
          priority 254;
          authentication-type simple;
          authentication-key booJUM;
        }
      }
    }
  }
}
```

On Router B

```
[edit interfaces]
ge-4/2/0 {
  unit 0 {
    family inet {
      address 192.168.1.24/24 {
```

Configuring One Router to Be the Master Virtual Router for the Group

```

    vrrp-group 27 {
        virtual-address 192.168.1.15;
        priority 200;
        authentication-type simple;
        authentication-key booJUM;
    }
}
}
}
}

```

```

[edit interfaces]
ge-0/0/0 {
    unit 0 {
        family inet {
            address 192.168.1.20/24 {
                vrrp-group 2 {
                    virtual-address 192.168.1.20;
                    priority 255;
                    advertise-interval 3;
                    preempt;
                }
                vrrp-group 10 {
                    virtual-address 192.168.1.55;
                    priority 201;
                    advertise-interval 3;
                }
                vrrp-group 1 {
                    virtual-address 192.168.1.54;
                    priority 22;
                    advertise-interval 4;
                }
            }
        }
    }
}
}
}

```

Configuring VRRP and MAC Source Address Filtering

The VRRP group number is the decimal equivalent of the last byte of the virtual MAC address.

```

[edit interfaces]
ge-5/2/0 {
    gigether-options {
        source-filtering;
        source-address-filter {
            00:00:5e:00:01:0a; # Virtual MAC address
        }
    }
    unit 0 {
        family inet {
            address 192.168.1.10/24 {
                vrrp-group 10; # VRRP group number
                virtual-address 192.168.1.10;
                priority 255;
                preempt;
            }
        }
    }
}

```

```

    }
  }
}

```

Related Documentation

- [Understanding VRRP on page 2110](#)
- [VRRP Configuration Hierarchy](#)
- [VRRP for IPv6 Configuration Hierarchy](#)
- [Example: Configuring VRRP for IPv6 on page 2239](#)
- [Example: Configuring VRRP Route Tracking on page 2240](#)

Example: Configuring VRRP for IPv6

Configure VRRP properties for IPv6 in one master (Router A) and one backup (Router B).

On Router A

```

[edit interfaces]
ge-1/0/0 {
  unit 0 {
    family inet6 {
      address fe80::5:0:0:6/64;
      address fec0::5:0:0:6/64 {
        vrrp-inet6-group 3; # VRRP inet6 group number
        virtual-inet6-address fec0::5:0:0:7;
        virtual-link-local-address fe80::5:0:0:7;
        priority 200;
        preempt;
      }
    }
  }
}

[edit protocols]
router-advertisement {
  interface ge-1/0/0.0 {
    prefix fec0::/64;
    max-advertisement-interval 4;
  }
}

```

On Router B

```

[edit interfaces]
ge-1/0/0 {
  unit 0 {
    family inet6 {
      address fe80::5:0:0:8/64;
      address fec0::5:0:0:8/64 {
        vrrp-inet6-group 3; # VRRP inet6 group number
        virtual-inet6-address fec0::5:0:0:7;
        virtual-link-local-address fe80::5:0:0:7;
        priority 100;
        preempt;
      }
    }
  }
}

```

```
}

[edit protocols]
router-advertisement {
  interface ge-1/0/0.0 {
    prefix fec0::/64;
    max-advertisement-interval 4;
  }
}
```

Related Documentation

- [Understanding VRRP on page 2110](#)
- [VRRP Configuration Hierarchy](#)
- [VRRP for IPv6 Configuration Hierarchy](#)
- [Example: Configuring VRRP on page 2237](#)
- [Example: Configuring VRRP Route Tracking on page 2240](#)

[Example: Configuring VRRP Route Tracking](#)

Configure Routers R1 and R2 to run VRRP. Configure static routes and a policy for exporting the static routes on Router R3. The VRRP routing instances on R2 track the routes that are advertised by R3.

On Router R1

```
[edit interfaces]
ge-1/0/3 {
  unit 0 {
    vlan-id 1;
    family inet {
      address 200.100.50.2/24 {
        vrrp-group 0 {
          virtual-address 200.100.50.101;
          priority 195;
        }
      }
    }
  }
}
```

On Router R2

```
[edit interfaces]
ge-1/0/1 {
  unit 0 {
    vlan-id 1;
    family inet {
      address 200.100.50.1/24 {
        vrrp-group 0 {
          virtual-address 200.100.50.101;
          priority 200;
          track {
            route 59.0.58.153/32 routing-instance default priority-cost 5;
            route 59.0.58.154/32 routing-instance default priority-cost 5;
            route 59.0.58.155/32 routing-instance default priority-cost 5;
          }
        }
      }
    }
  }
}
```

```

    }
  }
}

```

On Router R3

```

[edit]
policy-options {
  policy-statement static-policy {
    term term1 {
      then accept;
    }
  }
}
protocols {
  ospf {
    export static-policy;
    reference-bandwidth 4g;
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
  }
}
routing-options {
  static {
    route 59.0.0.153/32 next-hop 45.45.45.46;
    route 59.0.0.154/32 next-hop 45.45.45.46;
    route 59.0.0.155/32 next-hop 45.45.45.46;
  }
}

```

Related Documentation

- [Understanding VRRP on page 2110](#)
- [VRRP Configuration Hierarchy](#)
- [VRRP for IPv6 Configuration Hierarchy](#)
- [Configuring a Route to Be Tracked on page 2231](#)
- [Example: Configuring VRRP on page 2237](#)
- [Example: Configuring VRRP for IPv6 on page 2239](#)

Tracing VRRP Operations

To trace VRRP operations, include the **traceoptions** statement at the **[edit protocols vrrp]** hierarchy level.

By default, VRRP logs the error, data carrier detect (DCD) configuration, and routing socket events in a file in the **/var/log** directory. By default, this file is named **/var/log/vrrpd**. The default file size is 1 megabyte (MB), and three files are created before the first one gets overwritten.

To change the configuration of the logging file, include the **traceoptions** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
traceoptions {
  file filename <files number> <match regular-expression> <microsecond-stamp>
    <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
flag flag;
```

You can specify the following VRRP tracing flags:

- **all**—Trace all VRRP operations.
- **database**—Trace all database changes.
- **general**—Trace all general events.
- **interfaces**—Trace all interface changes.
- **normal**—Trace all normal events.
- **packets**—Trace all packets sent and received.
- **state**—Trace all state transitions.
- **timer**—Trace all timer events.

- Related Documentation**
- [Understanding VRRP on page 2110](#)
 - [VRRP Configuration Hierarchy](#)

Configuration Statements: VRRP

- [\[edit protocols vrrp\] Hierarchy Level on page 2242](#)

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

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```
protocols {
  vrrp {
    asymmetric-hold-time;
    delegate-processing;
    failover-delay milliseconds;
    global-advertisements-threshold advertisement-value;
    skew-timer-disable;
    startup-silent-period seconds;
    traceoptions {
      file <filename> <files number> <match regular-expression> <microsecond-stamp>
        <size maximum-file-size> <world-readable | no-world-readable>;
      flag flag;
      no-remote-trace;
    }
  }
}
```




```
        version-3;  
    }  
}
```


**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*
- *Junos OS Hierarchy and RFC Reference*
- *Junos® OS Ethernet Interfaces*
- *Junos® OS Network Interfaces*

accept-data

Syntax	(accept-data no-accept-data);
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	<p>In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not a router that is acting as the master router accepts all packets destined for the virtual IP address.</p> <ul style="list-style-type: none"> • accept-data—Enable the master router to accept all packets destined for the virtual IP address. • no-accept-data—Prevent the master router from accepting packets other than the ARP packets destined for the virtual IP address.
Default	<p>If the router acting as the master router is the IP address owner or has its priority set to 255, the master router, by default, responds to all packets sent to the virtual IP address. However, if the router acting as the master router does not own the IP address or has its priority set to a value less than 255, the master router responds only to ARP requests.</p>
<div>  <p>NOTE:</p> <ul style="list-style-type: none"> • If you want to restrict the incoming IP packets to ICMP packets only, you must configure firewall filters to accept only ICMP packets. • If you include the accept-data statement, your routing platform configuration does not comply with RFC 3768 (see section 6.4.3 of RFC 3768, <i>Virtual Router Redundancy Protocol (VRRP)</i>). </div>	
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring an Interface to Accept Packets Destined for the Virtual IP Address on page 2228


advertise-interval

Syntax	<code>advertise-interval seconds;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
Description	Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv4 advertisement packets. All routers in the VRRP group must use the same advertisement interval.
<div>  <p>NOTE: When VRRPv3 is enabled, the <code>advertise-interval</code> statement cannot be used to configure advertisement intervals. Instead, use the <code>fast-interval</code> statement to configure advertisement intervals.</p> </div>	
Options	<i>seconds</i> —Interval between advertisement packets. Range: 1 through 255 seconds Default: 1 second
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Advertisement Interval for the VRRP Master Router on page 2224 • fast-interval on page 2251 • inet6-advertise-interval on page 2254 • version-3 on page 2265

asymmetric-hold-time

Syntax	asymmetric-hold-time;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Enable the VRRP master router to switch over to the backup router immediately, without waiting for the priority hold time to expire, when a route goes down. However, when the route comes back online, the backup router that is acting as the master waits for the priority hold time to expire before switching the mastership back to the original master VRRP router.
Default	asymmetric-hold-time is disabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Asymmetric Hold Time for VRRP Routers on page 2227

authentication-key

Syntax	<code>authentication-key key;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	<p>Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 authentication key. You also must specify a VRRP authentication scheme by including the authentication-type statement.</p> <p>All routers in the VRRP group must use the same authentication scheme and password.</p>
	<div>  <p>NOTE: When VRRPv3 is enabled, the authentication-type and authentication-key statements cannot be configured for any VRRP groups.</p> </div>
Options	<p>key—Authentication password. For simple authentication, it can be 1 through 8 characters long. For Message Digest 5 (MD5) authentication, it can be 1 through 16 characters long. If you include spaces, enclose all characters in quotation marks (" ").</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring VRRP Authentication (IPv4 Only) on page 2222 • Configuring VRRP Authentication (IPv4 Only) • authentication-type on page 2248 • version-3 on page 2265

authentication-type

Syntax	<code>authentication-type <i>authentication</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
Description	Enable Virtual Router Redundancy Protocol (VRRP) IPv4 authentication and specify the authentication scheme for the VRRP group. If you enable authentication, you must specify a password by including the authentication-key statement. All routers in the VRRP group must use the same authentication scheme and password.



NOTE: When VRRPv3 is enabled, the **authentication-type** and **authentication-key** statements cannot be configured for any VRRP groups.

Options	<i>authentication</i> —Authentication scheme: <ul style="list-style-type: none">• simple—Use a simple password. The password is included in the transmitted packet, so this method of authentication is relatively insecure.• md5—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing platform uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme. Default: none (no authentication is performed).
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring VRRP Authentication (IPv4 Only) on page 2222• Configuring VRRP Authentication (IPv4 Only)• authentication-key on page 2247• version-3 on page 2265

bandwidth-threshold

Syntax	<code>bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> track interface <i>interface-name</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> track interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> track interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> track interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	Specify the bandwidth threshold for Virtual Router Redundancy Protocol (VRRP) logical interface tracking.
Options	<p><i>bits-per-second</i>—Bandwidth threshold for the tracked interface. When the bandwidth of the tracked interface drops below the specified value, the VRRP group uses the bandwidth threshold priority cost value. You can include up to five bandwidth threshold statements for each interface you track.</p> <p>Range: 1 through 10000000000000 bits per second</p> <p><i>priority-cost priority</i>—The value subtracted from the configured VRRP priority when the tracked interface or route is down to force a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Logical Interface to Be Tracked on page 2229 • Configuring a Logical Interface to Be Tracked


delegate-processing (VRRP)

Syntax	delegate-processing;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 9.6.
Description	Configure the distributed ppmmd process to send VRRP advertisements.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration
Related Documentation	<ul style="list-style-type: none">• Enabling the Distributed Periodic Packet Management Process for VRRP on page 2235

failover-delay

Syntax	failover-delay <i>milliseconds</i> ;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 9.4.
Description	Configure the failover delay for VRRP and VRRP for IPv6 operations.
Options	<i>milliseconds</i> —Specify the failover delay time, in milliseconds. Range: 50 through 2000
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring VRRP and VRRP for IPv6

fast-interval

Syntax	<code>fast-interval milliseconds;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	<p>Configure the interval, in milliseconds, between Virtual Router Redundancy Protocol (VRRP) advertisement packets.</p> <p>All routers in the VRRP group must use the same advertisement interval.</p>
Options	<p><i>milliseconds</i>—Interval between advertisement packets.</p> <p>Range: 10 through 40,950 milliseconds (range extended from 100–999 to 10–40,950 in Junos OS Release 12.2).</p>
<div>  <p>NOTE: When configuring VRRP for IPv4, if you have chosen not to enable VRRPv3, you cannot set a value less than 100 for <i>fast-interval</i>. Commit check fails if a value less than 100 is configured.</p> </div>	
Default: 1 second	
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Advertisement Interval for the VRRP Master Router on page 2224 • Configuring the Advertisement Interval for the VRRP Master • advertise-interval on page 2245 • advertise-interval on page 2245 • inet6-advertise-interval on page 2254 • version-3 on page 2265

global-advertisements-threshold

Syntax	<code>global-advertisements-threshold <i>advertisement-value</i>;</code>
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 12.2.
Description	Configure the number of fast advertisements that can be missed by a backup router before the master router is declared as down.



NOTE:


- The advertisement value configured using the `global-advertisements-threshold` statement is applicable to all the Virtual Router Redundancy Protocol (VRRP) groups in the system.
 - Setting the advertisement value of the `global-advertisements-threshold` configuration to 1 is not recommended for a scaled configuration with an aggressive advertisement interval. For example, if you have 1000 VRRP groups with an advertisement interval of 100 ms, then do not set the `global-advertisements-threshold` value to 1.
 - Changing the advertisement value of the `global-advertisements-threshold` configuration during runtime can result in unpredictable behavior by the VRRP state machine. For example, momentary ownership change from the master router to the backup router and vice versa. Therefore, avoid changing the advertisement value of the `global-advertisements-threshold` statement during runtime.
-

Options	<i>advertisement-value</i> —Number of VRRP advertisements missed before the master router is declared as down. Range: 1 through 15 Default: 3
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Improving the Convergence Time for VRRP on page 2114• Configuring VRRP to Improve Convergence Time on page 2236

hold-time (VRRP)

Syntax	<code>hold-time seconds;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id preempt</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id preempt</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id preempt</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id preempt</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	In a Virtual Router Redundancy Protocol (VRRP) configuration, set the hold time before a higher-priority backup router preempts the master router.
Default	VRRP preemption is not timed.
Options	<p>seconds—Hold-time period.</p> <p>Range: 0 through 3600 seconds</p> <p>Default: 0 seconds (VRRP preemption is not timed.)</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Backup Router to Preempt the Master Router on page 2226 • Configuring VRRP Preemption and Hold Time


inet6-advertise-interval

Syntax	<code>inet6-advertise-interval <i>milliseconds</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i>]
Release Information	Statement introduced in Junos OS Release 8.4R2.
Description	<p>Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv6 advertisement packets.</p> <p>All routers in the VRRP group must use the same advertisement interval.</p>
	<div><p>NOTE: When VRRPv3 is enabled, the <code>inet6-advertise-interval</code> statement cannot be used to configure advertisement intervals. Instead, use the <code>fast-interval</code> statement to configure advertisement intervals.</p></div>
Options	<p><i>milliseconds</i>—Interval, in milliseconds, between advertisement packets.</p> <p>Range: 100 to 40,000 milliseconds (ms)</p> <p>Default: 1 second</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring the Advertisement Interval for the VRRP Master Router on page 2224• advertise-interval on page 2245• fast-interval on page 2251• version-3 on page 2265

interface (VRRP Group)

Syntax	<pre>interface <i>interface-name</i> { bandwidth-threshold <i>bits-per-second</i> <i>priority-cost</i> <i>priority</i>; priority-cost <i>priority</i>; }</pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> track],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i> track],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> vrrp-group <i>group-id</i> track],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> vrrp-inet6-group <i>group-id</i> track]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>bandwidth-threshold statement added in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	Enable logical interface tracking for a Virtual Router Redundancy Protocol (VRRP) group.
Options	<p><i>interface-name</i>—Interface to be tracked for this VRRP group.</p> <p>Range: 1 through 10 interfaces</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Logical Interface to Be Tracked on page 2229 • <i>Configuring a Logical Interface to Be Tracked</i> • <i>Junos Services Interfaces Configuration Release 11.2</i>

preempt (VRRP)

Syntax	(preempt no-preempt) { hold-time seconds; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
Description	In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not a backup router can preempt a master router: <ul style="list-style-type: none"> • preempt—Allow the master router to be preempted. <div style="margin-top: 10px;">  <p>NOTE: By default, a higher-priority backup router can preempt a lower-priority master router.</p> </div> <ul style="list-style-type: none"> • no-preempt—Prohibit the preemption of the master router. When no-preempt is configured, the backup router cannot preempt the master router even if the backup router has a higher priority. <p>The remaining statement is explained separately.</p>
Default	By default the preempt statement is enabled, and a higher-priority backup router preempts a lower-priority master router even if the preempt statement is not explicitly configured.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring a Backup Router to Preempt the Master Router on page 2226 • Configuring VRRP Preemption and Hold Time


priority (Protocols VRRP)

Syntax	<code>priority <i>priority</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	Configure a Virtual Router Redundancy Protocol (VRRP) router's priority for becoming the master default router. The router with the highest priority within the group becomes the master.
Options	<p>priority—Router's priority for being elected to be the master router in the VRRP group. A larger value indicates a higher priority for being elected.</p> <p>Range: 1 through 255</p> <p>Default: 100 (for backup routers)</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Basic VRRP Support on page 2220 • Configuring Basic VRRP Support

priority-cost (VRRP)

Syntax	<code>priority-cost priority;</code>
Hierarchy Level	<code>[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id track interface interface-name],</code> <code>[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id track interface interface-name],</code> <code>[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id track interface interface-name],</code> <code>[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id track interface interface-name]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.2 for ACX2000 Universal Access Routers.
Description	Configure a Virtual Router Redundancy Protocol (VRRP) router's priority cost for becoming the master default router. The router with the highest priority within the group becomes the master.
Options	priority —The value subtracted from the configured VRRP priority when the tracked interface or route is down to force a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group. Range: 1 through 254
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Logical Interface to Be Tracked on page 2229• Configuring a Logical Interface to Be Tracked

priority-hold-time

Syntax	<code>priority-hold-time seconds;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id track</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id track</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id track</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id track</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	<p>Configure a Virtual Router Redundancy Protocol (VRRP) router's priority hold time to define the minimum length of time that must elapse between dynamic priority changes. If the dynamic priority changes because of a tracking event, the priority hold timer begins running. If another tracking event or manual configuration change occurs while the timer is running, the new dynamic priority update is postponed until the timer expires.</p>
	<div>  <p>NOTE: When the track feature is configured, and if VRRP should pre-empt due to the tracking interface or route transition, any configured pre-empt hold time will be ignored. VRRP master will pre-empt according to the configuration of the priority-hold time.</p> </div>
Options	<p>seconds—Minimum length of time that must elapse between dynamic priority changes.</p> <p>Range: 0through 3600 seconds</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Logical Interface to Be Tracked on page 2229 • <i>Configuring a Logical Interface to Be Tracked</i>

route (Interfaces)

Syntax	<code>route <i>prefix</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id track</i>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id track</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id track</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id track</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS 11.3 for QFX Series. Statement introduced in Junos OS 12.1 for EX Series switches.
Description	Enable route tracking for a Virtual Router Redundancy Protocol (VRRP) group.
Options	<p><i>prefix</i>—Route to be tracked for this VRRP group.</p> <p><i>priority-cost priority</i>—The value subtracted from the configured VRRP priority when the tracked interface or route is down, forcing a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.</p> <p><i>routing-instance instance-name</i>—Routing instance in which the route is to be tracked. If the route is in the default, or global, routing instance, the value for <i>instance-name</i> must be default.</p>
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Route to Be Tracked on page 2231• <i>Configuring a Route to Be Tracked</i>

skew-timer-disable

Syntax	skew-timer-disable;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 12.2.
Description	Disable the skew timer, thereby reducing the time required to transition from the backup state to the master state.
	<div>  <p>NOTE: The <code>skew-timer-disable</code> statement is used when there is only one master router and one backup router in the network.</p> </div>
Default	By default, the skew timer is enabled for all the VRRP groups.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Improving the Convergence Time for VRRP on page 2114 • Configuring VRRP to Improve Convergence Time on page 2236

startup-silent-period

Syntax	startup-silent-period <i>seconds</i> ;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
Description	Instruct the system to ignore the Master Down Event when an interface transitions from the down state to the up state. This statement is used to avoid incorrect error alarms caused by the delay or interruption of incoming Virtual Router Redundancy Protocol (VRRP) advertisement packets during the interface startup phase.
Options	<p><i>seconds</i>—Number of seconds for the startup period.</p> <p>Default: 4 seconds</p> <p>Range: 1 through 2000 seconds</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Startup Period for VRRP Operations on page 2219 • Configuring the Startup Period for VRRP Operations

traceoptions (Protocols VRRP)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i>> <microsecond-stamp> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; no-remote-trace; }</pre>
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Define tracing operations for the Virtual Router Redundancy Protocol (VRRP) process.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p> <p>By default, VRRP logs the error, dcd configuration, and routing socket events in a file in the directory /var/log.</p>
Default	If you do not include this statement, no VRRP-specific tracing operations are performed.
Options	<p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. By default, VRRP tracing output is placed in the file vrrpd.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. When the maximum number is reached, the oldest trace file is overwritten.</p> <p>Range: 0 through 4,294,967,296 files</p> <p>Default: 3 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. These are the VRRP-specific tracing options:</p> <ul style="list-style-type: none">• all—All VRRP tracing operations• database—Database changes• general—General events• interfaces—Interface changes• normal—Normal events• packets—Packets sent and received

- **state**—State transitions

- **timer**—Timer events

match *regular-expression*—(Optional) Refine the output to include only those lines that match the given regular expression.

microsecond-stamp—(Optional) Provide a timestamp with microsecond granularity.

no-world-readable—(Optional) Restrict users from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes, megabytes, or gigabytes. When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the ***trace-file*** again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your routing platform

Default: 1 MB

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

world-readable—(Optional) Allow users to read the log file.

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

Related Documentation	• Tracing VRRP Operations on page 2241
------------------------------	--

track (VRRP)

Syntax	<pre>track { interface <i>interface-name</i> { bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>; priority-cost <i>priority</i>; } priority-hold-time <i>seconds</i>; route <i>prefix/prefix-length</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>; }</pre>
Hierarchy Level	<pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>priority-hold-time statement added in Junos OS Release 8.1.</p> <p>route statement added in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	Enable logical interface tracking, route tracking, or both, for a Virtual Router Redundancy Protocol (VRRP) group.
Options	The remaining statements are described separately.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring a Logical Interface to Be Tracked on page 2229• Configuring a Route to Be Tracked on page 2231• <i>Configuring a Logical Interface to Be Tracked</i>• <i>Configuring a Route to Be Tracked</i>

version-3

Syntax	version-3;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in Junos OS Release 12.2.
Description	Enable Virtual Router Redundancy Protocol version 3 (VRRPv3).



NOTE:

- Even though the version-3 statement can be configured only at the [edit protocols vrrp] hierarchy level, VRRPv3 is enabled on all the configured logical systems as well.
 - When enabling VRRPv3, you must ensure that VRRPv3 is enabled on all the VRRP routers in the network. This is because VRRPv3 does not interoperate with the previous versions of VRRP.
-

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Junos OS Support for VRRPv3 on page 2111 • <i>VRRP Configuration Hierarchy</i> • <i>VRRP for IPv6 Configuration Hierarchy</i>


virtual-address

Syntax	<code>virtual-address [<i>addresses</i>];</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
Description	Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv4 or IPv6 group. You can configure up to eight addresses.
Options	<i>addresses</i> —Addresses of one or more virtual routers. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the master virtual router for the group.
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Basic VRRP Support on page 2220• Configuring Basic VRRP Support

virtual-inet6-address

Syntax	<code>virtual-inet6-address [<i>addresses</i>];</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You can configure up to eight addresses.
Options	<i>addresses</i> —Addresses of one or more virtual routers. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the master virtual router for the group.
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Basic VRRP Support on page 2220

virtual-link-local-address

Syntax	<code>virtual-link-local-address <i>ipv6-address</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]
Release Information	Statement introduced in Junos OS Release 8.4. Statement introduced in Junos OS 11.3 for the QFX Series.
Description	Configure a virtual link-local address for a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You must explicitly define a virtual link-local address for each VRRP for IPv6 group. The virtual link-local address must be in the same subnet as the physical interface address.
	<div>  <p>NOTE: You do <i>not</i> need to configure link-local addresses and virtual link-local addresses when configuring VRRP for IPv6. Junos OS automatically generates link-local addresses and virtual link-local addresses. However, if link local addresses and virtual link-local addresses are configured, Junos OS considers the configured addresses.</p> </div>
Options	<i>ipv6-address</i> —virtual link-local IPv6 address for VRRP for an IPv6 group. Range: 0 through 255
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Basic VRRP Support on page 2220 • Junos OS Support for VRRPv3 on page 2111

vrrp-group

Syntax	<pre> vrrp-group group-id { (accept-data no-accept-data); advertise-interval seconds; advertisements-threshold number; authentication-key key; authentication-type authentication; fast-interval milliseconds; (preempt no-preempt) { hold-time seconds; } priority number; track { interface interface-name { bandwidth-threshold bits-per-second priority-cost priority; priority-cost priority; } priority-hold-time seconds; route prefix/prefix-length routing-instance instance-name priority-cost priority; } virtual-address [addresses]; vrrp-inherit-from vrrp-group; } </pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
Description	Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 group.
Options	<p>group-id—VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the source-address-filter statement. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 2338. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.</p> <p>Range: 0 through 255</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Basic VRRP Support on page 2220 • Example: Configuring VRRP on page 2237 • Configuring Basic VRRP Support

- *Example: Configuring VRRP for Load Sharing*
- [vrrp-inet6-group on page 2269](#)

vrrp-inet6-group

Syntax	<pre> vrrp-inet6-group group-id { (accept-data no-accept-data); advertisements-threshold number; fast-interval milliseconds; inet6-advertise-interval seconds; (preempt no-preempt) { hold-time seconds; } priority number; track { interface interface-name { bandwidth-threshold bits-per-second priority-cost priority; priority-cost priority; } priority-hold-time seconds; route prefix/prefix-length routing-instance instance-name priority-cost priority; } virtual-inet6-address [addresses]; virtual-link-local-address ipv6-address; vrrp-inherit-from vrrp-group; } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure a Virtual Router Redundancy Protocol (VRRP) IPv6 group.
Options	<p>group-id—VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the source-address-filter statement. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 2338. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.</p> <p>Range: 0 through 255</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Basic VRRP Support on page 2220 • VRRP for IPv6 Configuration Hierarchy

Administration

- [Routine Monitoring on page 2270](#)
- [Operational Commands on page 2274](#)

Routine Monitoring

- [Resetting Local Statistics on page 2270](#)
- [Tracing Restart Signaling-Based Helper Mode Events for OSPF Graceful Restart on page 2270](#)
- [Verifying Graceful Restart Operation on page 2271](#)

Resetting Local Statistics

When you enable graceful Routing Engine switchover, the master Routing Engine configuration is copied and loaded to the backup Routing Engine. User files, accounting information, and trace options information are not replicated to the backup Routing Engine.

When a graceful Routing Engine switchover occurs, local statistics such as process statistics and networking statistics are displayed as a cumulative value from the time the process first came online. Because processes on the master Routing Engine can start at different times from the processes on the backup Routing Engine, the statistics on the two Routing Engines for the same process might differ. After a graceful Routing Engine switchover, we recommend that you issue the **clear interface statistics** (*interface-name* | **all**) command to reset the cumulative values for local statistics. Forwarding statistics are not affected by graceful Routing Engine switchover.

For information about how to use the **clear** command to clear statistics and protocol database information, see the *Junos OS Operational Mode Commands*.



.....

NOTE: The **clear firewall** command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for graceful Routing Engine switchover.

.....

Related Documentation

- [Understanding Graceful Routing Engine Switchover in the Junos OS on page 2061](#)
- [Configuring Graceful Routing Engine Switchover on page 2117](#)

Tracing Restart Signaling-Based Helper Mode Events for OSPF Graceful Restart

Junos OS provides a tracing option to log restart signaling-based helper mode events for OSPF graceful restart. To enable tracing for restart signaling-based helper mode events, include the **traceoptions flag restart-signaling** statement at the **[edit protocols ospf]** hierarchy level.

To enable tracing for restart signaling-based events:

1. Create a log file for saving the log.

```
[edit protocols ospf]
user@host# set traceoptions file ospf-log
```

where *ospf-log* is the name of the log file.

2. Enable tracing for restart signaling-based helper mode events.

```
[edit protocols ospf]
user@host# set traceoptions flag restart-signaling
```

3. Commit the configuration.

```
[edit protocols ospf]
user@host# commit
```

The logs are saved to the *ospf-log* file in the */var/log* folder.

Viewing the Log File

To view the restart signaling-based events from the log file, type:

```
user@host> file show /var/log/ospf-log | match "restart signaling"
Jun 25 14:44:08.890216 OSPF Restart Signaling: Start helper mode for nbr ip
14.19.3.2 id 10.10.10.1
Jun 25 14:44:11.358636 OSPF restart signaling: Received DBD with R bit set from
nbr ip=14.19.3.2 id=10.10.10.1. Start oob-resync.
Jun 25 14:44:11.380198 OSPF restart signaling: Received DBD with LR bit on from
nbr ip=14.19.3.2 id=10.10.10.1. Save its oob-resync capability 1
Jun 25 14:44:11.467200 OSPF restart signaling: nbr fsm for nbr ip=14.19.3.2
id=10.10.10.1 moving to state Full. Reset oob-resync parameters.
```

Related Documentation

- [Understanding Restart Signaling-Based Helper Mode Support for OSPF Graceful Restart](#)
- [Example: Managing Helper Modes for OSPF Graceful Restart on page 2164](#)
- [helper-disable \(OSPF\)](#)

Verifying Graceful Restart Operation

This topic contains the following sections:

- [Graceful Restart Operational Mode Commands on page 2271](#)
- [Verifying BGP Graceful Restart on page 2272](#)
- [Verifying IS-IS and OSPF Graceful Restart on page 2273](#)
- [Verifying CCC and TCC Graceful Restart on page 2273](#)

Graceful Restart Operational Mode Commands

To verify proper operation of graceful restart, use the following commands:

- **show bgp neighbor** (for BGP graceful restart)
- **show log** (for IS-IS and OSPF/OSPFv3 graceful restart)
- **show (ospf | ospfv3) overview** (for OSPF/OSPFv3 graceful restart)

- **show rsvp neighbor detail** (for RSVP graceful restart—helper router)
- **show rsvp version** (for RSVP graceful restart—restarting router)
- **show ldp session detail** (for LDP graceful restart)
- **show connections** (for CCC and TCC graceful restart)
- **show route instance detail** (for Layer 3 VPN graceful restart and for any protocols using graceful restart in a routing instance)
- **show route protocol l2vpn** (for Layer 2 VPN graceful restart)

For more information about these commands and a description of their output fields, see the *Junos OS Operational Mode Commands*.

Verifying BGP Graceful Restart

To view graceful restart information for BGP sessions, use the **show bgp neighbor** command:

```
user@PE1> show bgp neighbor 192.255.10.1
Peer: 192.255.10.1+179 AS 64595 Local: 192.255.5.1+1106 AS 64595
  Type: Internal    State: Established    Flags: <>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Export: [ static ]
  Options:<Preference LocalAddress HoldTime GracefulRestart Damping PeerAS Refresh>

Local Address: 192.255.5.1 Holdtime: 90 Preference: 170
IPSec SA Name: hope
Number of flaps: 0
Peer ID: 192.255.10.1    Local ID: 192.255.5.1    Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 180
Stale routes from peer are kept for: 180
Restart time requested by this peer: 300
NLRI that peer supports restart for: inet-unicast
NLRI that peer saved forwarding for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Table inet.0 Bit: 10000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 0
  Received prefixes: 0
  Suppressed due to damping: 0
Last traffic (seconds): Received 19    Sent 19    Checked 19
Input messages: Total 2    Updates 1    Refreshes 0    Octets 42
Output messages: Total 3    Updates 0    Refreshes 0    Octets 116
Output Queue[0]: 0
```

Verifying IS-IS and OSPF Graceful Restart

To view graceful restart information for IS-IS and OSPF, configure traceoptions (see [“Tracking Graceful Restart Events” on page 2134](#)).

Here is the output of a traceoptions log from an OSPF restarting router:

```
Oct  8 05:20:12 Restart mode - sending grace lsas
Oct  8 05:20:12 Restart mode - estimated restart duration timer triggered
Oct  8 05:20:13 Restart mode - Sending more grace lsas
```

Here is the output of a traceoptions log from an OSPF helper router:

```
Oct  8 05:20:14 Helper mode for neighbor 192.255.5.1
Oct  8 05:20:14 Received multiple grace lsa from 192.255.5.1
```

Verifying CCC and TCC Graceful Restart

To view graceful restart information for CCC and TCC connections, use the **show connections** command. The following example assumes four remote interface CCC connections between CE1 and CE2:

```
user@PE1> show connections
CCC and TCC connections [Link Monitoring On]
Legend for status (St)           Legend for connection types
UN -- uninitialized             if-sw: interface switching
NP -- not present               rmt-if: remote interface switching
WE -- wrong encapsulation       lsp-sw: LSP switching
DS -- disabled
Dn -- down
-> -- only outbound conn is up  Legend for circuit types
<- -- only inbound conn is up  intf -- interface
Up -- operational              tlsp -- transmit LSP
RmtDn -- remote CCC down       rlsp -- receive LSP
Restart -- restarting
```

CCC Graceful restart : Restarting

Connection/Circuit	Type	St	Time last up	# Up trans
CE1-CE2-0	rmt-if	Restart	----	0
fe-1/1/0.0	intf	Up		
PE1-PE2-0	tlsp	Up		
PE2-PE1-0	rlsp	Up		
CE1-CE2-1	rmt-if	Restart	----	0
fe-1/1/0.1	intf	Up		
PE1-PE2-1	tlsp	Up		
PE2-PE1-1	rlsp	Up		
CE1-CE2-2	rmt-if	Restart	----	0
fe-1/1/0.2	intf	Up		
PE1-PE2-2	tlsp	Up		
PE2-PE1-2	rlsp	Up		
CE1-CE2-3	rmt-if	Restart	----	0
fe-1/1/0.3	intf	Up		
PE1-PE2-3	tlsp	Up		
PE2-PE1-3	rlsp	Up		

- Related Documentation**
- [Graceful Restart Concepts on page 2085](#)
 - [Configuring Graceful Restart for QFabric Systems](#)

Operational Commands

show system switchover

Syntax	show system switchover
Syntax (TX Matrix Router)	show system switchover <all-chassis all-lcc lcc <i>number</i> scc>
Syntax (TX Matrix Plus Router)	show system switchover <all-chassis all-lcc lcc <i>number</i> sfc <i>number</i> >
Syntax (MX Series Router)	show system switchover <all-members> <local> <member <i>member-id</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6.
Description	Display whether graceful Routing Engine switchover is configured, the state of the kernel replication (ready or synchronizing), any replication errors, and whether the primary and standby Routing Engines are using compatible versions of the kernel database.



NOTE: Issue the `show system switchover` command *only* on the backup Routing Engine. This command is *not* supported on the master Routing Engine, because the kernel-replication process daemon does not run on the master Routing Engine. This process runs only on the backup Routing Engine.

Beginning Junos OS Release 9.6, the `show system switchover` command has been deprecated on the master Routing Engine on all routers other than a TX Matrix (switch-card chassis) or a TX Matrix Plus (switch-fabric chassis) router.

However, in a routing matrix, if you issue the `show system switchover` command on the master Routing Engine of the TX Matrix router (or switch-card chassis), the CLI displays graceful switchover information for the master Routing Engine of the T640 routers (or line-card chassis) in the routing matrix. Likewise, if you issue the `show system switchover` command on the master Routing Engine of a TX Matrix Plus router (or switch-fabric chassis), the CLI displays output for the master Routing Engine of T1600 routers (or line-card chassis) in the routing matrix.

Options **all-chassis**—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for all Routing Engines on the TX Matrix router and the T640 routers configured in the routing matrix. On a TX Matrix Plus router, display graceful Routing Engine switchover information for all

Routing Engines on the TX Matrix Plus router and the T1600 routers configured in the routing matrix.

all-lcc—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display graceful Routing Engine switchover information for all T1600 routers (or line-card chassis) connected to the TX Matrix Plus router.

all-members—(MX Series routers only) (Optional) Display graceful Routing Engine switchover information for all Routing Engines on all members of the Virtual Chassis configuration.

lcc *number*—(TX Matrix and TX Matrix Plus router only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for a specific T640 router (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display graceful Routing Engine switchover information for a specific T1600 router (or line-card chassis) connected to the TX Matrix Plus router. Replace ***number*** with 0.

local—(MX Series routers only) (Optional) Display graceful Routing Engines switchover information for all Routing Engines on the local Virtual Chassis member.

member *member-id*—(MX Series routers only) (Optional) Display graceful Routing Engine switchover information for all Routing Engines on the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value of 0 or 1.

scc—(TX Matrix router only) (Optional) Display graceful Routing Engine switchover information for the TX Matrix router (or switch-card chassis).

sfc—(TX Matrix Plus router only) (Optional) Display graceful Routing Engine switchover information for the TX Matrix Plus router (or switch-fabric chassis).

Additional Information If you issue the **show system switchover** command on a TX Matrix backup Routing Engine, the command is broadcast to all the T640 backup Routing Engines that are connected to it.

Likewise, if you issue the **show system switchover** command on a TX Matrix Plus backup Routing Engine, the command is broadcast to all the T1600 backup Routing Engines that are connected to it.

If you issue the **show system switchover** command on the active Routing Engine in the master router of an MX Series Virtual Chassis, the router displays an error message that graceful Routing Engine switchover (GRES) is not enabled on this member.

Required Privilege Level

view

List of Sample Output

[show system switchover \(Backup Routing Engine\) on page 2277](#)
[show system switchover all-lcc \(Routing Matrix\) on page 2277](#)

Output Fields Table 166 on page 2277 describes the output fields for the **show system switchover** command. Output fields are listed in the approximate order in which they appear.

Table 166: show system switchover Output Fields

Field Name	Field Description
Graceful switchover	Display graceful Routing Engine switchover status: <ul style="list-style-type: none"> • On—Indicates graceful-switchover is specified for the routing-options configuration command. • Off—Indicates graceful-switchover is not specified for the routing-options configuration command.
Configuration database	State of the configuration database: <ul style="list-style-type: none"> • Ready—Configuration database has synchronized. • Synchronizing—Configuration database is synchronizing. Displayed when there are updates within the last 5 seconds. • Synchronize failed—Configuration database synchronize process failed.
Kernel database	State of the kernel database: <ul style="list-style-type: none"> • Ready—Kernel database has synchronized. • Synchronizing—Kernel database is synchronizing. Displayed when there are updates within the last 5 seconds. • Version incompatible—The primary and standby Routing Engines are running incompatible kernel database versions. • Replication error—An error occurred when the state was replicated from the primary Routing Engine. Inspect <code>/var/log/ksyncd</code> for possible causes, or notify Juniper Networks customer support.
Peer state	Routing Engine peer state: <ul style="list-style-type: none"> • Steady State—Peer completed switchover transition. • Peer Connected—Peer in switchover transition.

Sample Output

show system switchover (Backup Routing Engine)

```
user@host> show system switchover
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
```

show system switchover all-lcc (Routing Matrix)

```
user@host> show system switchover all-lcc
```

```
lcc0-re0:
```

```
-----
Multichassis replication: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
lcc2-re0:
```

```
-----
Multichassis replication: On
```

Configuration database: Ready
Kernel database: Ready
Peer state: Steady State

Troubleshooting

- [Troubleshooting Unified ISSU on page 2278](#)

Troubleshooting Unified ISSU

- [Troubleshooting Unified ISSU Problems on page 2278](#)

Troubleshooting Unified ISSU Problems

If the unified ISSU procedure stops progressing, complete the following steps:

1. Open a new session on the master Routing Engine and issue the **request system software abort in-service-upgrade** command.
2. Check the existing router session to verify that the upgrade has been aborted.

An “ISSU: aborted!” message is provided. Additional system messages provide you with information about where the upgrade stopped and recommendations for the next step to take.

For more information about the **request system software abort in-service-upgrade** command, see the *Junos OS Operational Mode Commands*.

Related Documentation

- [Unified ISSU Concepts on page 2092](#)
- [Unified ISSU Process on the TX Matrix Router](#)
- [Unified ISSU System Requirements on page 2097](#)
- [Best Practices on page 2197](#)
- [Performing a Unified ISSU on page 2200](#)
- [Verifying a Unified ISSU on page 2212](#)
- [Managing and Tracing BFD Sessions During Unified ISSU Procedures on page 2212](#)

CHAPTER 10

Interfaces

- [Overview on page 2279](#)
- [Configuration on page 2285](#)
- [Administration on page 2374](#)

Overview

- [Aggregated Ethernet Overview on page 2279](#)
- [IP-Directed Broadcast Overview on page 2281](#)
- [Reverse Path Forwarding on page 2283](#)

Aggregated Ethernet Overview

- [Aggregated Ethernet Interfaces Overview on page 2279](#)
- [Load Balancing and Ethernet Link Aggregation Overview on page 2281](#)

Aggregated Ethernet Interfaces Overview

Link aggregation of Ethernet interfaces is defined in the IEEE 802.3ad standard. The Junos implementation of 802.3ad balances traffic across the member links within an aggregated Ethernet bundle based on the Layer 3 information carried in the packet. This implementation uses the same load-balancing algorithm used for per-flow load balancing.



NOTE: For information about configuring circuit cross-connects over aggregated Ethernet, see *Circuit and Translational Cross-Connects Overview*.

Platform Support for Aggregated Ethernet Interfaces

You configure an aggregated Ethernet virtual link by specifying the link number as a physical device and then associating a set of ports that have the same speed and are in full-duplex mode. The physical interfaces can be Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ, Gigabit Ethernet IQ2 and IQ2-E, or 10-Gigabit Ethernet IQ2 and IQ2-E. Generally, you cannot use a combination of these interfaces within the same aggregated link; however, you can combine Gigabit Ethernet and Gigabit Ethernet IQ interfaces in a single aggregated Ethernet bundle.

The following routers support a maximum of 16 physical interfaces per single aggregated Ethernet bundle:

- M120
- M320
- All MX Series 3D Universal Edge Routers
- All T Series routers

All other routers support a maximum of 8 physical interfaces per aggregated Ethernet bundle.

On M Series and T Series routers, you can create a maximum of 1024 logical interfaces on an aggregated Ethernet interface.

Aggregated Ethernet interfaces can use interfaces from different FPCs, DPCs, PICs, or MPCs.

Configuration Guidelines for Aggregated Ethernet Interfaces

Simple filters are not supported for interfaces in aggregated Ethernet bundles:

- On M Series routers, simple filters are supported in Gigabit Ethernet Enhanced Intelligent Queuing interfaces only, except when the interface is part of an aggregated Ethernet bundle.
- On MX Series routers, simple filters are supported in Enhanced Queuing Dense Port Concentrator (EQ DPC) interfaces only, except when the interface is part of an aggregated Ethernet bundle.

For more information about simple filters, see the *Junos OS Class of Service Configuration Guide*.

On the aggregated bundle, no IQ-specific capabilities such as MAC accounting, VLAN rewrites, and VLAN queuing are available. For more information about IQ-specific capabilities, see *Gigabit Ethernet Accounting and Policing Overview*.

Use the **show interfaces aggregate-interface extensive** and **show interfaces aggregate.logical-interface** commands to show the bandwidth of the aggregate. Also, the SNMP object identifier **ifSpeed/ifHighSpeed** shows the corresponding bandwidth on the aggregate logical interface if it is configured properly.

Aggregated Ethernet interfaces can be either tagged or untagged, with LACP enabled or disabled. Aggregated Ethernet interfaces on MX Series routers support the configuration of **flexible-vlan-tagging**, **native-vlan-id**, and on dual-tagged frames, which consist of the following configuration statements:

- **inner-tag-protocol-id**
- **inner-vlan-id**
- **pop-pop**
- **pop-swap**

- [push-push](#)
- [swap-push](#)
- [swap-swap](#)

In all cases, you must set the number of aggregated Ethernet interfaces on the chassis. You can also set the link speed and the minimum links in a bundle.

**Related
Documentation**

- [inner-tag-protocol-id on page 1891](#)
- [inner-vlan-id on page 1892](#)
- [pop-pop on page 1916](#)
- [pop-swap on page 1917](#)
- [push-push on page 1920](#)
- [swap-push on page 1926](#)
- [swap-swap on page 1927](#)
- *Gigabit Ethernet Accounting and Policing Overview*
- *Junos® OS Ethernet Interfaces*

Load Balancing and Ethernet Link Aggregation Overview

You can create a link aggregation group (LAG) for a group of Ethernet ports. Layer 2 bridging traffic is load balanced across the member links of this group, making the configuration attractive for congestion concerns as well as for redundancy. You can configure up to 480 LAG bundles on a Juniper Networks MX Series Ethernet Services Router. Each LAG bundle contains up to 16 links.

By default, the hash key mechanism to load-balance frames across LAG interfaces is based on Layer 2 fields (such as frame source and destination address) as well as the input logical interface (unit). The default LAG algorithm is optimized for Layer 2 switching. You can also configure the load balancing hash key for Layer 2 traffic to use fields in the Layer 3 and Layer 4 headers using the **payload** statement, see “[Configuring Load Balancing on a LAG Link](#)” on page 2304. In a Layer 2 switch, one link is overutilized and other links are underutilized.

**Related
Documentation**

- [Configuring Load Balancing on a LAG Link on page 2304](#)
- [Load Balancing on a LAG Link on page 2305](#)
- *payload*

IP-Directed Broadcast Overview

- [Understanding Targeted Broadcast on page 2282](#)

Understanding Targeted Broadcast

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

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

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

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



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

Related Documentation

- [Configuring Targeted Broadcast on page 2313](#)
- [targeted-broadcast on page 2368](#)

Reverse Path Forwarding

- [Understanding Unicast Reverse Path Forwarding on page 2283](#)
- [Understanding Multicast Reverse Path Forwarding on page 2283](#)

Understanding Unicast Reverse Path Forwarding

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

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



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

Related Documentation

- [Example: Configuring Unicast Reverse-Path-Forwarding Check](#)

Understanding Multicast Reverse Path Forwarding

Unicast forwarding decisions are typically based on the destination address of the packet arriving at a router. The unicast routing table is organized by destination subnet and mainly set up to forward the packet toward the destination.

In multicast, the router or switch forwards the packet away from the source to make progress along the distribution tree and prevent routing loops. The router's or switch's multicast forwarding state runs more logically by organizing tables based on the reverse path, from the receiver back to the root of the distribution tree. This process is known as *reverse-path forwarding (RPF)*.

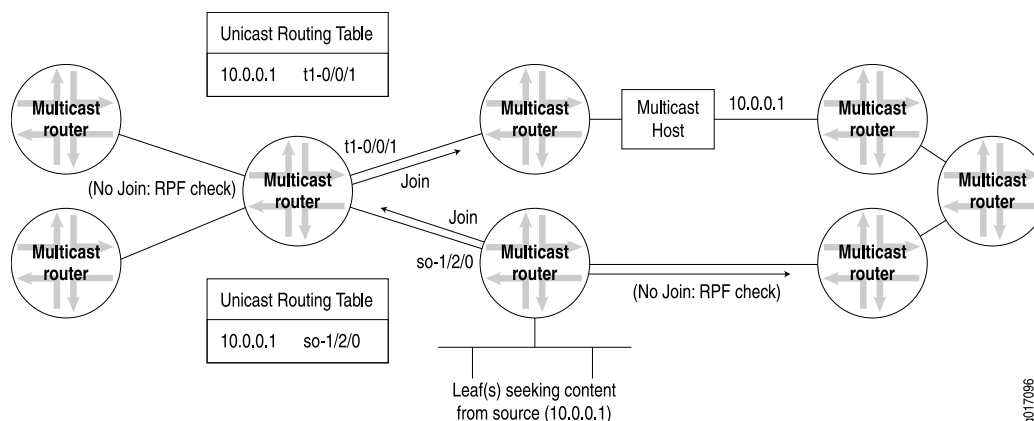
The router or switch adds a branch to a distribution tree depending on whether the request for traffic from a multicast group passes the reverse-path-forwarding check (RPF check). Every multicast packet received must pass an RPF check before it is eligible to be replicated or forwarded on any interface.

The RPF check is essential for every router's multicast implementation. When a multicast packet is received on an interface, the router or switch interprets the source address in the multicast IP packet as the destination address for a unicast IP packet. The source multicast address is found in the unicast routing table, and the outgoing interface is

determined. If the outgoing interface found in the unicast routing table is the same as the interface that the multicast packet was received on, the packet passes the RPF check. Multicast packets that fail the RPF check are dropped because the incoming interface is not on the *shortest path* back to the source.

Figure 27 on page 2284 shows how multicast routers can use the unicast routing table to perform an RPF check and how the results obtained at each router determine where join messages are sent.

Figure 27: Multicast Routers and the RPF Check



Routers and switches can build and maintain separate tables for RPF purposes. The router must have some way to determine its RPF interface for the group, which is the interface topologically closest to the root. For greatest efficiency, the distribution tree follows the shortest-path tree topology. The RPF check helps to construct this tree.

RPF Table

The RPF table plays the key role in the multicast router or switch. The RPF table is consulted for every RPF check, which is performed at intervals on multicast packets entering the multicast router. Distribution trees of all types rely on the RPF table to form properly, and the multicast forwarding state also depends on the RPF table.

RPF checks are performed only on unicast addresses to find the upstream interface for the multicast source or RP.

The routing table used for RPF checks can be the same routing table used to forward unicast IP packets, or it can be a separate routing table used only for multicast RPF checks. In either case, the RPF table contains only unicast routes, because the RPF check is performed on the source address of the multicast packet, not the multicast group destination address, and a multicast address is forbidden from appearing in the source address field of an IP packet header. The unicast address can be used for RPF checks because there is only one source host for a particular stream of IP multicast content for a multicast group address, although the same content could be available from multiple sources.

If the same routing table used to forward unicast packets is also used for the RPF checks, the routing table is populated and maintained by the traditional unicast routing protocols such as BGP, IS-IS, OSPF, and the Routing Information Protocol (RIP). If a dedicated

multicast RPF table is used, this table must be populated by some other method. Some multicast routing protocols (such as the Distance Vector Multicast Routing Protocol [DVMRP]) essentially duplicate the operation of a unicast routing protocol and populate a dedicated RPF table. Others, such as PIM, do not duplicate routing protocol functions and must rely on some other routing protocol to set up this table, which is why PIM is *protocol independent*.

Some traditional routing protocols such as BGP and IS-IS now have extensions to differentiate between different sets of routing information sent between routers and switches for unicast and multicast. For example, there is multiprotocol BGP (MBGP) and multitopology routing in IS-IS (M-IS-IS). IS-IS routes can be added to the RPF table even when special features such as traffic engineering and “shortcuts” are turned on. Multicast Open Shortest Path First (MOSPF) also extends OSPF for multicast use, but goes further than MBGP or M-IS-IS and makes MOSPF into a complete multicast routing protocol on its own. When these routing protocols are used, routes can be tagged as multicast RPF routers and used by the receiving router differently than the unicast routing information.

Using the main unicast routing table for RPF checks provides simplicity. A dedicated routing table for RPF checks allows a network administrator to set up separate paths and routing policies for unicast and multicast traffic, allowing the multicast network to function more independently of the unicast network.

By default, PIM uses **inet.0** as its reverse-path forwarding (RPF) routing table group. PIM uses an RPF routing table group to resolve its RPF neighbor for a particular multicast source address and for the RP address. PIM can optionally use **inet.2** as its RPF routing table group. The **inet.2** routing table is organized more efficiently for RPF checks.

Once configured, the RPF routing table must be applied to a PIM as a routing table group.

Configuration

- [Configuration on page 2285](#)
- [Configuration Statements on page 2320](#)

Configuration

- [Configuring a Layer 2 Virtual Switch on page 2286](#)
- [Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port on page 2287](#)
- [Understanding Layer 2 Virtual Switches Instances on page 2287](#)
- [Configuring VLAN Encapsulation on page 2288](#)
- [Rewriting the Inner and Outer VLAN Tags on page 2289](#)
- [Rewriting the VLAN Tag on Tagged Frames on page 2290](#)
- [Binding VLAN IDs to Logical Interfaces on page 2291](#)
- [Configuring a Logical Interface for Access Mode on page 2296](#)
- [Configuring Junos OS for Supporting Aggregated Devices on page 2296](#)
- [Configuring an Aggregated Ethernet Interface on page 2299](#)
- [Deleting an Aggregated Ethernet Interface on page 2300](#)

- [Configuring the Number of Aggregated Ethernet Interfaces on the Device \(Enhanced Layer 2 Software CLI Procedure\) on page 2300](#)
- [Example: Configuring Aggregated Ethernet Interfaces on page 2301](#)
- [Configuring Tagged Aggregated Ethernet Interfaces on page 2302](#)
- [Configuring Untagged Aggregated Ethernet Interfaces on page 2303](#)
- [Configuring Aggregated Ethernet Minimum Links on page 2303](#)
- [Configuring Load Balancing on a LAG Link on page 2304](#)
- [Example: Configuring Load Balancing on a LAG Link on page 2305](#)
- [Configuring Multichassis Link Aggregation on page 2305](#)
- [Configuring Aggregated Ethernet LACP on page 2307](#)
- [Configuring Targeted Broadcast on page 2313](#)
- [Configuring Unicast RPF on page 2314](#)

Configuring a Layer 2 Virtual Switch

A Layer 2 virtual switch, which isolates a LAN segment with its spanning-tree protocol instance and separates its VLAN ID space, filters and forwards traffic only at the data link layer. Each VLAN consists of a set of logical ports that participate in Layer 2 learning and forwarding. A virtual switch represents a Layer 2 network.

Two main types of interfaces are used in virtual switch hierarchies:

- Layer 2 logical interface—This type of interface uses the VLAN-ID as a virtual circuit identifier and the scope of the VLAN-ID is local to the interface port. This type of interface is often used in service-provider-centric applications.
- Access or trunk interface—This type of interface uses a VLAN-ID with global significance. The access or trunk interface is implicitly associated with VLANs based on VLAN membership. Access or trunk interfaces are typically used in enterprise-centric applications.



NOTE: The difference between access interfaces and trunk interfaces is that access interfaces can be part of one VLAN only and the interface is normally attached to an end-user device (packets are implicitly associated with the configured VLAN). In contrast, trunk interfaces multiplex traffic from multiple VLANs and usually interconnect switches.

To configure a Layer 2 virtual switch, include the following statements:

```
[edit]
routing-instances {
  routing-instance-name (
    instance-type virtual-switch;
    vlans vlan-name{
      vlan-id (all | none | number);
      [...configure optional VLAN parameters]
    }
  }
}
```

```
}
}
```

To enable a virtual switch, you must specify **virtual-switch** as the **instance-type**.

The VLANs that are specified with the **vlan-id** statement are included in the virtual switch.

You can configure other optional VLAN parameters in the virtual switch.

Related Documentation

- [Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port on page 1848](#)

Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port

You can associate one or more Layer 2 trunk interfaces with a virtual switch.

A virtual switch configured with a Layer 2 trunk port also supports IRB within a VLAN. IRB provides simultaneous support for Layer 2 bridging and Layer 3 IP routing on the same interface. Only an interface configured with the **interface-mode (access | trunk)** statement can be associated with a virtual switch. An access interface enables you to accept packets with no VLAN identifier.

In addition, you can configure Layer 2 learning and forwarding properties for the virtual switch.

To configure a virtual switch with a Layer 2 trunk interface, include the following statements:

```
[edit]
routing-instances {
  routing-instance-name {
    instance-type virtual-switch;
    interface interface-name;
    vlans name {
      vlan-id (all | none | number);
      [...configure optional VLAN parameters]
    }
  }
}
```

Related Documentation

- [Configuring a Layer 2 Virtual Switch on page 1847](#)

Understanding Layer 2 Virtual Switches Instances

At Layer 2, you can group one or more VLANs into a single routing instance to form a virtual switch instance. A virtual switch instance is composed of VLANs. The virtual switch instance isolates a LAN segment and contains most Layer 2 functions, such as spanning-tree protocol instances and VLAN ID spaces, into its own smaller, logical network. Splitting Layer 2 traffic using virtual switch instances allows you to more logically organize your Layer 2 traffic into multiple “virtual” Layer 2 networks.

A default virtual switch, called default-switch, is automatically created when a virtual switch is configured. All Layer 2 traffic not assigned to a VLAN in a virtual switch automatically becomes part of the default virtual switch.

You can configure a virtual switch to participate only in Layer 2 bridging and optionally to perform Layer 3 routing. In addition, you can configure spanning-tree protocols (STPs) within the virtual switch to prevent forwarding loops. For more information about how to configure Layer 2 logical ports on an interface, see the *Junos® OS Network Interfaces*.

You can associate one or more logical interfaces configured as trunk interfaces with a virtual switch. A trunk interface, or Layer 2 trunk port, enables you to configure a logical interface to represent multiple VLANs on the physical interface. For more information about how to configure trunk interfaces, see the *Junos® OS Network Interfaces*.

You can also configure Layer 2 forwarding and learning properties for the virtual switch.

Related Documentation

- [Configuring a Layer 2 Virtual Switch on page 1847](#)
- [Configuring a Layer 2 Virtual Switch with a Layer 2 Trunk Port on page 1848](#)
- [Configuring a Layer 2 Control Protocol Routing Instance](#)

Configuring VLAN Encapsulation

To configure encapsulation on an interface, enter the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]  
encapsulation type;
```

The following list contains important notes regarding encapsulation:

- Ethernet interfaces in VLAN mode can have multiple logical interfaces. In CCC and VPLS modes, VLAN IDs from 1 through 511 are reserved for normal VLANs, and VLAN IDs 512 through 4094 are reserved for CCC or VPLS VLANs. For 4-port Fast Ethernet interfaces, you can use VLAN IDs 512 through 1024 for CCC or VPLS VLANs.
- For encapsulation type **flexible-ethernet-services**, all VLAN IDs are valid.
- For some encapsulation types, including flexible Ethernet services, Ethernet VLAN CCC, and VLAN VPLS, you can also configure the encapsulation type that is used inside the VLAN circuit itself. To do this, include the **encapsulation** statement:

```
encapsulation (vlan-ccc | vlan-tcc | vlan-vpls);
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**
- You cannot configure a logical interface with VLAN CCC or VLAN VPLS encapsulation unless you also configure the physical device with the same encapsulation or with flexible Ethernet services encapsulation. In general, the logical interface must have a VLAN ID of 512 or higher; if the VLAN ID is 511 or lower, it will be subject to the normal destination filter lookups in addition to source address filtering. However if you configure flexible Ethernet services encapsulation, this VLAN ID restriction is removed.

In general, you configure an interface's encapsulation at the `[edit interfaces interface-name]` hierarchy level.

Example: Configuring VLAN Encapsulation on a Gigabit Ethernet Interface

Configure VLAN CCC encapsulation on a Gigabit Ethernet interface:

```
interfaces ge-2/1/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  unit 0 {
    encapsulation vlan-ccc;
    vlan-id 600;
  }
}
```

Example: Configuring VLAN Encapsulation on an Aggregated Ethernet Interface

Configure VLAN CCC encapsulation on an aggregated Gigabit Ethernet interface:

```
interfaces ae0 {
  vlan-tagging;
  encapsulation vlan-vpls;
  unit 0 {
    vlan-id 100;
  }
}
```

- Related Documentation**
- [802.1Q VLANs Overview](#)
 - [Junos® OS Ethernet Interfaces](#)

Rewriting the Inner and Outer VLAN Tags

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to replace both the inner and the outer VLAN tags of the incoming frame with a user-specified VLAN tag value, include the **swap-swap** statement in the input VLAN map or output VLAN map:

```
swap-swap;
```

You can include this statement at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number input-vlan-map]`
- `[edit interfaces interface-name unit logical-unit-number output-vlan-map]`
- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number input-vlan-map]`
- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number output-vlan-map]`

See *Configuring Inner and Outer TPIDs and VLAN IDs* and *Configuring Inner and Outer TPIDs and VLAN IDs* for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

- Related Documentation
- [input-vlan-map on page 1893](#)
 - [output-vlan-map on page 1911](#)
 - [swap-swap on page 1927](#)
 - *unit*
 - *Junos® OS Ethernet Interfaces*

Rewriting the VLAN Tag on Tagged Frames

To rewrite the VLAN tag on all tagged frames entering the interface to a specified VLAN ID and TPID, include the **swap**, **tag-protocol-id**, and **vlan-id** statements in the input VLAN map:

```
input-vlan-map {  
  swap;  
  vlan-id number;  
  tag-protocol-id tpid;  
}
```

To rewrite the VLAN tag on all tagged frames exiting the interface to a specified VLAN ID and TPID, include the **swap** and **tag-protocol-id** statements in the output VLAN map:

```
output-vlan-map {  
  swap;  
  vlan-id number;  
  tag-protocol-id tpid;  
}
```

You can include these statements at the following hierarchy levels:

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

You cannot include both the **swap** statement and the **vlan-id** statement in the output VLAN map configuration. If you include the **swap** statement in the configuration, the VLAN ID in outgoing frames is rewritten to the VLAN ID bound to the logical interface. For more information about binding a VLAN ID to the logical interface, see *802.1Q VLANs Overview*.

The swap operation works on the outer tag only, whether or not you include the **stacked-vlan-tagging** statement in the configuration. For more information, see *Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags*.

- Related Documentation
- [input-vlan-map on page 1893](#)
 - [output-vlan-map on page 1911](#)

- [swap on page 1925](#)
- [vlan-id on page 1934](#)
- [tag-protocol-id on page 1929](#)
- *unit*
- For more information about binding a VLAN ID to the logical interface, see *802.1Q VLANs Overview*.
- For more information about the swap operation, see *Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags*.
- *Junos® OS Ethernet Interfaces*

Binding VLAN IDs to Logical Interfaces

The following sections describe how to configure logical interfaces to receive and forward VLAN-tagged frames:

- [Binding VLAN IDs to Logical Interfaces Overview on page 2291](#)
- [Binding a VLAN ID to a Logical Interface on page 2292](#)
- [Binding a Range of VLAN IDs to a Logical Interface on page 2292](#)
- [Binding a List of VLAN IDs to a Logical Interface on page 2294](#)

Binding VLAN IDs to Logical Interfaces Overview

To configure a logical interface to receive and forward VLAN-tagged frames, you must bind a VLAN ID, a range of VLAN IDs, or a list of VLAN IDs to the logical interface.

[Table 167 on page 2291](#) lists the configuration statements you use to bind VLAN IDs to logical interfaces, organized by scope of the VLAN IDs used to match incoming packets:

Table 167: Configuration Statements Used to Bind VLAN IDs to Logical Interfaces

Scope of VLAN ID Matching	Type of VLAN Framing Supported on the Logical Interface	
	Single-Tag Framing	Dual-Tag Framing
VLAN ID	<code>vlan-id <i>vlan-id</i>;</code>	<code>vlan-tags outer <i>tpid</i>.<<i>vlan-id</i>> inner <i>tpid</i><i>vlan-id</i>;</code>
VLAN ID Range	<code>vlan-id-range <i>vlan-id-vlan-id</i>;</code>	<code>vlan-tags outer <i>tpid.vlan-id</i> inner-range <i>tpid.vlan-id-vlan-id</i>;</code>
VLAN ID List	<code>vlan-id-list [<i>vlan-id</i> <i>vlan-id-vlan-id</i>];</code>	<code>vlan-tags outer <<i>tpid</i>.><i>vlan-id</i> inner-list [<i>vlan-id</i> <i>vlan-id-vlan-id</i>];</code>

You can include all of the statements at the following hierarchy levels:

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



NOTE: The inner-list option of the `vlan-tags` statement does not support Tag Protocol ID (TPID) values.

Binding a VLAN ID to a Logical Interface

A logical interface that you have associated (bound) to a particular VLAN ID will receive and forward incoming frames that contain a matching VLAN ID.

Binding a VLAN ID to a Single-Tag Logical Interface

To bind a VLAN ID to a single-tag logical interface, include the `vlan-id` statement:

```
vlan-id vlan-id;
```

You can include the statement at the following hierarchy levels:

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

To configure an Ethernet interface to support single-tag logical interfaces, include the `vlan-tagging` statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the `flexible-vlan-tagging` statement instead.

Binding a VLAN ID to a Dual-Tag Logical Interface

To bind a VLAN ID to a dual-tag logical interface, include the `vlan-tags` statement:

```
vlan-tags inner <tpid>vlan-id outer <tpid>vlan-id;
```

You can include the statement at the following hierarchy levels:

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

To configure an Ethernet interface to support dual-tag logical interfaces, include the `stacked-vlan-tagging` statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the `flexible-vlan-tagging` statement instead.

Binding a Range of VLAN IDs to a Logical Interface

A VLAN range can be used by service providers to interconnect multiple VLANs belonging to a particular customer over multiple sites. Using a VLAN ID range conserves switch resources and simplifies configuration.

Binding a Range of VLAN IDs to a Single-Tag Logical Interface

To bind a range of VLAN IDs to a single-tag logical interface, include the `vlan-id-range` statement:

```
vlan-id-range vlan-id-vlan-id;
```

You can include the statement at the following hierarchy levels:

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

To configure an Ethernet interface to support single-tag logical interfaces, include the **vlan-tagging** statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

Binding a Range of VLAN IDs to a Dual-Tag Logical Interface

To bind a range of VLAN IDs to a dual-tag logical interface, include the **vlan-tags** statement. Use the **inner-list** option to specify the VLAN IDs as an inclusive range by separating the starting VLAN ID and ending VLAN ID with a hyphen.

```
vlan-tags inner-list [ vlan-id vlan-id-vlan-id ] outer <tpid.>vlan-id;
```

You can include the statement at the following hierarchy levels:

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

To configure an Ethernet interface to support dual-tag logical interfaces, include the **stacked-vlan-tagging** statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

Example: Binding Ranges VLAN IDs to Logical Interfaces

The following example configures two different ranges of VLAN IDs on two different logical ports:

```
[edit interfaces]
ge-3/0/0 {
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 500-600;
  }
}
ge-3/0/1 {
  flexible-vlan-tagging;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 200-300;
  }
  unit 1 {
    encapsulation vlan-bridge;
    vlan-tags outer 1000 inner-range 100-200;
  }
}
```

Binding a List of VLAN IDs to a Logical Interface

In Junos OS Release 9.5 and later, on MX Series routers and in Junos OS Release 12.2R2 and later on EX Series switches, you can bind a list of VLAN IDs to a single logical interface, eliminating the need to configure a separate logical interface for every VLAN or VLAN range. A logical interface that accepts packets tagged with any VLAN ID specified in a VLAN ID list is called a *VLAN-bundled* logical interface.

You can use VLAN-bundled logical interfaces to configure circuit cross-connects between Layer 2 VPN routing instances or Layer 2 circuits. Using VLAN-bundled logical interfaces simplifies configuration and reduces use of system resources such as logical interfaces, next hops, and circuits.

As an alternative to configuring multiple logical interfaces (one for each VLAN ID and one for each range of VLAN IDs), you can configure a single VLAN-bundled logical interface based on a list of VLAN IDs.



NOTE: The `vlan-id` option is not supported to achieve VLAN normalization on VPLS instances that are configured with `vlan-id-list`. However, you can use the `vlan-maps` option to achieve VLAN normalization.

Binding a List of VLAN IDs to a Single-Tag Logical Interface

To bind a list of VLAN IDs to a single-tag logical interface, include the `vlan-id-list` statement. Specify the VLAN IDs in the list individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both.

```
vlan-id-list [ vlan-id vlan-id–vlan-id ];
```

You can include the statement at the following hierarchy levels:

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

To configure an Ethernet interface to support single-tag logical interfaces, include the `vlan-tagging` statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the `flexible-vlan-tagging` statement instead.

Binding a List of VLAN IDs to a Dual-Tag Logical Interface

To bind a list of VLAN IDs to a dual-tag logical interface, include the `vlan-tags` statement. Use the `inner-list` option to specify the VLAN IDs individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both:

```
vlan-tags inner-list [ vlan-id vlan-id–vlan-id ] outer <tpid>vlan-id;
```



NOTE: The `inner-list` option of the `vlan-tags` statement does not support Tag Protocol ID (TPID) values.

You can include the statement at the following hierarchy levels:

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

To configure an Ethernet interface to support dual-tag logical interfaces, include the **stacked-vlan-tagging** statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the **flexible-vlan-tagging** statement instead.

Example: Binding Lists of VLAN IDs to Logical Interfaces

The following example configures two different lists of VLAN IDs on two different logical ports:

```
[edit interfaces]
ge-1/1/0 {
  vlan-tagging; # Only for single-tagging
  encapsulation flexible-ethernet-services;
  unit 10 {
    encapsulation vlan-ccc;
    vlan-id-list [20 30–40 45];
  }
}
ge-1/1/1 {
  flexible-vlan-tagging; # Only for mixed tagging
  encapsulation flexible-ethernet-services;
  unit 10 {
    encapsulation vlan-ccc;
    vlan-id-list [1 10 20 30–40];
  }
  unit 20 {
    encapsulation vlan-ccc;
    vlan-tags outer 200 inner-list [50–60 80 90–100];
  }
}
```

In the example configuration above, **ge-1/1/0** supports single-tag logical interfaces, and **ge-1/1/1** supports mixed tagging. The single-tag logical interfaces **ge-1/1/0.10** and **ge-1/1/1.20** each bundle lists of VLAN IDs. The dual-tag logical interface **ge-1/1/1.20** bundles lists of inner VLAN IDs.



TIP: You can group a range of identical interfaces into an interface range and then apply a common configuration to that interface range. For example, in the above example configuration, both interfaces **ge-1/1/0** and **ge-1/1/1** have the same physical encapsulation type of **flexible-ethernet-services**. Thus you can define an interface range with the interfaces **ge-1/1/0** and **ge-1/1/1** as its members and apply the encapsulation type **flexible-ethernet-services** to that defined interface range. For more information about interface ranges, see *Configuring Interface Ranges*.

- Related Documentation**
- [802.1Q VLANs Overview](#)
 - [Configuring Interface Ranges](#)
 - [Junos® OS Ethernet Interfaces](#)

Configuring a Logical Interface for Access Mode

Enterprise network administrators can configure a single logical interface to accept untagged packets and forward the packets within a specified VLAN. A logical interface configured to accept untagged packets is called an *access interface* or *access port*.

`interface-mode access;`

You can include this statement at the following hierarchy levels:

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

When an untagged or tagged packet is received on an access interface, the packet is accepted, the VLAN ID is added to the packet, and the packet is forwarded within the VLAN that is configured with the matching VLAN ID.

The following example configures a logical interface as an access port with a VLAN ID of 20 on routers and switches that support the enhanced Layer 2 software:

```
[edit interfaces ge-1/2/0]
unit 1 {
  family ethernet-switching {
    interface-mode access;
    vlan members 20;
  }
}
```

- Related Documentation**
- [802.1Q VLANs Overview](#)
 - [Junos® OS Ethernet Interfaces](#)

Configuring Junos OS for Supporting Aggregated Devices

Junos OS supports the aggregation of physical devices into defined virtual links, such as the link aggregation of Ethernet interfaces defined by the IEEE 802.3ad standard.

Tasks for configuring aggregated devices are:

- [Configuring Virtual Links for Aggregated Devices on page 2297](#)
- [Configuring LACP Link Protection at the Chassis Level on page 2297](#)
- [Enabling LACP Link Protection on page 2298](#)
- [Configuring System Priority on page 2298](#)
- [Configuring the Maximum Links Limit on page 2298](#)

Configuring Virtual Links for Aggregated Devices

To define virtual links, you need to specify the associations between physical and logical devices within the **[edit interfaces]** hierarchy, and assign the correct number of logical devices by including the **device-count** statement at the **[edit chassis aggregated-devices ethernet]** and **[edit chassis aggregated-devices sonet]** hierarchy levels:

```
[edit chassis]
aggregated-devices {
  ethernet {
    device-count number;
  }
  sonet {
    device-count number;
  }
}
```

The maximum number of Ethernet logical interfaces that you can configure is 128. On M Series and T Series routers, you can configure a maximum number of 128 aggregated interfaces. On MX Series routers, you can configure a maximum of 480 aggregated interfaces. The aggregated interfaces are numbered from **ae0** through **ae127** for M Series and T Series routers, and the aggregated interfaces (LAG bundles) are numbered from **ae0** through **ae479** on MX Series routers. The maximum number of SONET/SDH logical interfaces is 16. The aggregated SONET/SDH interfaces are numbered from **as0** through **as15**.

Configuring LACP Link Protection at the Chassis Level

Link Aggregation Control Protocol (LACP) is one method of bundling several physical interfaces to form one logical interface. You can configure both VLAN-tagged and untagged aggregated Ethernet with or without LACP enabled. LACP exchanges are made between actors and partners. An actor is the local interface in an LACP exchange. A partner is the remote interface in an LACP exchange.

LACP link protection enables you to force active and standby links within an aggregated Ethernet. You configure LACP link protection by using the **link-protection** and **system-priority** statements at either the chassis or interface level and by configuring port priority at the interface level using the **system-priority** statement. Configuring LACP parameters at the chassis level results in all aggregated Ethernet interfaces using the defined values unless overridden by the LACP configuration on a specific interface.

```
[edit chassis]
aggregated-devices {
  ethernet {
    lacp {
      link-protection {
        non-revertive;
      }
      system-priority priority;
    }
  }
}
```



NOTE: LACP link protection also uses port priority. You can configure port priority at the Ethernet interface `[gigether-options]` hierarchy level using the `port-priority` statement. If you choose not to configure port priority, LACP link protection uses the default value for port priority (127).

Enabling LACP Link Protection

To enable LACP link protection for aggregated Ethernet interfaces on the chassis, use the `link-protection` statement at the `[edit chassis aggregated-devices ethernet lacp]` hierarchy level:

```
[edit chassis aggregated-devices ethernet lacp]
link-protection {
  non-revertive;
}
```

By default, LACP link protection reverts to a higher-priority (lower-numbered) link when that higher-priority link becomes operational or a link is added to the aggregator that is determined to be higher in priority. However, you can suppress link calculation by adding the `non-revertive` statement to the LACP link protection configuration. In nonrevertive mode, after a link is active and collecting and distributing packets, the subsequent addition of a higher-priority (better) link does not result in a switch, and the current link remains active.



CAUTION: If both ends of an aggregator have LACP link protection enabled, make sure to configure both ends of the aggregator to use the same mode. Mismatching LACP link protection modes can result in lost traffic.

Configuring System Priority

To configure LACP system priority for aggregated Ethernet interfaces on the chassis, use the `system-priority` statement at the `[edit chassis aggregated-devices ethernet lacp]` hierarchy level:

```
[edit chassis aggregated-devices ethernet lacp]
system-priority priority;
```

The system priority is a 2-octet binary value that is part of the LACP system ID. The LACP system ID consists of the system priority as the two most-significant octets and the interface MAC address as the six least-significant octets. The system with the numerically lower value for system priority has the higher priority. By default, system priority is 127, with a range of 0 through 65,535.

Configuring the Maximum Links Limit

To configure the maximum links limit, use the `maximum-links` statement at the `[edit chassis aggregated-devices]` hierarchy level:

```
[edit chassis aggregated-devices]
maximum-links maximum-links-limit;
```


Related Documentation

- [Configuring an Aggregated Ethernet Interface](#)
- [Junos® OS Ethernet Interfaces](#)
- [Configuring Aggregated Ethernet Interfaces on PTX Series Packet Transport Switches](#)

Configuring an Aggregated Ethernet Interface

You can associate a physical interface with an aggregated Ethernet interface.

To configure an aggregated Ethernet interface:

1. Specify that you want to configure the link aggregation group interface.

```
user@host# edit interfaces interface-name
```

2. Configure the aggregated Ethernet interface.

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

You specify the interface instance number *x* to complete the link association; *x* can be from 0 through 480, for a total of 480 aggregated interfaces on MX Series routers or EX9200 switches. You must also include a statement defining **aex** at the **[edit interfaces]** hierarchy level. You can optionally specify other physical properties that apply specifically to the aggregated Ethernet interfaces; for details, see *Ethernet Interfaces Overview*.



NOTE: In general, aggregated Ethernet bundles support the features available on all supported interfaces that can become a member link within the bundle. As an exception, Gigabit Ethernet IQ features and some newer Gigabit Ethernet features are not supported in aggregated Ethernet bundles.

Gigabit Ethernet IQ and SFP interfaces can be member links, but IQ- and SFP-specific features are not supported on the aggregated Ethernet bundle even if all the member links individually support those features.

You need to configure the correct link speed for the aggregated Ethernet interface to eliminate any warning message.



NOTE: Before you commit an aggregated Ethernet configuration, ensure that link mode is not configured on any member interface of the aggregated Ethernet bundle; otherwise, the configuration commit check fails.

Related Documentation

- [Configuring an Aggregated Ethernet Interface on page 2299](#)
- [Configuring the Number of Aggregated Ethernet Interfaces on the Device \(Enhanced Layer 2 Software CLI Procedure\) on page 2300](#)
- [Deleting an Aggregated Ethernet Interface on page 2300](#)
- [Aggregated Ethernet Interfaces Overview on page 2279](#)

- *Junos® OS Ethernet Interfaces*

Deleting an Aggregated Ethernet Interface

There are two approaches to deleting an aggregated Ethernet interface:

- You can delete an aggregated Ethernet interface from the interface configuration. The Junos OS removes the configuration statements related to **aex** and sets this interface to down state.
- You can also permanently remove the aggregated Ethernet interface from the device configuration by deleting it from the device-count on the routing device.

To delete an aggregated Ethernet interface:

1. Delete the aggregated Ethernet configuration.

This step changes the interface state to down and removing the configuration statements related to **aex**.

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

2. Delete the interface from the device count.

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

Related Documentation

- *Configuring an Aggregated Ethernet Interface*
- *Configuring the Number of Aggregated Ethernet Interfaces on the Device*
- [Aggregated Ethernet Interfaces Overview on page 2279](#)
- *Junos® OS Ethernet Interfaces*

Configuring the Number of Aggregated Ethernet Interfaces on the Device (Enhanced Layer 2 Software CLI Procedure)

By default, no aggregated Ethernet interfaces are created. You must set the number of aggregated Ethernet interfaces on the routing device before you can configure them.

On MX Series routers and EX9200 switches, you can configure a maximum of 480 aggregated interfaces. The aggregated interfaces (LAG bundles) are numbered from **ae0** through **ae479** on MX Series routers and EX9200 switches.

1. Specify that you want to access the aggregated Ethernet configuration on the device.

```
user@host# edit chassis aggregated-devices ethernet
```

2. Set the number of aggregated Ethernet interfaces.

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

You must also specify the constituent physical links by including the **802.3ad** statement at the `[edit interfaces interface-name ether-options]` or `[edit interfaces interface-name ether-options]` hierarchy level.

Related Documentation

- For information about physical links, see [Configuring an Aggregated Ethernet Interface on page 2299](#)
- *Junos® OS Ethernet Interfaces*
- For information about configuring aggregated devices, see the *Junos OS System Basics Configuration Guide*.

Example: Configuring Aggregated Ethernet Interfaces

Aggregated Ethernet interfaces can use interfaces from different FPCs, DPCs, or PICs. The following configuration is sufficient to get an aggregated Gigabit Ethernet interface up and running.

```
[edit chassis]
aggregated-devices {
  ethernet {
    device-count 15;
  }
}

[edit interfaces]
ge-1/3/0 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-2/0/1 {
  gigether-options {
    802.3ad ae0;
  }
}
ae0 {
  aggregated-ether-options {
    link-speed 1g;
    minimum-links 1;
  }
}
vlan-tagging;
unit 0 {
  vlan-id 1;
  family inet {
    address 14.0.100.50/24;
  }
}
unit 1 {
  vlan-id 1024;
  family inet {
    address 14.0.101.50/24;
  }
}
```

```
unit 2 {  
    vlan-id 1025;  
    family inet {  
        address 14.0.102.50/24;  
    }  
}  
unit 3 {  
    vlan-id 4094;  
    family inet {  
        address 14.0.103.50/24;  
    }  
}
```

- Related Documentation**
- [Junos® OS Ethernet Interfaces](#)
 - [Configure 'link-speed' for Gigabit Ethernet based Aggregate Ethernet interface bundles](#)

Configuring Tagged Aggregated Ethernet Interfaces

To specify aggregated Ethernet interfaces, include the **vlan-tagging** statement at the **[edit interfaces aex]** hierarchy level:

```
[edit interfaces aex]  
vlan-tagging;
```

You must also include the **vlan-id** statement:

```
vlan-id number;
```

You can include this statement at the following hierarchy levels:

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

For more information about the **vlan-tagging** and **vlan-id** statements, see *802.1Q VLANs Overview*.

- Related Documentation**
- [vlan-id](#)
 - [vlan-tagging on page 2373](#)

Configuring Untagged Aggregated Ethernet Interfaces

When you configure an untagged Aggregated Ethernet interface, the existing rules for untagged interfaces apply. These rules are as follows:

- You can configure only one logical interface (unit 0) on the port. The logical unit 0 is used to send and receive LACP or marker protocol data units (PDUs) to and from the individual links.
- You cannot include the **vlan-id** statement in the configuration of the logical interface.

Configure an untagged aggregated Ethernet interface by omitting the **vlan-tagging** and **vlan-id** statements from the configuration:

```
[edit interfaces]
ge-1/1/1 {
  ether-options {
    802.3ad ae0;
  }
}
ae0 {
  # vlan-tagging; OMIT FOR UNTAGGED AE CONFIGURATIONS
  unit 0 {
    # vlan-id 100; OMIT FOR UNTAGGED AE CONFIGURATIONS
    family inet {
      address 13.1.1.2/24 {
        vrrp-group 0 {
          virtual-address 13.1.1.4;
          priority 200;
        }
      }
    }
  }
}
```

Related Documentation

- For more information about configuring LACP, see [Configuring Aggregated Ethernet LACP on page 2307](#).
- *Junos® OS Ethernet Interfaces*

Configuring Aggregated Ethernet Minimum Links

On aggregated Ethernet interfaces, you can configure the minimum number of links that must be up for the bundle as a whole to be labeled **up**. By default, only one link must be up for the bundle to be labeled **up**.

To configure the minimum number of links:

1. Specify that you want to configure the aggregated Ethernet options.

```
user@host# edit interfaces interface-name aggregated-ether-options
```
2. Configure the minimum number of links.

```
[edit interfaces interface-name aggregated-ether-options]
```

`user@host# set minimum-links number`

On M120, M320, MX Series, T Series, and TX Matrix routers with Ethernet interfaces, and EX 9200 switches, the valid range for **minimum-links *number*** is 1 through 16. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

On all other routers and on EX Series switches, other than EX8200 switches, the range of valid values for **minimum-links *number*** is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

On EX8200 switches, the range of valid values for **minimum-links *number*** is 1 through 12. When the maximum value (12) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

If the number of links configured in an aggregated Ethernet interface is less than the minimum link value configured under the **aggregated-ether-options** statement, the configuration commit fails and an error message is displayed.

**Related
Documentation**

- [aggregated-ether-options](#)
- [minimum-links](#)
- [Junos® OS Ethernet Interfaces](#)

Configuring Load Balancing on a LAG Link

You can configure the load balancing hash key for Layer 2 traffic to use fields in the Layer 3 and Layer 4 headers inside the frame payload for load-balancing purposes using the **payload** statement. You can configure the statement to look at **layer-3** (and **source-ip-only** or **destination-ip-only** packet header fields) or **layer-4** fields. You configure this statement at the **[edit forwarding-options hash-key family multiservice]** hierarchy level.

You can configure Layer 3 or Layer 4 options, or both. The **source-ip-only** or **destination-ip-only** options are mutually exclusive. The **layer-3-only** statement is not available on MX Series routers.



NOTE: For Dense Port Concentrators (DPC), any change in the hash key configuration requires a reboot for the changes to take effect. For Modular Port Concentrators (MPC) the reboot is not required.

For more information about link aggregation group (LAG) configuration, see the *Junos® OS Network Interfaces*.

**Related
Documentation**

- [Load Balancing and Ethernet Link Aggregation on page 2281](#)
- [Load Balancing on a LAG Link on page 2305](#)

Example: Configuring Load Balancing on a LAG Link

This example configures the load-balancing hash key to use the source Layer 3 IP address option and Layer 4 header fields as well as the source and destination MAC addresses for load balancing on a link aggregation group (LAG) link:

```
[edit]
forwarding-options {
  hash-key {
    family multiservice {
      source-mac;
      destination-mac;
      payload {
        ip {
          layer-3 {
            source-ip-only;
          }
          layer-4;
        }
      }
    }
  }
}
```

Related Documentation

- [Load Balancing and Ethernet Link Aggregation on page 2281](#)
- [Configuring Load Balancing on a LAG Link on page 2304](#)

Configuring Multichassis Link Aggregation

On MX Series routers and EX Series switches, multichassis link aggregation (MC-LAG) enables a device to form a logical LAG interface with two or more other devices. MC-LAG provides additional benefits over traditional LAG in terms of node level redundancy, multi-homing support, and loop-free Layer 2 network without running Spanning Tree Protocol (STP). MC-LAG can be configured for VPLS routing instance, CCC application, and Layer 2 circuit encapsulation types.

The MC-LAG devices use Inter-Chassis Communication Protocol (ICCP) to exchange the control information between two MC-LAG network devices.

On one end of MC-LAG is a MC-LAG client device that has one or more physical links in a link aggregation group (LAG). This client device does not need to be aware of MC-LAG. On the other side of MC-LAG are two MC-LAG network devices. Each of these network devices has one or more physical links connected to a single client device. The network devices coordinate with each other to ensure that data traffic is forwarded properly.

MC-LAG includes the following functionality:

- Active standby mode is supported using Link Aggregation Control Protocol (LACP)
- MC-LAG operates only between two chassis.
- Layer 2 circuit functions are supported with **ether-ccc** encapsulation.
- VPLS functions are supported with **ether-vpls** and **vlan-vpls**.



NOTE: Ethernet connectivity fault management (CFM) specified in IEEE 802.1ag standard for Operation, Administration, and Management (OAM) is *not* supported on MC-LAG interfaces.

To enable MC-LAG, include the **mc-ae** statement at the **[edit interfaces aeX aggregated-ether-options]** hierarchy level along with either the **ethernet-bridge**, **encapsulation ethernet-ccc**, **encapsulation ethernet-vpls**, or **flexible-ethernet-services** statement at the **[edit interfaces aeX]** hierarchy level. You also need to configure the **lACP** statement and the **admin-key** and **system-id** statements at the **[edit interfaces aeX aggregated-ether-options]** hierarchy level:

```
[edit interfaces aeX]
encapsulation (ethernet-bridge | ethernet-ccc | ethernet-vpls | flexible-ethernet-services);
aggregated-ether-options {
  lACP {
    active;
    admin-key number;
    system-id mac-address;
    system-priority number;
  }
  mc-ae {
    chassis-id chassis-id;
    events {
      iccp-peer-down {
        force-icl-down;
        prefer-status-control-active;
      }
    }
    mc-ae-id mc-ae-id;
    mode (active-active | active-standby);
    redundancy-group group-id;
    status-control (active | standby);
  }
}
```



NOTE: When you configure the **prefer-status-control-active** statement, you must also configure the **status-control active** statement. If you configure the **status-control standby** statement with the **prefer-status-control-active** statement, the system issues a warning.

To delete a MC-LAG interface from the configuration, issue the **delete interfaces aeX aggregated-ether-options mc-ae** command at the **[edit]** hierarchy level in configuration mode:

```
[edit]
user@host# delete interfaces aeX aggregated-ether-options mc-ae
```

Related Documentation

- [Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview](#)
- [Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation](#)

- `show interfaces mc-ae`
- *Junos® OS Ethernet Interfaces*

Configuring Aggregated Ethernet LACP

For aggregated Ethernet interfaces, you can configure the Link Aggregation Control Protocol (LACP). LACP is one method of bundling several physical interfaces to form one logical interface. You can configure both VLAN-tagged and untagged aggregated Ethernet with or without LACP enabled.

For Multichassis Link Aggregation (MC-LAG), you must specify the **system-id** and **admin key**. MC-LAG peers use the same **system-id** while sending the LACP messages. The **system-id** can be configured on the MC-LAG network device and synchronized between peers for validation.

LACP exchanges are made between actors and partners. An actor is the local interface in an LACP exchange. A partner is the remote interface in an LACP exchange.

LACP is defined in IEEE 802.3ad, *Aggregation of Multiple Link Segments*.

LACP was designed to achieve the following:

- Automatic addition and deletion of individual links to the aggregate bundle without user intervention
- Link monitoring to check whether both ends of the bundle are connected to the correct group

The Junos OS implementation of LACP provides link monitoring but not automatic addition and deletion of links.

The LACP mode can be active or passive. If the actor and partner are both in passive mode, they do not exchange LACP packets, which results in the aggregated Ethernet links not coming up. If either the actor or partner is active, they do exchange LACP packets. By default, LACP is turned off on aggregated Ethernet interfaces. If LACP is configured, it is in passive mode by default. To initiate transmission of LACP packets and response to LACP packets, you must configure LACP in active mode.

To enable LACP active mode, include the **lACP** statement at the **[edit interfaces interface-name aggregated-ether-options]** hierarchy level, and specify the **active** option:

```
[edit interfaces interface-name aggregated-ether-options]
lACP {
  active;
}
```



NOTE: The LACP process exists in the system only if you configure the system in either active or passive LACP mode.

To restore the default behavior, include the **lACP** statement at the **[edit interfaces *interface-name* aggregated-ether-options]** hierarchy level, and specify the **passive** option:

```
[edit interfaces interface-name aggregated-ether-options]
lACP {
  passive;
}
```

Starting with Junos OS release 12.2, you can also configure LACP to override the IEEE 802.3ad standard and to allow the standby link always to receive traffic. Overriding the default behavior facilitates subsecond failover.

To override the IEEE 802.3ad standard and facilitate subsecond failover, include the **fast-failover** statement at the **[edit interfaces *interface-name* aggregated-ether-options lACP]** hierarchy level.

For more information, see the following sections:

- [Configuring the LACP Interval on page 2308](#)
- [Configuring LACP Link Protection on page 2309](#)
- [Configuring LACP System Priority on page 2310](#)
- [Configuring LACP System Identifier on page 2310](#)
- [Configuring LACP administrative Key on page 2310](#)
- [Configuring LACP Port Priority on page 2311](#)
- [Tracing LACP Operations on page 2311](#)
- [LACP Limitations on page 2312](#)
- [Example: Configuring Aggregated Ethernet LACP on page 2312](#)

Configuring the LACP Interval

By default, the actor and partner send LACP packets every second. You can configure the interval at which the interfaces send LACP packets by including the **periodic** statement at the **[edit interfaces *interface-name* aggregated-ether-options lACP]** hierarchy level:

```
[edit interfaces interface-name aggregated-ether-options lACP]
periodic interval;
```

The interval can be fast (every second) or slow (every 30 seconds). You can configure different periodic rates on active and passive interfaces. When you configure the active and passive interfaces at different rates, the transmitter honors the receiver's rate.



NOTE: Source address filtering does not work when LACP is enabled.

Percentage policers are not supported on aggregated Ethernet interfaces with the CCC protocol family configured. For more information about percentage policers, see the *Routing Policy Configuration Guide*.

Generally, LACP is supported on all untagged aggregated Ethernet interfaces. For more information, see *Configuring Untagged Aggregated Ethernet Interfaces*.

Configuring LACP Link Protection



NOTE: When using LACP link protection, you can configure only two member links to an aggregated Ethernet interface: one active and one standby.

To force active and standby links within an aggregated Ethernet, you can configure LACP link protection and system priority at the aggregated Ethernet interface level using the **link-protection** and **system-priority** statements. Configuring values at this level results in only the configured interfaces using the defined configuration. LACP interface configuration also enables you to override global (chassis) LACP settings.

LACP link protection also uses port priority. You can configure port priority at the Ethernet interface **[ether-options]** hierarchy level using the **port-priority** statement. If you choose not to configure port priority, LACP link protection uses the default value for port priority (127).



NOTE: LACP link protection supports per-unit scheduling configuration on aggregated Ethernet interfaces.

To enable LACP link protection for an aggregated Ethernet interfaces, use the **link-protection** statement at the **[edit interfaces aeX aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
link-protection;
  disable;
  revertive;
  non-revertive;
}
```

By default, LACP link protection reverts to a higher-priority (lower-numbered) link when that higher-priority link becomes operational or a link is added to the aggregator that is determined to be higher in priority. However, you can suppress link calculation by adding the **non-revertive** statement to the LACP link protection configuration. In nonrevertive mode, once a link is active and collecting and distributing packets, the subsequent addition of a higher-priority (better) link does not result in a switch and the current link remains active.

If LACP link protection is configured to be nonrevertive at the global (**[edit chassis]** hierarchy) level, you can add the **revertive** statement to the LACP link protection configuration to override the nonrevertive setting for the interface. In revertive mode, the addition of a higher-priority link to the aggregator results in LACP performing a priority recalculation and switching from the current active link to the new active link.



CAUTION: If both ends of an aggregator have LACP link protection enabled, make sure to configure both ends of the aggregator to use the same mode. Mismatching LACP link protection modes can result in lost traffic.

We strongly recommend you to use LACP on both ends of the aggregator, when you connect an aggregated Ethernet interface with two member interfaces to any other vendor device. Otherwise, the vendor device (say a Layer 2 switch, or a router), will not be able to manage the traffic coming from the two link aggregated Ethernet bundle. As a result, you might observe the vendor device sending back the traffic to the backup member link of the aggregated Ethernet interface.

Currently, MX-MPC2-3D, MX-MPC2-3D-Q, MX-MPC2-3D-EQ, MX-MPC1-3D, MX-MPC1-3D-Q, and MPC-3D-16XGE-SFP do not drop traffic coming back to the backup link, whereas DPCE-R-Q-20GE-2XGE, DPCE-R-Q-20GE-SFP, DPCE-R-Q-40GE-SFP, DPCE-R-Q-4XGE-XFP, DPCE-X-Q-40GE-SFP, and DPCE-X-Q-4XGE-XFP drop traffic coming to the backup link.

Configuring LACP System Priority

To configure LACP system priority for aggregated Ethernet interfaces on the interface, use the **system-priority** statement at the **[edit interfaces aeX aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
system-priority;
```

The system priority is a 2-octet binary value that is part of the LACP system ID. The LACP system ID consists of the system priority as the two most-significant octets and the interface MAC address as the six least-significant octets. The system with the numerically lower value for system priority has the higher priority. By default, system priority is 127, with a range of 0 to 65,535.

Configuring LACP System Identifier

To configure the LACP system identifier for aggregated Ethernet interfaces, use the **system-id** statement at the **[edit interfaces aeX aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
system-id system-id;
```

The user-defined system identifier in LACP enables two ports from two separate devices to act as though they were part of the same aggregate group.

The system identifier is a 48-bit (6-byte) globally unique field. It is used in combination with a 16-bit system-priority value, which results in a unique LACP system identifier.

Configuring LACP administrative Key

To configure an administrative key for LACP, include the **admin-key number** statement at the **edit interfaces aex aggregated-ether-options lacp** hierarchy level:

```
[edit interfaces ae x aggregated-ether-options-lacp]
admin-key number;
```



NOTE: You must configure MC-LAG to configure the `admin-key` statement. For more information about MC-LAG, see [“Configuring Multichassis Link Aggregation” on page 2305](#).

Configuring LACP Port Priority

To configure LACP port priority for aggregated Ethernet interfaces, use the `port-priority` statement at the `[edit interfaces interface-name ether-options 802.3ad aeX lacp]` or `[edit interfaces interface-name ether-options 802.3ad aeX lacp]` hierarchy levels:

```
[edit interfaces interface-name ether-options 802.3ad aeX lacp]
port-priority priority;
```

The port priority is a 2-octet field that is part of the LACP port ID. The LACP port ID consists of the port priority as the two most-significant octets and the port number as the two least-significant octets. The system with the numerically lower value for port priority has the higher priority. By default, port priority is 127, with a range of 0 to 65,535.

Port aggregation selection is made by each system based on the highest port priority and are assigned by the system with the highest priority. Ports are selected and assigned starting with the highest priority port of the highest priority system and working down in priority from there.



NOTE: Port aggregation selection (discussed above) is performed for the active link when LACP link protection is enabled. Without LACP link protection, port priority is not used in port aggregation selection.

Tracing LACP Operations

To trace the operations of the LACP process, include the `traceoptions` statement at the `[edit protocols lacp]` hierarchy level:

```
[edit protocols lacp]
traceoptions {
  file <filename> <files number> <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
```

You can specify the following flags in the `protocols lacp traceoptions` statement:

- **all**—All LACP tracing operations
- **configuration**—Configuration code
- **packet**—Packets sent and received
- **process**—LACP process events
- **protocol**—LACP protocol state machine
- **routing-socket**—Routing socket events
- **startup**—Process startup events

For general information about tracing, see the tracing and logging information in the *Junos OS System Basics Configuration Guide*.

LACP Limitations

LACP can link together multiple different physical interfaces, but only features that are supported across all of the linked devices will be supported in the resulting link aggregation group (LAG) bundle. For example, different PICs can support a different number of forwarding classes. If you use link aggregation to link together the ports of a PIC that supports up to 16 forwarding classes with a PIC that supports up to 8 forwarding classes, the resulting LAG bundle will only support up to 8 forwarding classes. Similarly, linking together a PIC that supports WRED with a PIC that does not support it will result in a LAG bundle that does not support WRED.

Example: Configuring Aggregated Ethernet LACP

Configure aggregated Ethernet LACP over a VLAN-tagged interface:

```
LACP with      [edit interfaces]
VLAN-Tagged    ge--1/1/1 {
Aggregated Ethernet ether-options {
                  802.3ad ae0;
                  }
                }
                ae0 {
                  aggregated-ether-options {
                    lACP {
                      active;
                    }
                  }
                  vlan-tagging;
                  unit 0 {
                    vlan-id 100;
                    family inet {
                      address 10.1.1.2/24 {
                        vrrp-group 0 {
                          virtual-address 10.1.1.4;
                          priority 200;
                        }
                      }
                    }
                  }
                }
              }
```

```
}

```

Configure aggregated Ethernet LACP over an untagged interface:

LACP with Untagged Aggregated Ethernet

```
[edit interfaces]
ge-1/1/1 {
  ether-options-options {
    802.3ad ae0;
  }
}
ae0 {
  aggregated-ether-options {
    lacp {
      active;
    }
  }
  unit 0 {
    family inet {
      address 10.1.1.2/24 {
        vrrp-group 0 {
          virtual-address 10.1.1.4;
          priority 200;
        }
      }
    }
  }
}
}
```

Related Documentation

- [lacp on page 2353](#)
- [link-protection on page 2355](#)
- *traceoptions*
- *Junos® OS Ethernet Interfaces*

Configuring Targeted Broadcast

You can configure targeted broadcast with different options to forward the IP packets destined for a Layer 3 broadcast address to an egress interface and the Routing Engine or to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.

To enable targeted broadcast:

1. Configure the physical interface:

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

2. Configure the logical unit number:

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

3. Configure the protocol family *inet*:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@host# edit family inet
```

4. Configure targeted broadcast:

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

5. Specify one of the following options:

- To send packets to the egress interface and to the Routing Engine:

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

- To send packets to only the egress interface:

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

6. Verify the configuration. The following example configures targeted broadcast to both the egress interface and the Routing Engine:

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

Related Documentation

- [targeted-broadcast on page 2368](#)
- [Understanding Targeted Broadcast on page 2282](#)

Configuring Unicast RPF

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



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

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

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

The following sections describe unicast RPF in detail:

- [Configuring Unicast RPF Strict Mode on page 2315](#)
- [Configuring Unicast RPF Loose Mode on page 2316](#)
- [Unicast RPF and Default Routes on page 2317](#)
- [Unicast RPF with Routing Asymmetry on page 2318](#)
- [Configuring Unicast RPF on a VPN on page 2318](#)
- [Example: Configuring Unicast RPF on page 2319](#)

Configuring Unicast RPF Strict Mode

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

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

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

When unicast RPF is enabled on an interface, Bootstrap Protocol (BOOTP) packets and Dynamic Host Configuration Protocol (DHCP) packets are not accepted on the interface. To allow the interface to accept BOOTP packets and DHCP packets, you must apply a fail filter that accepts all packets with a source address of **0.0.0.0** and a destination address of **255.255.255.255**. For a configuration example, see [“Example: Configuring Unicast RPF” on page 2319](#).

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

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

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

You can include this statement at the following hierarchy levels:

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

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

- RPF fail filters are evaluated after input filters and before output filters.
- If you configure a filter counter for packets dropped by an input filter, and you want to know the total number of packets dropped, you must also configure a filter counter for packets dropped by the RPF check.

- To count packets that fail the RPF check and are accepted by the RPF fail filter, you must configure a filter counter.
- If an input filter forwards packets anywhere other than the **inet.0** or **inet6.0** routing tables, the unicast RPF check is not performed.
- If an input filter forwards packets anywhere other than the routing instance the input interface is configured for, the unicast RPF check is not performed.

Configuring Unicast RPF Loose Mode

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

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

```
mode loose;
```

You can include this statement at the following hierarchy levels:

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

Configuring Unicast RPF Loose Mode with Ability to Discard Packets

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

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

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

```
}
```

Unicast RPF and Default Routes

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

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

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

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

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

- [Unicast RPF Behavior with a Default Route on page 2317](#)
- [Unicast RPF Behavior Without a Default Route on page 2317](#)

Unicast RPF Behavior with a Default Route

If you configure a default route that uses an interface configured with unicast RPF, unicast RPF behaves as follows:

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

The packet is not accepted when either of the following is true:

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

Unicast RPF Behavior Without a Default Route

If you do not configure a default route, or if the default route does not use an interface configured with unicast RPF, unicast RPF behaves as described in “[Configuring Unicast RPF Strict Mode](#)” on page 2315 and “[Configuring Unicast RPF Loose Mode](#)” on page 2316. To summarize, unicast RPF without a default route behaves as follows:


- Strict mode—The packet is not accepted when either of the following is true:

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

Unicast RPF with Routing Asymmetry

In general, we recommend that you not enable unicast RPF on interfaces that are internal to the network because internal interfaces are likely to have *routing asymmetry*. Routing asymmetry means that a packet's outgoing and return paths are different. Routers in the core of the network are more likely to have asymmetric reverse paths than routers at the customer or provider edge. [Figure 28 on page 2318](#) shows unicast RPF in an environment with routing asymmetry.

Figure 28: Unicast RPF with Routing Asymmetry



In [Figure 28 on page 2318](#), if you enable unicast RPF on interface `so-0/0/0`, traffic destined for Router A is not rejected. If you enable unicast RPF on interface `so-1/0/1`, traffic from Router A is rejected.

If you need to enable unicast RPF in an asymmetric routing environment, you can use fail filters to allow the router to accept incoming packets that are known to be arriving by specific paths. For an example of a fail filter that accepts packets with a specific source and destination address, see [“Example: Configuring Unicast RPF” on page 2319](#).

Configuring Unicast RPF on a VPN

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

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

- For Layer 3 VPNs, unicast RPF is supported on the CE router interface.
- Unicast RPF is not supported on core-facing interfaces.
- For virtual-router routing instances, unicast RPF is supported on all interfaces you specify in the routing instance.
- If an input filter forwards packets anywhere other than the routing instance the input interface is configured for, the unicast RPF check is not performed.

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

Example: Configuring Unicast RPF on a VPN

Configure unicast RPF on a Layer 3 VPN interface:

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

Example: Configuring Unicast RPF

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

```
[edit firewall]
filter rpf-special-case-dhcp-bootp {
  term allow-dhcp-bootp {
    from {
      source-address {
        0.0.0.0/32;
      }
      address {
        255.255.255.255/32;
      }
    }
    then {
      count rpf-dhcp-bootp-traffic;
      accept;
    }
  }
  term default {
    then {
      log;
      reject;
    }
  }
}
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        rpf-check fail-filter rpf-special-case-dhcp-bootp;
      }
    }
  }
}
```

- Related Documentation**
- [unicast-reverse-path on page 2371](#)
 - *Example: Configuring Unicast Reverse-Path-Forwarding Check*

Configuration Statements

- [\[edit chassis\] Hierarchy Level on page 2320](#)
- [\[edit dynamic-profiles\] Hierarchy Level on page 2328](#)
- [\[edit forwarding-options rpf-loose-mode-discard\] Hierarchy Level on page 2329](#)
- [\[edit interfaces\] Hierarchy Level on page 2329](#)
- [\[edit multi-chassis\] Hierarchy Level on page 2340](#)
- [\[edit protocols isis\] Hierarchy Level on page 2340](#)
- [Layer 2 Routing Instances Configuration Hierarchy on page 2343](#)

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

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
      lacp {
        link-protection {
          non-revertive;
        }
        system-priority;
      }
    }
    sonet {
      device-count number;
    }
    maximum-links maximum-links-limit;
  }
  alarm {
    ds1 {
      ais (ignore | red | yellow);
      ylw (ignore | red | yellow);
    }
    ethernet {
      link-down (ignore | red | yellow);
    }
    integrated-services {
      failure (ignore | red | yellow);
    }
    management-ethernet {
      link-down (ignore | red | yellow);
    }
  }
  relay
  input {
    port port-number {
      mode (close | open);
      trigger (ignore | red | yellow);
    }
  }
}
```

```

    }
    output {
        port port-number {
            input-relay input-relay;
            mode (close | open);
            temperature;
        }
    }
    serial {
        cts-absent (ignore | red | yellow);
        dcd-absent (ignore | red | yellow);
        dsr-absent (ignore | red | yellow);
        loss-of-rx-clock (ignore | red | yellow);
        loss-of-tx-clock (ignore | red | yellow);
        tm-absent (ignore | red | yellow);
    }
    services {
        hw-down (ignore | red | yellow);
        linkdown (ignore | red | yellow);
        pic-hold-reset (ignore | red | yellow);
        pic-reset (ignore | red | yellow);
        rx-errors (ignore | red | yellow);
        sw-down (ignore | red | yellow);
        tx-errors (ignore | red | yellow);
    }
    sonet {
        (ais-l | ais-p | ber-sd | ber-sf | locd | lol | lop-p | los | pll | plm-p | rfi-l | rfl-p | uneq-p)
        (ignore | red | yellow);
    }
    t3 {
        (ais | exz | ferf | idle | lcv | lof | los | pll | ylw) (ignore | red | yellow);
    }
}
cluster {
    control-link-recovery;
    control-ports {
        fpc slot-number port port-number;
    }
    heartbeat-interval milliseconds;
    heartbeat-threshold number;
    redundancy-group {
        ... the redundancy-group subhierarchy appears at the end of the [edit chassis cluster]
        hierarchy ...
    }
    reth-count number;
    traceoptions {
        file <filename> <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
        flag flag;
        level severity;
        no-remote-trace;
    }
    redundancy-group group-number {
        gratuitous-arp-count number;
        hold-down-interval seconds;
        interface-monitor {

```

```

        interface-name weight number;
    }
    ip-monitoring {
        family {
            inet {
                ipv4-address {
                    interface rethindex.logical-unit-number secondary-ip-address ipv4-address;
                    weight number;
                }
            }
        }
        global-threshold number;
        global-weight number;
        retry-count count;
        retry-interval interval;
    }
    node node-number priority priority-number;
    preempt;
}
config-button {
    no-clear;
    no-rescue;
}
container-devices {
    device-count number;
}
craft-lockout;
disable-power-management;
disk-partition partition-name (/config | /var) {
    level (full | high) {
        free-space threshold-value (mb | percent);
    }
}
enhanced-policer;
extended-statistics;
fabric {
    degraded {
        action-fpc-restart-disable;
        degraded-fabric-detection-enable
        degraded-fpc-bad-plane-threshold number-bad-planes;
    }
    redundancy-mode (increased-bandwidth | redundant);
}
filter;
fpc slot-number {
    ... the fpc subhierarchy appears after the main [edit chassis] hierarchy ...
}
fpc-feb-connectivity {
    fpc slot-number feb (slot-number | none);
}
fpc-resync;
fru-poweron-sequence sequence;
lcc index {
    ... the lcc subhierarchy appears after the main [edit chassis] hierarchy ...
}
maximum-ecmp value;

```



```

memory-enhanced {
    filter;
    route;
    vpn-label;
}
network-services (ethernet | enhanced-ethernet | ip | enhanced-ip);
(packet-scheduling | no-packet-scheduling);
pem {
    minimum number;
}
policer-drop-probability-low;
ppp-subscriber-services (disable | enable);
redundancy {
    cfeb slot (always | preferred);
    failover {
        on-disk-failure;
        on-loss-of-keepalives;
    }
    feb {
        redundancy-group group-name {
            description description;
            feb slot-number <backup | primary>;
            no-auto-failover;
        }
    }
    graceful-switchover;
    keepalive-time seconds;
    routing-engine slot-number (backup | disabled | master);
    sfm slot-number (always | preferred);
    ssb slot-number (always | preferred);
}
route-memory-enhanced;
route-localization {
    inet (chassis);
    inet6;
}
routing-engine {
    bios {
        no-auto-upgrade;
    }
    on-disk-failure disk-failure-action (halt | reboot);
}
sfm slot-number {
    power off;
}
sib {
    minimum number;
}
(source-route | no-source-route);
state [
    cb-upgrade [on | off];
]
synchronization { # for M Series and T Series routers
    primary (external-a | external-b);
    secondary (external-a | external-b);
    signal-type (e1 | t1);
}

```

```
switching-mode (non-revertive | revertive);
transmitter-enable;
validation-interval seconds;
y-cable-line-termination;
}
synchronization { # for MX80 and MX240 routers
  clock-mode (auto-select | free-run);
  esmc-transmit {
    interfaces (all | interface-name);
  }
  hold-interval {
    configuration-change seconds;
    restart seconds;
    switchover seconds;
  }
  network-option (option-1 | option-2);
  quality-mode-enable;
  selection-mode (configured-quality|received-quality);
  source {
    (external-a | external-b) {
      priority number;
      quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
      request (force-switch | lockout);
    }
    interfaces interface-name {
      priority number;
      quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
      request (force-switch | lockout);
      wait-to-restore minutes;
    }
  }
  switchover-mode (revertive | non-revertive);
}
synchronization { # for ACX Series routers
  clock-mode (auto-select | free-run);
  esmc-transmit {
    interfaces (all | interface-name);
  }
  hold-interval {
    configuration-change seconds;
    restart seconds;
    switchover seconds;
  }
  network-option (option-1 | option-2);
  quality-mode-enable;
  selection-mode (configured-quality | received-quality);
  source {
    (bits | gps) {
      priority number;
      quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
      request (force-switch | lockout);
    }
    interfaces interface-name {
      priority number;
      quality-level (prc | prs |sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc);
      request (force-switch | lockout);
    }
  }
}
```

```

        wait-to-restore minutes;
    }
}
switchover-mode(non-revertive | revertive);
}
system-domains {
    protected-system-domains psdnumerical-index {
        control-plane-bandwidth-percent percent;
        control-slot-numbers [ slot-numbers ];
        control-system-id control-system-id;
        description description;
        fpcs [ slot-numbers ];
    }
    root-domain-id root-domain-id;
}
vrf-mtu-check;
}

chassis {
    fpc slot-number {
        number-of-ports active-ports;
        offline;
        pic slot-number {
            ... the pic subhierarchy appears after the main [edit chassis fpc slot-number] hierarchy
            ...
        }
        port-mirror-instance port-mirror-instance-name;
        power (off | on);
        sampling-instance instance-name;
    }

    fpc slot-number {
        pic slot-number {
            adaptive-services {
                service-package (layer-2 | layer-3 | ...the following extension-provider subhierarchy
                ...);
            }
            extension-provider {
                control-cores number;
                data-cores number;
                data-flow-affinity {
                    hash-key (layer-3 | layer-4);
                }
                channelization;
                forwarding-db-size megabytes;
                object-cache-size megabytes;
                package package-name;
                policy-db-size megabytes;
                syslog {
                    facility {
                        severity;
                        destination (pic-console | routing-engine);
                    }
                }
                wired-process-mem-size megabytes;
            }
        }
    }
}

```

```
aggregated-devices {
  ima {
    device-count number;
  }
}
aggregate-ports;
atm-cell-relay-accumulation;
atm-l2circuit-mode (aal5 | cell | trunk trunk);
cel {
  el port-number {
    channel-group group-number timeslots slot-number;
  }
}
ct3 {
  port port-number {
    t1 link-number {
      channel-group group-number timeslots slot-number;
    }
  }
}
ethernet {
  pic-mode (enhanced-switching | routing | switching);
}
fibre-channel {
  port port-number;
  port-range port-range-low port-range-high
}
egress-policer-overhead bytes;
forwarding-mode {
  sa-multicast;
  vlan-steering {
    vlan-rule (high-low | odd-even);
  }
}
framing (e1 | e3 | sdh | sonet | t1 | t3);
idle-cell-format {
  itu-t;
  payload-pattern payload-pattern-byte;
}
ingress-policer-overhead bytes;
inline-services {
  bandwidth (1g | 10g);
}
linerate-mode;
max-queues-per-interface (4 | 8);
mlfr-uni-nni-bundles number;
no-concatenate;
no-multi-rate;
port port-number {
  framing (e1 | e3 | sdh | sonet | t1 | t3);
  forwarding-mode {
    sa-multicast;
  }
  speed ( oc3-stm1 | oc12-stm4 | oc48-stm16);
}
port-mirror-instance port-mirror-instance-name;
```

```

    q-pic-large-buffer {
        (large-scale | small-scale);
    }
    red-buffer-occupancy {
        weighted-averaged <instant-usage-weight-exponent weight-value>;
    }
    shdsl {
        pic-mode (1-port-atm | 2-port-atm);
    }
    sparse-dlcis;
    traffic-manager {
        egress-shaping-overhead number;
        ingress-shaping-overhead number;
        mode {
            egress-only;
            ingress-and-egress;
            session-shaping;
        }
    }
    tunnel-queuing;
    tunnel-services {
        bandwidth (1g | 10g | 20g | 40g);
        tunnel-only;
    }
    vtmapping (itu-t | klm);
}

}

chassis {
    lcc index {
        fpc slot-number {
            ... the fpc subhierarchy appears after the main [edit chassis lcc index] hierarchy ...
        }
        offline;
        online-expected;
    }
}

lcc index {
    fpc slot-number {
        pic slot-number {
            ... the pic subhierarchy appears after the main [edit chassis lcc index fpc slot-number] hierarchy ...
        }
        power (off | on);
        sampling-instance instance-name;
    }

    fpc slot-number {
        pic slot-number {
            aggregate-ports;
            atm-cell-relay-accumulation;
            atm-l2circuit-mode (aal5 | cell | trunk trunk);
            framing (e1 | e3 | sdh | sonet | t1 | t3);
            idle-cell-format {

```

```
    itu-t;
    payload-pattern payload-pattern-byte;
  }
  linerate-mode;
  max-queues-per-interface (4 | 8);
  no-concatenate;
  no-mcast-replication;
  no-pre-classifier;
  port port-number {
    framing (e1 | e3 | sdh | sonet | t1 | t3);
  }
  q-pic-large-buffer {
    (large-scale | small-scale);
  }
  red-buffer-occupancy {
    weighted-averaged <instant-usage-weight-exponent weight-value>;
  }
  shdsl {
    pic-mode (1-port-atm | 2-port-atm);
  }
  traffic-manager {
    egress-shaping-overhead bytes;
    ingress-shaping-overhead bytes;
    mode {
      egress-only;
      ingress-and-egress;
    }
  }
}
}
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

[\[edit dynamic-profiles\] Hierarchy Level](#)

```
dynamic-profiles {
  profile-name {
    class-of-service
    ... statements from those in [edit class-of-service] Hierarchy Level.
    firewall
    ... statements from those in [edit firewall] Hierarchy Level.
    interfaces
    ... statements from those in [edit interfaces] Hierarchy Level.
    policy-options
    ... statements from those in [edit policy-options] Hierarchy Level.
    predefined-variable-defaults variable-name default-value
    profile-variable-set variable-set-name dynamic-variable-name substitute-variable-name
    protocols
    ... statements from those in [edit protocols] Hierarchy Level.
    routing-instances
    ... statements from those in [edit routing-instances] Hierarchy Level.
    routing-options
    ... statements from those in [edit routing-options] Hierarchy Level.
    services
```

```

... statements from those in [edit services] Hierarchy Level.
variables {
  variable-name {
    default-value default-value;
    equals expression;
    mandatory;
  }
  uid;
  uid-reference;
}
}

```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit dynamic-profiles routing-instances] Hierarchy Level*
- *[edit dynamic-profiles routing-options] Hierarchy Level*
- *[edit dynamic-profiles variables] Hierarchy Level*

[\[edit forwarding-options rpf-loose-mode-discard\] Hierarchy Level](#)

```

rpf-loose-mode-discard {
  family {
    inet;
    inet6;
  }
}

```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- [\[edit forwarding-options\] Hierarchy Level on page 333](#)

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

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```

interfaces {
  interface-name {
    ... the "interface-name" subhierarchy appears after the main [edit interfaces] hierarchy
    level ...
  }
  interface-set interface-set-name {
    interface interface-name {
      (unit unit-number | vlan-tags-outer vlan-tag);
    }
  }
  irb (Interfaces) {
    accounting-profile name;
    description text;
    disable;

    (gratuitous-arp-reply | no-gratuitous-arp-reply);

```

```
hold-time up milliseconds down milliseconds;  
mtu bytes;  
no-gratuitous-arp-request;  
  
traceoptions {  
    flag flag;  
}  
(traps | no-traps);  
unit logical-unit-number {  
    accounting-profile name;  
    bandwidth rate;  
    description text;  
    disable;  
    encapsulation type;  
    family inet {  
        accounting {  
            destination-class-usage;  
            source-class-usage {  
                input;  
                output;  
            }  
        }  
    }  
    address ipv4-address {  
        arp ip-address (mac | multicast-mac) mac-address <publish>;  
        broadcast address;  
        preferred;  
        primary;  
        vrrp-group group-id {  
            (accept-data | no-accept-data);  
            advertise-interval seconds;  
            advertisements-threshold number;  
            authentication-key key;  
            authentication-type authentication;  
            fast-interval milliseconds;  
            (preempt | no-preempt) {  
                hold-time seconds;  
            }  
            priority number;  
            track {  
                interface interface-name {  
                    bandwidth-threshold bits-per-second priority-cost priority;  
                    priority-cost priority;  
                }  
                priority-hold-time seconds;  
                route prefix/prefix-length routing-instance instance-name priority-cost priority;  
            }  
            virtual-address [ addresses ];  
            vrrp-inherit-from vrrp-group;  
        }  
    }  
    filter {  
        input filter-name;  
        output filter-name;  
    }  
    mtu bytes;
```



```

no-neighbor-learn;
no-redirects;
primary;
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bandwidth priority-cost number;
                priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-inet6-address [addresses];
        virtual-link-local-address ipv6-address;
        vrrp-inherit-from {
            active-group group-number;
            active-interface interface-name;
        }
    }
}
}
(dad-disable | no-dad-disable);
filter {

```

```

        input filter-name;
        output filter-name;
    }
    mtu bytes;
    nd6-stale-time seconds;
    no-neighbor-learn;
    no-redirects;
    policer {
        input policer-name;
        output policer-name;
    }
    rpf-check {
        fail-filter filter-name;
        mode {
            loose;
        }
    }
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    filter {
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    policer {
        input policer-name;
        output policer-name;
    }
}
native-inner-vlan-id vlan-id;
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id-list [vlan-id's];
vlan-id-range [vlan-id-range];
}
}
traceoptions {
    file <filename> <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag <disable>;
    no-remote-trace;
}
}

interfaces {
    interface-name {
        disable;
        accounting-profile name;
        aggregated-ether-options {
            ethernet-switch-profile {
                tag-protocol-id [ hexadecimal-identifiers ];
            }
        }
    }
}

```

```

(flow-control | no-flow-control);
lACP {
    (active | passive);
    admin-key key;
    fast-failover;
    link-protection {
        disable;
        (revertive | non-revertive);
    }
    periodic (fast | slow);
    system-id mac-address;
    system-priority priority;
}
(link-protection | no-link-protection);
link-speed (100m | 1g | 8g | 10g | 40g | 50g | 80g | 100g | oc192);
logical-interface-fpc-redundancy;
(loopback | no-loopback);
mc-ae {
    chassis-id chassis-id;
    events {
        iccp-peer-down {
            force-icl-down;
            prefer-status-control-active;
        }
    }
    mc-ae-id mc-ae-id;
    mode (active-active | active-standby);
    redundancy-group group-id;
    status-control (active | standby);
}
minimum-links number;
rebalance-periodic {
    start-time time;
    interval number;
}
source-address-filter {
    mac-address;
}
(source-filtering | no-source-filtering);
}
auto-configure {
    remove-when-no-subscribers;
    stacked-vlan-ranges {
        access-profile profile-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                mac-address;
                option-82 ( circuit-id | remote-id);
                radius-realm radius-realm-string;
                user-prefix user-prefix-string;
            }
        }
    }
}

```

```

    }
    dynamic-profile profile-name {
        accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
        ranges (any | low-tag-high-tag), (any | low-tag-high-tag);
    }
}
vlan-ranges {
    access-profile profile-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-name;
            mac-address;
            option-82;
            radius-realm radius-realm-string;
            user-prefix user-prefix-string;
        }
    }
    dynamic-profile profile-name {
        accept (any | dhcp-v4 | dhcp-v6 | inet | inet6);
        ranges (any | low-tag)—(any | high-tag);
    }
}
override tag vlan-tag dynamic-profile profile name;
}
encapsulation (ethernet-bridge | ethernet-vpls | extended-vlan-bridge |
    extended-vlan-vpls | flexible-ethernet-services | vlan-vpls);
ether-options {
    802.3ad {
        aex;
        (backup | primary);
        lacp {
            force-up;
            port-priority
        }
    }
}
asynchronous-notification;
(auto-negotiation | no-auto-negotiation);
ethernet-switch-profile {
    ethernet-policer-profile {
        input-priority-map {
            ieee802.1p premium [ values ];
        }
        output-priority-map {
            classifier {
                premium {
                    forwarding-class class-name {
                        loss-priority (high | low);
                    }
                }
            }
        }
    }
    policer cos-policer-name {

```

```

        aggregate {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
        premium {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
    }
    tag-protocol-id;
}
(mac-learn-enable | no-mac-learn-enable);
}
(flow-control | no-flow-control);
ignore-l3-incompletes;
link-mode (automatic | full-duplex | half-duplex);
(loopback | no-loopback);
keepalives <interval seconds> <down-count number> <up-count number>;
speed (1g | 10m | 100m | 10m-100m | auto-negotiation);
source-address-filter {
    mac-address;
}
source-filtering | no-source-filtering;
}
flexible-vlan-tagging;
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time (up milliseconds | down milliseconds);
interface-transmit-statistics;
(keepalives <down-count number> <interval seconds> <up-count number> |
no-keepalives);
layer2-policer {
    apply-groups [ group-names ];
    apply-groups-except [ group-names ];
}
link-mode (automatic | full-duplex);
mac mac-address;
mtu bytes;
multi-chassis-protection peer-ip-address {
    interface interface-name;
}
native-vlan-id number;
no-gratuitous-arp-request;
optics-options {
    alarm low-light-alarm {
        (link-down | syslog);
    }
    warning low-light-warning {
        (link-down | syslog);
    }
}
wavelength nm;
}
passive-monitor-mode;
per-unit-scheduler;
speed (10m | 100m | 1g | auto | oc3 | oc12 | oc48);
stacked-vlan-tagging;
traceoptions {

```

```

    flag flag;
  }
  transmit-bucket {
    overflow discard;
    rate percentage;
    threshold bytes;
  }
  (traps | no-traps);
  unidirectional;
  vlan-tagging;
}

interface-name {
  unit logical-unit-number {
    disable;
    accept-source-mac {
      mac-address mac-address {
        policer {
          input policer-name;
          output policer-name;
        }
      }
    }
  }
  account-layer2-overhead (Interface Level) {
    value;
    egress bytes;
    ingress bytes;
  }
  accounting-profile name;
  advisory-options {
    downstream-rate rate;
    upstream-rate rate;
  }
  arp-resp (restricted|unrestricted);
  bandwidth rate;
  clear-dont-fragment-bit;
  copy-tos-to-outer-ip-header;
  demux-destination family;
  encapsulation (vlan-bridge | vlan-vpls);
  epd-threshold cells plp1 cells;
  filter filter-name;
  inner-vlan-id-range start start-id end end-id;
  input-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
  }
  interface-shared-with psdnumerical-index;
  layer2-policer {
    input-hierarchical-policer policer-name;
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
  }
}

```

```

    output-three-color policer-name;
}
multi-chassis-protection peer-ip-address {
    interface interface-name;
}
native-inner-vlan-id number;
output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
peer-interface interface-name;
peer-unit unit-number;
plp-to-clp;
proxy-arp <restricted | unrestricted>;
rpm {
    (client | server);
    twamp-server;
}
swap-by-poppush;
vlan-id number;
vlan-id-list [ vlan-id vlan-id-vlan-id ];
vlan-id-range number-number;
vlan-tags (inner <tpid.>vlan-id | inner-list [ vlan-id vlan-id-vlan-id ] |
    inner-range <tpid.>vlan-id-vlan-id) outer <tpid.>vlan-id;
}

unit logical-unit-number {
    family ethernet-switching {
        filter {
            group filter-group-number;
            (input filter-name | input-list [ filter-names ]);
            (output filter-name | output-list [ filter-names ]);
            (inner-vlan-id-list [ vlan-ids ] | vlan-id number | vlan-id-list [ number
                number-number ]);
            interface-mode (access | trunk);
        }
        policer {
            input policer-name;
            output policer-name;
        }
        vlan-rewrite {
            translate old-vlan-id new-vlan-id;
        }
        vlan {
            members [ all vlan-identifiers ];
        }
    }
}
family inet {
    filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
    input-hierarchical-policer policer-name;
}

```

```
mac-validate (loose | strict);
mtu bytes;
no-neighbor-learn;
no-redirects;
policer {
    arp policer-template-name;
    input policer-name;
    output policer-name;
}
primary;
receive-options-packets;
receive-ttl-exceeded;
rpf-check {
    fail-filter filter-name;
    mode loose;
}
sampling {
    (input | output | input output);
}
simple-filter {
    input filter-name;
}
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
unnumbered-address interface-name <destination address>
    <destination-profile profile-name> <preferred-source-address address>;
}

family inet6 {
    address ipv6-address {
        destination destination-address;
        eui-64;
        ndp ipv6-address <l2-interface interface-name> <(mac mac-address |
            multicast-mac multicast-mac-address) <publish>>;
        preferred;
        primary;
        vrrp-inet6-group group-number {
            (accept-data | no-accept-data);
            fast-interval milliseconds;
            inet6-advertise-interval seconds;
            (no-preempt; | ... the following preempt statement ...)
            preempt {
                hold-time seconds;
            }
            priority number;
            track {
                interface interface-name {
                    bandwidth-threshold bits-per-second priority-cost priority;
                    priority-cost priority;
                }
                priority-hold-time seconds;
                route ip-address-prefix/prefix-length routing-instance instance-name
                    priority-cost priority;
            }
        }
    }
}
```



```

    }
    virtual-inet6-address [ addresses ];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
    }
}
}
(dad-disable | no-dad-disable);
filter {
    group filter-group-number;
    (input filter-name | input-list [ filter-names ]);
    (output filter-name | output-list [ filter-names ]);
}
input-hierarchical-policer policer-name;
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode loose;
}
sampling {
    (input | output | input output);
}
unnumbered-address interface-name preferred-source-address address;
}

family iso {
    address iso-address;
    mtu bytes;
}

family mlfr-end-to-end {
    bundle logical-interface-name;
}

family mpls {
    filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
    input-hierarchical-policer policer-name;
    maximum-labels maximum-labels;
    mtu bytes;
    policer {
        input policer-name;
    }
}

```

```
        output policer-name;
    }
}

family vpls {
    core-facing;
    filter {
        group filter-group-number;
        (input filter-name | input-list [ filter-names ]);
        (output filter-name | output-list [ filter-names ]);
    }
    policer {
        input policer-name;
        output policer-name;
    }
}
}
```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

[edit multi-chassis] Hierarchy Level

```
multi-chassis {
    multi-chassis-protection ipv4-address {
        interface interface-name;
    }
}
```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

[edit protocols isis] Hierarchy Level

The following statement hierarchy can also be included at the **[edit protocols isis]** hierarchy level.

```
protocols {
    isis {
        disable;
        clns-routing;
        context-identifier ip-address</prefix> {
            level (1 | 2) <disable>;
        }
        export [ policy-names ];
        graceful-restart {
            disable;
            helper-disable;
            restart-duration seconds;
        }
        ignore-attached-bit;
    }
}
```

```

interface interface-name {
    ... the interface subhierarchy appears after the main [edit protocols isis] hierarchy ...
}
label-switched-path name level level metric metric;
level (1 | 2) {
    disable;
    authentication-key key;
    authentication-type authentication;
    external-preference preference;
    no-csnp-authentication;
    no-hello-authentication;
    no-psnp-authentication;
    preference preference;
    prefix-export-limit number;
    wide-metrics-only;
}
loose-authentication-check;
lsp-lifetime seconds;
max-areas number;
no-adjacency-holddown;
no-authentication-check;
no-ipv4-routing;
no-ipv6-routing;
overload {
    advertise-high-metrics;
    timeout seconds;
}
reference-bandwidth reference-bandwidth;
rib-group {
    inet group-name;
    inet6 group-name;
}
spf-options {
    delay milliseconds;
    holddown milliseconds;
    rapid-runs number;
}
topologies {
    ipv4-multicast;
    ipv6-multicast;
    ipv6-unicast;
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
traffic-engineering {
    disable;
    family inet {
        shortcuts {
            multicast-rpf-routes;
        }
    }
    family inet6 {
        shortcuts;
    }
}

```

```
    }
  }
  ignore-lsp-metrics;
}

isis {
  interface interface-name {
    disable;
    bfd-liveness-detection {
      authentication {
        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
          meticulous-keyed-sha-1 | simple-password);
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      version (1 | automatic);
    }
    checksum;
    csnp-interval (seconds | disable);
    hello-padding (adaptive | loose | strict);
    ldp-synchronization {
      disable;
      hold-time seconds;
    }
    level (1 | 2) {
      disable;
      hello-authentication-key key;
      hello-authentication-type authentication;
      hello-interval seconds;
      hold-time seconds;
      ipv4-multicast-metric number;
      ipv6-multicast-metric number;
      ipv6-unicast-metric number;
      metric metric;
      passive;
      priority number;
      te-metric metric;
    }
    link-protection;
    lsp-interval milliseconds;
    mesh-group (value | blocked);
    no-adjacency-down-notification;
    no-eligible-backup;
    no-ipv4-multicast;
    no-ipv6-multicast;
```

```

no-ipv6-unicast;
no-unicast-topology;
node-link-protection;
passive;
point-to-point;
}
}
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

Layer 2 Routing Instances Configuration Hierarchy

Use the **vpls** routing instance type for point-to-multipoint LAN implementations between a set of sites in a VPN.

To configure routing instances for Layer 2 networks, include the following statements:

```

routing-instances {
  routing-instance-name {
    access {
      address-assignment {
        ... same statements as in the address-assignment subhierarchy in [edit access]
        Hierarchy Level ...
      }
      address-protection;
      description text;
      egress-protection {
        context-identifier context-id;
      }
      forwarding-options {
        ...forwarding-options...
      }
      instance-role role;
      instance-type type;
      interface interface-name;
      l2-domain-id-for-l3 id;
      l2vpn-id community;
      layer3-domain-identifier identifier;
      multicast-snooping-options {
        ... same statements as in [edit multicast-snooping-options] Hierarchy Level EXCEPT
        FOR ...
      }
      traceoptions {...} # NOT valid at this level
    }
    no-irb-layer-2-copy;
    no-local-switching;
    no-vrf-advertise;
    no-vrf-propagate-ttl;
    pbb-options {
      default-bvlan bvlan;
      peer-instance instance;
      vlan-id vlan-id isid-list [ isid-numbers ]
    }
  }
}

```

```
protocols {
  ... the protocols subhierarchy appears after the main [edit routing-instances
    routing-instance-name] hierarchy ...
}
provider-tunnel {
  ... the provider-tunnel subhierarchy appears after the main [edit routing-instances
    routing-instance-name] hierarchy ...
}
route-distinguisher (as-number:number | ip-address:number);
routing-interface interface;
routing-options {
  ... the routing-options subhierarchy appears after the main [edit routing-instances
    routing-instance-name] hierarchy ...
}
service-groups {
  service-group-name {
    pbb-service-options {
      default-isid isid-number;
      isid isid-number vlan-id-list [ vlan-ids ];
      mac-address mac-address;
    }
    service-type type;
  }
}
services {
  mobile-ip {
    ... same statements as in [edit services mobile-ip] Hierarchy Level ...
  }
}
switch-options {
  ... same statements as in [edit switch-options] Hierarchy Level ...
}
vlan-id (id | all | none);
vlan-model one-to-one;
vlan-tags outer <tpid.>vlan-id inner <tpid.>vlan-id;
[edit vlans] Hierarchy Level on page 445 {
  ... same statements as in [edit vlans] Hierarchy Level ...
}
vrf-advertise-selective {
  family {
    inet-mvpn;
    inet6-mvpn;
  }
}
vrf-export [ policy-names ];
vrf-import [ policy-names ];
vrf-propagate-ttl;
vrf-table-label;
vrf-target {
  export community-name;
  import community-name;
}
protocols {
  ... protocols-configuration ...
}
routing-options {
```

```

...routing-options-configuration ...
}
bridge-domains {
  bridge-domain-name {
    domain-type bridge;
    interface interface-name;
    routing-interface routing-interface-name;
    vlan-id (Bridge Domain or VLAN) (none | all | number);
    vlan-tags outer number inner number;
    bridge-options {
      interface-mac-limit limit {
        packet-action drop;
      }
      interface interface-name {
        interface-mac-limit limit {
          packet-action drop;
        }
      }
      mac-statistics;
      mac-table-size limit {
        packet-action drop;
      }
      no-mac-learning;
      static-mac mac-address;
    }
  }
}
}
}

```

With the exception of the **instance-type virtual-switch** statement (which configures a virtual-switch routing instance), you can include the statements at the following hierarchy levels:

- **[edit]**
- **[edit logical-systems *logical-system-name*]**

The **instance-type virtual-switch** statement is not supported at the **[edit logical-systems *logical-system-name*]** hierarchy level.

Related Documentation

- [Routing Instances Overview](#)
- [Layer 2 Routing Instance Types](#)
- [Configuring a Layer 2 Virtual Switch on page 1847](#)
- [Configuring a Layer 2 Control Protocol Routing Instance](#)

bandwidth (Interfaces)

Syntax	<code>bandwidth rate;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure an informational-only bandwidth value for an interface. This statement is valid for all logical interface types except multilink and aggregated interfaces.



NOTE: We recommend that you be careful when setting this value. Any interface bandwidth value that you configure using the `bandwidth` statement affects how the interface cost is calculated for a dynamic routing protocol, such as OSPF. By default, the interface cost for a dynamic routing protocol is calculated using the following formula:

$$\text{cost} = \text{reference-bandwidth} / \text{bandwidth},$$


where bandwidth is the physical interface speed. However, if you specify a value for bandwidth using the `bandwidth` statement, that value is used to calculate the interface cost, rather than the actual physical interface bandwidth.

Options	rate —Peak rate, in bits per second (bps) or cells per second (cps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). You can also specify a value in cells per second by entering a decimal number followed by the abbreviation c ; values expressed in cells per second are converted to bits per second by means of the formula 1 cps = 384 bps. Range: Not limited.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Configuring the Interface Bandwidth

chassis-id


Syntax	<code>chassis-id <i>chassis-id</i>;</code>
Hierarchy Level	[edit]interfaces aggregated-ether-options mc-ae]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify the chassis ID of the multichassis aggregated Ethernet interface device. LACP uses the chassis ID to calculate the port number of the multichassis link aggregation group (MC-LAG) physical member links.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

filter

Syntax	<pre>filter { group <i>filter-group-number</i>; input <i>filter-name</i>; input-list [<i>filter-names</i>]; output <i>filter-name</i>; output-list [<i>filter-names</i>]; }</pre>
Hierarchy Level	<pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	<p> NOTE: On EX Series switches, the <code>group</code>, <code>input-list</code>, <code>output-filter</code> statements are not supported under the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>inet</i>]</code>, <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>inet6</i>]</code>, and <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>mpls</i>]</code> hierarchies.</p> <p>Apply a filter to an interface. You can also use filters for encrypted traffic. When you configure filters, you can configure them under the family <code>ethernet-switching</code>, <code>inet</code>, <code>inet6</code>, <code>mpls</code>, or <code>vpls</code> only.</p>
Options	<p>group <i>filter-group-number</i>—Define an interface to be part of a filter group. Range: 1 through 255</p> <p>input <i>filter-name</i>—Name of one filter to evaluate when packets are received on the interface.</p> <p>output <i>filter-name</i>—Name of one filter to evaluate when packets are transmitted on the interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Applying a Filter to an Interface</i> • <i>Junos Services Interfaces Configuration Release 11.2</i> • <i>Routing Policy Configuration Guide</i> • <i>Junos OS System Basics Configuration Guide</i> • <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i>

- *Configuring Firewall Filters (CLI Procedure)*
- *family*

flow-control

Syntax	(flow-control no-flow-control);
Hierarchy Level	[edit interfaces <i>interface-name</i> aggregated-ether-options], [edit interfaces <i>interface-name</i> ether-options], [edit interfaces <i>interface-name</i> fastether-options], [edit interfaces <i>interface-name</i> ggether-options], [edit interfaces <i>interface-name</i> multiservice-options], [edit interfaces interface-range <i>name</i> aggregated-ether-options], [edit interfaces interface-range <i>name</i> ether-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 in EX Series switches. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Description	For aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, explicitly enable flow control, which regulates the flow of packets from the router or switch to the remote side of the connection. Enabling flow control is useful when the remote device is a Gigabit Ethernet switch. Flow control is not supported on the 4-port Fast Ethernet PIC.
	<div>  <p>NOTE: On the Type 5 FPC, to prioritize control packets in case of ingress oversubscription, you must ensure that the neighboring peers support MAC flow control. If the peers do not support MAC flow control, then you must disable flow control.</p> </div>
Default	Flow control is enabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Flow Control</i> • <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i>

forward-and-send-to-re

Syntax	forward-and-send-to-re;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet targeted-broadcast], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet targeted-broadcast]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	Specify that IP packets destined for a Layer 3 broadcast address be forwarded to an egress interface and the Routing Engine. The packets are broadcast only if the egress interface is a LAN interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Targeted Broadcast on page 2313• targeted-broadcast on page 2368• Understanding Targeted Broadcast on page 2282

forward-only

Syntax	forward-only;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet targeted-broadcast], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet targeted-broadcast]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	Specify that IP packets destined for a Layer 3 broadcast address be forwarded to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Targeted Broadcast on page 2313• targeted-broadcast on page 2368• Understanding Targeted Broadcast on page 2282

gratuitous-arp-reply

Syntax	(gratuitous-arp-reply no-gratuitous-arp-reply);
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 in EX Series switches. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Description	For Ethernet interfaces, enable updating of the ARP cache for replies received in response to gratuitous ARP requests.
Default	Updating of the ARP cache is disabled on all Ethernet interfaces.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Gratuitous ARP</i>• no-gratuitous-arp-request on page 2358


group (RPF Selection)

Syntax	<pre>group group-address{ source source-address { next-hop next-hop-address; } wildcard-source { next-hop next-hop-address; } }</pre>
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> edit protocols pim rpf-selection]
Release Information	Statement introduced in JUNOS Release 10.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure the PIM group address for which you configure RPF selection group (RPF Selection) .
Default	By default, PIM RPF selection is not configured.
Options	group-address —PIM group address for which you configure RPF selection.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring PIM RPF Selection

l2-domain-id-for-l3

Syntax	<pre>l2-domain-id-for-l3 id;</pre>
Hierarchy Level	[edit routing-instances <i>instance-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3R2.
Description	Specify a Layer 2 domain ID within a routing instance.
Options	id —Layer 2 identification number.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Layer 2 Virtual Switch on page 1847

lACP (Aggregated Ethernet)

Syntax	<pre>lACP { (active passive); admin-key key; fast-failover; link-protection { disable; (revertive non-revertive); } periodic interval; system-id mac-address; system-priority priority; }</pre>
Hierarchy Level	[edit interfaces aex aggregated-ether-options]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>fast-failover option introduced in Junos OS Release 12.2.</p>
Description	For aggregated Ethernet interfaces only, configure Link Aggregation Control Protocol (LACP).
Default	If you do not specify LACP as either active or <i>passive</i> , LACP remains passive.
Options	<p>active—Initiate transmission of LACP packets.</p> <p>admin-key number—Specify an administrative key for the router or switch.</p>
	<div>  <p>NOTE: You must also configure Multichassis Link Aggregation (MC-LAG) when you configure the admin-key.</p> </div>
	<p>passive—Respond to LACP packets.</p> <p>fast-failover—Specify to override the IEEE 802.3ad standard and allow the standby link to receive traffic. Overriding the default behavior facilitates subsecond failover.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring Aggregated Ethernet LACP Configuring Aggregated Ethernet LACP (CLI Procedure) Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch

layer3-domain-identifier

Syntax	layer3-domain-identifier <i>identifier</i> ;
Hierarchy Level	[edit routing-instances <i>instance-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3R2.
Description	Specify a Layer 3 domain ID within a routing instance.
Options	id —Layer 3 Identification number.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Layer 2 Virtual Switch on page 1847

link-protection

Syntax	<pre>link-protection { disable; (revertive non-revertive); }</pre>
Hierarchy Level	<p>[edit interfaces aex aggregated-ether-options]</p> <p>[edit interfaces aex aggregated-ether-options <i>lACP</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for disable, revertive, and non-revertive statements added in Junos OS Release 9.3.</p>
Description	<p>On the router, for aggregated Ethernet interfaces only, configure link protection. In addition to enabling link protection, a primary and a secondary (backup) link must be configured to specify what links egress traffic should traverse. To configure primary and secondary links on the router, include the primary and backup statements at the [edit interfaces <i>ge-fpc/pic/port</i> gigether-options 802.3ad aex] hierarchy level or the [edit interfaces <i>fe-fpc/pic/port</i> fastether-options 802.3ad aex] hierarchy level.</p> <p>On the switch, you can configure either Junos OS link protection for aggregated Ethernet interfaces or the LACP standards link protection for aggregated Ethernet interfaces.</p> <p>For Junos OS link protection, specify link-protection at the following hierarchy levels:</p> <ul style="list-style-type: none"> • [edit interfaces <i>ge-fpc/pic/port</i> ether-options 802.3ad aex] • [edit interfaces <i>xe-fpc/pic/port</i> ether-options 802.3ad aex] <p>For LACP standards link protection, specify link-protection at the following hierarchy levels:</p> <ul style="list-style-type: none"> • For global LACP link protection, specify at [edit chassis aggregated-devices ethernet lACP] • For a specific aggregated Ethernet interface, specify at [edit interfaces aeX aggregated-ether-options lACP]
Options	The statements are explained separately.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Aggregated Ethernet Link Protection</i> • <i>Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)</i>

mc-ae-id

Syntax	<code>mc-ae-id mc-ae-id;</code>
Hierarchy Level	[edit interfaces aggregated-ether-options mc-ae]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify the multichassis aggregated Ethernet (MC-AE) identification number of the MC-AE that a given aggregated Ethernet interface belongs to. The two peers that host a given multichassis link aggregation group (MC-LAG) must have the same multichassis aggregated Ethernet ID.
Options	Range: 1 through 65535.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

mode (QFX Series)

Syntax	<code>mode active-active ;</code>
Hierarchy Level	[edit interfaces aggregated-ether-options mc-ae]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the multichassis link aggregation group (MC-LAG) to be in active-active mode. In active-active mode, all of the members of the MC-LAG will be active on both routing or switching devices. Only active-active mode is supported at this time.
Options	active-active —Specify that all of the members of the MC-LAG will be active on both routing or switching devices.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	

mode (Interfaces)

Syntax	mode loose;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) rpf-check], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) rpf-check]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Check whether the packet has a source address with a corresponding prefix in the routing table. If a corresponding prefix is not found, unicast reverse path forwarding (RPF) loose mode does not accept the packet. Unlike strict mode, loose mode does not check whether the interface expects to receive a packet with a specific source address prefix.
Default	If you do not include this statement, unicast RPF is in strict mode.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Unicast RPF Strict Mode on page 2315

multicast-rpf-routes

Syntax	multicast-rpf-routes;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis traffic-engineering family inet shortcuts], [edit logical-systems <i>logical-system-name</i> routing-instances traffic-engineering family inet shortcuts], [edit protocols isis traffic-engineering family inet shortcuts], [edit routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering family inet shortcuts]
Release Information	Statement introduced in Junos OS Release 9.3.
Description	Install unicast IPv4 routes into the multicast routing table (inet.2) for multicast reverse-path-forwarding (RPF) checks. Traffic engineering shortcuts must be enabled. IPv4 multicast topology must not be enabled. Label-switched paths (LSPs) must not be advertised into IS-IS.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Enabling IS-IS Traffic Engineering Support

next-hop (PIM RPF Selection)

Syntax	<code>next-hop <i>next-hop-address</i>;</code>
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> source <i>source-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> wildcard-source], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> source <i>source-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> wildcard-source]
Release Information	Statement introduced in JUNOS Release 10.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure the specific next-hop address for the PIM group source.
Options	<i>next-hop-address</i> —Specific next-hop address for the PIM group source.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring PIM RPF Selection</i>


no-gratuitous-arp-request

Syntax	<code>no-gratuitous-arp-request;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Description	For Ethernet interfaces, do not respond to gratuitous ARP requests.
Default	Gratuitous ARP responses are enabled on all Ethernet interfaces.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Gratuitous ARP</i>• gratuitous-arp-reply on page 2351

no-local-switching

Syntax	no-local-switching;
Hierarchy Level	[edit routing-instances <i>instance-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3R2.
Description	Specify that access ports in this routing instance do not forward packets to each other.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Layer 2 Virtual Switch on page 1847

policer (MAC)

Syntax	<pre>policer { input <i>cos-policer-name</i>; output <i>cos-policer-name</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> accept-source-mac mac-address <i>mac-address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> accept-source-mac mac-address <i>mac-address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), configure MAC policing.
<div> NOTE: On MX Series routers with Gigabit Ethernet or Fast Ethernet PICs, the following considerations apply:</div> <ul style="list-style-type: none">• Interface counters do not count the 7-byte preamble and 1-byte frame delimiter in Ethernet frames.• In MAC statistics, the frame size includes MAC header and CRC before any VLAN rewrite/imposition rules are applied.• In traffic statistics, the frame size encompasses the L2 header without CRC after any VLAN rewrite/imposition rule.	
Options	<p>input <i>cos-policer-name</i>—Name of one policer to specify the premium bandwidth and aggregate bandwidth.</p> <p>output <i>cos-policer-name</i>—Name of one policer to specify the premium bandwidth and aggregate bandwidth.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring MAC Address Filtering</i>

prefix-list (PIM RPF Selection)

Syntax	<pre> prefix-list <i>prefix-list-addresses</i> { source <i>source-address</i> { next-hop <i>next-hop-address</i>; } wildcard-source { next-hop <i>next-hop-address</i>; } } </pre>
Hierarchy Level	<p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> source <i>source-address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> wildcard-source],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> source <i>source-address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> wildcard-source]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.4.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	(Optional) Configure a list of prefixes (addresses) for multiple PIM groups.
Options	<p><i>prefix-list-addresses</i>—List of prefixes (addresses) for multiple PIM groups.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>view-level—To view this statement in the configuration.</p> <p>control-level—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring PIM RPF Selection</i>

redundancy-group

Syntax	<code>redundancy-group <i>group-id</i>;</code>
Hierarchy Level	<code>[edit interfaces aggregated-ether-options mc-ae]</code>
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify the redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate the routing or switching devices contained in a redundancy group.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Multichassis Link Aggregation on page 2305

rpf-check (Dynamic Profiles)

Syntax	<pre>rpf-check { mode loose; }</pre>
Hierarchy Level	<code>[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.6.
Description	Check whether traffic is arriving on an expected path. You can include this statement with the inet protocol family only. The remaining statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Unicast RPF Strict Mode on page 2315

rpf-check (interfaces)

Syntax	<pre>rpf-check { fail-filter <i>filter-name</i>; mode loose; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Check whether traffic is arriving on an expected path. You can include this statement with the inet or inet6 protocol family only.</p> <p>The mode statement is explained separately.</p>
Options	fail-filter —A filter to evaluate when packets are received on the interface. If the RPF check fails, this optional filter is evaluated. If the fail filter is not configured, the default action is to silently discard the packet.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Unicast RPF Strict Mode on page 2315 • Configuring Unicast RPF Loose Mode on page 2316 • Example: Configuring Unicast Reverse-Path-Forwarding Check

rpf-check-policy (Routing Options RPF)

Syntax	<code>rpf-check-policy [<i>policy-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Apply policies for disabling RPF checks on arriving multicast packets. The policies must be correctly configured.
Options	<i>policy-names</i> —Name of one or more multicast RPF check policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring RPF Policies</i>

rpf-loose-mode-discard

Syntax	<pre>rpf-loose-mode-discard { family { inet; inet6; } }</pre>
Hierarchy Level	[edit forwarding-options]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure unicast reverse path forwarding (unicast RPF) loose mode with the ability to discard packets with the source address pointing to the discard next hop.
Options	inet —IPv4 address family. inet6 —IPv6 address family.
Required Privilege Level	interface-control —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Unicast RPF on page 2314

rpf-selection

Syntax	<pre>rpf-selection { group group-address { source source-address { next-hop next-hop-address; } wildcard-source { next-hop next-hop-address; } } prefix-list prefix-list-addresses { source source-address { next-hop next-hop-address; } wildcard-source { next-hop next-hop-address; } } }</pre>
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> protocols pim]
Release Information	Statement introduced in JUNOS Release 10.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure the PIM RPF next-hop neighbor for a specific group and source for a VRF routing instance. The remaining statements are explained separately.
Default	If you omit the rpf-selection statement, PIM RPF checks typically choose the best path determined by the unicast protocol for all multicast flows.
Options	source-address —Specific source address for the PIM group.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring PIM RPF Selection</i>

source (PIM RPF Selection)

Syntax	<code>source source-address { next-hop next-hop-address; }</code>
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i>]
Release Information	Statement introduced in JUNOS Release 10.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure the source address for the PIM group.
Options	source-address —Specific source address for the PIM group. The remaining statements are explained separately.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring PIM RPF Selection</i>

status-control

Syntax	<code>status-control (active standby);</code>
Hierarchy Level	[edit interfaces aggregated-ether-options mc-ae]
Release Information	Statement introduced in Junos OS Release 12.2 for the QFX Series. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify whether a peer hosting a multichassis link aggregation group (MC-LAG) is primary or secondary. Primary is considered active, and secondary is considered standby.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

targeted-broadcast

Syntax	targeted-broadcast { forward-and-send-to-re; forward-only; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	<p>Specify the IP packets destined for a Layer 3 broadcast address to be forwarded to both an egress interface and the Routing Engine, or to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.</p> <p>The statements are explained separately.</p>
Default	When this statement is not included, broadcast packets are sent to the Routing Engine only.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Targeted Broadcast on page 2313• Understanding Targeted Broadcast on page 2282

traceoptions (Individual Interfaces)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>name</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; match; } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p>
Description	<p>Define tracing operations for individual interfaces.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p> <p>The interfaces traceoptions statement does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system syslog file in the directory /var/log.</p>
Default	If you do not include this statement, no interface-specific tracing operations are performed.
Options	<p>file name—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. By default, interface process tracing output is placed in the file files number—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.dcd.</p> <p>match—(Optional) Regular expression for lines to be traced.</p> <p>no-world-readable—(Optional) Prevent any user from reading the log file.</p> <p>world-readable—(Optional) Allow any user to read the log file.</p> <p>size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches this size, it is renamed trace-file.0. When the trace-file again reaches its maximum size, trace-file.0 is renamed trace-file.1 and trace-file is renamed trace-file.0. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. The following are the interface-specific tracing options.</p> <ul style="list-style-type: none"> • all—All interface tracing operations • event—Interface events

- **ipc**—Interface interprocess communication (IPC) messages
- **media**—Interface media changes
- **q921**—Trace ISDN Q.921 frames
- **q931**—Trace ISDN Q.931 frames

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

Related Documentation

- *Tracing Operations of an Individual Router or Switch Interface*

traps

Syntax (traps | no-traps);

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description Enable or disable the sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.

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


Related Documentation

- *Enabling or Disabling SNMP Notifications on Physical Interfaces*
- *Enabling or Disabling SNMP Notifications on Logical Interfaces*

unicast-reverse-path

Syntax	unicast-reverse-path (active-paths feasible-paths);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options forwarding-table], [edit routing-instances <i>routing-instance-name</i> instance-type <i>name</i> routing-options forwarding-table], [edit routing-options forwarding-table]
Release Information	Statement introduced before Junos OS Release 7.4. Support for routing instances added in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Control the operation of unicast reverse-path-forwarding check. This statement enables the RPF check to be used when routing is asymmetrical.
Options	active-paths —Consider only active paths during the unicast reverse-path check. feasible-paths —Consider all feasible paths during the unicast reverse-path check. Default: If you omit the unicast-reverse-path statement, only the active paths to a particular destination are considered.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Unicast Reverse-Path-Forwarding Check</i> • <i>Enabling Unicast Reverse-Path Forwarding Check for VPNs</i>

unidirectional

Syntax	unidirectional;
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Create two new, unidirectional (transmit-only and receive-only) physical interfaces subordinate to the original parent interface. Unidirectional links are currently supported only on 10-Gigabit Ethernet interfaces on the following hardware: <div data-bbox="472 741 542 810"></div> <div data-bbox="581 783 1099 816">NOTE: Which interfaces and pic for EX9200?</div> <ul style="list-style-type: none">• 4-port 10-Gigabit Ethernet DPC on the MX960 router• 10-Gigabit Ethernet IQ2 PIC and 10-Gigabit Ethernet IQ2E PIC on the T Series router
Default	Disabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding Unidirectional Traffic Flow on Physical Interfaces</i>• <i>Enabling Unidirectional Traffic Flow on Physical Interfaces</i>

vlan-tagging

Syntax	vlan-tagging;
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Description	For Fast Ethernet and Gigabit Ethernet interfaces and aggregated Ethernet interfaces configured for VPLS, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch</i> • <i>Example: Configuring BGP Autodiscovery for LDP VPLS</i> • <i>Configuring a Layer 3 Subinterface (CLI Procedure)</i> • Configuring Tagged Aggregated Ethernet Interfaces on page 2302 • Configuring Interfaces for VPLS Routing on page 5336 • Enabling VLAN Tagging on page 5339 • <i>802.1Q VLANs Overview</i> • <i>vlan-id</i>

wildcard-source (PIM RPF Selection)

Syntax	wildcard-source { next-hop next-hop-address; }
Hierarchy Level	[edit routing-instances routing-instance-name protocols pim rpf-selection group group-address], [edit routing-instances routing-instance-name protocols pim rpf-selection prefix-list prefix-list-addresses]
Release Information	Statement introduced in Junos OS Release 10.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Use a wildcard for the multicast source instead of (or in addition to) a specific multicast source. The remaining statements are explained separately.
Required Privilege Level	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring PIM RPF Selection

Administration

- [Routine Monitoring on page 2374](#)
- [Operational Commands on page 2376](#)

Routine Monitoring

- [Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface on page 2374](#)

Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface

Purpose	To monitor statistics for a Fast Ethernet or Gigabit Ethernet interface, use the following Junos OS CLI operational mode command:
Action	user@host> monitor interface (fe-fpc/pic/port ge-fpc/pic/port)



CAUTION: We recommend that you use the monitor interface *fe-fpc/pic/port* or monitor interface *ge-fpc/pic/port* command only for diagnostic purposes. Do not leave these commands on during normal router or switch operations because real-time monitoring of traffic consumes additional CPU and memory resources.

Sample Output

The following sample output is for a Fast Ethernet interface:

```

user@host> monitor interface fe-2/1/0
Interface: fe-2/1/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 100mbps
Traffic statistics:
  Input bytes:          282556864218 (14208 bps)      [40815]
  Output bytes:         42320313078 (384 bps)        [890]
  Input packets:        739373897 (11 pps)           [145]
  Output packets:       124798688 (1 pps)            [14]
Error statistics:
  Input errors:          0                          [0]
  Input drops:           0                          [0]
  Input framing errors:  0                          [0]
  Policed discards:     6625892                     [6]
  L3 incompletes:        75                         [0]
  L2 channel errors:    0                          [0]
  L2 mismatch timeouts: 0                          [0]
  Carrier transitions:   1                          [0]
  Output errors:         0                          [0]
  Output drops:         0                          [0]
  Aged packets:         0                          [0]
Active alarms : None
Active defects: None
Input MAC/Filter statistics:
  Unicast packets       464751787                   [154]
  Packet error count    0                          [0]

```

Meaning Use the information from this command to help narrow down possible causes of an interface problem.



NOTE: If you are accessing the router or switch from the console connection, make sure you set the CLI terminal type using the `set cli terminal` command.

The statistics in the second column are the cumulative statistics since the last time they were cleared using the `clear interfaces statistics interface-name` command. The statistics in the third column are the cumulative statistics since the `monitor interface interface-name` command was executed.

If the input errors are increasing, verify the following:

1. Check the cabling to the router and have the carrier verify the integrity of the line. To verify the integrity of the cabling, make sure that you have the correct cables for the interface port. Make sure you have single-mode fiber cable for a single-mode interface and multimode fiber cable for a multimode interface.
2. For a fiber-optic connection, measure the received light level at the receiver end and make sure that it is within the receiver specification of the Ethernet interface. See *Fiber-Optic Ethernet Interface Specifications* for the fiber-optic Ethernet interface specifications.

3. Measure the transmit light level on the Tx port to verify that it is within specification. See *Fiber-Optic Ethernet Interface Specifications* for the optical specifications.

Operational Commands

- [Common Output Fields Description on page 2376](#)

Common Output Fields Description

This chapter explains the content of the output fields, which appear in the output of most **show interfaces** commands.

Destination Class Field

For the logical interface, the **Destination class** field provides the names of destination class usage (DCU) counters per family and per class for a particular interface. The counters display packets and bytes arriving from designated user-selected prefixes. For example:

Destination class	Packets (packet-per-second)	Bytes (bits-per-second)
gold	1928095	161959980
	(889)	(597762)
bronze	0	0
	(0)	(0)
silver	0	0
	(0)	(0)

Enabled Field

For the physical interface, the **Enabled** field provides information about the state of the interface, displaying one or more of the following values:

- **Administratively down, Physical link is Down**—The interface is turned off, and the physical link is inoperable and cannot pass packets even when it is enabled. To change the interface state to **Enabled**, use the following command:

```
user@host# set interfaces interface enable
```

Manually verify the connections to bring the physical link up.

- **Administratively down, Physical link is Up**—The interface is turned off, but the physical link is operational and can pass packets when it is enabled. To change the interface state to **Enabled**, use the following command:

```
user@host# set interfaces interface enable
```

- **Enabled, Physical link is Down**—The interface is turned on, but the physical link is inoperable and cannot pass packets. Manually verify the connections to bring the physical link up.
- **Enabled, Physical link is Up**—The interface is turned on, and the physical link is operational and can pass packets.

Filters Field

For the logical interface, the **Filters** field provides the name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is **Filters: Input: *filter-name* and Filters: Output: *filter-name***. For example:

```
Filters: Input: sample-all
Filters: Output: cp-ftp
```

Flags Fields

The following sections provide information about flags that are specific to interfaces:

- [Addresses, Flags Field on page 2377](#)
- [Device Flags Field on page 2377](#)
- [Family Flags Field on page 2378](#)
- [Interface Flags Field on page 2379](#)
- [Link Flags Field on page 2379](#)
- [Logical Interface Flags Field on page 2380](#)

Addresses, Flags Field

The **Addresses, Flags** field provides information about the addresses configured for the protocol family on the logical interface and displays one or more of the following values:

- **Dest-route-down**—The routing process detected that the link was not operational and changed the interface routes to nonforwarding status
- **Is-Default**—The default address of the router used as the source address by SNMP, ping, traceroute, and other network utilities.
- **Is-Preferred**—The default local address for packets originating from the local router and sent to destinations on the subnet.
- **Is-Primary**—The default local address for broadcast and multicast packets originated locally and sent out the interface.
- **Preferred**—This address is a candidate to become the preferred address.
- **Primary**—This address is a candidate to become the primary address.

Device Flags Field

The **Device flags** field provides information about the physical device and displays one or more of the following values:

- **Down**—Device has been administratively disabled.
- **Hear-Own-Xmit**—Device receives its own transmissions.
- **Link-Layer-Down**—The link-layer protocol has failed to connect with the remote endpoint.
- **Loopback**—Device is in physical loopback.

- **Loop-Detected**—The link layer has received frames that it sent, thereby detecting a physical loopback.
- **No-Carrier**—On media that support carrier recognition, no carrier is currently detected.
- **No-Multicast**—Device does not support multicast traffic.
- **Present**—Device is physically present and recognized.
- **Promiscuous**—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media.
- **Quench**—Transmission on the device is quenched because the output buffer is overflowing
- **Recv-All-Multicasts**—Device is in multicast promiscuous mode and therefore provides no multicast filtering.
- **Running**—Device is active and enabled.

Family Flags Field

The **Family flags** field provides information about the protocol family on the logical interface and displays one or more of the following values:

- **DCU**—Destination class usage is enabled.
- **Dest-route-down**—The software detected that the link is down and has stopped forwarding the link's interface routes.
- **Down**—Protocol is inactive.
- **Is-Primary**—Interface is the primary one for the protocol.
- **Mac-Validate-Loose**—Interface is enabled with loose MAC address validation.
- **Mac-Validate-Strict**—Interface is enabled with strict MAC address validation.
- **Maximum labels**—Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.
- **MTU-Protocol-Adjusted**—The effective MTU is not the configured value in the software.
- **No-Redirects**—Protocol redirects are disabled.
- **Primary**—Interface can be considered for selection as the primary family address.
- **Protocol-Down**—Protocol failed to negotiate correctly.
- **SCU-in**—Interface is configured for source class usage input.
- **SCU-out**—Interface is configured for source class usage output.
- **send-bcast-packet-to-re**—Interface is configured to forward IPv4 broadcast packets to the Routing Engine.
- **targeted-broadcast**—Interface is configured to forward IPv4 broadcast packets to the LAN interface and the Routing Engine.

- **Unnumbered**—Protocol family is configured for unnumbered Ethernet. An unnumbered Ethernet interface borrows an IPv4 address from another interface, which is referred to as the donor interface.
- **Up**—Protocol is configured and operational.
- **uRPF**—Unicast Reverse Path Forwarding is enabled.

Interface Flags Field

The **Interface flags** field provides information about the physical interface and displays one or more of the following values:

- **Admin-Test**—Interface is in test mode and some sanity checking, such as loop detection, is disabled.
- **Disabled**—Interface is administratively disabled.
- **Down**—A hardware failure has occurred.
- **Hardware-Down**—Interface is nonfunctional or incorrectly connected.
- **Link-Layer-Down**—Interface keepalives have indicated that the link is incomplete.
- **No-Multicast**—Interface does not support multicast traffic.
- **No-receive No-transmit**—Passive monitor mode is configured on the interface.
- **Point-To-Point**—Interface is point-to-point.
- **Pop all MPLS labels from packets of depth**—MPLS labels are removed as packets arrive on an interface that has the **pop-all-labels** statement configured. The depth value can be one of the following:
 - **1**—Takes effect for incoming packets with one label only.
 - **2**—Takes effect for incoming packets with two labels only.
 - **[1 2]**—Takes effect for incoming packets with either one or two labels.
- **Promiscuous**—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses.
- **Recv-All-Multicasts**—Interface is in multicast promiscuous mode and provides no multicast filtering.
- **SNMP-Traps**—SNMP trap notifications are enabled.
- **Up**—Interface is enabled and operational.

Link Flags Field

The **Link flags** field provides information about the physical link and displays one or more of the following values:

- **ACFC**—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option.
- **Give-Up**—Link protocol does not continue connection attempts after repeated failures.

- **Loose-LCP**—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational.
- **Loose-LMI**—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational.
- **Loose-NCP**—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational.
- **No-Keepalives**—Link protocol keepalives are disabled.
- **PFC**—Protocol field compression is configured. The PPP session negotiates the PFC option.

Logical Interface Flags Field

The **Logical interface flags** field provides information about the logical interface and displays one or more of the following values:

- **ACFC Encapsulation**—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer).
- **Device-down**—Device has been administratively disabled.
- **Disabled**—Interface is administratively disabled.
- **Down**—A hardware failure has occurred.
- **Clear-DF-Bit**—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit.
- **Hardware-Down**—Interface protocol initialization failed to complete successfully.
- **PFC**—Protocol field compression is enabled for the PPP session.
- **Point-To-Point**—Interface is point-to-point.
- **SNMP-Traps**—SNMP trap notifications are enabled.
- **Up**—Interface is enabled and operational.

Label-Switched Interface Traffic Statistics Field

When you use the **vrf-table-label** statement to configure a VRF routing table, a label-switched interface (LSI) logical interface label is created and mapped to the VRF routing table.

Any routes present in a VRF routing table and configured with the **vrf-table-label** statement are advertised with the LSI logical interface label allocated for the VRF routing table. When packets for this VPN arrive on a core-facing interface, they are treated as if the enclosed IP packet arrived on the LSI interface and are then forwarded and filtered based on the correct table. For more information on the **vrf-table-label** statement, including a list of supported interfaces, see the *Junos VPNs Configuration Guide*.

If you configure the **family mpls** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level and you also configure the **vrf-table-label** statement at the **[edit routing-instances *routing-instance-name*]** hierarchy level, the output for the

show interface *interface-name* extensive command includes the following output fields about the LSI traffic statistics:

- **Input bytes**—Number of bytes entering the LSI and the current throughput rate in bits per second (bps).
- **Input packets**—Number of packets entering the LSI and the current throughput rate in packets per second (pps).



NOTE: If LSI interfaces are used with VPLS when **no-tunnel-services** is configured or L3VPN when **vrf-table-label** configuration is applied inside the routing-instance, the **Input packets** field associated with the core-facing interfaces may not display the correct value. Only the Input counter is affected because the LSI is used to receive traffic from the remote PEs. Traffic that arrives on an LSI interface might not be counted at both the Traffic Statistics and the Label-switched interface (LSI) traffic statistics levels.

This note applies to the following platforms:

- M Series routers with -E3 FPC model numbers or configured with an Enhanced CFEB (CFEB-E), and M120 routers
- MX Series routers with DPC or ADPC only

The following example shows the LSI traffic statistics that you might see as part of the output of the **show interface *interface-name* extensive** command:

Label-switched interface (LSI) traffic statistics:

Input bytes:	0	0 bps
Input packets:	0	0 pps

Policer Field

For the logical interface, the **Policer** field provides the policers that are to be evaluated when packets are received or transmitted on the interface. The format is **Policer: Input: *type-fpc/picport-in-policer*, Output: *type-fpc/pic/port-out-policer***. For example:

Policer: Input: at-1/2/0-in-policer, Output: at-2/4/0-out-policer

Protocol Field

For the logical interface, the **Protocol** field indicates the protocol family or families that are configured on the interface, displaying one or more of the following values:

- **aenet**—Aggregated Ethernet. Displayed on Fast Ethernet interfaces that are part of an aggregated Ethernet bundle.
- **ccc**—Circuit cross-connect (CCC). Configured on the logical interface of CCC physical interfaces.
- **inet**—IP version 4 (IPv4). Configured on the logical interface for IPv4 protocol traffic, including Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), Internet Control Message Protocol (ICMP), and Internet Protocol Control Protocol (IPCP).

- **inet6**—IP version 6 (IPv6). Configured on the logical interface for IPv6 protocol traffic, including Routing Information Protocol for IPv6 (RIPng), Intermediate System-to-Intermediate System (IS-IS), and BGP.
- **iso**—International Organization for Standardization (ISO). Configured on the logical interface for IS-IS traffic.
- **mlfr-uni-nni**—Multilink Frame Relay (MLFR) FRF.16 user-to-network network-to-network (UNI NNI). Configured on the logical interface for link services bundling.
- **mlfr-end-to-end**—Multilink Frame Relay end-to-end. Configured on the logical interface for multilink bundling.
- **mlppp**—Multilink Point-to-Point Protocol (MLPPP). Configured on the logical interface for multilink bundling.
- **mpls**—Multiprotocol Label Switching (MPLS). Configured on the logical interface for participation in an MPLS path.
- **pppoe**—Point-to-Point Protocol over Ethernet (PPPoE). Configured on Ethernet interfaces enabled to support multiple protocol families.
- **tcc**—Translational cross-connect (TCC). Configured on the logical interface of TCC physical interfaces.
- **tnp**—Trivial Network Protocol (TNP). Used to communicate between the Routing Engine and the router's packet forwarding components. The Junos OS automatically configures this protocol family on the router's internal interfaces only.
- **vpls**—Virtual private LAN service (VPLS). Configured on the logical interface on which you configure VPLS.

RPF Failures Field

For the logical interface, the **RPF Failures** field provides information about the amount of incoming traffic (in packets and bytes) that failed a unicast reverse path forwarding (RPF) check on a particular interface. The format is **RPF Failures: Packets: xx,Bytes: yy**. For example:

RPF Failures: Packets: 0, Bytes:0

Source Class Field

For the logical interface, the **Source class** field provides the names of source class usage (SCU) counters per family and per class for a particular interface. The counters display packets and bytes arriving from designated user-selected prefixes. For example:

Source class	Packets (packet-per-second)	Bytes (bits-per-second)
gold	1928095	161959980
(889)	(597762)
bronze	0	0
(0)	(0)
silver	0	0
(0)	(0)

clear interfaces statistics

Syntax	clear interfaces statistics (all <i>interface-name</i>)
Release Information	Command introduced before Junos OS Release 7.4.
Description	Set interface statistics to zero. If you issue the clear interfaces statistics <i>interface-name</i> command and then perform a graceful Routing Engine switchover, the interface statistics are not cleared on the new master. Reissue the command to clear the interface statistics again.
Options	all —Set statistics on all interfaces to zero. <i>interface-name</i> —Set statistics on a particular interface to zero.
Required Privilege Level	clear
List of Sample Output	clear interfaces statistics on page 2383
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear interfaces statistics

```
user@host> clear interfaces statistics
```

show interfaces (10-Gigabit Ethernet)

Syntax	<code>show interfaces <i>xe-fpc/pic/port</i></code> <code><brief detail extensive terse></code> <code><descriptions></code> <code><media></code> <code><snmp-index <i>snmp-index</i>></code> <code><statistics></code>
Release Information	Command introduced in Junos OS Release 8.0.
Description	(M320, M120, MX Series, and T Series routers and EX Series switches only) Display status information about the specified 10-Gigabit Ethernet interface.
Options	<p><code><i>xe-fpc/pic/port</i></code>—Display standard information about the specified 10-Gigabit Ethernet interface.</p> <p><code>brief detail extensive terse</code>—(Optional) Display the specified level of output.</p> <p><code>descriptions</code>—(Optional) Display interface description strings.</p> <p><code>media</code>—(Optional) Display media-specific information about network interfaces.</p> <p><code>snmp-index <i>snmp-index</i></code>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><code>statistics</code>—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
List of Sample Output	<p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2) on page 2399</p> <p>show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode) on page 2402</p> <p>show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC) on page 2404</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode) on page 2406</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only) on page 2406</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only) on page 2407</p>
Output Fields	See Table 135 on page 1957 for the output fields for the show interfaces (10-Gigabit Ethernet) command.

Table 168: show interfaces Gigabit Ethernet Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 2376 .	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
WAN-PHY mode	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: Enabled or Disabled for parent interface; Rx-only or Tx-only for child interfaces.	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under “Common Output Fields Description” on page 2376 .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 2376 .	All levels

Table 168: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Link flags	Information about the link. Possible values are described in the “Links Flags” section under “ Common Output Fields Description ” on page 2376.	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	detail extensive none
Schedulers	(Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified
Output Rate	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Egress accounting overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	detail extensive
Ingress accounting overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	detail extensive

Table 168: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Output bytes—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 135 on page 1957.</p>	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the <code>ignore-l3-incompletes</code> statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 168: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 168: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 168: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 168: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 136 on page 1971 • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—Number of frames that exceed 1518 octets. • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive
OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	extensive
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	extensive

Table 168: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 168: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 168: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 168: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 168: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under "Common Output Fields Description" on page 2376 .	All levels

Table 168: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> push—An outer VLAN tag is pushed in front of the existing VLAN tag. pop—The outer VLAN tag of the incoming frame is removed. swap—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. push—An outer VLAN tag is pushed in front of the existing VLAN tag. push-push—Two VLAN tags are pushed in from the incoming frame. swap-push—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value. pop-swap—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed. 	brief detail extensive none
Demux:	<p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> Source Family Inet Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under “ Common Output Fields Description ” on page 2376.	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> Input bytes, Output bytes—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the routing device.	extensive

Table 168: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the Output bytes and Output packets interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under “ Common Output Fields Description ” on page 2376.	detail extensive
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	detail extensive none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	detail extensive
Output Filters	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	detail extensive
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under “ Common Output Fields Description ” on page 2376.	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about address flag (possible values are described in the “Addresses Flags” section under “ Common Output Fields Description ” on page 2376.	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interlace.	detail extensive none

Table 168: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

For Gigabit Ethernet IQ PICs, traffic and MAC statistics output varies. [Table 136 on page 1971](#) describes the traffic and MAC statistics for two sample interfaces, each of which is sending traffic in packets of 500 bytes (including 478 bytes for the Layer 3 packet, 18 bytes for the Layer 2 VLAN traffic header, and 4 bytes for cyclic redundancy check [CRC] information). In [Table 136 on page 1971](#), the **ge-0/3/0** interface is the inbound physical interface, and the **ge-0/0/0** interface is the outbound physical interface. On both interfaces, traffic is carried on logical unit .50 (VLAN 50).

Table 169: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Inbound physical interface	show interfaces ge-0/3/0 extensive	Traffic statistics: Input bytes: 496 bytes per packet, representing the Layer 2 packet MAC statistics: Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes	The additional 4 bytes are for the CRC.
Inbound logical interface	show interfaces ge-0/3/0.50 extensive	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	
Outbound physical interface	show interfaces ge-0/0/0 extensive	Traffic statistics: Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes MAC statistics: Received octets: 478 bytes per packet, representing the Layer 3 packet	For input bytes, the additional 12 bytes includes 6 bytes for the destination MAC address + 4 bytes for VLAN + 2 bytes for the Ethernet type.
Outbound logical interface	show interfaces ge-0/0/0.50 extensive	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	

Sample Output

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2)

```

user@host> show interfaces xe-5/0/0 extensive
Physical interface: xe-5/0/0, Enabled, Physical link is Up
  Interface index: 177, SNMP ifIndex: 99, Generation: 178
  Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Speed: 10Gbps, Loopback:

```

```

None, Source filtering: Enabled,
Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags : None
CoS queues : 8 supported, 4 maximum usable queues
Schedulers : 1024
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:14:f6:b9:f1:f6, Hardware address: 00:14:f6:b9:f1:f6
Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes : 6970332384 0 bps
Output bytes : 0 0 bps
Input packets: 81050506 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Ingress traffic statistics at Packet Forwarding Engine:
Input bytes : 6970299398 0 bps
Input packets: 81049992 0 pps
Drop bytes : 0 0 bps
Drop packets: 0 0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0,
MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets

0 best-effort 81049992 81049992 0
1 expedited-fo 0 0 0
2 assured-forw 0 0 0
3 network-cont 0 0 0

Egress queues: 4 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets

0 best-effort 0 0 0
1 expedited-fo 0 0 0
2 assured-forw 0 0 0
3 network-cont 0 0 0

Active alarms : None
Active defects : None
PCS statistics Seconds
Bit errors 0
Errored blocks 0

```

```

MAC statistics:
Total octets          6970332384
Total packets        81050506
Unicast packets      81050000
Broadcast packets    506
Multicast packets    0
CRC/Align errors     0
FIFO errors          0
MAC control frames   0
MAC pause frames     0
Oversized frames     0
Jabber frames        0
Fragment frames      0
VLAN tagged frames   0
Code violations       0

Filter statistics:
Input packet count    81050506
Input packet rejects  506
Input DA rejects      0
Input SA rejects      0
Output packet count   0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0

Packet Forwarding Engine configuration:
Destination slot: 5

CoS information:
Direction : Output
CoS transmit queue   Bandwidth      Buffer Priority Limit
                        %      bps      %      usec
0 best-effort        95    950000000  95      0      low  none
3 network-control    5     50000000   5      0      low  none

Direction : Input
CoS transmit queue   Bandwidth      Buffer Priority Limit
                        %      bps      %      usec
0 best-effort        95    950000000  95      0      low  none
3 network-control    5     50000000   5      0      low  none

Logical interface xe-5/0/0.0 (Index 71) (SNMP ifIndex 95) (Generation 195)
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ] Encapsulation: ENET2
Egress accounting overhead: 100
Ingress accounting overhead: 90

Traffic statistics:
Input bytes : 0
Output bytes : 46
Input packets: 0
Output packets: 1

IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Local statistics:
Input bytes : 0
Output bytes : 46
Input packets: 0
Output packets: 1

Transit statistics:
Input bytes : 0
Output bytes : 0

```

```

Input packets:                0                0 pps
Output packets:               0                0 pps
IPv6 transit statistics:
  Input bytes :                0
  Output bytes :               0
  Input packets:              0
  Output packets:             0
Protocol inet, MTU: 1500, Generation: 253, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 192.1.1/24, Local: 192.1.1.1, Broadcast: 192.1.1.255,
Generation: 265
Protocol multiservice, MTU: Unlimited, Generation: 254, Route table: 0
  Flags: None
  Policer: Input: __default_arp_policer__

```

show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode)

```

user@host> show interfaces xe-1/0/0 extensive
Physical interface: xe-1/0/0, Enabled, Physical link is Up
  Interface index: 141, SNMP ifIndex: 34, Generation: 47
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Loopback: Disabled
  WAN-PHY mode
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Link flags     : None
  CoS queues     : 4 supported
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:a2:10:9d, Hardware address: 00:05:85:a2:10:9d
  Last flapped   : 2005-07-07 11:22:34 PDT (3d 12:28 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes :                0                0 bps
    Output bytes :               0                0 bps
    Input packets:              0                0 pps
    Output packets:             0                0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    HS Link CRC errors: 0, HS Link FIFO overflows: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0,
    Aged packets: 0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
    Resource errors: 0
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets
    0 best-effort   0                0                0
    1 expedited-fo  0                0                0
    2 assured-forw  0                0                0
    3 network-cont  0                0                0
  Active alarms : LOL, LOS, LBL
  Active defects: LOL, LOS, LBL, SEF, AIS-L, AIS-P
  PCS statistics
    Seconds  Count
    Bit errors  0        0
    Errored blocks  0        0
  MAC statistics:
    Receive  Transmit
    Total octets  0        0
    Total packets  0        0
    Unicast packets  0        0
    Broadcast packets  0        0
    Multicast packets  0        0

```



```

CRC/Align errors                0          0
FIFO errors                     0          0
MAC control frames              0          0
MAC pause frames                0          0
Oversized frames               0
Jabber frames                  0
Fragment frames                0
VLAN tagged frames             0
Code violations                 0
Filter statistics:
  Input packet count            0
  Input packet rejects          0
  Input DA rejects              0
  Input SA rejects              0
  Output packet count           0
  Output packet pad count       0
  Output packet error count     0
CAM destination filters: 0, CAM source filters: 0
PMA PHY:
  Seconds      Count  State
  PLL lock     0      0 OK
  PHY light    63159  1 Light Missing
WIS section:
  BIP-B1        0      0
  SEF           434430  434438 Defect Active
  LOS           434430  1 Defect Active
  LOF           434430  1 Defect Active
  ES-S          434430
  SES-S         434430
  SEFS-S        434430
WIS line:
  BIP-B2        0      0
  REI-L         0      0
  RDI-L         0      0 OK
  AIS-L         434430  1 Defect Active
  BERR-SF       0      0 OK
  BERR-SD       0      0 OK
  ES-L          434430
  SES-L         434430
  UAS-L         434420
  ES-LFE        0
  SES-LFE       0
  UAS-LFE       0
WIS path:
  BIP-B3        0      0
  REI-P         0      0
  LOP-P         0      0 OK
  AIS-P         434430  1 Defect Active
  RDI-P         0      0 OK
  UNEQ-P        0      0 OK
  PLM-P         0      0 OK
  ES-P          434430
  SES-P         434430
  UAS-P         434420
  ES-PFE        0
  SES-PFE       0
  UAS-PFE       0
Received path trace:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted path trace: orissa so-1/0/0
6f 72 69 73 73 61 20 73 6f 2d 31 2f 30 2f 30 00 orissa so-1/0/0.
Packet Forwarding Engine configuration:

```

```

Destination slot: 1
CoS information:
  CoS transmit queue      Bandwidth      Buffer      Priority  Limit
                           %      bps      %      bytes
  0 best-effort           95      950000000  95         0      low      none
  3 network-control       5       50000000   5         0      low      none

```

show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC)

```

user@host> show interfaces ge-7/0/0 extensive
Physical interface: ge-7/0/0, Enabled, Physical link is Down
Interface index: 143, SNMP ifIndex: 508, Generation: 208
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Enabled
Device flags   : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
Link flags     : None
Wavelength     : 1550.12 nm, Frequency: 193.40 THz
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:05:85:70:2b:72, Hardware address: 00:05:85:70:2b:72
Last flapped   : 2011-04-20 15:48:54 PDT (18:39:49 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes   : 0          0 bps
Output bytes  : 0          0 bps
Input packets : 0          0 pps
Output packets: 0          0 pps
IPv6 transit statistics:
Input bytes   : 0
Output bytes  : 0
Input packets : 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 2, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort           0              0              0

  1 expedited-fo         0              0              0

  2 assured-forw         0              0              0

  3 network-cont
Queue number:      Mapped forwarding classes
  0                best-effort
  1                expedited-forwarding
  2                assured-forwarding
  3                network-control
Active alarms  : LINK
Active defects : LINK
MAC statistics:
Total octets      Receive      Transmit
Total packets     0              0

```

```

Unicast packets                0                0
Broadcast packets              0                0
Multicast packets              0                0
CRC/Align errors               0                0
FIFO errors                    0                0
MAC control frames             0                0
MAC pause frames               0                0
Oversized frames               0
Jabber frames                  0
Fragment frames                0
VLAN tagged frames             0
Code violations                 0
Total octets                   0                0
Total packets                  0                0
Unicast packets                0                0
Broadcast packets              0                0
Multicast packets              0                0
CRC/Align errors               0                0
FIFO errors                    0                0
MAC control frames             0                0
MAC pause frames               0                0
Oversized frames               0
Jabber frames                  0
Fragment frames                0
VLAN tagged frames             0
Code violations                 0
OTN alarms                     :   None
OTN defects                    :   None
OTN FEC Mode                   :   GFEC
OTN Rate                       :   Fixed Stuff Bytes 11.0957Gbps
OTN Line Loopback              :   Enabled
OTN FEC statistics :
  Corrected Errors              0
  Corrected Error Ratio (      0 sec average) 0e-0
OTN FEC alarms:                Seconds    Count  State
  FEC Degrade                   0          0    OK
  FEC Excessive                 0          0    OK
OTN OC:                        Seconds    Count  State
  LOS                           2          1    OK
  LOF                           67164      2  Defect Active
  LOM                           67164      71  Defect Active
  Wavelength Lock               0          0    OK
OTN OTU:
  AIS                           0          0    OK
  BDI                           65919     4814  Defect Active
  IAE                           67158      1  Defect Active
  TTIM                          7          1    OK
  SF                             67164      2  Defect Active
  SD                             67164      3  Defect Active
  TCA-ES                        0          0    OK
  TCA-SES                       0          0    OK
  TCA-UAS                       80         40    OK
  TCA-BBE                       0          0    OK
  BIP                           0          0    OK
  BBE                           0          0    OK
  ES                            0          0    OK
  SES                           0          0    OK
  UAS                           587         0    OK
Received DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Received SAPI:

```

```

00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
OTN Received Overhead Bytes:
  APS/PCC0: 0x02, APS/PCC1: 0x42, APS/PCC2: 0xa2, APS/PCC3: 0x48
  Payload Type: 0x03
OTN Transmitted Overhead Bytes:
  APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
  Payload Type: 0x03
Filter statistics:
  Input packet count                0
  Input packet rejects              0
  Input DA rejects                  0
  Input SA rejects                  0
  Output packet count                0
  Output packet pad count            0
  Output packet error count          0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 7
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
      0 best-effort        95      9500000000    95      0      low
none
      3 network-control    5       500000000    5       0      low
none
...

```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode)

```

user@host> show interfaces xe-7/0/0 extensive
Physical interface: xe-7/0/0, Enabled, Physical link is Up
  Interface index: 173, SNMP ifIndex: 212, Generation: 174
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
  Unidirectional: Enabled,
  Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
...

```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only)

```

user@host> show interfaces xe-7/0/0-tx extensive
Physical interface: xe-7/0/0-tx, Enabled, Physical link is Up
  Interface index: 176, SNMP ifIndex: 137, Generation: 177
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
  Unidirectional: Tx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:73:e4:83, Hardware address: 00:05:85:73:e4:83
  Last flapped   : 2007-06-01 09:08:19 PDT (3d 02:31 ago)
  Statistics last cleared: Never
Traffic statistics:
  Input bytes   :                0                0 bps

```

```

Output bytes :      322891152287160      9627472888 bps
Input packets:              0              0 pps
Output packets:    328809727380      1225492 pps

...

Filter statistics:
  Output packet count      328810554250
  Output packet pad count      0
  Output packet error count    0
...

Logical interface xe-7/0/0-tx.0 (Index 73) (SNMP ifIndex 138) (Generation 139)

Flags: SNMP-Traps Encapsulation: ENET2
Egress accounting overhead: 100
Ingress accounting overhead: 90
Traffic statistics:
  Input bytes :              0
  Output bytes :    322891152287160
  Input packets:              0
  Output packets:    328809727380
IPv6 transit statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:            0
Local statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:            0
Transit statistics:
  Input bytes :              0              0 bps
  Output bytes :    322891152287160      9627472888 bps
  Input packets:              0              0 pps
  Output packets:    328809727380      1225492 pps
IPv6 transit statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:            0
Protocol inet, MTU: 1500, Generation: 147, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.11.12/24, Local: 10.11.12.13, Broadcast: 10.11.12.255,
Generation: 141
Protocol multiservice, MTU: Unlimited, Generation: 148, Route table: 0
  Flags: None
  Policer: Input: __default_arp_policer__

```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only)

```

user@host> show interfaces xe-7/0/0-rx extensive
Physical interface: xe-7/0/0-rx, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 118, Generation: 175
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Rx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues

```

```

Hold-times      : Up 0 ms, Down 0 ms
Current address: 00:05:85:73:e4:83, Hardware address: 00:05:85:73:e4:83
Last flapped   : 2007-06-01 09:08:22 PDT (3d 02:31 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes :      322857456303482      9627496104 bps
Output bytes :              0          0 bps
Input packets:      328775413751      1225495 pps
Output packets:              0          0 pps

...

Filter statistics:
Input packet count      328775015056
Input packet rejects    1
Input DA rejects        0

...

Logical interface xe-7/0/0-rx.0 (Index 72) (SNMP ifIndex 120) (Generation 138)

Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
Input bytes :      322857456303482
Output bytes :              0
Input packets:      328775413751
Output packets:              0
IPv6 transit statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:              0
Local statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:              0
Transit statistics:
Input bytes :      322857456303482      9627496104 bps
Output bytes :              0          0 bps
Input packets:      328775413751      1225495 pps
Output packets:              0          0 pps
IPv6 transit statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:              0
Protocol inet, MTU: 1500, Generation: 145, Route table: 0
Addresses, Flags: Is-Preferred Is-Primary
Destination: 192.1.1/24, Local: 192.1.1.1, Broadcast: 192.1.1.255,
Generation: 139
Protocol multiservice, MTU: Unlimited, Generation: 146, Route table: 0
Flags: None
Policer: Input: __default_arp_policer__

```

show interfaces (Discard)

Syntax	<pre>show interfaces dsc <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display status information about the specified discard interface.
Options	<p>dsc—Display standard information about the specified discard interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—This option is not relevant for the discard interface and always shows a value of 0.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) This option is not relevant for the discard interface and always shows a value of 0.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> <i>show interfaces (ATM)</i> <i>show interfaces routing</i>
List of Sample Output	show interfaces dsc on page 2412 show interfaces dsc brief on page 2412 show interfaces dsc detail on page 2412 show interfaces dsc extensive on page 2413
Output Fields	Table 170 on page 2409 lists the output fields for the show interfaces (discard) command. Output fields are listed in the approximate order in which they appear.

Table 170: Discard show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface, whether the interface is enabled, and the state of the physical interface: Up or Down .	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	detail extensive none

Table 170: Discard show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Type	Type of interface. Software-Pseudo indicates a standard software interface with no associated hardware device.	All levels
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	MTU size on the physical interface.	All levels
Clocking	Reference clock source. It can be Internal or External .	brief detail extensive
Speed	Speed at which the interface is running.	brief detail extensive
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under " Common Output Fields Description " on page 2376.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under " Common Output Fields Description " on page 2376.	All levels
Link type	Encapsulation being used on the physical interface.	detail extensive
Link flags	Information about the link. Possible values are described in the "Link Flags" section under " Common Output Fields Description " on page 2376.	detail extensive
Physical info	Information about the physical interface.	detail extensive
Hold-times	Current interface hold-time up and hold-time down. Value is in milliseconds.	detail extensive
Current address, Hardware address	Configured MAC address and hardware MAC address.	detail extensive
Alternate link address	Backup address of the link.	detail extensive
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive

Table 170: Discard show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. 	detail extensive
Input errors	<p>Input errors on the interface:</p> <ul style="list-style-type: none"> • Errors—Sum of incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Giants—Number of frames received that are larger than the giant threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • Resource errors—Sum of transmit drops. 	detail extensive
Output errors	<p>(Extensive only) Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	detail extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	detail extensive
SNMP ifIndex	Logical interface SNMP interface index number.	detail extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Table 170: Discard show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under “ Common Output Fields Description ” on page 2376.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family configured on the logical interface, such as iso , inet6 , or mpls .	All levels
MTU	MTU size on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0 .	detail extensive

Sample Output

show interfaces dsc

```

user@host> show interfaces dsc
Physical interface: dsc, Enabled, Physical link is Up
  Interface index: 5, SNMP ifIndex: 5
  Type: Software-Pseudo, MTU: Unlimited
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : None
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface dsc.0 (Index 66) (SNMP ifIndex 235)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Unspecified
    Protocol inet, MTU: Unlimited
    Flags: None

```

show interfaces dsc brief

```

user@host> show interfaces dsc brief
Physical interface: dsc, Enabled, Physical link is Up
  Type: Software-Pseudo, Link-level type: Unspecified, MTU: Unlimited, Clocking:
  Unspecified, Speed: Unspecified
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps

  Logical interface dsc.0
    Flags: Point-To-Point SNMP-Traps Encapsulation: Unspecified
    inet

```

show interfaces dsc detail

```

user@host> show interfaces dsc detail
Physical interface: dsc, Enabled, Physical link is Up
  Interface index: 5, SNMP ifIndex: 5, Generation: 9
  Type: Software-Pseudo, Link-level type: Unspecified, MTU: Unlimited, Clocking:

```

```

Unspecified, Speed: Unspecified
Device flags   : Present Running
Interface flags: Point-To-Point SNMP-Traps
Link type      : Unspecified
Link flags     : None
Physical info  : Unspecified
Hold-times    : Up 0 ms, Down 0 ms
Current address: Unspecified, Hardware address: Unspecified
Alternate link address: Unspecified
Last flapped   : Never
Statistics last cleared: Never
Traffic statistics:
  Input bytes   : 0
  Output bytes  : 0
  Input packets: 0
  Output packets: 0

Logical interface dsc.0 (Index 66) (SNMP ifIndex 235) (Generation 6)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Unspecified
  Protocol inet, MTU: Unlimited, Generation: 14, Route table: 0
  Flags: None

```

show interfaces dsc extensive

```

user@host> show interfaces dsc extensive
Physical interface: dsc, Enabled, Physical link is Up
  Interface index: 5, SNMP ifIndex: 5, Generation: 9
  Type: Software-Pseudo, Link-level type: Unspecified, MTU: Unlimited, Clocking:
Unspecified, Speed: Unspecified
Device flags   : Present Running
Interface flags: Point-To-Point SNMP-Traps
Link type      : Unspecified
Link flags     : None
Physical info  : Unspecified
Hold-times    : Up 0 ms, Down 0 ms
Current address: Unspecified, Hardware address: Unspecified
Alternate link address: Unspecified
Last flapped   : Never
Statistics last cleared: Never
Traffic statistics:
  Input bytes   : 0
  Output bytes  : 0
  Input packets: 0
  Output packets: 0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
  Policed discards: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
  Resource errors: 0
Logical interface dsc.0 (Index 66) (SNMP ifIndex 235) (Generation 6)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Unspecified
  Protocol inet, MTU: Unlimited, Generation: 14, Route table: 0

```

show interfaces (Gigabit Ethernet)

Syntax	<code>show interfaces <i>ge-fpc/pic/port</i></code> <code><brief detail extensive terse></code> <code><descriptions></code> <code><media></code> <code><snmp-index <i>snmp-index</i>></code> <code><statistics></code>
Release Information	Command introduced before Junos OS Release 7.4.
Description	(M Series, T Series, and MX Series routers and EX Series switches only) Display status information about the specified Gigabit Ethernet interface.
Options	<p><code><i>ge-fpc/pic/port</i></code>—Display standard information about the specified Gigabit Ethernet interface.</p> <p><code>brief detail extensive terse</code>—(Optional) Display the specified level of output.</p> <p><code>descriptions</code>—(Optional) Display interface description strings.</p> <p><code>media</code>—(Optional) Display media-specific information about network interfaces.</p> <p><code>snmp-index <i>snmp-index</i></code>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><code>statistics</code>—(Optional) Display static interface statistics.</p>
Additional Information	In a logical system, this command displays information only about the logical interfaces and not about the physical interfaces.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• <i>Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration</i>
List of Sample Output	<p>show interfaces (Gigabit Ethernet) on page 2429</p> <p>show interfaces (Gigabit Ethernet on MX Series Routers) on page 2429</p> <p>show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration) on page 2430</p> <p>show interfaces brief (Gigabit Ethernet) on page 2430</p> <p>show interfaces detail (Gigabit Ethernet) on page 2431</p> <p>show interfaces extensive (Gigabit Ethernet IQ2) on page 2432</p> <p>show interfaces (Gigabit Ethernet Unnumbered Interface) on page 2435</p> <p>show interfaces (ACI Interface Set Configured) on page 2435</p>
Output Fields	Table 137 on page 1982 describes the output fields for the show interfaces (Gigabit Ethernet) command. Output fields are listed in the approximate order in which they appear. For Gigabit Ethernet IQ and IQE PICs, the traffic and MAC statistics vary by interface type. For more information, see Table 138 on page 1995 .

Table 171: show interfaces Gigabit Ethernet Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 2376 .	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
WAN-PHY mode	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: Enabled or Disabled for parent interface; Rx-only or Tx-only for child interfaces.	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under “Common Output Fields Description” on page 2376 .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 2376 .	All levels

Table 171: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Link flags	Information about the link. Possible values are described in the “Links Flags” section under “ Common Output Fields Description ” on page 2376.	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	detail extensive none
Schedulers	(Gigabit Ethernet intelligent queuing 2 [IQ2] interfaces only) Number of CoS schedulers configured.	extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds (ms).	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None
Output Rate	Output rate in bps and pps.	None
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Egress accounting overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	detail extensive
Ingress accounting overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	detail extensive

Table 171: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Output bytes—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the show interfaces (10-Gigabit Ethernet) command.</p>	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 171: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 171: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
Interface transmit statistics	<p>(On MX Series devices) Status of the interface-transmit-statistics configuration: Enabled or Disabled.</p> <ul style="list-style-type: none"> • Enabled—When the interface-transmit-statistics statement is included in the configuration. If this is configured, the interface statistics show the actual transmitted load on the interface. • Disabled—When the interface-transmit-statistics statement is not included in the configuration. If this is not configured, the interface statistics show the offered load on the interface. 	detail extensive
OTN FEC statistics	<p>The forward error correction (FEC) counters provide the following statistics:</p> <ul style="list-style-type: none"> • Corrected Errors—The count of corrected errors in the last second. • Corrected Error Ratio—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits. 	detail extensive
PCS statistics	<p>(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.</p> <ul style="list-style-type: none"> • Bit errors—High bit error rate. Indicates the number of bit errors when the PCS receiver is operating in normal mode. • Errored blocks—Loss of block lock. The number of errored blocks when the PCS receiver is operating in normal mode. 	detail extensive

Table 171: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the show interfaces (10-Gigabit Ethernet) command. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—There are two possible conditions regarding the number of oversized frames: <ul style="list-style-type: none"> • Packet length exceeds 1518 octets, or • Packet length exceeds MRU • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. <p>NOTE: The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the VLAN tagged frames field displays 0 when the show interfaces command is executed on a 20-port Gigabit Ethernet MIC. In other words, the number of VLAN tagged frames cannot be determined for the 20-port Gigabit Ethernet MIC.</p> <ul style="list-style-type: none"> • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive
OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	extensive
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	extensive

Table 171: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 171: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 171: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 171: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 171: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
CoS information	<p>Information about the CoS queue for the physical interface.</p> <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under " Common Output Fields Description " on page 2376.	All levels

Table 171: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • push-push—Two VLAN tags are pushed in from the incoming frame. • swap-push—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value. • pop-swap—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed. 	brief detail extensive none
Demux	<p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> • Source Family Inet • Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels
ACI VLAN: Dynamic Profile	Name of the dynamic profile that defines the agent circuit identifier (ACI) interface set. If configured, the ACI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information.	brief detail extensive none
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under “ Common Output Fields Description ” on page 2376.	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Dynamic Profile	(MX Series routers with Trio MPCs only) Name of the dynamic profile that was used to create this interface configured with a Point-to-Point Protocol over Ethernet (PPPoE) family.	detail extensive none
Service Name Table	(MX Series routers with Trio MPCs only) Name of the service name table for the interface configured with a PPPoE family.	detail extensive none
Max Sessions	(MX Series routers with Trio MPCs only) Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.	detail extensive none

Table 171: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Duplicate Protection	(MX Series routers with Trio MPCs only) State of PPPoE duplicate protection: On or Off . When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the router.	extensive
Transit statistics	<p>Number and rate of bytes and packets transiting the switch.</p> <p>NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the Output bytes and Output packets interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.</p>	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. Possible values are described in the "Family Flags" section under " Common Output Fields Description " on page 2376.	detail extensive
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	detail extensive none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	detail extensive

Table 171: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output Filters	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	detail extensive
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 2376 .	detail extensive none
protocol-family	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about the address flag. Possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 2376 .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Table 172: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Inbound physical interface	show interfaces ge-0/3/0 extensive	<p>Traffic statistics:</p> <p>Input bytes: 496 bytes per packet, representing the Layer 2 packet</p> <p>MAC statistics:</p> <p>Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes</p>	The additional 4 bytes are for the CRC.
Inbound logical interface	show interfaces ge-0/3/0.50 extensive	<p>Traffic statistics:</p> <p>Input bytes: 478 bytes per packet, representing the Layer 3 packet</p>	

Table 172: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type (*continued*)

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Outbound physical interface	show interfaces ge-0/0/0 extensive	<p>Traffic statistics:</p> <p>Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes</p> <p>MAC statistics:</p> <p>Received octets: 478 bytes per packet, representing the Layer 3 packet</p>	For input bytes, the additional 12 bytes include 6 bytes for the destination MAC address plus 4 bytes for VLAN plus 2 bytes for the Ethernet type.
Outbound logical interface	show interfaces ge-0/0/0.50 extensive	<p>Traffic statistics:</p> <p>Input bytes: 478 bytes per packet, representing the Layer 3 packet</p>	

Sample Output

show interfaces (Gigabit Ethernet)

```

user@host> show interfaces ge-3/0/2
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:05:85:4a:e9:7c, Hardware address: 00:05:85:4a:e9:7c
  Last flapped   : 2006-08-10 17:25:10 PDT (00:01:08 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69)
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
  0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  Egress accounting overhead: 100
  Ingress accounting overhead: 90
  Input packets : 0
  Output packets: 0
  Protocol ccc, MTU: 1522
  Flags: Is-Primary

```

show interfaces (Gigabit Ethernet on MX Series Routers)

```

user@host> show interfaces ge-2/2/2
Physical interface: ge-2/2/2, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 188
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, MAC-REWRITE Error: None,
  Loopback: Disabled,

```

```
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags     : None
CoS queues     : 8 supported, 4 maximum usable queues
Schedulers    : 0
Current address: 00:1f:12:b7:d7:c0, Hardware address: 00:1f:12:b7:d6:76
Last flapped   : 2008-09-05 16:44:30 PDT (3d 01:04 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
Active alarms  : None
Active defects : None
Logical interface ge-2/2/2.0 (Index 82) (SNMP ifIndex 219)
  Flags: SNMP-Traps 0x20000000 Encapsulation: Ethernet-Bridge
  Egress accounting overhead: 100
  Ingress accounting overhead: 90
  Input packets : 0
  Output packets: 0
  Protocol aenet, AE bundle: ae0.0    Link Index: 4
```

show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration)

```
user@host> show interfaces ge-2/1/2 extensive | match "output|interface"
Physical interface: ge-2/1/2, Enabled, Physical link is Up
Interface index: 151, SNMP ifIndex: 530, Generation: 154
Interface flags: SNMP-Traps Internal: 0x4000
Output bytes   :      240614363944      772721536 bps
Output packets:      3538446506      1420444 pps
Direction : Output
Interface transmit statistics: Enabled

Logical interface ge-2/1/2.0 (Index 331) (SNMP ifIndex 955) (Generation 146)
Output bytes   :      195560312716      522726272 bps
Output packets:      4251311146      1420451 pps
```

show interfaces brief (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2 brief
Physical interface: ge-3/0/2, Enabled, Physical link is Up
Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags     : None

Logical interface ge-3/0/2.0
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  ccc

Logical interface ge-3/0/2.32767
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2
```

show interfaces detail (Gigabit Ethernet)

```

user@host> show interfaces ge-3/0/2 detail
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35, Generation: 177
  Link-level type: 52, MTU: 1522, Speed: 1000Mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:4a:e9:7c, Hardware address: 00:05:85:4a:e9:7c
  Last flapped   : 2006-08-09 17:17:00 PDT (01:31:33 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes  :                0                0 bps
    Output bytes :                0                0 bps
    Input packets:                0                0 pps
    Output packets:              0                0 pps
  Ingress traffic statistics at Packet Forwarding Engine:
    Input bytes  :                0                0 bps
    Input packets:                0                0 pps
    Drop bytes   :                0                0 bps
    Drop packets:                0                0 pps
  Ingress queues: 4 supported, 4 in use
  Queue counters:


|                | Queued packets | Transmitted packets | Dropped packets |
|----------------|----------------|---------------------|-----------------|
| 0 best-effort  | 0              | 0                   | 0               |
| 1 expedited-fo | 0              | 0                   | 0               |
| 2 assured-forw | 0              | 0                   | 0               |
| 3 network-cont | 0              | 0                   | 0               |


  Egress queues: 4 supported, 4 in use
  Queue counters:


|                | Queued packets | Transmitted packets | Dropped packets |
|----------------|----------------|---------------------|-----------------|
| 0 best-effort  | 0              | 0                   | 0               |
| 1 expedited-fo | 0              | 0                   | 0               |
| 2 assured-forw | 0              | 0                   | 0               |
| 3 network-cont | 0              | 0                   | 0               |


  Active alarms : None
  Active defects : None

  Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69) (Generation 140)
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530)
  Out(swap-push 0x8100.512 0x8100.513)
    Encapsulation: VLAN-CCC
    Egress accounting overhead: 100
    Ingress accounting overhead: 90
    Traffic statistics:
      Input bytes  :                0
      Output bytes :                0

```

```

Input packets:          0
Output packets:         0
Local statistics:
Input bytes :           0
Output bytes :           0
Input packets:          0
Output packets:         0
Transit statistics:
Input bytes :           0          0 bps
Output bytes :           0          0 bps
Input packets:          0          0 pps
Output packets:         0          0 pps
Protocol ccc, MTU: 1522, Generation: 149, Route table: 0
Flags: Is-Primary

```

```

Logical interface ge-3/0/2.32767 (Index 71) (SNMP ifIndex 70)
(Generation 139)
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2
Traffic statistics:
Input bytes :           0
Output bytes :           0
Input packets:          0
Output packets:         0
Local statistics:
Input bytes :           0
Output bytes :           0
Input packets:          0
Output packets:         0
Transit statistics:
Input bytes :           0          0 bps
Output bytes :           0          0 bps
Input packets:          0          0 pps
Output packets:         0          0 pps

```

show interfaces extensive (Gigabit Ethernet IQ2)

```

user@host> show interfaces ge-7/1/3 extensive
Physical interface: ge-7/1/3, Enabled, Physical link is Up
Interface index: 170, SNMP ifIndex: 70, Generation: 171
Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4004000
Link flags : None
CoS queues : 8 supported, 4 maximum usable queues
Schedulers : 256
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:14:f6:30:5e:74, Hardware address: 00:14:f6:30:5e:74
Last flapped : 2007-11-07 21:31:41 PST (02:03:33 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes :          38910844056          7952 bps
Output bytes :           7174605          8464 bps
Input packets:        418398473          11 pps
Output packets:         78903          12 pps
IPv6 transit statistics:
Input bytes :           0
Output bytes :           0
Input packets:          0
Output packets:         0

```

Ingress traffic statistics at Packet Forwarding Engine:

```

Input bytes :          38910799145          7952 bps
Input packets:         418397956           11 pps
Drop bytes :              0              0 bps
Drop packets:           0              0 pps

```

Input errors:

```

Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0

```

Output errors:

```

Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

```

```

FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0

```

Ingress queues: 4 supported, 4 in use

Queue counters:	Queued packets	Transmitted packets	Dropped packets
0 best-effort	418390823	418390823	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	7133	7133	0

Egress queues: 4 supported, 4 in use

Queue counters:	Queued packets	Transmitted packets	Dropped packets
0 best-effort	1031	1031	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	77872	77872	0

```

Active alarms : None

```

```

Active defects : None

```

MAC statistics:

	Receive	Transmit
Total octets	38910844056	7174605
Total packets	418398473	78903
Unicast packets	408021893366	1026
Broadcast packets	10	12
Multicast packets	418398217	77865
CRC/Align errors	0	0
FIFO errors	0	0
MAC control frames	0	0
MAC pause frames	0	0
Oversized frames	0	
Jabber frames	0	
Fragment frames	0	
VLAN tagged frames	0	
Code violations	0	OTN Received Overhead Bytes:
APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58		
Payload Type: 0x08		

OTN Transmitted Overhead Bytes:

```

APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
Payload Type: 0x08

```

Filter statistics:

Input packet count	418398473
Input packet rejects	479
Input DA rejects	479

```

Input SA rejects                                0
Output packet count                             78903
Output packet pad count                         0
Output packet error count                       0
CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link partner:
Link mode: Full-duplex, Flow control: Symmetric/Asymmetric,
Remote fault: OK
Local resolution:
Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
Destination slot: 7
CoS information:
Direction : Output
CoS transmit queue      Bandwidth      Buffer      Priority      Limit
                        %      bps      %      usec
0 best-effort           95      950000000  95          0
low none
3 network-control       5      500000000   5          0
low none
Direction : Input
CoS transmit queue      Bandwidth      Buffer      Priority      Limit
                        %      bps      %      usec
0 best-effort           95      950000000  95          0
low none
3 network-control       5      500000000   5          0
low none

Logical interface ge-7/1/3.0 (Index 70) (SNMP ifIndex 85) (Generation 150)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
Input bytes :      812400
Output bytes :    1349206
Input packets:      9429
Output packets:    9449
IPv6 transit statistics:
Input bytes :      0
Output bytes :      0
Input packets:      0
Output packets:     0
Local statistics:
Input bytes :      812400
Output bytes :    1349206
Input packets:      9429
Output packets:    9449
Transit statistics:
Input bytes :      0      7440 bps
Output bytes :      0      7888 bps
Input packets:      0      10 pps
Output packets:      0      11 pps
IPv6 transit statistics:
Input bytes :      0
Output bytes :      0
Input packets:      0
Output packets:     0
Protocol inet, MTU: 1500, Generation: 169, Route table: 0
Flags: Is-Primary, Mac-Validate-Strict
Mac-Validate Failures: Packets: 0, Bytes: 0
Addresses, Flags: Is-Preferred Is-Primary

```



```

Input Filters: F1-ge-3/0/1.0-in, F3-ge-3/0/1.0-in
Output Filters: F2-ge-3/0/1.0-out (53)
Destination: 10.74.2/24, Local: 10.74.2.2, Broadcast: 10.74.2.255,
Generation: 196
Protocol multiservice, MTU: Unlimited, Generation: 170, Route table: 0
Flags: Is-Primary
Policer: Input: __default_arp_policer__

```

NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics displayed in the **show interfaces** command output might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the interface counters. For detailed information, see the description of the logical interface **Transit statistics** fields in [Table 137 on page 1982](#).

show interfaces (Gigabit Ethernet Unnumbered Interface)

```

user@host> show interfaces ge-3/2/0
Physical interface: ge-3/2/0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 50
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Current address: 00:14:f6:11:26:f8, Hardware address: 00:14:f6:11:26:f8
  Last flapped   : 2006-10-27 04:42:23 PDT (08:01:52 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 624 bps (1 pps)
  Active alarms  : None
  Active defects : None

Logical interface ge-3/2/0.0 (Index 67) (SNMP ifIndex 85)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 0
  Output packets: 6
  Protocol inet, MTU: 1500
  Flags: Unnumbered
  Donor interface: lo0.0 (Index 64)
  Preferred source address: 22.22.22.22

```

show interfaces (ACI Interface Set Configured)

```

user@host> show interfaces ge-1/0/0.4001
Logical interface ge-1/0/0.4001 (Index 340) (SNMP ifIndex 548)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.4001 ] Encapsulation: PPP-over-

Ethernet
ACI VLAN:
  Dynamic Profile: aci-vlan-set-profile
  PPPoE:
    Dynamic Profile: aci-vlan-pppoe-profile,
    Service Name Table: None,
    Max Sessions: 32000, Max Sessions VSA Ignore: Off,
    Duplicate Protection: On, Short Cycle Protection: Off,
    AC Name: nbc
  Input packets : 9

```

Output packets: 8
Protocol multiservice, MTU: Unlimited

show interfaces (Serial)

Syntax	show interfaces <i>interface-type</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i> > <statistics>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display status information about serial interfaces, including RS-232, RS-422/449, EIA-530, X.21, and V.35.
Options	<p><i>interface-type</i>—On M Series and T Series routers, the interface type is <i>se-fpc/pic/port</i>. On the J Series routers, the interface type is <i>se-pim/O/port</i>.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
List of Sample Output	show interfaces (Serial, EIA-530) on page 2443 show interfaces brief (Serial, EIA-530) on page 2443 show interfaces detail (Serial, EIA-530) on page 2444 show interfaces extensive (Serial, EIA-530) on page 2444 show interfaces (Serial, V.35) on page 2445 show interfaces brief (Serial, V.35) on page 2446 show interfaces detail (Serial, V.35) on page 2446 show interfaces extensive (Serial, V.35) on page 2447 show interfaces statistics detail (RS 449) on page 2448
Output Fields	Table 173 on page 2437 lists the output fields for the show interfaces (Serial) command. Output fields are listed in the approximate order in which they appear.

Table 173: show interfaces (Serial) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels

Table 173: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 2376 .	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Type	Type of interface.	All levels
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit (MTU) size on the physical interface.	All levels
Maximum speed	Maximum speed. The nonconfigurable value is 16,384 kbps.	detail extensive none
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under “Common Output Fields Description” on page 2376 .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 2376 .	All levels
Link flags	Information about the link. Possible values are described in the “Link Flags” section under “Common Output Fields Description” on page 2376 .	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Keepalive settings	(PPP and HDLC) Configured settings for keepalive packets. <ul style="list-style-type: none"> Interval <i>seconds</i>—Time between successive keepalive requests. The range of values, in seconds, is 10 to 32,767. The default value is 10. Up-count <i>number</i>—Number of keepalive packets a destination must receive to change a link's status from down to up. The range of values is 1 to 255. The default value is 1. Down-count <i>number</i>—Number of keepalive packets a destination must fail to receive before the network takes a link down. The range is 1 to 255. The default value is 3. 	All levels
Keepalive	(PPP and HDLC) Information about keepalive packets. <ul style="list-style-type: none"> Input: <i>number (hh:mm:ss ago)</i>—Number of keepalive packets received by PPP and the time since the last keepalive packet was received. Output: <i>number (hh:mm:ss ago)</i>—Number of keepalive packets sent by PPP and the time since the last keepalive packet was sent. 	brief none

Table 173: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Keepalive statistics	(PPP and HDLC) Information about keepalive packets. <ul style="list-style-type: none"> • Input: <i>number (last seen hh:mm:ss ago)</i>—Number of keepalive packets received by PPP and the time since the last keepalive packet was received. • Output: <i>number (last seen hh:mm:ss ago)</i>—Number of keepalive packets sent by PPP and the time since the last keepalive packet was sent. 	detail extensive
LCP state	(PPP) Link Control Protocol state. <ul style="list-style-type: none"> • Conf-ack-received—Acknowledgement was received. • Conf-ack-sent—Acknowledgement was sent. • Conf-req-sent—Request was sent. • Down—LCP negotiation is incomplete (not yet completed or has failed). • Not-configured—LCP is not configured on the interface. • Opened—LCP negotiation is successful. 	detail extensive none
NCP state	(PPP) Network Control Protocol state. <ul style="list-style-type: none"> • Conf-ack-received—Acknowledgement was received. • Conf-ack-sent—Acknowledgement was sent. • Conf-req-sent—Request was sent. • Down—NCP negotiation is incomplete (not yet completed or has failed). • Not-configured—NCP is not configured on the interface. • Opened—NCP negotiation is successful. 	detail extensive none
CHAP state	(PPP) Displays the state of the Challenge Handshake Authentication Protocol (CHAP) during its transaction. <ul style="list-style-type: none"> • Chap-Chal-received—Challenge was received but response not yet sent. • Chap-Chal-sent—Challenge was sent. • Chap-Resp-received—Response was received for the challenge sent, but CHAP has not yet moved into the Success state. (Most likely with RADIUS authentication.) • Chap-Resp-sent—Response was sent for the challenge received. • Closed—CHAP authentication is incomplete. • Failure—CHAP authentication failed. • Not-configured—CHAP is not configured on the interface. • Success—CHAP authentication was successful. 	detail extensive none
CoS queues	Number of CoS queues configured.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified

Table 173: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Giants—Number of frames received that are larger than the giant threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • Resource errors—Sum of transmit drops. 	extensive
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • MTU errors—Number of packets whose size exceeds the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Egress queues supported	Total number of egress queues supported on the specified interface. Displayed with the statistics option.	detail extensive
Egress queues in use	Total number of egress queues in use on the specified interface. Displayed with the statistics option.	detail extensive

Table 173: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Queue counters	CoS queue number and its associated user-configured forwarding class name. Displayed with the statistics option. <ul style="list-style-type: none"> Queued packets—Number of queued packets. Transmitted packets—Number of transmitted packets. Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive
Serial media information	Information about the physical media: <ul style="list-style-type: none"> Line protocol—eia530, eia530a, rs232, rs449, v.35, or x.21. Resync history—Information about resynchronization events: <ul style="list-style-type: none"> Sync loss count—Number of times the synchronization was lost. Data signal—(X.21 and V.35) Information about the data signal: <ul style="list-style-type: none"> Rx Clock—Receive clock status: OK (DTE is receiving the receive clock signal) or Not detected (receive clock signal is not being received). Control signals—Information about modem control signals: <ul style="list-style-type: none"> Local mode:DCE (data communication equipment) or DTE (data terminal equipment) To DCE—Control signals that the Serial PIC sent to the DCE: DTR (Data Terminal Ready:up or down) or RTS (Request To Send: up or down.) From DC—Control signals that the Serial PIC received from the DCE: CTS (Clear To Send: up or down), DCD (Data Carrier Detect: up or down), DSR (Data Set Ready: up or down), or TM (Test Mode: up or down). Clocking mode—Clocking used for the transmit clock: <ul style="list-style-type: none"> dte—Transmit clock is generated by DTE. dce—Transmit clock is generated by the DCE and is looped back as the transmit clock. loop-timed—Receive clock from the DCE is looped back as the transmit clock. Clock rate—Rate, in megahertz (MHz), at which the clock is configured. Loopback—Configured loopback mode for the interface: dce-remote, dce-local, liu, local, or none. Tx clock—Clocking phase of the transmit clock: invert (transmit clock polarity is inverted) or non-invert (transmit clock polarity is not inverted). Line encoding—Type of line encoding used: nrz (nonreturn to zero) or nrzi (return to zero inverted). 	detail extensive
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> Destination slot—FPC slot number. PLP byte—Packet Level Protocol byte. 	extensive

Table 173: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface: <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	Logical interface SNMP interface index number.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under “Common Output Fields Description” on page 2376 .	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the source and destination address are also displayed.	brief
Protocol	Protocol family configured on the logical interface, such as iso , inet6 , mpls .	detail extensive none
MTU	MTU size on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0 .	detail extensive
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under “Common Output Fields Description” on page 2376 .	detail extensive

Table 173: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 2376 .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

show interfaces (Serial, EIA-530)

```

user@host> show interfaces se-5/0/1
Physical interface: se-5/0/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 41
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags   : Present Running
  Interface flags: Point-To-Point Internal: 0x4000
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 32 (00:00:10 ago), Output: 31 (00:00:07 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Closed
  CoS queues    : 8 supported, 8 maximum usable queues
  Last flapped  : 2006-04-26 15:10:18 PDT (00:05:22 ago)
  Input rate    : 0 bps (0 pps)
  Output rate   : 0 bps (0 pps)

Logical interface se-5/0/1.0 (Index 71) (SNMP ifIndex 45)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
  Protocol inet, MTU: 1500
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 12.0.0.0/30, Local: 12.0.0.1, Broadcast: 12.0.0.3

```

show interfaces brief (Serial, EIA-530)

```

user@host> show interfaces se-5/0/1 brief
Physical interface: se-5/0/1, Enabled, Physical link is Up
  Type: Serial, Link-level type: PPP, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point Internal: 0x4000
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 235 (00:00:10 ago), Output: 234 (00:00:00 ago)

Logical interface se-5/0/1.0

```

```
Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
inet 12.0.0.1/30
```

show interfaces detail (Serial, EIA-530)

```
user@host> show interfaces se-5/0/1 detail
Physical interface: se-5/0/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 41, Generation: 25
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags   : Present Running
  Interface flags: Point-To-Point Internal: 0x4000
  Link flags     : Keepalives
  Hold-times     : Up 0 ms, Down 0 ms
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive statistics:
    Input : 37 (last seen 00:00:06 ago)
    Output: 35 (last sent 00:00:01 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Closed
  CoS queues   : 8 supported, 8 maximum usable queues
  Last flapped : 2006-04-26 15:10:18 PDT (00:06:02 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes :          928          40 bps
    Output bytes:         1023          48 bps
    Input packets:           76           0 pps
    Output packets:          77           0 pps
  Serial media information:
    Line protocol: eia530
    Resync history:
      Sync loss count: 0
    Data signal:
      Rx Clock: OK
    Control signals:
      Local mode: DTE
      To DCE: DTR: up, RTS: up
      From DCE: CTS: up, DCD: up, DSR: up
    Clocking mode: loop-timed
    Clock rate: 8.0 MHz
    Loopback: none
    Tx clock: non-invert
    Line encoding: nrz

  Logical interface se-5/0/1.0 (Index 71) (SNMP ifIndex 45) (Generation 9)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
    Protocol inet, MTU: 1500, Generation: 15, Route table: 0
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 12.0.0.0/30, Local: 12.0.0.1, Broadcast: 12.0.0.3,
      Generation: 23
```

show interfaces extensive (Serial, EIA-530)

```
user@host> show interfaces se-5/0/1 extensive
Physical interface: se-5/0/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 41, Generation: 25
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags   : Present Running
  Interface flags: Point-To-Point Internal: 0x4000
```

```

Link flags      : Keepalives
Hold-times     : Up 0 ms, Down 0 ms
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive statistics:
  Input : 40 (last seen 00:00:00 ago)
  Output: 37 (last sent 00:00:09 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Closed
CoS queues      : 8 supported, 8 maximum usable queues
Last flapped    : 2006-04-26 15:10:18 PDT (00:06:28 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes :          988          40 bps
  Output bytes :         1088          48 bps
  Input packets:           81           0 pps
  Output packets:          82           0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 2, Runts: 0, Giants: 0,
  Policed discards: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, MTU errors: 0,
  Resource errors: 0
Serial media information:
  Line protocol: eia530
  Resync history:
    Sync loss count: 0
  Data signal:
    Rx Clock: OK
  Control signals:
    Local mode: DTE
    To DCE: DTR: up, RTS: up
    From DCE: CTS: up, DCD: up, DSR: up
  Clocking mode: loop-timed
  Clock rate: 8.0 MHz
  Loopback: none
  Tx clock: non-invert
  Line encoding: nrz
Packet Forwarding Engine configuration:
  Destination slot: 5, PLP byte: 1 (0x00)
CoS information:
  CoS transmit queue      Bandwidth      Buffer      Priority      Limit
                           %          bps          %          usec
  0 best-effort           95        15564800    95           0          low      none
  3 network-control        5          819200      5           0          low      none

Logical interface se-5/0/1.0 (Index 71) (SNMP ifIndex 45) (Generation 9)
Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
Protocol inet, MTU: 1500, Generation: 15, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 12.0.0.0/30, Local: 12.0.0.1, Broadcast: 12.0.0.3,
  Generation: 23

```

show interfaces (Serial, V.35)

```

user@host> show interfaces se-5/0/0
Physical interface: se-5/0/0, Enabled, Physical link is Down
Interface index: 150, SNMP ifIndex: 39
Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps

```

```
Device flags      : Present Running Down
Interface flags: Hardware-Down Point-To-Point Internal: 0x4000
Link flags       : Loose-NCP
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive: Input: 0 (never), Output: 0 (never)
LCP state: Down
NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
mpls: Not-configured
CHAP state: Closed
CoS queues       : 8 supported, 8 maximum usable queues
Last flapped    : 2006-04-26 14:51:27 PDT (01:02:23 ago)
Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)

Logical interface se-5/0/0.0 (Index 73) (SNMP ifIndex 27)
  Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps
  Encapsulation: PPP
  Protocol inet, MTU: 1500
  Flags: Protocol-Down
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 13.0.0.0/30, Local: 13.0.0.2, Broadcast: 13.0.0.3
```

show interfaces brief (Serial, V.35)

```
user@host> show interfaces se-5/0/0 brief
Physical interface: se-5/0/0, Enabled, Physical link is Down
  Type: Serial, Link-level type: PPP, MTU: 1504
  Device flags      : Present Running Down
  Interface flags: Hardware-Down Point-To-Point Internal: 0x4000
  Link flags       : Loose-NCP
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 0 (never), Output: 0 (never)

Logical interface se-5/0/0.0
  Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps
  Encapsulation: PPP
  inet 13.0.0.2/30
```

show interfaces detail (Serial, V.35)

```
user@host> show interfaces se-5/0/0 detail
Physical interface: se-5/0/0, Enabled, Physical link is Down
  Interface index: 150, SNMP ifIndex: 39, Generation: 31
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags      : Present Running Down
  Interface flags: Hardware-Down Point-To-Point Internal: 0x4000
  Link flags       : Loose-NCP
  Hold-times       : Up 0 ms, Down 0 ms
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive statistics:
    Input : 0 (last seen: never)
    Output: 0 (last sent: never)
  LCP state: Down
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
mpls: Not-configured
  CHAP state: Closed
  CoS queues       : 8 supported, 8 maximum usable queues
  Last flapped    : 2006-04-26 14:51:27 PDT (01:03:15 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes :                                0                0 bps
```

```

Output bytes :                0                0 bps
Input packets:                0                0 pps
Output packets:               0                0 pps
Serial media information:
  Line protocol: v.35
  Resync history:
    Sync loss count: 0
  Data signal:
    Rx Clock: Not Detected
  Control signals:
    Local mode: DCE
    To DTE: CTS: down, DCD: down, DSR: up
    From DTE: DTR: down, RTS: down
  DCE loopback override: Off
  Clocking mode: internal
  Clock rate: 38.4 KHz
  Loopback: none
  Tx clock: non-invert
  Line encoding: nrz

Logical interface se-5/0/0.0 (Index 73) (SNMP ifIndex 27) (Generation 12)
  Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps
  Encapsulation: PPP
  Protocol inet, MTU: 1500, Generation: 17, Route table: 0
  Flags: Protocol-Down
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 13.0.0.0/30, Local: 13.0.0.2, Broadcast: 13.0.0.3,
    Generation: 23

```

show interfaces extensive (Serial, V.35)

```

user@host> show interfaces se-5/0/0 extensive
Physical interface: se-5/0/0, Enabled, Physical link is Down
  Interface index: 150, SNMP ifIndex: 39, Generation: 31
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point Internal: 0x4000
  Link flags     : Loose-NCP
  Hold-times     : Up 0 ms, Down 0 ms
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive statistics:
    Input : 0 (last seen: never)
    Output: 0 (last sent: never)
  LCP state: Down
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Closed
  CoS queues   : 8 supported, 8 maximum usable queues
  Last flapped : 2006-04-26 14:51:27 PDT (01:04:17 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes :                0                0 bps
    Output bytes :                0                0 bps
    Input packets:                0                0 pps
    Output packets:               0                0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

```

```

Serial media information:
  Line protocol: v.35
  Resync history:
    Sync loss count: 0
  Data signal:
    Rx Clock: Not Detected
  Control signals:
    Local mode: DCE
    To DTE: CTS: down, DCD: down, DSR: up
    From DTE: DTR: down, RTS: down
  DCE loopback override: Off
  Clocking mode: internal
  Clock rate: 38.4 KHz
  Loopback: none
  Tx clock: non-invert
  Line encoding: nrz
Packet Forwarding Engine configuration:
  Destination slot: 5, PLP byte: 1 (0x00)
CoS information:
  CoS transmit queue      Bandwidth      Buffer      Priority  Limit
                           %      bps      %      usec
  0 best-effort           95      15564800  95        0      low  none
  3 network-control        5       819200   5         0      low  none

Logical interface se-5/0/0.0 (Index 73) (SNMP ifIndex 27) (Generation 12)
Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps
Encapsulation: PPP
Protocol inet, MTU: 1500, Generation: 17, Route table: 0
Flags: Protocol-Down
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
  Destination: 13.0.0.0/30, Local: 13.0.0.2, Broadcast: 13.0.0.3,
  Generation: 23

```

show interfaces statistics detail (RS 449)

```

user@host> show interfaces se-6/0/0 statistics detail
Interface index: 149, SNMP ifIndex: 59, Generation: 150
Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 8mbps
Device flags : Present Running
Interface flags: Point-To-Point Internal: 0x4000
Link flags : No-Keepalives Loose-NCP
Hold-times : Up 0 ms, Down 0 ms
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Closed
PAP state: Closed
CoS queues : 8 supported, 8 maximum usable queues
Last flapped : 2007-11-28 19:38:36 PST (00:14:06 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes : 744 0 bps
  Output bytes : 5978 0 bps
  Input packets: 33 0 pps
  Output packets: 129 0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
0,
  Resource errors: 0
Output errors:
  Carrier transitions: 13, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:

```

0

Egress queues: 8 supported, 5 in use

Queue counters:	Queued packets	Transmitted packets	Dropped packets
0 best-effort	24	24	0
1 expedited-fo	0	0	0
2 bulk	0	0	0
3 assured-forw	105	105	0
4 voip	0	0	0

Serial media information:

Line protocol: rs449

Resync history:

Sync loss count: 0

Data signal:

Rx Clock: OK

Control signals:

Local mode: DTE

To DCE: DTR: up, RTS: up

From DCE: CTS: up, DCD: up, DSR: up

Clocking mode: internal

Loopback: none

Tx clock: non-invert

Line encoding: nrz

Logical interface se-6/0/0.0 (Index 75) (SNMP ifIndex 69) (Generation 141)

Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP

Protocol inet, MTU: 256, Generation: 145, Route table: 0

Flags: None

Addresses, Flags: Is-Preferred Is-Primary

Destination: 11.11.11/24, Local: 11.11.11.2, Broadcast: 11.11.11.255,

Generation: 157

show interfaces extensive

Syntax show interfaces extensive

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches.

Description Display extensive information about all interfaces configured on the router.



NOTE:

- At some times, the cumulative byte counters displayed with the `show interfaces extensive` command on the 10-Gigabit Ethernet MPC with SFP+ is not always increasing and cumulative and does not give the correct results. There is a time lag in collecting these statistics, during which the display might decrease or go from a nonzero number to zero. Eventually, the counter will display the correct result.
 - When the `show interfaces extensive` command is executed on a router with an MPC or a T4000 Type 5 FPC, the *Input packet rejects* counter of the *Filter statistics* field also displays statistics related to the following packet errors:
 - Invalid VLAN range
 - Tagged packet received on an untagged interface
 - When the `show interfaces extensive` command is executed on an interface that is configured on a T4000 Type 5 FPC, the *IPv6 transit statistics* field displays:
 - Total statistics (sum of transit and local statistics) at the physical interface level
 - Transit statistics at the logical interface level
 - When the `show interfaces extensive` command is executed on an aggregate interface in a T1600 Core Router, the *IPv6 Input bytes* is displayed for an aggregate interface. However, the *IPv6 Input bytes* is always zero on a member link of an aggregated bundle even when there are IPv6 transit traffic on the member link. This is because the logical interface index of the aggregate logical interface is updated but not the logical interface of the member links in the channel lookup table.
-

Options This command has no options.

Required Privilege Level view

List of Sample Output [show interfaces extensive \(Circuit Emulation\) on page 2451](#)

[show interfaces extensive \(Fast Ethernet\) on page 2451](#)
[show interfaces extensive \(Gigabit Ethernet\) on page 2453](#)
[show interfaces extensive \(10-Gigabit Ethernet\) on page 2454](#)
[show interfaces extensive \(IQ2 and IQ2E\) on page 2456](#)
[show interfaces extensive \(100-Gigabit Ethernet\) on page 2459](#)
[show interfaces extensive \(PTX5000 Packet Transport Switch\) on page 2461](#)
[show interfaces extensive \(T4000 Routers with Type 5 FPCs\) on page 2463](#)
[show interfaces extensive \(T4000 Routers with 24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC\) on page 2465](#)
[show interfaces extensive \(Aggregated Ethernet\) on page 2466](#)

Output Fields For more information, see the output fields table for the particular interface type in which you are interested. For information about destination class and source class statistics, see the “Destination Class Field” section and the “Source Class Field” section under [“Common Output Fields Description” on page 2376](#). For sample output for specific interfaces, see the other topics in this collection.

Sample Output

show interfaces extensive (Circuit Emulation)

If a Circuit Emulation (CE) PIC is configured for SAToP pseudowire, then pseudowire statistics are displayed in the CE information section of the show interface extensive output. If SAToP pseudowire is not configured on the CE PIC, then all the CE information counters will be displayed as 0 (zero).

```

user@host> show interface t1-0/0/0 extensive
Physical interface :t1-0/0/0, Enabled, Physical Link : Up
    Interface index:61441
    Speed : 1.54 Mbps, Loopback: Disabled
    Operational state : Enabled,      Encapsulation : Trans
    Encoding : b8zs,      Framing : unframe,  Build-out : 0-30
    Inversion : enable,  Clock source : master
    Description :
    Traffic statistics:
    T1 media:           Seconds
    ES                  1643
    SES                 1643

    CE Info            Packets           Bytes
    CE Rx              : 2395529         306627712
    CE Tx              : 2396259         306721152
    CE Rx Drop:        0                 0
    CE Tx Drop:        0                 0

    CE Overrun  Events: 0
    CE Underrun Events: 0

```

Sample Output

show interfaces extensive (Fast Ethernet)

```

user@host> show interfaces fe-0/2/1 extensive
Physical interface: fe-0/2/0, Enabled, Physical link is Up
    Interface index: 129, SNMP ifIndex: 23, Generation: 130
    Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
    Source filtering: Disabled, Flow control: Enabled

```

```

Device flags      : Present Running
Interface flags:  SNMP-Traps Internal: 0x4000
CoS queues       : 4 supported, 4 maximum usable queues
Hold-times       : Up 0 ms, Down 0 ms
Current address:  00:90:69:91:c4:3e, Hardware address: 00:90:69:91:c4:3e
Last flapped     : 2006-04-16 23:00:41 PDT (02:08:05 ago)
Statistics last cleared: 2006-04-16 21:42:00 PDT (03:26:46 ago)
Traffic statistics:
Input bytes      :          17539          152 bps
Output bytes     :          92968          224 bps
Input packets    :           348           0 pps
Output packets   :          1349           0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort          66              66              0
  1 expedited-fo         0              0              0
  2 assured-forw         0              0              0
  3 network-cont       1283             1283             0

Active alarms : None
Active defects : None
MAC statistics:
Total octets          24721          105982
Total packets         348           1349
Unicast packets       347           430
Broadcast packets      1            37
Multicast packets      0           882
CRC/Align errors      0            0
FIFO errors            0            0
MAC control frames     0            0
MAC pause frames       0            0
Oversized frames       0
Jabber frames          0
Fragment frames        0
VLAN tagged frames     0
Code violations         0
Filter statistics:
Input packet count    348
Input packet rejects   0
Input DA rejects      0
Input SA rejects      0
Output packet count    1349
Output packet pad count 0
Output packet error count 0
CAM destination filters: 3, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link partner:
  Link mode: Full-duplex, Flow control: None, Remote fault: OK

```

Packet Forwarding Engine configuration:

Destination slot: 0

CoS information:

CoS transmit queue		Bandwidth	Buffer	Priority	Limit
	%	bps	% usec		
0 best-effort	95	95000000	95 0	low	none
3 network-control	5	5000000	5 0	low	none

Logical interface fe-0/2/0.0 (Index 66) (SNMP ifIndex 46) (Generation 133)
 Flags: SNMP-Traps Encapsulation: ENET2
 Protocol inet, MTU: 1500, Generation: 142, Route table: 0
 Flags: DCU, SCU-out

Destination class	Packets (packet-per-second)	Bytes (bits-per-second)
silv1_new	0	0
(0)	0)
silv2_new	0	0
(0)	0)
silv_misc	0	0
(0)	0)
silver0	0	0
(0)	0)
silver2	0	0
(0)	0)
silver3	0	0
(0)	0)
silver4	0	0
(0)	0)
silver5	0	0
(0)	0)
silver6	0	0
(0)	0)
silver7	0	0
(0)	0)
silver9	0	0
(0)	0)

Source class	Packets (packet-per-second)	Bytes (bits-per-second)
gold1	0	0
(0)	0)
gold2	16600	1062400
(0)	0)
gold3	0	0
(0)	0)

Addresses, Flags: Is-Preferred Is-Primary

Destination: 12.1.1/24, Local: 12.1.1.1, Broadcast: 12.1.1.255,
Generation: 150

Sample Output

show interfaces extensive (Gigabit Ethernet)

user@host> show interfaces ge-5/0/0.0 extensive

Logical interface ge-5/0/0.0 (Index 71) (SNMP ifIndex 1930) (Generation 139)
 Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
 Traffic statistics:

Input bytes :	0
Output bytes :	42
Input packets:	0
Output packets:	1

```

Local statistics:
  Input bytes :          0
  Output bytes :         42
  Input packets:         0
  Output packets:        1
Transit statistics:
  Input bytes :          0          0 bps
  Output bytes :          0          0 bps
  Input packets:         0          0 pps
  Output packets:        0          0 pps
Output Filters: f-any
Protocol inet, MTU: 1500, Generation: 155, Route table: 0
  Output Filters: f-inet,
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.11.1/24, Local: 10.11.1.1, Broadcast: 10.11.1.255,
    Generation: 170
Protocol multiservice, MTU: Unlimited, Generation: 156, Route table: 0
  Flags: Is-Primary
  Policer: Input: __default_arp_policer__

```

Sample Output

show interfaces extensive (10-Gigabit Ethernet)

```

user@host> show interfaces xe-2/1/0 extensive

Physical interface: xe-2/1/0, Enabled, Physical link is Up
  Interface index: 258, SNMP ifIndex: 762, Generation: 2046
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
  None, Loopback: None, Source filtering: Disabled,
  Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f8:6d:eb, Hardware address: 00:1d:b5:f8:6d:eb
  Last flapped   : 2011-12-17 00:19:02 PST (07:36:37 ago)
  Statistics last cleared: 2011-12-17 07:55:24 PST (00:00:15 ago)
Traffic statistics:
  Input bytes :          110000          0 bps
  Output bytes :           0          0 bps
  Input packets:          1000          0 pps
  Output packets:           0          0 pps
IPv6 transit statistics:
  Input bytes :          110000
  Output bytes :           0
  Input packets:          1000
  Output packets:           0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
  L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
  FIFO errors: 0, HS link CRC errors: 0,
  MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

```

```

0 best-effort          0          0          0
1 expedited-fo        0          0          0
2 assured-forw        0          0          0
3 network-cont        0          0          0

Queue number:         Mapped forwarding classes
0                     best-effort
1                     expedited-forwarding
2                     assured-forwarding
3                     network-control

Active alarms : None
Active defects : None
PCS statistics                Seconds
  Bit errors                0
  Errored blocks            0
MAC statistics:              Receive      Transmit
  Total octets              128000        0
  Total packets             1000         0
  Unicast packets           1000         0
  Broadcast packets         0           0
  Multicast packets         0           0
  CRC/Align errors          0           0
  FIFO errors               0           0
  MAC control frames        0           0
  MAC pause frames          0           0
  Oversized frames          0
  Jabber frames             0
  Fragment frames           0
  VLAN tagged frames        0
  Code violations           0
Filter statistics:
  Input packet count        1000
  Input packet rejects      0
  Input DA rejects          0
  Input SA rejects          0
  Output packet count       0
  Output packet pad count   0
  Output packet error count 0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 2
CoS information:
  Direction : Output
  CoS transmit queue        Bandwidth      Buffer Priority
Limit
                                %      bps      %      usec
0 best-effort                95    9500000000  95      0      low
none
3 network-control            5     500000000   5       0      low
none
Interface transmit statistics: Disabled

Logical interface xe-2/1/0.0 (Index 83) (SNMP ifIndex 1677) (Generation 10082)

Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
Traffic statistics:
  Input bytes :             110000
  Output bytes :              0

```

```

Input packets:          1000
Output packets:         0
IPv6 transit statistics:
  Input bytes :          55000
  Output bytes :         0
  Input packets:         500
  Output packets:        0
Local statistics:
  Input bytes :          55000
  Output bytes :         0
  Input packets:         500
  Output packets:        0
Transit statistics:
  Input bytes :          55000          0 bps
  Output bytes :         0          0 bps
  Input packets:         500          0 pps
  Output packets:        0          0 pps
IPv6 transit statistics:
  Input bytes :          55000
  Output bytes :         0
  Input packets:         500
  Output packets:        0
Protocol inet6, MTU: 1500, Generation: 23739, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 2001:1000:abcd:2312:1432:abcd:1234:0/112, Local:
2001:1000:abcd:2312:1432:abcd:1234:1234
  Generation: 506
    Addresses, Flags: Is-Preferred
      Destination: fe80::/64, Local: fe80::21d:b5ff:fef8:6deb
Protocol multiservice, MTU: Unlimited, Generation: 508
Generation: 23740, Route table: 0
  Policer: Input: __default_arp_policer__

```

Sample Output

show interfaces extensive (IQ2 and IQ2E)

```

user@host> show interfaces ge-3/2/2 extensive
Physical interface: ge-3/2/2, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 548, Generation: 159
  Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues : 8 supported, 8 maximum usable queues
  Schedulers : 128
  Hold-times : Up 0 ms, Down 0 ms
  Current address: 00:14:f6:12:86:fa, Hardware address: 00:14:f6:12:86:fa
  Last flapped : 2010-03-17 04:03:11 PDT (00:45:30 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes :          1716096          0 bps
    Output bytes :          1716448          0 bps
    Input packets:          13407          0 pps
    Output packets:          13411          0 pps
  IPv6 total statistics:
    Input bytes :          1716096
    Output bytes :          1716096
    Input packets:          13407
    Output packets:          13407

```

```

Ingress traffic statistics at Packet Forwarding Engine:
Input bytes :          1716096          0 bps
Input packets:         13407          0 pps
Drop bytes :           0          0 bps
Drop packets:          0          0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 1, L2 mismatch timeouts: 0, FIFO errors:
0,
  Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets:
0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Ingress queues: 8 supported, 4 in use
  Queue counters:      Queued packets  Transmitted packets      Dropped
packets
    0 best-effort          13407          13407
0
    1 expedited-fo           0           0
0
    2 assured-forw          0           0
0
    3 network-cont          0           0
0
  Egress queues: 8 supported, 4 in use
  Queue counters:      Queued packets  Transmitted packets      Dropped
packets
    0 best-effort          13407          13407
0
    1 expedited-fo           0           0
0
    2 assured-forw          0           0
0
    3 network-cont          4           4
0
Active alarms : None
Active defects : None
MAC statistics:
Total octets          1716096          1716448
Total packets          13407          13411
Unicast packets          13407          13407
Broadcast packets           0           0
Multicast packets         0           4
CRC/Align errors         0           0
FIFO errors               0           0
MAC control frames        0           0
MAC pause frames          0           0
Oversized frames          0
Jabber frames             0
Fragment frames           0
VLAN tagged frames        0
Code violations           0
Filter statistics:
Input packet count      13407
Input packet rejects     0
Input DA rejects         0
Input SA rejects         0
Output packet count      13411
Output packet pad count   0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0

```

Autonegotiation information:

Negotiation status: Complete

Link partner:

Link mode: Full-duplex, Flow control: None, Remote fault: OK

Local resolution:

Flow control: Symmetric, Remote fault: Link OK

Packet Forwarding Engine configuration:

Destination slot: 3

CoS information:

Direction : Output

CoS transmit queue	Bandwidth		Buffer Priority	
Limit	%	bps	%	usec
0 best-effort	95	950000000	95	0
none				
3 network-control	5	50000000	5	0
none				
Direction : Input				
CoS transmit queue	Bandwidth		Buffer Priority	
Limit	%	bps	%	usec
0 best-effort	95	950000000	95	0
none				
3 network-control	5	50000000	5	0
none				

Logical interface ge-3/2/2.0 (Index 83) (SNMP ifIndex 6080) (Generation

148)

Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.100] Encapsulation: ENET2

Traffic statistics:

Input bytes : 0
Output bytes : 336
Input packets: 0
Output packets: 4

IPv6 total statistics:

Input bytes : 1716096
Output bytes : 1716096
Input packets: 13407
Output packets: 13407

Local statistics:

Input bytes : 0
Output bytes : 336
Input packets: 0
Output packets: 4

Transit statistics:

Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps

IPv6 total statistics:

Input bytes : 1716096
Output bytes : 1716096
Input packets: 13407
Output packets: 13407

Protocol inet6, MTU: 1500, Generation: 159, Route table: 0

Flags: Is-Primary

Addresses, Flags: Is-Default Is-Primary

Destination: Unspecified, Local: 2000::2

Generation: 146

Addresses, Flags: Is-Preferred

Destination: fe80::/64, Local: fe80::214:f600:6412:86fa


```

Protocol multiservice, MTU: Unlimited, Generation: 148
Generation: 160, Route table: 0
Policer: Input: __default_arp_policer__

Logical interface ge-3/2/2.32767 (Index 84) (SNMP ifIndex 6081) (Generation
149)
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
Protocol multiservice, MTU: Unlimited, Generation: 161, Route table: 0
Flags: None
Policer: Input: __default_arp_policer__

```

Sample Output

show interfaces extensive (100-Gigabit Ethernet)

```

user@host> show interfaces et-0/0/0:0 extensive
Physical interface: et-0/0/0:0, Enabled, Physical link is Down
  Interface index: 156, SNMP ifIndex: 516, Generation: 163
  Link-level type: Ethernet, MTU: 9192, Speed: 50000mbps, BPDU Error: None,
  MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled
  Device flags : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags : None
  CoS queues : 8 supported, 8 maximum usable queues
  Hold-times : Up 0 ms, Down 0 ms
  Current address: 00:aa:aa:aa:aa:00, Hardware address: 00:21:59:5c:48:00
  Last flapped : 2010-01-07 16:36:49 PST (18:02:35 ago)
  Statistics last cleared: Never
Traffic statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0,
  L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors:
0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

```

```

FIFO errors: 0,
  HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 8 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 DEFAULT, NC-                0                0                0
  1 REALTIME                    0                0                0
  2 PRIVATE, NC-                0                0                0
  3 CONTROL                     1253             1253             0
  4 BC-H, CLASS_                0                0                0
  5 BC-M, CLASS_                0                0                0
  6 IA, CLASS_V_                0                0                0
  7 CLASS_S_OUTP                0                0                0

Queue      Mapped Forwarding Class
0          DEFAULT, NC-Q0
1          REALTIME
2          PRIVATE, NC-Q1
3          CONTROL
4          BC-H, CLASS-Q4
5          BC-M, CLASS-Q5
6          IA, CLASS_V_OUTPUT
7          CLASS_S_OUTPUT

Active alarms : None
Active defects : None
MAC statistics:
Total octets      Receive      Transmit
Total packets    0          0
Unicast packets  0          0
Broadcast packets 0          0
Multicast packets 0          0
CRC/Align errors 0          0
FIFO errors       0          0
MAC control frames 0          0
MAC pause frames  0          0
Oversized frames  0
Jabber frames     0
Fragment frames   0
VLAN tagged frames 0
Code violations    0

Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority Limit

                                %      bps      %      usec
0 best-effort             95    47500000000    95      0    low none
3 network-control         5     2500000000     5      0    low none

Logical interface et-0/0/0:0.0 (Index 68) (SNMP ifIndex 546) (Generation 161)
Flags: Deviet-Down SNMP-Traps Encapsulation: ENET2

```

```

Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
Protocol inet, MTU: 9178, Generation: 220, Route table: 0
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 210.160.0/24, Local: 210.160.0.1, Broadcast: 210.160.0.255,
  Generation: 192
Protocol mpls, MTU: 9166, Maximum labels: 3, Generation: 221, Route table: 0

Protocol multiservice, MTU: Unlimited, Generation: 222, Route table: 0
  Policer: Input: __default_arp_policer

```

Sample Output

show interfaces extensive (PTX5000 Packet Transport Switch)

```

user@host> show interfaces et-7/0/0 extensive
Physical interface: et-7/0/0, Enabled, Physical link is Up
  Interface index: 168, SNMP ifIndex: 501, Generation: 171
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None,
  MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled
  Device flags : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags : None
  CoS queues : 8 supported, 8 maximum usable queues
  Hold-times : Up 0 ms, Down 0 ms
  Current address: 88:e0:f3:3b:de:43, Hardware address: 88:e0:f3:3b:de:43
  Last flapped : 2012-01-18 11:48:24 PST (01:47:08 ago)
  Statistics last cleared: Never
Traffic statistics:
  Input bytes : 3583014 0 bps
  Output bytes : 758050 0 bps
  Input packets: 17740 0 pps
  Output packets: 3418 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
  incompletes: 0,
  L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors:
  0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
  FIFO errors: 0,
  HS link CRC errors: 0, MTU errors: 0, Resource errors: 0

```

```

Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort      252                252                0
  1 expedited-fo      0                  0                  0
  2 assured-forw      0                  0                  0
  3 network-cont      6196               6196               0

Queue number:      Mapped forwarding classes
  0                best-effort
  1                expedited-forwarding
  2                assured-forwarding
  3                network-control

Active alarms : None
Active defects : None
MAC statistics:
  Total octets      4108825      Transmit 1159686
  Total packets      21166      6448
  Unicast packets      14824      3255
  Broadcast packets      3          0
  Multicast packets      6339      3193
  CRC/Align errors      0          0
  FIFO errors          0          0
  MAC control frames      0          0
  MAC pause frames      0          0
  Oversized frames      0
  Jabber frames          0
  Fragment frames        0
  VLAN tagged frames      16091
  Code violations        0
Filter statistics:
  Input packet count      9
  Input packet rejects      9
  Input DA rejects        9
  Input SA rejects        0
  Output packet count      0
  Output packet pad count  0
  Output packet error count 0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Incomplete
Packet Forwarding Engine configuration:
  Destination slot: 7
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
  %      bps      %      usec
  0 best-effort      95      9500000000      95      0      low
none
  3 network-control      5      500000000      5      0      low
none
Interface transmit statistics: Disabled

```

Sample Output

show interfaces extensive (T4000 Routers with Type 5 FPCs)

```

user@host> show interfaces xe-4/0/0 extensive
Physical interface: xe-4/0/0, Enabled, Physical link is Up
  Interface index: 170, SNMP ifIndex: 859, Generation: 173
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags      : Present Running
  Interface flags:  SNMP-Traps Internal: 0x4000
  Link flags       : None
  CoS queues       : 8 supported, 8 maximum usable queues
  Hold-times       : Up 0 ms, Down 0 ms
  Current address:  00:12:1e:37:53:f8, Hardware address: 00:12:1e:37:53:f8
  Last flapped     : 2012-06-06 02:25:56 PDT (10:11:58 ago)
  Statistics last cleared: 2012-06-06 12:36:59 PDT (00:00:55 ago)
  Traffic statistics:
    Input bytes  :                0                0 bps
    Output bytes :                0                0 bps
    Input packets:                0                0 pps
    Output packets:              0                0 pps
  IPv6 transit statistics:
    Input bytes  :                0
    Output bytes :                0
    Input packets:                0
    Output packets:              0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
    Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:      Queued packets  Transmitted packets      Dropped packets

    0 best-effort              0                0                0
    1 expedited-fo              0                0                0
    2 assured-forw              0                0                0
    3 network-cont              0                0                0

  Queue number:      Mapped forwarding classes
    0                best-effort
    1                expedited-forwarding
    2                assured-forwarding
    3                network-control
  Active alarms  : None
  Active defects : None
  PCS statistics                Seconds
    Bit errors                0
    Errored blocks            0
  MAC statistics:      Receive      Transmit
    Total octets          0          0
    Total packets         0          0
    Unicast packets       0          0
    Broadcast packets     0          0

```

```

Multicast packets          0          0
CRC/Align errors          0          0
FIFO errors                0          0
MAC control frames        0          0
MAC pause frames          0          0
Oversized frames          0
Jabber frames             0
Fragment frames           0
VLAN tagged frames        0
Code violations            0
Filter statistics:
  Input packet count       0
  Input packet rejects     0
  Input DA rejects         0
  Input SA rejects         0
  Output packet count      0
  Output packet pad count  0
  Output packet error count 0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 4
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
      %      bps      %      usec
0 best-effort      95      4750000000    95      0      low
none
3 network-control  5      2500000000    5      0      low
none
Interface transmit statistics: Disabled

Logical interface xe-4/0/0.0 (Index 93) (SNMP ifIndex 834) (Generation 158)
Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0      0 bps
  Output bytes : 0      0 bps
  Input packets: 0      0 pps
  Output packets: 0      0 pps
Protocol inet, MTU: 1500, Generation: 192, Route table: 0
Flags: Sendbcst-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 34.1.1/24, Local: 34.1.1.2, Broadcast: 34.1.1.255, Generation:
157
Protocol multiservice, MTU: Unlimited, Generation: 193, Route table: 0
Policer: Input: __default_arp_policer__

```

Sample Output

show interfaces extensive (T4000 Routers with 24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC)

```

user@host> show interfaces xe-3/1/0 extensive
Physical interface: xe-3/1/0, Enabled, Physical link is Up
  Interface index: 160, SNMP ifIndex: 1285, Generation: 163
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
None, Loopback: None,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 2c:6b:f5:e1:cb:39, Hardware address: 2c:6b:f5:e1:cb:39
  Last flapped   : 2012-05-09 07:15:54 UTC (03:39:52 ago)
  Statistics last cleared: Never
Traffic statistics:
  Input bytes   : 0                      0 bps
  Output bytes  : 0                      0 bps
  Input packets : 0                      0 pps
  Output packets: 0                      0 pps
IPv6 transit statistics:
  Input bytes   : 0
  Output bytes  : 0
  Input packets : 0
  Output packets: 0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
  L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0,
  HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort      0                0                0

  1 ay_q1            0                0                0

  2 assured-forw     0                0                0

  3 network-cont     0                0                0

Queue number:      Mapped forwarding classes
  0                best-effort
  1                ay_q1
  2                assured-forwarding
  3                network-control
Active alarms  : None
Active defects : None
PCS statistics      Seconds
  Bit errors        0
  Errored blocks    0
MAC statistics:      Receive      Transmit
  Total octets      0            0
  Total packets     0            0
  Unicast packets   0            0

```

```

Broadcast packets          0          0
Multicast packets         0          0
CRC/Align errors          0          0
FIFO errors               0          0
MAC control frames        0          0
MAC pause frames          0          0
Oversized frames          0
Jabber frames             0
Fragment frames           0
VLAN tagged frames        0
Code violations            0
Filter statistics:
  Input packet count       0
  Input packet rejects     0
  Input DA rejects         0
  Input SA rejects         0
  Output packet count      0
  Output packet pad count  0
  Output packet error count 0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 3
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority  Limit

                                %      bps      %      usec
0 best-effort              95    9500000000    95      0      low    none
3 network-control          5     5000000000     5      0      low    none

Preclassifier statistics:
Traffic Class      Received Packets  Transmitted Packets  Dropped Packets

network-control      0                  0                  0
best-effort          0                  0                  0
Interface transmit statistics: Disabled

```

Sample Output

show interfaces extensive (Aggregated Ethernet)

```

user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Up
Interface index: 199, SNMP ifIndex: 570, Generation: 202
Link-level type: Ethernet, MTU: 1514, Speed: 2Gbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth needed: 0
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Current address: 2c:6b:f5:d1:0f:c0, Hardware address: 2c:6b:f5:d1:0f:c0
Last flapped : 2012-06-06 23:33:03 PDT (00:00:58 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes :          18532          1984 bps
  Output bytes :           0          0 bps
  Input packets:          158          2 pps
  Output packets:          0          0 pps
IPv6 transit statistics:
  Input bytes :           0
  Output bytes :           0

```



```

    Input packets:          0
    Output packets:         0
Dropped traffic statistics due to STP State:
    Input bytes :          0
    Output bytes :         0
    Input packets:         0
    Output packets:        0
Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
0,
    Resource errors: 0
Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
Ingress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

    0 best-effort          0              0              0
    1 expedited-fo        0              0              0
    2 assured-forw        0              0              0
    3 network-cont        0              0              0

Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

    0 best-effort          57             57             0
    1 expedited-fo         0              0              0
    2 assured-forw         0              0              0
    3 network-cont       63605           63605           0

Queue number:      Mapped forwarding classes
    0              best-effort
    1              expedited-forwarding
    2              assured-forwarding
    3              network-control

Logical interface ae0.0 (Index 331) (SNMP ifIndex 583) (Generation 142)
Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
Statistics      Packets      pps      Bytes      bps
Bundle:
    Input :      149          2      17416      1984
    Output:       0          0         0         0
Link:
    ge-3/2/5.0
        Input :      90          1      10100      992
        Output:       0          0         0         0
    ge-3/3/9.0
        Input :      59          1       7316      992
        Output:       0          0         0         0
LACP info:      Role      System      System      Port
Port  Port
          priority      identifier  priority      number
key
ge-3/2/5.0  Actor      100    00:00:00:00:00:01      127      1
1

```

```

    ge-3/2/5.0  Partner      127  00:24:dc:98:67:c0      127      1      1
    ge-3/3/9.0  Actor       100  00:00:00:00:00:01      127      2
1   ge-3/3/9.0  Partner      127  00:24:dc:98:67:c0      127      2      1

LACP Statistics:      LACP Rx      LACP Tx      Unknown Rx      Illegal Rx
ge-3/2/5.0            38          137           0              0
ge-3/3/9.0            36          139           0              0
Marker Statistics:    Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
ge-3/2/5.0            0           0            0              0
ge-3/3/9.0            0           0            0              0
Protocol inet, MTU: 1500, Generation: 169, Route table: 0
  Flags: Sendbcst-pkt-to-re
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 1.1.1/24, Local: 1.1.1.2, Broadcast: 1.1.1.255, Generation:
153
Protocol multiservice, MTU: Unlimited, Generation: 170, Route table: 0
  Flags: Is-Primary
  Policers: Input: __default_arp_policer__
```

show interfaces queue

Syntax	<pre>show interfaces queue <aggregate remaining-traffic> <both-ingress-egress> <egress> <forwarding-class forwarding-class> <ingress> <interface-name interface-name> <l2-statistics> <remaining-traffic></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>both-ingress-egress, egress, and ingress options introduced in Junos OS Release 7.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>l2-statistics option introduced in Junos OS Release 12.1.</p>
Description	Display class-of-service (CoS) queue information for physical interfaces.
Options	<p>none—Show detailed CoS queue statistics for all physical interfaces.</p> <p>aggregate—(Optional) Display the aggregated queuing statistics of all logical interfaces that have traffic-control profiles configured. (Not on the QFX Series.)</p> <p>both-ingress-egress—(Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics. (Not on the QFX Series.)</p> <p>egress—(Optional) Display egress queue statistics.</p> <p>forwarding-class forwarding-class—(Optional) Forwarding class name for this queue. Shows detailed CoS statistics for the queue associated with the specified forwarding class.</p> <p>ingress—(Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics. (Not on the QFX Series.)</p> <p>interface-name interface-name—(Optional) Show detailed CoS queue statistics for the specified interface.</p> <p>l2-statistics—(optional) Display layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles</p>

Overhead for Layer 2 Statistics

Transmitted packets and transmitted byte counts are displayed for the layer 2 level with the addition of encapsulation overheads applied for fragmentation, as shown in [Table 140 on page 2012](#). Others counters, such as packets and bytes queued (input) and drop counters, are displayed at the layer 3 level. In the case of link fragmentation and interleaving (LFI) for which not fragmentation is applied, corresponding layer 2 overheads are added, as shown in [Table 140 on page 2012](#).

Table 174: Layer 2 Overhead, Transmitted Packets/Bytes

Protocol	Fragmentation		LFI
	First fragmentation	Second to n fragmentations	
	Bytes	Bytes	
MLPPP (Long)	13	12	8
MLPPP (short)	11	10	8
MLFR (FRF15)	12	10	8
MFR (FRF16)	10	8	-
MCMLPPP(Long)	13	12	-
MCMLPPP(Short)	11	10	-

Layer 2 Statistics - Fragmentation Overhead Calculation

MLPPP/MC-MLPPP Overhead details:

=====

Fragment 1:

```

Outer PPP header           : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
Inner PPP header           : 1 byte
HDLC flag and FCS bytes    : 4 bytes

```

Fragments 2 .. n :

```

Outer PPP header           : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
HDLC flag and FCS bytes    : 4 bytes

```

MLFR (FRF15) Overhead details:

=====

Fragment 1:

```

Framereelay header        : 2 bytes
Control,NLPID             : 2 bytes
Fragmentaion header       : 2 bytes
Inner proto               : 2 bytes
HDLC flag and FCS         : 4 bytes

```

Fragments 2 ...n :

```

Framereelay header        : 2 bytes
Control,NLPID             : 2 bytes
Fragmentaion header       : 2 bytes
HDLC flag and FCS         : 4 bytes

```

MFR (FRF16) Overhead details:

=====

```

Fragment 1:
  Fragmentation header : 2 bytes
  Framereelay header   : 2 bytes
  Inner proto          : 2 bytes
  HDLC flag and FCS    : 4 bytes

Fragments 2 ...n :
  Fragmentation header : 2 bytes
  Framereelay header   : 2 bytes
  HDLC flag and FCS    : 4 bytes

```

Overhead with LFI

```

MLPPP(Long & short sequence):
=====
  Outer PPP header : 4 bytes
  HDLC flag and FCS : 4 bytes

MLFR (FRF15):
=====
  Framereelay header : 2 bytes
  Control,NLPID      : 2 bytes
  HDLC flag and FCS  : 4 bytes

```

The following examples show overhead for different cases:

- A 1000-byte packet is sent to a mlppp bundle without any fragmentation. At the layer 2 level, bytes transmitted is 1013 in 1 packet. This overhead is for MLPPP long sequence encap.
- A 1000-byte packet is sent to a mlppp bundle with a fragment threshold of 250byte. At the layer 2 level, bytes transmitted is 1061 bytes in 5 packets.
- A 1000-byte LFI packet is sent to an mlppp bundle. At the layer 2 level, bytes transmitted is 1008 in 1 packet.

remaining-traffic—(Optional) Display the queuing statistics of all logical interfaces that do not have traffic-control profiles configured. (Not on the QFX Series.)

Additional Information On M Series routers (except for the M320 and M120 routers), this command is valid only for a PIC installed on an enhanced Flexible PIC Concentrator (FPC).

Queue statistics for aggregated interfaces are supported on the M Series and T Series routers only. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can view the statistics for a child interface by using the **show interfaces statistics** command for that child interface.

When you configure tricolor marking on a 10-port 1-Gigabit Ethernet PIC, for queues 6 and 7 only, the output does not display the number of queued bytes and packets, or the number of bytes and packets dropped because of RED. If you do not configure tricolor marking on the interface, these statistics are available for all queues.

For the 4-port Channelized OC12 IQE PIC and 1-port Channelized OC48 IQE PIC, the **Packet Forwarding Engine Chassis Queues** field represents traffic bound for a particular

physical interface on the PIC. For all other PICs, the **Packet Forwarding Engine Chassis Queues** field represents the total traffic bound for the PIC.

For Gigabit Ethernet IQ2 PICs, the **show interfaces queue** command output does not display the number of tail-dropped packets. This limitation does not apply to Packet Forwarding Engine chassis queues.

When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (under the **Packet Forwarding Engine Chassis Queues** field) shows the prefragmentation values.

The behavior of the **egress** queues for the **Routing Engine-Generated Traffic** is not same as the configured queue for MLPPP and MFR configurations.

For information about how to configure CoS, see the *Junos® OS Network Interfaces*. For related CoS operational mode commands, see the *Junos OS Operational Mode Commands*.

Required Privilege Level view

List of Sample Output

- [show interfaces queue \(Aggregated Ethernet on a T320 Router\) on page 2476](#)
- [show interfaces queue \(Fast Ethernet on a J4300 Router\) on page 2478](#)
- [show interfaces queue \(Gigabit Ethernet on a T640 Router\) on page 2479](#)
- [show interfaces queue aggregate \(Gigabit Ethernet Enhanced DPC\) on page 2479](#)
- [show interfaces queue \(Gigabit Ethernet IQ2 PIC\) on page 2483](#)
- [show interfaces queue both-ingress-egress \(Gigabit Ethernet IQ2 PIC\) on page 2486](#)
- [show interfaces queue ingress \(Gigabit Ethernet IQ2 PIC\) on page 2488](#)
- [show interfaces queue egress \(Gigabit Ethernet IQ2 PIC\) on page 2489](#)
- [show interfaces queue remaining-traffic \(Gigabit Ethernet Enhanced DPC\) on page 2490](#)
- [show interfaces queue \(Channelized OC12 IQE Type 3 PIC in SONET Mode\) on page 2493](#)
- [show interfaces queue \(QFX Series\) on page 2503](#)
- [show interfaces queue l2-statistics \(lsq interface\) on page 2504](#)

Output Fields [Table 141 on page 2014](#) lists the output fields for the **show interfaces queue** command. Output fields are listed in the approximate order in which they appear.

Table 175: show interfaces queue Output Fields

Field Name	Field Description
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in the "Enabled Field" section under " Common Output Fields Description " on page 2376.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Forwarding classes supported	Total number of forwarding classes supported on the specified interface.

Table 175: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
Forwarding classes in use	Total number of forwarding classes in use on the specified interface.
Ingress queues supported	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues supported on the specified interface.
Ingress queues in use	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues in use on the specified interface.
Output queues supported	Total number of output queues supported on the specified interface.
Output queues in use	Total number of output queues in use on the specified interface.
Egress queues supported	Total number of egress queues supported on the specified interface.
Egress queues in use	Total number of egress queues in use on the specified interface.
Queue	Queue number.
Queue counters (Ingress)	<p>CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism.
Burst size	(Logical interfaces on IQ PICs only) Maximum number of bytes up to which the logical interface can burst. The burst size is based on the shaping rate applied to the interface.
Forwarding classes	Forwarding class name.
Queued Packets	<p>Number of packets queued to this queue.</p> <p>NOTE: For Gigabit Ethernet IQ2 interfaces, the Queued Packets count is calculated by the Junos OS interpreting one frame buffer as one packet. If the queued packets are very large or very small, the calculation might not be completely accurate for transit traffic. The count is completely accurate for traffic terminated on the router.</p>
Queued Bytes	<p>Number of bytes queued to this queue. The byte counts vary by PIC type. For more information, see Table 142 on page 2017.</p>
Transmitted Packets	<p>Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the Packet Forwarding Engine Chassis Queues field) shows the prefragmentation values.</p> <p>NOTE: For layer 2 statistics, see “Overhead for Layer 2 Statistics” on page 2011</p>

Table 175: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
Transmitted Bytes	<p>Number of bytes transmitted by this queue. The byte counts vary by PIC type. For more information, see Table 142 on page 2017.</p> <p>NOTE: On MX Series routers, this number can be inaccurate when you issue the command for a physical interface repeatedly and in quick succession, because the statistics for the child nodes are collected infrequently. Wait ten seconds between successive iterations to avoid this situation.</p> <p>NOTE: For layer 2 statistics, see “Overhead for Layer 2 Statistics” on page 2011</p>
Tail-dropped packets	Number of packets dropped because of tail drop.
RED-dropped packets	<p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> • (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> • Low, non-TCP—Number of low-loss priority non-TCP packets dropped because of RED. • Low, TCP—Number of low-loss priority TCP packets dropped because of RED. • High, non-TCP—Number of high-loss priority non-TCP packets dropped because of RED. • High, TCP—Number of high-loss priority TCP packets dropped because of RED. • (J Series routers and MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> • Low—Number of low-loss priority packets dropped because of RED. • Medium-low—Number of medium-low loss priority packets dropped because of RED. • Medium-high—Number of medium-high loss priority packets dropped because of RED. • High—Number of high-loss priority packets dropped because of RED.
RED-dropped bytes	<p>Number of bytes dropped because of RED. The byte counts vary by PIC type. For more information, see Table 142 on page 2017.</p> <ul style="list-style-type: none"> • (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> • Low, non-TCP—Number of low-loss priority non-TCP bytes dropped because of RED. • Low, TCP—Number of low-loss priority TCP bytes dropped because of RED. • High, non-TCP—Number of high-loss priority non-TCP bytes dropped because of RED. • High, TCP—Number of high-loss priority TCP bytes dropped because of RED. • (J Series routers only) The output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> • Low—Number of low-loss priority bytes dropped because of RED. • Medium-low—Number of medium-low loss priority bytes dropped because of RED. • Medium-high—Number of medium-high loss priority bytes dropped because of RED. • High—Number of high-loss priority bytes dropped because of RED.

Byte counts vary by PIC type. [Table 142 on page 2017](#) shows how the byte counts on the outbound interfaces vary depending on the PIC type. [Table 142 on page 2017](#) is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.

Table 176: Byte Count by PIC Type

PIC Type	Output Level	Byte Count Includes	Comments
Gigabit Ethernet IQ and IQE PICs	Interface	<p>Queued: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>Transmitted: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>RED dropped: 496 bytes per packet representing 478 bytes of Layer 3 packet + 18 bytes</p>	<p>The 12 additional bytes include 6 bytes for the destination MAC address + 4 bytes for the VLAN + 2 bytes for the Ethernet type.</p> <p>For RED dropped, 6 bytes are added for the source MAC address.</p>
	Packet forwarding component	<p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p>	—
Non-IQ PIC	Interface	<p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> Queued: 478 bytes of Layer 3 packet. Transmitted: 478 bytes of Layer 3 packet. <p>T4000 routers with Type 5 FPCs :</p> <ul style="list-style-type: none"> Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Inter frame Gap. Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Interframe Gap. <p>M Series routers:</p> <ul style="list-style-type: none"> Queued: 478 bytes of Layer 3 packet. Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead. <p>PTX Series Packet Transport Switches:</p> <ul style="list-style-type: none"> Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes FCS + the full Layer 1 overhead of the MAC header DA + SA + EtherType (non-VLAN). Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead of the MAC header DA + SA + EtherType (non-VLAN). RED dropped: 478 bytes of Layer 3 packet + 22 bytes special header. To the TQ, this packet has 4 bytes more than queued or transmitted. 	<p>The Layer 2 overhead is 14 bytes for non-VLAN traffic and 18 bytes for VLAN traffic.</p>

Table 176: Byte Count by PIC Type (*continued*)

PIC Type	Output Level	Byte Count Includes	Comments
IQ and IQE PICs with a SONET/SDH interface	Interface	<p>Queued: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>Transmitted: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>RED dropped: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p>	The additional 4 bytes are for the Layer 2 Point-to-Point Protocol (PPP) header.
	Packet forwarding component	<p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 486 bytes per packet, representing 478 bytes of Layer 3 packet + 8 bytes</p>	For transmitted packets, the additional 8 bytes includes 4 bytes for the PPP header and 4 bytes for a cookie.
Non-IQ PIC with a SONET/SDH interface	Interface	<p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> Queued: 478 bytes of Layer 3 packet. Transmitted: 478 bytes of Layer 3 packet. <p>M Series routers:</p> <ul style="list-style-type: none"> Queued: 478 bytes of Layer 3 packet. Transmitted: 483 bytes per packet, representing 478 bytes of Layer 3 packet + 5 bytes RED dropped: 478 bytes per packet, representing 478 bytes of Layer 3 packet 	For transmitted packets, the additional 5 bytes includes 4 bytes for the PPP header and 1 byte for the packet loss priority (PLP).
Interfaces configured with Frame Relay Encapsulation	Interface	The default Frame Relay overhead is 7 bytes. If you configure the Frame Check Sequence (FCS) to 4 bytes, then the overhead increases to 10 bytes.	
1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs	Interface	<p>Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p> <p>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p>	The Layer 2 overhead is 18 bytes for non-VLAN traffic and 22 bytes for VLAN traffic.
4-port 1G IQ2 and IQ2-E PICs	Packet forwarding component	Queued: 478 bytes of Layer 3 packet.	—
8-port 1G IQ2 and IQ2-E PICs		Transmitted: 478 bytes of Layer 3 packet.	

Sample Output

show interfaces queue (Aggregated Ethernet on a T320 Router)

The following example shows that the aggregated Ethernet interface, **ae1**, has traffic on queues **af1** and **af12**:

```

user@host> show interfaces queue ae1
Physical interface: ae1, Enabled, Physical link is Up
Interface index: 158, SNMP ifIndex: 33 Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:
    Packets      :           5           0 pps
    Bytes        :          242           0 bps
  Transmitted:
    Packets      :           5           0 pps
    Bytes        :          242           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 1, Forwarding classes: af1
  Queued:
    Packets      :      42603765      595484 pps
    Bytes        :     5453281920     609776496 bps
  Transmitted:
    Packets      :      42603765      595484 pps
    Bytes        :     5453281920     609776496 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 2, Forwarding classes: ef1
  Queued:
    Packets      :           0           0 pps
    Bytes        :           0           0 bps
  Transmitted:
    Packets      :           0           0 pps
    Bytes        :           0           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 3, Forwarding classes: nc
  Queued:
    Packets      :           45           0 pps
    Bytes        :          3930           0 bps
  Transmitted:
    Packets      :           45           0 pps
    Bytes        :          3930           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 4, Forwarding classes: af11
  Queued:
    Packets      :           0           0 pps
    Bytes        :           0           0 bps
  Transmitted:
    Packets      :           0           0 pps
    Bytes        :           0           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 5, Forwarding classes: ef11
  Queued:
    Packets      :           0           0 pps
    Bytes        :           0           0 bps
  Transmitted:
    Packets      :           0           0 pps
    Bytes        :           0           0 bps

```

```

Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 6, Forwarding classes: af12
Queued:
Packets : 31296413 437436 pps
Bytes : 4005940864 447935200 bps
Transmitted:
Packets : 31296413 437436 pps
Bytes : 4005940864 447935200 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 7, Forwarding classes: nc2
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

show interfaces queue (Fast Ethernet on a J4300 Router)

```

user@host> show interfaces queue fe-4/0/0.0
Logical interface fe-4/0/0.0 (Index 71) (SNMP ifIndex 42)
Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
Queued:
Packets : 5240762 3404 pps
Bytes : 3020710354 15934544 bps
Transmitted:
Packets : 5240762 3404 pps
Bytes : 3020710354 15934544 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
Queue: 1, Forwarding classes: af1
Queued:
Packets : 2480391 1650 pps
Bytes : 1304685666 6945704 bps
Transmitted:
Packets : 2478740 1650 pps
Bytes : 1303817240 6945704 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 1651 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 1651 0 pps

```

RED-dropped bytes	:	868426	0 bps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	868426	0 pps

show interfaces queue (Gigabit Ethernet on a T640 Router)

```

user@host> show interfaces queue
Physical interface: ge-7/0/1, Enabled, Physical link is Up
Interface index: 150, SNMP ifIndex: 42
Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:
    Packets      :          13          0 pps
    Bytes        :         622          0 bps
  Transmitted:
    Packets      :          13          0 pps
    Bytes        :         622          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 1, Forwarding classes: af1
  Queued:
    Packets      :      1725947945      372178 pps
    Bytes        :    220921336960    381110432 bps
  Transmitted:
    Packets      :      1725947945      372178 pps
    Bytes        :    220921336960    381110432 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 2, Forwarding classes: ef1
  Queued:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
  Transmitted:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 3, Forwarding classes: nc
  Queued:
    Packets      :          571          0 pps
    Bytes        :         49318        336 bps
  Transmitted:
    Packets      :          571          0 pps
    Bytes        :         49318        336 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps

```

show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 aggregate
Physical interface: ge-2/2/9, Enabled, Physical link is Up
Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use

```

```

Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :      148450735      947295 pps
    Bytes        :      8016344944    409228848 bps
  Transmitted:
    Packets      :      76397439      487512 pps
    Bytes        :     4125461868    210602376 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :      72053285      459783 pps
    Low          :      72053285      459783 pps
    Medium-low   :           0         0 pps
    Medium-high  :           0         0 pps
    High         :           0         0 pps
  RED-dropped bytes  :     3890877444    198626472 bps
    Low          :     3890877444    198626472 bps
    Medium-low   :           0         0 bps
    Medium-high  :           0         0 bps
    High         :           0         0 bps

Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :           0         0 pps
    Bytes        :           0         0 bps
  Transmitted:
    Packets      :           0         0 pps
    Bytes        :           0         0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :           0         0 pps
    Low          :           0         0 pps
    Medium-low   :           0         0 pps
    Medium-high  :           0         0 pps
    High         :           0         0 pps
  RED-dropped bytes  :           0         0 bps
    Low          :           0         0 bps
    Medium-low   :           0         0 bps
    Medium-high  :           0         0 bps
    High         :           0         0 bps

Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      :      410278257      473940 pps
    Bytes        :     22156199518    204742296 bps
  Transmitted:
    Packets      :      4850003       4033 pps
    Bytes        :     261900162     1742256 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :     405425693     469907 pps
    Low          :     405425693     469907 pps
    Medium-low   :           0         0 pps
    Medium-high  :           0         0 pps
    High         :           0         0 pps
  RED-dropped bytes  :     21892988124    203000040 bps
    Low          :     21892988124    203000040 bps
    Medium-low   :           0         0 bps
    Medium-high  :           0         0 bps
    High         :           0         0 bps

Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      :           0         0 pps
    Bytes        :           0         0 bps
  Transmitted:
    Packets      :           0         0 pps
    Bytes        :           0         0 bps

```

```

Tail-dropped packets : Not Available
RED-dropped packets :                0                0 pps
  Low                  :                0                0 pps
  Medium-low           :                0                0 pps
  Medium-high          :                0                0 pps
  High                 :                0                0 pps
RED-dropped bytes    :                0                0 bps
  Low                  :                0                0 bps
  Medium-low           :                0                0 bps
  Medium-high          :                0                0 bps
  High                 :                0                0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets            :                76605230          485376 pps
    Bytes              :                5209211400        264044560 bps
  Transmitted:
    Packets            :                76444631          484336 pps
    Bytes              :                5198235612        263478800 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :                160475            1040 pps
    Low                  :                160475            1040 pps
    Medium-low           :                0                0 pps
    Medium-high          :                0                0 pps
    High                 :                0                0 pps
  RED-dropped bytes    :                10912300          565760 bps
    Low                  :                10912300          565760 bps
    Medium-low           :                0                0 bps
    Medium-high          :                0                0 bps
    High                 :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets            :                0                0 pps
    Bytes              :                0                0 bps
  Transmitted:
    Packets            :                0                0 pps
    Bytes              :                0                0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :                0                0 pps
    Low                  :                0                0 pps
    Medium-low           :                0                0 pps
    Medium-high          :                0                0 pps
    High                 :                0                0 pps
  RED-dropped bytes    :                0                0 bps
    Low                  :                0                0 bps
    Medium-low           :                0                0 bps
    Medium-high          :                0                0 bps
    High                 :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets            :                4836136           3912 pps
    Bytes              :                333402032        2139056 bps
  Transmitted:
    Packets            :                3600866           1459 pps
    Bytes              :                244858888        793696 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :                1225034           2450 pps
    Low                  :                1225034           2450 pps
    Medium-low           :                0                0 pps
    Medium-high          :                0                0 pps

```

```

      High : 0 0 pps
    RED-dropped bytes : 83302312 1333072 bps
      Low : 83302312 1333072 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
      High : 0 0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps

```

Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

```

  Queued:
    Packets : 77059796 486384 pps
    Bytes : 3544750624 178989576 bps
  Transmitted:
    Packets : 77059797 486381 pps
    Bytes : 3544750670 178988248 bps
  Tail-dropped packets : 0 0 pps
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps

```

Queue: 1, Forwarding classes: expedited-forwarding

```

  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Tail-dropped packets : 0 0 pps
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps

```



```

      Medium-high      :           0           0 bps
      High             :           0           0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets           :      4846580           3934 pps
    Bytes             :      222942680        1447768 bps
  Transmitted:
    Packets           :      4846580           3934 pps
    Bytes             :      222942680        1447768 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
      Low             :           0           0 pps
      Medium-low      :           0           0 pps
      Medium-high     :           0           0 pps
      High            :           0           0 pps
    RED-dropped bytes :           0           0 bps
      Low             :           0           0 bps
      Medium-low      :           0           0 bps
      Medium-high     :           0           0 bps
      High            :           0           0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets           :           0           0 pps
    Bytes             :           0           0 bps
  Transmitted:
    Packets           :           0           0 pps
    Bytes             :           0           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
      Low             :           0           0 pps
      Medium-low      :           0           0 pps
      Medium-high     :           0           0 pps
      High            :           0           0 pps
    RED-dropped bytes :           0           0 bps
      Low             :           0           0 bps
      Medium-low      :           0           0 bps
      Medium-high     :           0           0 bps
      High            :           0           0 bps

```

show interfaces queue (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-7/1/3
Physical interface: ge-7/1/3, Enabled, Physical link is Up
  Interface index: 170, SNMP ifIndex: 70 Forwarding classes: 16 supported, 4 in
  use Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets           :      418390039           10 pps
    Bytes             :      38910269752        7440 bps
  Transmitted:
    Packets           :      418390039           10 pps
    Bytes             :      38910269752        7440 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets           :           0           0 pps
    Bytes             :           0           0 bps
  Transmitted:
    Packets           :           0           0 pps

```

```

Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 7055 1 pps
Bytes : 451552 512 bps
Transmitted:
Packets : 7055 1 pps
Bytes : 451552 512 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Forwarding classes: 16 supported, 4 in use Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : 1031 0 pps
Bytes : 143292 0 bps
Transmitted:
Packets : 1031 0 pps
Bytes : 143292 0 bps
Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

Queue: 3, Forwarding classes: network-control

Queued:

Packets	:	77009	11 pps
Bytes	:	6894286	7888 bps

Transmitted:

Packets	:	77009	11 pps
Bytes	:	6894286	7888 bps

Tail-dropped packets : Not Available

RL-dropped packets	:	0	0 pps
--------------------	---	---	-------

RL-dropped bytes	:	0	0 bps
------------------	---	---	-------

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

Queued:

Packets	:	1031	0 pps
Bytes	:	147328	0 bps

Transmitted:

Packets	:	1031	0 pps
Bytes	:	147328	0 bps

Tail-dropped packets : 0 0 pps

RED-dropped packets : 0 0 pps

Low, non-TCP	:	0	0 pps
--------------	---	---	-------

Low, TCP	:	0	0 pps
----------	---	---	-------

High, non-TCP	:	0	0 pps
---------------	---	---	-------

High, TCP	:	0	0 pps
-----------	---	---	-------

RED-dropped bytes : 0 0 bps

Low, non-TCP	:	0	0 bps
--------------	---	---	-------

Low, TCP	:	0	0 bps
----------	---	---	-------

High, non-TCP	:	0	0 bps
---------------	---	---	-------

High, TCP	:	0	0 bps
-----------	---	---	-------

Queue: 1, Forwarding classes: expedited-forwarding

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Tail-dropped packets : 0 0 pps

RED-dropped packets : 0 0 pps

Low, non-TCP	:	0	0 pps
--------------	---	---	-------

Low, TCP	:	0	0 pps
----------	---	---	-------

High, non-TCP	:	0	0 pps
---------------	---	---	-------

High, TCP	:	0	0 pps
-----------	---	---	-------

RED-dropped bytes : 0 0 bps

Low, non-TCP	:	0	0 bps
--------------	---	---	-------

Low, TCP	:	0	0 bps
----------	---	---	-------

High, non-TCP	:	0	0 bps
---------------	---	---	-------

High, TCP	:	0	0 bps
-----------	---	---	-------

Queue: 2, Forwarding classes: assured-forwarding

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Tail-dropped packets : 0 0 pps

RED-dropped packets : 0 0 pps

Low, non-TCP	:	0	0 pps
--------------	---	---	-------

```

        Low, TCP           :           0           0 pps
        High, non-TCP      :           0           0 pps
        High, TCP          :           0           0 pps
        RED-dropped bytes  :           0           0 bps
        Low, non-TCP       :           0           0 bps
        Low, TCP           :           0           0 bps
        High, non-TCP      :           0           0 bps
        High, TCP          :           0           0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets           :           94386           12 pps
  Bytes             :          13756799          9568 bps
Transmitted:
  Packets           :           94386           12 pps
  Bytes             :          13756799          9568 bps
  Tail-dropped packets :           0           0 pps
  RED-dropped packets :           0           0 pps
  Low, non-TCP       :           0           0 pps
  Low, TCP           :           0           0 pps
  High, non-TCP      :           0           0 pps
  High, TCP          :           0           0 pps
  RED-dropped bytes  :           0           0 bps
  Low, non-TCP       :           0           0 bps
  Low, TCP           :           0           0 bps
  High, non-TCP      :           0           0 bps
  High, TCP          :           0           0 bps

```

show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 both-ingress-egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets           : Not Available
    Bytes             :           0           0 bps
  Transmitted:
    Packets           :           254           0 pps
    Bytes             :          16274          0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets           : Not Available
    Bytes             :           0           0 bps
  Transmitted:
    Packets           :           0           0 pps
    Bytes             :           0           0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets           : Not Available
    Bytes             :           0           0 bps
  Transmitted:
    Packets           :           0           0 pps
    Bytes             :           0           0 bps

```

```

Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 3 0 pps
Bytes : 126 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : 80564692 0 pps
Bytes : 3383717100 0 bps

```

```

Transmitted:
  Packets      :      80564692      0 pps
  Bytes        :      3383717100    0 bps
  Tail-dropped packets :      0      0 pps
  RED-dropped packets :      0      0 pps
  RED-dropped bytes  :      0      0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets      :      80564685      0 pps
  Bytes        :      3383716770    0 bps
Transmitted:
  Packets      :      80564685      0 pps
  Bytes        :      3383716770    0 bps
  Tail-dropped packets :      0      0 pps
  RED-dropped packets :      0      0 pps
  RED-dropped bytes  :      0      0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets      :      0      0 pps
  Bytes        :      0      0 bps
Transmitted:
  Packets      :      0      0 pps
  Bytes        :      0      0 bps
  Tail-dropped packets :      0      0 pps
  RED-dropped packets :      0      0 pps
  RED-dropped bytes  :      0      0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets      :      9397      0 pps
  Bytes        :      3809052      232 bps
Transmitted:
  Packets      :      9397      0 pps
  Bytes        :      3809052      232 bps
  Tail-dropped packets :      0      0 pps
  RED-dropped packets :      0      0 pps
  RED-dropped bytes  :      0      0 bps

```

show interfaces queue ingress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 ingress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets      : Not Available
  Bytes        :      0      0 bps
Transmitted:
  Packets      :      288      0 pps
  Bytes        :      18450    0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :      0      0 pps
  RED-dropped bytes  :      0      0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets      : Not Available
  Bytes        :      0      0 bps
Transmitted:
  Packets      :      0      0 pps
  Bytes        :      0      0 bps

```

```

Tail-dropped packets : Not Available
RED-dropped packets  : 0 0 pps
RED-dropped bytes    : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets              : Not Available
Bytes                : 0 0 bps
Transmitted:
Packets              : 0 0 pps
Bytes                : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets  : 0 0 pps
RED-dropped bytes    : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets              : Not Available
Bytes                : 0 0 bps
Transmitted:
Packets              : 0 0 pps
Bytes                : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets  : 0 0 pps
RED-dropped bytes    : 0 0 bps

```

show interfaces queue egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets              : Not Available
Bytes                : 0 0 bps
Transmitted:
Packets              : 3 0 pps
Bytes                : 126 0 bps
Tail-dropped packets : Not Available
RED-dropped packets  : 0 0 pps
RED-dropped bytes    : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets              : Not Available
Bytes                : 0 0 bps
Transmitted:
Packets              : 0 0 pps
Bytes                : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets  : 0 0 pps
RED-dropped bytes    : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets              : Not Available
Bytes                : 0 0 bps
Transmitted:
Packets              : 0 0 pps
Bytes                : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets  : 0 0 pps
RED-dropped bytes    : 0 0 bps

```

```

Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      : Not Available
    Bytes        :                      0          0 bps
  Transmitted:
    Packets      :                      0          0 pps
    Bytes        :                      0          0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                      0          0 pps
    RED-dropped bytes   :                      0          0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :          80564692          0 pps
    Bytes        :          3383717100        0 bps
  Transmitted:
    Packets      :          80564692          0 pps
    Bytes        :          3383717100        0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :          80564685          0 pps
    Bytes        :          3383716770        0 bps
  Transmitted:
    Packets      :          80564685          0 pps
    Bytes        :          3383716770        0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
  Transmitted:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      :          9538          0 pps
    Bytes        :          3819840          0 bps
  Transmitted:
    Packets      :          9538          0 pps
    Bytes        :          3819840          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps

```

show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 remaining-traffic
Physical interface: ge-2/2/9, Enabled, Physical link is Up
  Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort

```



```

Queued:
  Packets      :      110208969      472875 pps
  Bytes       :      5951284434    204282000 bps
Transmitted:
  Packets      :      110208969      472875 pps
  Bytes       :      5951284434    204282000 bps
Tail-dropped packets : Not Available
RED-dropped packets :      0      0 pps
  Low          :      0      0 pps
  Medium-low   :      0      0 pps
  Medium-high  :      0      0 pps
  High         :      0      0 pps
RED-dropped bytes  :      0      0 bps
  Low          :      0      0 bps
  Medium-low   :      0      0 bps
  Medium-high  :      0      0 bps
  High         :      0      0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets      :      0      0 pps
  Bytes       :      0      0 bps
Transmitted:
  Packets      :      0      0 pps
  Bytes       :      0      0 bps
Tail-dropped packets : Not Available
RED-dropped packets :      0      0 pps
  Low          :      0      0 pps
  Medium-low   :      0      0 pps
  Medium-high  :      0      0 pps
  High         :      0      0 pps
RED-dropped bytes  :      0      0 bps
  Low          :      0      0 bps
  Medium-low   :      0      0 bps
  Medium-high  :      0      0 bps
  High         :      0      0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets      :      0      0 pps
  Bytes       :      0      0 bps
Transmitted:
  Packets      :      0      0 pps
  Bytes       :      0      0 bps
Tail-dropped packets : Not Available
RED-dropped packets :      0      0 pps
  Low          :      0      0 pps
  Medium-low   :      0      0 pps
  Medium-high  :      0      0 pps
  High         :      0      0 pps
RED-dropped bytes  :      0      0 bps
  Low          :      0      0 bps
  Medium-low   :      0      0 bps
  Medium-high  :      0      0 bps
  High         :      0      0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets      :      0      0 pps
  Bytes       :      0      0 bps
Transmitted:
  Packets      :      0      0 pps
  Bytes       :      0      0 bps
Tail-dropped packets : Not Available

```

```

RED-dropped packets : 0 0 pps
  Low : 0 0 pps
  Medium-low : 0 0 pps
  Medium-high : 0 0 pps
  High : 0 0 pps
RED-dropped bytes : 0 0 bps
  Low : 0 0 bps
  Medium-low : 0 0 bps
  Medium-high : 0 0 bps
  High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets : 109355853 471736 pps
    Bytes : 7436199152 256627968 bps
  Transmitted:
    Packets : 109355852 471736 pps
    Bytes : 7436198640 256627968 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps

```

```

RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps

```

show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode)

```

user@host> show interfaces queue t3-1/1/0:7
Physical interface: t3-1/1/0:7, Enabled, Physical link is Up

Interface index: 192, SNMP ifIndex: 1948

Description: full T3 interface connect to 6ce13 t3-3/1/0:7 for FR testing -
Lam

Forwarding classes: 16 supported, 9 in use

Egress queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

Queued:

Packets : 214886 13449 pps
Bytes : 9884756 5164536 bps

Transmitted:

Packets : 214886 13449 pps
Bytes : 9884756 5164536 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps

```

High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 1, Forwarding classes: REALTIME

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 2, Forwarding classes: PRIVATE

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 3, Forwarding classes: CONTROL

Queued:

Packets	:	60	0 pps
Bytes	:	4560	0 bps

Transmitted:

Packets	:	60	0 pps
Bytes	:	4560	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 4, Forwarding classes: CLASS_B_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 5, Forwarding classes: CLASS_C_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps

RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 6, Forwarding classes: CLASS_V_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 7, Forwarding classes: CLASS_S_OUTPUT, GETS

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps

RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Packet Forwarding Engine Chassis Queues:

Queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

Queued:

Packets	:	371365	23620 pps
Bytes	:	15597330	7936368 bps

Transmitted:

Packets	:	371365	23620 pps
Bytes	:	15597330	7936368 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps

High	:	0	0 bps
Queue: 1, Forwarding classes: REALTIME			
Queued:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Transmitted:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps
Queue: 2, Forwarding classes: PRIVATE			
Queued:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Transmitted:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps

High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 3, Forwarding classes: CONTROL

Queued:

Packets	:	32843	0 pps
Bytes	:	2641754	56 bps

Transmitted:

Packets	:	32843	0 pps
Bytes	:	2641754	56 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 4, Forwarding classes: CLASS_B_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 5, Forwarding classes: CLASS_C_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 6, Forwarding classes: CLASS_V_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 7, Forwarding classes: CLASS_S_OUTPUT, GETS

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps

RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

show interfaces queue (QFX Series)

```

user@switch> show interfaces queue xe-0/0/15
Physical interface: xe-0/0/15, Enabled, Physical link is Up
Interface index: 49165, SNMP ifIndex: 539
Forwarding classes: 12 supported, 8 in use
Egress queues: 12 supported, 8 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      : 0          0 pps
    Bytes        : 0          0 bps
  Transmitted:
    Packets      : 0          0 pps
    Bytes        : 0          0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0          0 pps
    Total-dropped bytes  : 0          0 bps
Queue: 3, Forwarding classes: fcoe
  Queued:
    Packets      : 0          0 pps
    Bytes        : 0          0 bps
  Transmitted:
    Packets      : 0          0 pps
    Bytes        : 0          0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0          0 pps
    Total-dropped bytes  : 0          0 bps
0 bps
Queue: 4, Forwarding classes: no-loss
  Queued:
    Packets      : 0          0 pps
    Bytes        : 0          0 bps
  Transmitted:
    Packets      : 0          0 pps
    Bytes        : 0          0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0          0 pps
    Total-dropped bytes  : 0          0 bps
Queue: 7, Forwarding classes: network-control
  Queued:
    Packets      : 0          0 pps
    Bytes        : 0          0 bps
  Transmitted:
    Packets      : 0          0 pps
    Bytes        : 0          0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0          0 pps
    Total-dropped bytes  : 0          0 bps
Queue: 8, Forwarding classes: mcast
  Queued:

```

Packets	:	0	0 pps
Bytes	:	0	0 bps
Transmitted:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets : Not Available			
Total-dropped packets:		0	0 pps
Total-dropped bytes :		0	0 bps

show interfaces queue l2-statistics (lsq interface)

```

user@switch> show interfaces queue lsq-2/2/0.2 l2-statistics
Logical interface lsq-2/2/0.2 (Index 69) (SNMP ifIndex 1598)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
  Queued:
    Packets      :           1      0 pps
    Bytes        :        1001      0 bps
  Transmitted:
    Packets      :           5      0 pps
    Bytes        :        1062      0 bps
    Tail-dropped packets :           0      0 pps
    RED-dropped packets :           0      0 pps
    RED-dropped bytes  :           0      0 bps
Queue: 1, Forwarding classes: ef
  Queued:
    Packets      :           1      0 pps
    Bytes        :        1500      0 bps
  Transmitted:
    Packets      :           6      0 pps
    Bytes        :        1573      0 bps
    Tail-dropped packets :           0      0 pps
    RED-dropped packets :           0      0 pps
    RED-dropped bytes  :           0      0 bps
Queue: 2, Forwarding classes: af
  Queued:
    Packets      :           1      0 pps
    Bytes        :         512      0 bps
  Transmitted:
    Packets      :           3      0 pps
    Bytes        :         549      0 bps
    Tail-dropped packets :           0      0 pps
    RED-dropped packets :           0      0 pps
    RED-dropped bytes  :           0      0 bps
Queue: 3, Forwarding classes: nc
  Queued:
    Packets      :           0      0 pps
    Bytes        :           0      0 bps
  Transmitted:
    Packets      :           0      0 pps
    Bytes        :           0      0 bps
    Tail-dropped packets :           0      0 pps
    RED-dropped packets :           0      0 pps
    RED-dropped bytes  :           0      0 bps
=====

```

CHAPTER 11

Layer 3 Protocols

- [BGP on page 2505](#)
- [IS-IS on page 2675](#)
- [OSPF on page 2784](#)
- [RIP and RIPng on page 2915](#)
- [Routing Options on page 3193](#)

BGP

- [Overview on page 2505](#)
- [Configuration on page 2553](#)
- [Administration on page 2640](#)

Overview

- [Feature Support on page 2505](#)

Feature Support

- [EX Series Switch Software Features Overview on page 2505](#)

EX Series Switch Software Features Overview

This topic lists the Juniper Networks EX Series Ethernet Switches software features, the Juniper Networks Junos operating system (Junos OS) release in which they were introduced, and the first Junos OS release for each switch.



NOTE: For Virtual Chassis features, see *EX Series Virtual Chassis Software Features Overview*.



NOTE: In all tables in this topic, “N.S.” = “Not supported”, and “—” = “Not applicable”.

- [Table 3 on page 64](#)—First Junos OS Release for Each EX Series Switch
- [Table 4 on page 65](#)—Access Control Features

- [Table 5 on page 66](#)—Administration Features
- [Table 6 on page 67](#)—Class-of-Service (CoS) Features
- [Table 7 on page 68](#)—Class-of-Service (CoS) Features for EX9200 Switches
- [Table 8 on page 70](#)—Converged Networks (LAN and SAN) Features
- [Table 9 on page 70](#)—Device Security Features
- [Table 10 on page 71](#)—High Availability and Resiliency Features
- [Table 11 on page 73](#)—High Availability and Resiliency Features on EX9200 Switches
- [Table 12 on page 75](#)—Interfaces Features
- [Table 13 on page 76](#)—Interfaces Features on EX9200 Switches
- [Table 14 on page 77](#)—IP Address Management Features
- [Table 15 on page 77](#)—IP Address Management Features on EX9200 Switches
- [Table 16 on page 78](#)—IPv6 Features
- [Table 17 on page 81](#)—Layer 2 Network Protocols Features
- [Table 18 on page 82](#)—Layer 2 Network Protocols Features on EX9200 Switches
- [Table 19 on page 84](#)—Layer 3 Protocols Features
- [Table 20 on page 86](#)—Layer 3 Protocols Features for EX9200 Switches
- [Table 21 on page 89](#)—Logical Systems Features on EX9200 Switches
- [Table 22 on page 89](#)—MPLS Features
- [Table 23 on page 92](#)—MPLS Features on EX9200 Switches
- [Table 24 on page 94](#)—Multicast Features
- [Table 25 on page 95](#)—Multicast Features on EX9200 Switches
- [Table 26 on page 97](#)—Network Management and Monitoring Features
- [Table 27 on page 99](#)—Network Management and Monitoring Features on EX9200 Switches
- [Table 28 on page 100](#)—Port Security Features
- [Table 29 on page 102](#)—Power over Ethernet (PoE) Features
- [Table 30 on page 103](#)—Routing Policy and Packet Filtering Features
- [Table 31 on page 103](#)—Routing Policy and Packet Filtering Features on EX9200 Switches
- [Table 32 on page 105](#)—Spanning-Tree Protocols Features
- [Table 33 on page 106](#)—Spanning-Tree Protocols Features on EX9200 Switches
- [Table 34 on page 107](#)—System Management Features
- [Table 35 on page 107](#)—System Management Features on EX9200 Switches
- [Table 36 on page 108](#)—User Interface and Configuration Features on EX9200 Switches
- [Table 37 on page 108](#)—VPN Features on EX9200 Switches

The Junos OS release for software features on a switch cannot be earlier than the first Junos OS release for that switch.

Table 177: First Junos OS Release for Each EX Series Switch

Switch	Junos OS Release
EX2200 switches*	Junos OS Release 10.1R1 *EX2200-C models: Junos OS Release 11.3R1
EX3200 switches	Junos OS Release 9.0R1
EX3300 switches	Junos OS Release 11.3R1
EX4200 switches	Junos OS Release 9.0R1
EX4300 switches	Junos OS Release 13.2X50-D10
EX4500 switches**	Junos OS Release 10.2R1* **EX4500-C models: Junos OS Release 10.3R2
EX4550 switches	Junos OS Release 12.2R1
EX6200 switch	Junos OS Release 11.3R2
EX8208 switches	Junos OS Release 9.4R1
EX8216 switches	Junos OS Release 9.5R1
EX9200 switches	Junos OS Release 12.3R2

Table 178: Access Control Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
802.1X authentication (port-based, multiple supplicant)	10.1R1	9.0R2	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	10.2R1	N.S.
802.1X authentication with authentication bypass	10.1R1	9.0R2	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	10.2R1	N.S.
802.1X authentication with VLAN assignment, VoIP VLAN support	10.1R1	9.0R1	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	10.3R1	N.S.
802.1X user-based dynamic firewall filters	10.1R1	9.0R2	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	10.3R1	N.S.

Table 178: Access Control Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
802.1X user-based dynamic firewall filters on multiple-suplicant ports	10.1R1	9.5R2	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	10.3R1	N.S.
802.1X per-user statistics	10.1R1	9.2R1	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R1	10.3R1	N.S.
Authentication fallback	11.3R1	10.3R1	12.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	N.S.	N.S.
Captive portal authentication for Layer 3 interfaces	11.3R1	10.1R1	N.S.	13.2X50 - D10	12.1R1	12.2R1	N.S.	N.S.	N.S.
Captive portal authentication for Layer 2 interfaces	11.3R1	10.3R1	12.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	N.S.	N.S.
Energy Efficient Ethernet (EEE)	N.S.	N.S.	12.2R1	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.
Infranet Controller (IC) as an external captive-portal server	12.2R1	12.2R1	12.2R1	N.S.	12.2R1	12.2R1	N.S.	N.S.	N.S.
MAC RADIUS authentication	10.1R1	9.3R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	10.3R1	N.S.
NetBIOS snooping	11.3R5	11.1R1	11.3R5	N.S.	N.S.	N.S.	11.3R5	11.1R1	N.S.
Server fail fallback	10.1R1	9.3R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	N.S.	10.2R1	N.S.
TACACS+	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2

Table 179: Administration Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
System logging (syslog) over IPv4	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2
System logging (syslog) over IPv6	10.3R1	9.3R2	11.3R1	13.2X50 - D10	10.4R1	12.2R1	11.3R2	10.1R1	12.3R2
System snapshot	N.S.	10.0R1	N.S.	13.2X50 - D10	10.2R1	12.2R1	N.S.	10.0R1	12.3R2

Table 180: Class-of-Service (CoS) Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Class of service (CoS)—Class-based queuing with prioritization, Layer 2 and Layer 3 classification, rewrite, and queuing; strict priority queuing on egress	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	See Table 63 for a list of EX9200 CoS features
CoS—DSCP, IEEE 802.1p, and IP precedence packet rewrites on RVIs or IRB interfaces	N.S.	9.5R1	11.3R1	N.S.	10.2R1	12.2R1	11.3R2	9.4R1	
CoS—Interface-specific classifiers on RVIs or IRB interfaces	N.S.	9.4R1	11.3R1	13.2X50 - D10	11.3R1	12.2R1	N.S.	10.2R1	
CoS—Multidestination	—	—	—	—	—	—	N.S.	9.5R1	
CoS—Per-interface classification	N.S.	9.3R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	10.2R1	
CoS support on link aggregation groups (LAGs)	10.1R1	9.2R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	N.S.	9.4R1	
CoS support on RVIs or IRB interfaces	10.3R1	9.4R1	11.3R1	13.2X50 - D10 (classifiers only; rewrites N.S.)	10.4R1	12.2R1	N.S.	9.4R1	
DSCP setting on ingress interface	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	12.3R1	
Flexible CoS-outer 802.1p marking	N.S.	9.6R1	12.3R1	N.S.	12.1R1	12.2R1	N.S.	N.S.	
Interface-specific CoS rewrite rules	10.3R1	9.4R1	N.S.		11.2R1	12.2R1	N.S.	10.2R1	

Table 180: Class-of-Service (CoS) Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
				13.2X50 - D10 (for Layer 3 interfaces; IRB interfaces and Layer 3 subinterfaces N.S.)					
Junos EZQoS for CoS	10.1R1	9.3R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Port shaping and queue shaping	10.1R1	9.3R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	10.1R1	
Re-marking of bridged packets	11.2R1	9.4R1	N.S.	13.2X50 - D10	10.2R1	12.2R1	11.3R2	10.2R1	
Shaped-deficit weighted round-robin (SDWRR)	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Single-rate two-color marking	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	

Table 181: Class-of-Service (CoS) Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Assigning forwarding class and DSCP value for Routing Engine generated traffic	12.3R2
BA classification for VPLS based on IEEE 802.1p bits	12.3R2
Classification—Associate packets with CoS servicing levels. Types of classification: <ul style="list-style-type: none"> • Behavior aggregate (BA)—Operates on packets as they enter the switch • Multifield classification—Examines multiple fields in packets. • Fixed classification—Associate a forwarding class with a packet regardless of packet contents. 	12.3R2
Classification and DSCP marking of distributed protocol handler traffic	12.3R2
Classification of control-plane traffic	12.3R2
CoS classification and rewrite for IRB and Layer 2 interfaces and for other Layer 3 interfaces. Port-level queuing, scheduling, and shaping are supported.	12.3R2

Table 181: Class-of-Service (CoS) Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Egress GRE classification based on DSCP	12.3R2
IEEE 802.1p inheritance push and swap from transparent tags configuration	12.3R2
Elevated packet drops during oversubscription	12.3R2
Layer 2 policers for the ingress and egress interfaces. Policer types: <ul style="list-style-type: none"> • Single-rate two-color • Single-rate three-color (color-blind and color-aware) • Two-rate three-color (color-blind and color-aware) 	12.3R2
Independent values for DSCP and EXP bits	12.3R2
Apply CoS schedulers on ingress interfaces	12.3R2
Ingress DSCP bits for multicast traffic over Layer 3 VPNs	12.3R2
Layer 2 traffic policing	12.3R2
Policer support for aggregated Ethernet bundles (link aggregation groups, or LAGs)	12.3R2
Queuing support for logical tunnel interfaces	12.3R2
Rate-limit and excess rate or excess priority option	12.3R2
Re-marking of MVPN GRE encapsulation DSCP at ASBR	12.3R2
Scheduling	12.3R2
Set IPv6 DSCP and MPLS EXP independently	12.3R2
Set IPv6 DiffServ code point (DSCP) and MPLS EXP independently	12.3R2
Support for Layer 2 policers at the VLAN level	12.3R2
Support for applying a transmit rate limit to logical interfaces on Type 1, 2, or 3 Multiservices PICs	12.3R2
Support for configuring ToS rewrite rules so that DSCP bits of GRE packets are consistent with service providers' CoS policy	12.3R2
Support for copying the TOS bits to the outer IP header on GRE tunnel traffic sent by the Routing Engine	12.3R2
Support for setting the forwarding class and DSCP value for traffic generated by the Routing Engine	12.3R2
Unified command to display all CoS statistics	12.3R2

Table 182: Converged Networks (LAN and SAN) Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
NOTE: The EX4500 switch models that support Fibre Channel over Ethernet features must be Converged Enhanced Ethernet (CEE) capable. The CEE-capable EX4500 switch models have a “-C” in the hardware model number. See <i>EX4500 Switch Models</i> .									
Data Center Bridging Capability Exchange protocol (DCBX)	N.S.	N.S.	N.S.	N.S.	11.3R1	12.2R1	N.S.	N.S.	N.S.
DCBX application protocol TLV exchange	N.S.	N.S.	N.S.	N.S.	12.1R1	12.2R1	N.S.	N.S.	N.S.
FIP snooping	N.S.	N.S.	N.S.	N.S.	10.4R1	12.2R1	N.S.	N.S.	N.S.
Priority-based flow control (PFC)	N.S.	N.S.	N.S.	N.S.	10.4R1	12.2R1	N.S.	N.S.	N.S.

Table 183: Device Security Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Automatic recovery for port error disable conditions	10.1R1	9.6R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	10.0R1	N.S.
Storm control (broadcast and unicast)	10.1R1	9.1R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	N.S.
Storm control (multicast)	10.3R2	10.3R2	N.S.	13.2X50 - D10	10.3R2	12.2R1	11.3R2	10.3R2	N.S.
Unknown Layer 2 unicast forwarding	10.1R1	9.3R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	N.S.	10.0R1	12.3R2

Table 184: High Availability and Resiliency Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
NOTE: For complete lists of Virtual Chassis features, see <i>EX Series Virtual Chassis Software Features Overview</i> .									See Table 3 for a list of EX9200 HA features.
Graceful protocol restart for BGP	–	9.0R2	N.S.	13.2X50 - D10	–	–	11.3R2	9.4R1	
Graceful protocol restart for IS-IS	–	9.3R2	N.S.	13.2X50 - D10	–	–	11.3R2	9.4R1	
Graceful protocol restart for OSPF	–	9.0R2	N.S.	13.2X50 - D10	–	–	11.3R2	9.4R1	
Graceful protocol restart for RSVP and LDP	N.S.	N.S.	N.S.	13.2X50 - D10	N.S.	N.S.	N.S.	12.3R1	
GRES for ARP entries, forwarding database, and Layer 3 protocols	–	9.2R1 (Virtual Chassis only)	11.3R1	13.2X50 - D10	11.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	11.3R2	9.4R1	
GRES for IGMP snooping	–	11.3R1 (Virtual Chassis only)	12.1R1 (Virtual Chassis only)	13.2X50 - D10	11.4R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	N.S.	11.3R1	
GRES for LACP	N.S.	11.3R1	N.S.	13.2X50 - D10	11.3R1	12.2R1	N.S.	11.3R1	
GRES for Layer 2 and Layer 3 VPN LSPs	N.S.	N.S.	N.S.	13.2X50 - D10	N.S.	N.S.	N.S.	12.3R1	
GRES for port security (DHCP snooping, DAI, and IP source guard)	–	9.2R1 (Virtual Chassis only)	N.S.	13.2X50 - D10	–	–	11.3R2	9.6R1	
LACP support for dual-homing applications in data centers	10.1R1	10.0R1	N.S.	13.2X50 - D10	10.2R1	12.2R1	N.S.	10.0R1	
Link Aggregation Control Protocol (LACP)	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Link aggregation groups (LAGs)	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	

Table 184: High Availability and Resiliency Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Nonstop active routing (NSR) for BGP, IS-IS, IGMP with BFD, and RIP	—	11.1R1 (Virtual Chassis only)	12.1R1 (Virtual Chassis only)	N.S.	11.3R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	11.3R2	11.1R1	
Nonstop active routing (NSR) for IPv6 IS-IS, RIPng, and OSPFv3 with BFD	—	12.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	N.S.	12.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	N.S.	11.2R1	
Nonstop active routing (NSR) for OSPFv2	—	11.1R1	12.1R1 (Virtual Chassis only)	N.S.	11.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	11.3R2	10.4R1	
Nonstop active routing (NSR) for Protocol Independent Multicast (PIM)	N.S.	12.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	N.S.	12.2R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	N.S.	11.4R2	
Nonstop active routing (NSR) for RSVP and LDP	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	12.3R1	
Nonstop bridging (NSB) for LAGs and LACP	—	11.4R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	N.S.	11.4R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	12.1R1	11.3R1	
Nonstop bridging (NSB) for LLDP and LLDP-MED	—	11.3R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	N.S.	N.S.	N.S.	N.S.	11.3R1	
Nonstop bridging (NSB) for spanning-tree protocols	—	11.3R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	N.S.	12.1R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	12.1R1	11.3R1	
Nonstop software upgrade (NSSU)	—	12.1R1 (Virtual Chassis only)	12.2R1 (Virtual Chassis only)	N.S.	12.1R1 (Virtual Chassis only)	12.1R1 (Virtual Chassis only)	12.2R1	10.4R1	
Power budget management	—	—	N.S.	13.2X50 - D10	—	—	11.3R2	10.2R1	
	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	11.4R1	

Table 184: High Availability and Resiliency Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Virtual Router—Network Time Protocol (NTP), system logging, Simple Network Management Protocol (SNMP), RADIUS, and TACACS support in a virtual router									
Virtual Router Redundancy Protocol (VRRP)	12.3R1	9.0R2	12.1R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Virtual Router Redundancy Protocol (VRRP)—Support for multiple VRRP owners per physical interface	12.3R1	12.3R1	12.3R1	N.S.	12.3R1	12.3R1	N.S.	12.3R1	
Virtual Router Redundancy Protocol (VRRP) for IPv6 (except authentication type and authentication key)	N.S.	10.2R1	12.3R1	13.2X50 - D10	11.2R1	12.2R1	12.1R1	10.1R1	

Table 185: High Availability and Resiliency Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Graceful Routing Engine switchover (GRES)	12.3R2
Nonstop active routing (NSR) support for Protocol Independent Multicast (PIM) for IPv4 and IPv6	12.3R2
Nonstop active routing (NSR) support for VPLS and for LDP-based VPLS	12.3R2
Nonstop active routing (NSR) support for LDP OAM features	12.3R2
Nonstop active routing (NSR) support for Layer 2 VPNs	12.3R2
Unified ISSU (requires EX9200-40T or EX9200-40F line cards)	12.3R3

Table 185: High Availability and Resiliency Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Virtual Router Redundancy Protocol version 3 (VRRPv3)	12.3R2

Table 186: Interfaces Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Digital optical monitoring (DOM)	N.S.	10.0R1	11.3R1	13.2X50 - D10	N.S.	N.S.	N.S.	10.0R1	See Table 186 for a list of EX9200 interfaces features.
Interface ranges	10.1R1	10.0R1	11.3R1	13.2X50 - D10	10.2R	12.2R1	11.3R2	10.1R1	
IPv4 over generic routing encapsulation (GRE) tunnels—encapsulation support	N.S.	12.1R1	N.S.	N.S.	N.S.	N.S.	N.S.	12.1R1	
IPv4 over generic routing encapsulation (GRE) tunnels—de-encapsulation support	N.S.	12.1R1	N.S.	N.S.	N.S.	N.S.	N.S.	12.3R1	
IPv6 over generic routing encapsulation (GRE) tunnels using IPv4 transport—encapsulation support	N.S.	12.1R1	N.S.	N.S.	N.S.	N.S.	N.S.	12.1R1	
IPv6 over generic routing encapsulation (GRE) tunnels using IPv4 transport—de-encapsulation support	N.S.	12.1R1	N.S.	N.S.	N.S.	N.S.	N.S.	12.3R1	
IP directed broadcast	11.3R1	9.4R1	12.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Time domain reflectometry (TDR)	10.1R1	9.0R2	11.3R1	13.2X50 - D10	N.S.	N.S.	11.3R2	9.4R1	
Unicast reverse-path forwarding (RPF)	N.S.	9.3R2	12.3R1	13.2X50 - D10	11.2R1	12.2R1	11.3R2	10.1R1	
	N.S.	9.2R1	N.S.	13.2X50 - D10	11.2R1	12.2R1	11.3R2	9.4R1	

Table 186: Interfaces Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
VLAN-tagged Layer 3 subinterfaces									

Table 187: Interfaces Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
ICMP redirect	12.3R2
Clear MAC address information	12.3R2
IPv6 subnet support on loopback interfaces	12.3R2
IPv6 support for unnumbered Ethernet interfaces	12.3R2
Multichassis link aggregation (MC-LAG)	12.3R2
Nonstop active routing (NSR) support for Bidirectional Forwarding Detection (BFD)	12.3R2
Protection against distributed denial-of-service (DDOS) attacks	12.3R2
Software support for IPv4 to IPv6 transition	12.3R2
Static mapping for port forwarding	12.3R2
Support for active monitoring on logical systems	12.3R2
Support for VRF in Routing Engine-based sampling	12.3R2
Support for integrated routing and bridging (IRB) MAC synchronization in MC-LAG for aggregated Ethernet	12.3R2
Targeted broadcast for virtual routing and forwarding (VRF) instances	12.3R2
Trunk interface enhancements: <ul style="list-style-type: none"> Configure a single logical trunk interface to support a list of VLANs or to accept packets with no VLAN tag. Configure multiple logical trunk interfaces on a single physical interface. 	12.3R2
Unicast reverse-path forwarding (RPF) loose mode, with ability to discard packets with source addresses pointing to the discard interface	12.3R2
Unnumbered Ethernet—Configure IPv4 processing on interfaces without assigning explicit IP addresses to the interfaces.	12.3R2
VLAN rewrite operations on incoming and outgoing frames	12.3R2

Table 188: IP Address Management Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
DHCP server and relay with option 82 for Layer 2 VLANs	10.1R1	9.3R2	11.3R1	N.S.	10.2R1	12.2R1	11.3R2	9.4R1	See Table 189 for a list of EX9200 IP address management features.
DHCP server and relay with option 82 for Layer 3 interfaces	10.1R1	9.0R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
DNS for IPv6	N.S.	9.3R2	N.S.	13.2X50 - D10	N.S.	N.S.	N.S.	N.S.	
Local DHCP server	10.1R1	9.3R2	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	9.4R1	
Virtual router aware DHCP (VR-aware DHCP)	N.S.	12.3R1	N.S.	13.2X50 - D10	12.3R1	12.3R1	N.S.	12.3R1	
Virtual router aware DHCPv6 (VR-aware DHCPv6)	N.S.	N.S.	N.S.	13.2X50 - D10	12.3R1	12.3R1	12.3R1	N.S.	
Static addresses	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	N.S.	9.4R1	

Table 189: IP Address Management Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
DHCP server and relay	12.3R2
DHCPv6 local server	12.3R2
DHCPv6 support	12.3R2
Distinguishing DHCP duplicate clients by subinterface	12.3R2
Dynamic reconfiguration of extended DHCP and DHCPv6 local server clients	12.3R2
Dynamic IPv6 filters	12.3R2
Expression support for dynamic profiles	12.3R2
Extended DHCP relay proxy	12.3R2

Table 189: IP Address Management Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Optional disabling of automatic ARP table population	12.3R2
IPv6 address assignment pools	12.3R2
Overriding DHCP settings on specific interfaces	12.3R2
Per-interface DHCP tracing operations	12.3R2
S-VLAN-based shaping for dynamic profiles	12.3R2
Sending a DHCP relay and relay proxy release message	12.3R2
Specifying the DHCP source address used for IP packets	12.3R2
Support for MAC address validation	12.3R2
Support for address pool threshold traps	12.3R2
Address assignment pools	12.3R2
Per-interface DHCP lease limits	12.3R2

Table 190: IPv6 Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
NOTE: A separate software license is required for IPv6. See Understanding Software Licenses for EX Series Switches .									
Application identification (APPID) for IPv6 packets	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	12.3R2
BFD for IPv6	N.S.	12.3R1	N.S.	13.2X50 - D10	12.3R1	12.3R1	N.S.	12.3R1	12.3R2 (also for static routes)
BGP for IPv6	N.S.	9.4R1	12.3R1	13.2X50 - D10	11.1R1	12.2R1	12.1R1	10.1R1	12.3R2
IPv6 CoS (multifield classification and rewrite)	N.S.	10.2R1	12.3R1	13.2X50 - D10	12.1R1	12.2R1	12.1R1	10.4R1	12.3R2

Table 190: IPv6 Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
IPv6 management	10.3R1	9.3R2 (using loopback addresses only)	11.3R1	13.2X50 - D10	10.4R1	12.2R1	N.S.	10.1R1	12.3R2
IPv6 multicast protocols (PIM, MLDv1/v2)	N.S.	10.1R1	12.3R1	13.2X50 - D10	11.2R1	12.2R1	12.1R1	10.2R1	12.3R2
IPv6 neighbor redirect compliance with RFC 4861	12.3R1	12.3R1	12.3R1	13.2X50 - D10	12.3R1	12.3R1	12.3R1	12.3R1	12.3R2
IPv6 path MTU discovery	10.3R1	9.3R1	12.3R1	13.2X50 - D10	10.4R1	12.2R1	N.S.	10.2R1	12.3R2
IS-IS for IPv6	N.S.	9.4R1	N.S.	13.2X50 - D10	11.2R1	12.2R1	12.1R1	10.1R1	12.3R2
MBGP for IPv6	N.S.	9.3R1	12.3R1	13.2X50 - D10	N.S.	N.S.	12.1R1	10.1R1	12.3R2
OSPFv3	N.S.	9.3R1	12.3R1	13.2X50 - D10	11.1R1	12.2R1	12.1R1	10.1R1	12.3R2
RFC 4291 compliance	12.3R1 See note at end of table	12.3R1 See note at end of table	12.3R1 See note at end of table	N.S.	12.3R1 See note at end of table	12.3R1 See note at end of table	12.3R1 See note at end of table	12.3R1 See note at end of table	12.3R2 See note at end of table
RIPng	N.S.	9.3R1	12.3R1	13.2X50 - D10	11.1R1	12.2R1	12.1R1	10.1R1	12.3R2
VRRPv3 (RFC 5798 compliance, ability to send SNMP traps)	N.S.	12.3R1	N.S.	13.2X50 - D10	12.3R1	12.3R1	N.S.	12.3R1	12.3R2



NOTE: Compliance with RFC 4291

EX Series switches drop the following types of illegal IPv6 packets:

- Packets that have a link-local source or destination address. Because link-local addresses are intended to be used for addressing only on a single link, EX Series switches do not forward any packets with such addresses to other links.
- Packets with the IPv6 unspecified source address 0:0:0:0:0:0:0:0.
- Packets that are to be sent outside a node but have the IPv6 loopback address 0:0:0:0:0:0:0:1 as the source address. When IPv6 packets are received on an interface, EX Series switches drop packets that have the loopback address as the destination address.

EX Series switches do not support Subnet-Router Anycast address.

Table 191: Layer 2 Network Protocols Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
802.1Q VLAN tagging	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	See Table 82 for a list of EX9200 Layer 2 networking protocols features.
Edge virtual bridging (EVB) support with virtual Ethernet port aggregator (VEPA)	N.S.	N.S.	N.S.	N.S.	12.1R1	12.2R1	N.S.	12.1R1	
Ethernet ring protection switching (ERPS, G.8032/Y.1344)	12.1R1	12.1R1	12.3R1	13.2X50 - D10	12.3R1	12.3R1	N.S.	12.3R1	
Layer 2 protocol tunneling (L2PT)	11.1R1	10.0R1	12.3R1	N.S.	11.2R1	12.2R1	12.1R1	N.S.	
Link Layer Discovery Protocol (LLDP)	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Link Layer Discovery Protocol—Media Endpoint Discovery (LLDP-MED) with voice over IP (VoIP) integration	10.1R1	9.0R2	11.3R1	13.2X50 - D10	N.S.	N.S.	N.S.	N.S.	
MAC-based VLANs	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Multiple VLAN Registration Protocol (MVRP, IEEE 802.1ak)	11.3R1	10.0R1	12.3R1	13.2X50 - D10	11.2R1	12.2R1	12.1R1	10.0R1	
Private VLANs (PVLANS)	11.1R1	9.3R2	12.1R1	N.S.	11.2R1	12.2R1	11.3R2	10.1R1	
Private VLANs (PVLANS) support across switches	11.1R1	10.4R1	12.1R1	N.S.	11.2R1	12.2R1	11.3R2	11.2R1	
Proxy ARP—Restricted	10.1R1	10.0R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	N.S.	10.0R1	

Table 191: Layer 2 Network Protocols Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Proxy ARP—Unrestricted	10.1R1	9.6R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	12.1R1	10.0R1	
Proxy ARP per VLAN	10.1R1	10.1R1	N.S.	13.2X50 - D10	10.2R1	12.2R1	N.S.	10.1R1	
Q-in-Q tunneling	11.1R1	9.3R2	11.4R1	N.S.	11.2R1	12.2R1	12.1R1	11.1R1	
Q-in-Q VLAN extended support for multiple S-VLANs per access interface, firewall-filter-based VLAN assignment, and RVIs or IRB interfaces	N.S.	9.6R1	12.3R1	N.S.	11.2R1	12.2R1	12.1R1	11.1R1	
Redundant trunk groups	10.1R1	9.0R2	11.3R1	13.2X50 - D10	11.2R1	12.2R1	11.3R2	9.4R1	
Routed VLAN interfaces (RVIs)—Also known as integrated routing and bridging (IRB) interfaces	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
VLAN ID translation	11.1R1	10.0R1	N.S.	N.S.	11.2R1	12.2R1	N.S.	11.1R1	
VLAN ranges	10.1R1	9.2R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	

Table 192: Layer 2 Networking Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
VLANs and virtual switching	12.3R2
DHCP support for integrated routing and bridging (IRB)	12.3R2
MC-LAG support for IGMP snooping in IRB	12.3R2
Hash-key load-balancing support for Layer 3 and Layer 4 fields	12.3R2

Table 192: Layer 2 Networking Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
IP multicast over Layer 2 trunk port support	12.3R2
Integrated routing and bridging (IRB)	12.3R2
Layer 2 Ethernet OAM: <ul style="list-style-type: none"> Distributed periodic packet management process (ppmd) for improved scaling Graceful Routing Engine switchover (GRES) Remote defect indication (RDI) Configuration of action profiles 	12.3R2
Layer 2 address learning in logical systems	12.3R2
Layer 2 forwarding support for bridging and VPLS	12.3R2
Layer 2 policer statistics MIB	12.3R2
Firewall filter match conditions for Layer 2 bridging and VPLS	12.3R2
Next-hop groups using either IP addresses or Layer 2 addresses for the next hop	12.3R2
Unicast reverse-path forwarding (RPF) loose mode, with ability to discard packets with source addresses pointing to the discard interface	12.3R2
Spanning-tree protocols support for Layer 2 bridging and VPLS	12.3R2
VLAN rewrite operations on incoming and outgoing frames	12.3R2
STP root guard (root protection)	12.3R2
Support for Layer 2 and Layer 2.5 features: <ul style="list-style-type: none"> Extensive set of Layer 2 label-manipulation capabilities, Q-in-Q support MC-LAG active / standby and MC-LAG active / active xSTP protocol support Integrated Routing and Bridging (IRB) interface support IGMP snooping for multichassis link aggregation group (MC-LAG) interfaces Configurable label block sizes for VPLS Connectivity fault management process flooding to interfaces based on mesh groups Layer 2 address learning in logical systems Virtual switch support, providing virtual Layer 2 switch instances with separate Layer 2 learning domains, isolated 4K VLAN ID spaces, and STP instances Ethernet Ring Protocol (ERP) for multiple ring instances on the same physical ring Transit and bypass static label-switched paths (LSPs) Layer 2 Gigabit Ethernet logical interface policing Static LSP statistics Multiple VLAN Registration Protocol (MVRP)—IEEE 802.1ak-2007 	12.3R2

Table 192: Layer 2 Networking Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
VPLS root protection topology change-action control	12.3R2
VLAN ranges	12.3R2
Q-in-Q VLAN extended support for multiple S-VLANs per access interface, firewall-filter-based VLAN assignment, and RVIs or IRB interfaces	12.3R2
Q-in-Q tunneling	12.3R2
Proxy ARP—Unrestricted and restricted	12.3R2
MAC-based VLANs	12.3R2
Link Layer Discovery Protocol (LLDP)	12.3R2
Layer 2 protocol tunneling (L2PT)	12.3R2

Table 193: Layer 3 Protocols Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Bidirectional Forwarding Detection (BFD)	11.3R1	9.0R2	12.1R1	13.2X50 - D10	10.2R1	12.2R1	12.1R1	9.4R1	See Table 193 for a list of EX9200 Layer 3 protocols features.
Border Gateway Protocol (BGP)	N.S.	9.0R2	12.1R1	13.2X50 - D10	11.1R1	12.2R1	11.3R2	9.4R1	
Multiprotocol Border Gateway Protocol (MBGP)	N.S.	9.3R1	12.3R1	13.2X50 - D10	11.2R1	12.2R1	12.1R1	9.4R1	

A separate software license is required for BGP and MBGP. See [Understanding Software Licenses for EX Series Switches](#).

Table 193: Layer 3 Protocols Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Distributed periodic packet management (PPM) with BFD	N.S.	10.4R1	N.S.	13.2X50 - D10	N.S.	N.S.	12.1R1	10.4R1	
Distributed periodic packet management (PPM) with LACP	N.S.	10.2R1	N.S.	13.2X50 - D10	11.1R1	12.2R1	11.3R2	10.2R1	
Filter-based forwarding	N.S.	9.4R1	12.3R1	13.2X50 - D10	11.2R1	12.2R1	11.3R2	9.6R1	
Filter-based forwarding over IPv6	N.S.	10.1R1	N.S.	N.S.	N.S.	N.S.	N.S.	10.3R1	
Intermediate System-to-Intermediate System (IS-IS)	N.S.	9.0R2	N.S.	13.2X50 - D10	11.1R1	12.2R1	11.3R2	9.4R1	
A separate software license is required for IS-IS. See Understanding Software Licenses for EX Series Switches .									
IPv6 Layer 3 multicast protocols	N.S.	10.1R1	N.S.	13.2X50 - D10	N.S.	N.S.	N.S.	10.2R1	
Jumbo frames on RVIs or IRB interfaces	N.S.	9.4R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	N.S.	9.4R1	
OSPF Multitopology Routing (MT-OSPF)	N.S.	9.5R1	N.S.	13.2X50 - D10	N.S.	N.S.	N.S.	N.S.	
See the Junos OS Routing Protocols Configuration Guide .									

Table 193: Layer 3 Protocols Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
OSPFv2	11.1R1	9.0R2	11.4R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
OSPFv3 IPsec support	N.S.	10.3R1	N.S.	13.2X50 - D10	N.S.	N.S.	N.S.	N.S.	
Routing Information Protocol version 1 (RIPv1) and RIPv2	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Static routes	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Virtual routing and forwarding (VRF) with IPv4—Virtual routing instances	12.3R1	9.2R1	12.3R1	13.2X50 - D10	11.1R1	12.2R1	11.3R2	9.6R1	
VRF with IPv4—Virtual routing instances for PIM and IGMP	N.S.	10.0R1	N.S.	13.2X50 - D10	11.1R1	12.2R1	11.3R2	10.0R1	
VRF with IPv4—Virtual routing instances for IGMP snooping	N.S.	11.4R1	N.S.	13.2X50 - D10	12.1R1	12.2R1	N.S.	11.3R1	
VRF with IPv6—Virtual routing instances for multicast traffic	N.S.	10.1R1	N.S.	13.2X50 - D10	N.S.	N.S.	N.S.	10.1R1	
VRF with IPv6—Virtual routing instances for unicast traffic	N.S.	10.1R1	12.3R1	13.2X50 - D10	N.S.	N.S.	N.S.	10.1R1	

Table 194: Layer 3 Protocols Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Accumulated IGP attribute for BGP	12.3R2
Advertisement of the best external BGP path to internal peers	12.3R2

Table 194: Layer 3 Protocols Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Alias support for local autonomous system numbers for BGP	12.3R2
BFD liveness detection	12.3R2
BFD protocol support for OSPFv3	12.3R2
BGP remote next-hop support for single-hop EBGP peers	12.3R2
BGP support for 4-byte autonomous system numbers	12.3R2
BGP support for MDT-SAFI updates without a route target	12.3R2
Behavior change for BGP-independent autonomous system (AS) domains	12.3R2
Bidirectional Forwarding Detection (BFD) hold-down timer	12.3R2
Distributed periodic packet management support for aggregate interfaces	12.3R2
Egress filtering PIMv4/v6 join messages	12.3R2
For internal BGP (IBGP), advertise multiple paths to a destination	12.3R2
Frequent BGP keepalive messages and short BGP hold time	12.3R2
Hitless authentication key rollover for IS-IS	12.3R2
Hub-and-spoke support for multiprotocol BGP-based multicast VPNs with PIM-SSM GRE S-PMSI transport	12.3R2
IPv4 subnet support on loopback interfaces	12.3R2
IS-IS hold-down timer for subsequent SPF calculations	12.3R2
Keepalive support for GRE interfaces	12.3R2
Multitopology routing (MTR)	12.3R2
Nonstop active routing (NSR) support for the Routing Information Protocol (RIP) and RIP next generation (RIPng)	12.3R2
Nonstop active routing (NSR) support	12.3R2
OSPF graceful restart enhancement	12.3R2
OSPF hold-down timer for subsequent SPF calculations	12.3R2
Only the system log notes failure to add routes	12.3R2

Table 194: Layer 3 Protocols Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Origin validation for BGP	12.3R2
PIM join suppression support	12.3R2
Priority assignment for prefixes in OSPF import policies	12.3R2
Reduction in flooding of self-originated OSPF LSAs	12.3R2
Support for BFD over multihop static routes	12.3R2
Support for BFD on logical switches	12.3R2
Support for IPSec authentication for OSPFv2	12.3R2
Support for OSPF database protection for OSPF and OSPFv3	12.3R2
Support for OSPF export and import policies for network-summary LSAs	12.3R2
Support for alternate loop-free routes for IS-IS and OSPF	12.3R2
Support for clearing the VPN tag	12.3R2
Support for disabling the attribute set messages on independent AS domains for BGP loop detection	12.3R2
Support for disabling traps for passive OSPFv2 interfaces	12.3R2
Support for display of flood next-hop branch overflow condition	12.3R2
Support for dropping and ignoring path attributes during BGP neighbor updates	12.3R2
Support for the algorithm that determines the single best path to skip the step that evaluates an AS path	12.3R2
Support for limiting the number of prefixes accepted from a BGP peer	12.3R2
Support for multiarea adjacency in OSPFv2	12.3R2
Support for multiple address families in OSPFv3	12.3R2
Support for route leaking when the switch is in overload mode	12.3R2
Support for route-filter-based BGP outbound route filtering	12.3R2
Support for the BGP Monitoring Protocol	12.3R2
Support to hold down BGP peering sessions after a nonstop active routing (NSR) switchover Timer to delay MED updates for routes advertised by BGP groups or peers configured with the metric-out igp statement Virtual Router Redundancy Protocol (VRRP)	12.3R2

Table 194: Layer 3 Protocols Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Timer to delay MED updates for routes advertised by BGP groups or peers configured with the metric-out igp statement	12.3R2
Virtual Router Redundancy Protocol (VRRP)	12.3R2

Table 195: Logical Systems Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
A separate software license is required for logical systems. See Understanding Software Licenses for EX Series Switches .	
Logical systems	
Layer 2 address learning in logical systems	12.3R2

Table 196: MPLS Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
A separate software license is required for MPLS. See Understanding Software Licenses for EX Series Switches .									

Table 196: MPLS Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Aggregated Ethernet interfaces (LAGs) on circuit cross-connects (CCCs)	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	See Table 192 for a list of EX9200 MPLS features.
BFD for an LDP-based LSP	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	12.2R1	
BFD for an RSVP-based LSP	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	12.2R1	
CCC between 2 interfaces in the same switch	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
Interior gateway protocol (IGP) IS-IS and OSPF shortcuts	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
IP over MPLS	N.S.	10.1R1	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
					See Note at end of table	See Note at end of table			
IPv6 over MPLS label-switched paths (LSPs)	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	12.1R1	
					See Note at end of table	See Note at end of table			
LDP-based MPLS	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
LDP tunneling (LDP over RSVP)	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
MPLS-based circuit cross-connects (CCC)	N.S.	9.5R1	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
MPLS label-switched router (LSR) support	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
	N.S.	9.5R1	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	

Table 196: MPLS Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
MPLS Layer 2 CCC on Ethernet- encapsulated interfaces (RFC 6624)									
MPLS Layer 2 CCC on VLAN-encapsulated interfaces (RFC 4905)	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
MPLS Layer 2 VLAN CCC on Ethernet- encapsulated interfaces (RFC 6624)	N.S.	9.5R1	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.3R1	
MPLS Layer 2 VLAN CCC on VLAN-encapsulated interfaces (RFC 4905)	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.3R1	
MPLS Layer 2 VPN over CCC	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
MPLS Layer 2 VPN over VLAN CCC	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.3R1	
MPLS OAM-LSP ping	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	11.1R1	
MPLS over untagged Layer 3 interfaces	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
MPLS with class of service (CoS)	N.S.	9.5R1	N.S.	N.S.	12.2R5	12.2R5	N.S.	12.1R1	
MPLS Layer 3 VPNs	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
MPLS with RSVP-based label-switched paths (LSPs)	N.S.	9.5R1	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
	N.S.	N.S.	N.S.	N.S.			N.S.	12.1R1	

Table 196: MPLS Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Layer 3 subinterfaces as MPLS core interfaces					12.2R1 See Note at end of table.	12.2R1 See Note at end of table.			
Routed VLAN interfaces (RVIs) as MPLS core interfaces	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	12.1R1	
Path maximum transmission unit (MTU) and unicast reverse-path forwarding (RPF) checks for VPNs	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
Resource Reservation Protocol—traffic engineering (RSVP-TE)	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
Standby secondary path protection	N.S.	12.1R1	N.S.	N.S.	12.2R1	12.2R1	N.S.	11.1R1	
Static label-switched paths (LSPs)	N.S.	N.S.	N.S.	N.S.	12.2R1	12.2R1	N.S.	12.1R1	



NOTE: Layer 3 subinterfaces as MPLS core interfaces—For EX4500 and EX4550 switches to support Layer 3 subinterfaces as MPLS core interfaces, the peer switch that the Layer 3 subinterfaces connect to, must be an EX8200 switch.

IP over MPLS—The EX4500 and EX4550 switches do not support IP over MPLS (single MPLS label in the packet) when the switch is positioned as a non-penultimate-hop popping (non-PHP) switch.

Table 197: MPLS Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
---------	------------------

A separate software license is required for MPLS. See [Understanding Software Licenses for EX Series Switches](#).

Table 197: MPLS Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Bypass static LSPs	12.3R2
LDP LSP action based on a BFD failure event	12.3R2
LDP downstream on demand	12.3R2
LDP, BGP, and VPLS interworking	12.3R2
P2MP LSP traceroute	12.3R2
Static LSP: <ul style="list-style-type: none"> • Revert timer • Statistics • Traceoptions • At the ingress switch • At the transit switch 	12.3R2
Statistics for P2MP LSPs	12.3R2

Table 198: Multicast Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
IGMP (Internet Group Management Protocol) version 1 (IGMPv1) and IGMPv2	11.1R1	9.0R2	12.1R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	See Table 195 for a list of EX9200 multicast features
IGMP filtering	11.3R1	9.5R1	12.3R1	13.2X50 - D10	11.3R1	11.3R1	11.3R1	9.5R1	
IGMP snooping with RVIs or IRB interfaces	10.1R1	9.2R1	12.1R1	13.2X50 - D10	10.2R1	12.2R1	N.S.	9.4R1	
IGMPv3	11.1R1	9.3R2	12.1R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.6R1	
IGMPv1 and IGMPv2 snooping	10.1R1	9.1R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
IGMPv3 snooping	10.1R1	9.6R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.6R1	
Multicast Listener Discovery version 1 and 2 (MLDv1 and MLDv2)	N.S.	10.1R1	N.S.	13.2X50 - D10	11.2R1	12.2R1	12.1R1	10.2R1	
Multicast Listener Discovery version 1 (MLDv1) snooping (MLDv1 snooping)	12.1R1	12.1R1	12.1R1	N.S.	12.1R1	12.2R1	12.1R1	12.1R1	
Multicast Listener Discovery version 2 (MLDv2) snooping (MLDv2 snooping)	12.1R1	12.1R1	12.1R1	N.S.	12.1R1	12.2R1	12.1R1	12.1R1	
Multicast Source Discovery Protocol (MSDP)	N.S.	9.4R1	12.3R1	13.2X50 - D10	10.2R1	12.2R1	12.1R1	9.4R1	
See the Junos OS Multicast Protocols Configuration Guide .									
Multicast VLAN registration (MVR)	11.3R1	9.6R1	12.1R1	N.S.	N.S.	N.S.	N.S.	N.S.	
	11.1R1	9.2R1	12.1R1		11.2R1	12.2R1	11.3R2	9.4R1	

Table 198: Multicast Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Protocol Independent Multicast dense mode (PIM DM)				13.2X50 - D10					
See the Junos OS Multicast Protocols Configuration Guide .									
Protocol Independent Multicast sparse mode (PIM SM)	11.1R1	9.0R2	12.1R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
See the Junos OS Multicast Protocols Configuration Guide .									
Protocol Independent Multicast source-specific multicast (PIM SSM)	11.1R1	9.3R1	12.1R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
See the Junos OS Multicast Protocols Configuration Guide .									
Single-source multicast	N.S.	9.0R2	N.S.	13.2X50 - D10	N.S.	N.S.	N.S.	9.4R1	

Table 199: Multicast Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
BFD for PIM—IPv6	12.3R2
BFD support for ECMP LSPs signaled using LDP	12.3R2
Bidirectional PIM (RFC 5015)	12.3R2
Control of PIM resources for multicast VPNs	12.3R2
Disable PIM for IPv6 only	12.3R2
Dynamic reuse of data multicast distribution tree (MDT) group addresses	12.3R2
Flexible configuration for IGMP or MLD static-join	12.3R2
IGMPv3 and MLDv2 full support	12.3R2

Table 199: Multicast Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
IGMP and MLD enhancements— <ul style="list-style-type: none"> • immediate-leave (IGMP and MLD) • promiscuous-mode (IGMP only) 	12.3R2
IGMP and PIM support for unnumbered interfaces	12.3R2
IGMP join and leave recording for system or for specific interfaces	12.3R2
IGMP and MLD source or group access lists and MLD join and leave recording	12.3R2
IGMP and MLD support for dynamic interfaces	12.3R2
Independently configurable loopback addresses for VRF VPNs	12.3R2
Internet multicast using ingress replication provider tunnels	12.3R2
Software support for configuring accept any-source multicast (ASM) join messages (*;G) for group addresses	12.3R2
Software support for configuring a provider network to operate in source-specific multicast (SSM) mode	12.3R2
LDP signaling for point-to-multipoint LSPs in next-generation MBGP multicast VPNs	12.3R2
Load-balancing PIM join messages on multicast VPNs	12.3R2
Multicast flow maps	12.3R2
Nonstop active routing (NSR) PIM for Draft-Rosen VPNs	12.3R2
PIM automatic make-before-break (MBB) join load balancing	12.3R2
PIM join load balancing	12.3R2
Source-specific multicast (SSM)-map definition for different groups to different sources	12.3R2
Support for filtering unwanted PIM neighbors	12.3R2
Support for multicast output interface (OIF) mapping	12.3R2
Translation of PIM join/prune messages to IGMP or MLD report/leave messages	12.3R2
Turn off spanning-tree interface state in multicast snooping	12.3R2

Table 200: Network Management and Monitoring Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
802.1ag Ethernet OAM connectivity fault management (CFM)	11.2R1	10.2R1	12.3R1	13.2X50 - D10	12.2R1	12.2R1	N.S.	11.4R1	See Table 200-99 for a list of EX9200 network management and monitoring features.
Ethernet frame delay measurement (ETH-DM, Y.1731)	N.S.	11.4R1 (EX4200 only)	N.S.	13.2X50 - D10	11.4R1	12.2R1	N.S.	11.4R1	
Ethernet OAM link fault management (LFM—also known as Ethernet in the First Mile, EFM)	11.1R1	9.4R1	12.2R1	13.2X50 - D10	12.2R1	12.2R1	N.S.	10.0R1	
Port mirroring	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Port mirroring enhancements <ul style="list-style-type: none"> • Layer 3 interface support • Multiple VLAN support 	N.S.	9.5R1	N.S.	—	N.S.	N.S.	N.S.	9.5R1	
Port mirroring enhancements <ul style="list-style-type: none"> • For remote port mirroring, ingress and egress options on VLAN member interfaces on the intermediate (transit) switch to avoid flooding mirrored traffic to those interfaces 	N.S.	10.0R1	N.S.	—	N.S.	N.S.	N.S.	N.S.	
	N.S.	N.S.	N.S.	13.2X50 - D10	11.2R1	12.2R1	N.S.	N.S.	

Table 200: Network Management and Monitoring Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Port mirroring support for multiple analyzers per session									
Real-time performance monitoring (RPM)	10.1R1	9.3R2	12.2R1	13.2X50 - D10	10.2R1	12.2R1	12.1R1	10.1R1	
Real-time performance monitoring (RPM)—hardware timestamps with RVIs or IRB interfaces	N.S.	10.3R1	12.2R1	13.2X50 - D10	10.2R1	12.2R1	12.1R1	10.3R1	
Real-time performance monitoring (RPM)—client and server on same interface	10.3R1	10.3R1	12.2R1	13.2X50 - D10	11.1R1	12.2R1	N.S.	10.3R1	
Routing Engine Software Development Kit (SDK)	N.S.	12.2R1	12.2R1 (EX4200 only)	N.S.	12.2R1	12.2R1	N.S.	12.2R1	
RMON	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
sFlow monitoring technology	11.1R1	9.3R2	12.1R1	13.2X50 - D10	11.2R1	12.2R1	12.1R1	10.0R1	
sFlow monitoring technology—Persistent IP addresses for agent IDs and use in datagrams	11.1R1	10.2R1	N.S.	13.2X50 - D10	N.S.	N.S.	12.1R1	10.2R1	
Simple Network Management Protocol version 1 (SNMPv1), SNMPv2, and SNMPv3	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
	11.1R1	11.1R1	12.1R2		11.1R1	12.2R1	N.S.	12.1R1	

Table 200: Network Management and Monitoring Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Uplink failure detection				13.2X50 - D10					

Table 201: Network Management and Monitoring Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Junos OS XML API and scripting—NETCONF Java toolkit for rapid development of Java applications to manage devices running Junos OS	12.3R2
Junos OS XML API and scripting—NETCONF Perl client installation—Supports loading prerequisites from Comprehensive Perl Archive Network (CPAN) global repository	12.3R2
Junos OS XML API and scripting—NETCONF tracing operations	12.3R2
Junos OS XML API and scripting: <ul style="list-style-type: none"> • Dedicated directory for user script library • Global variable provided to Junos OS automation scripts • References to a correlating event in a policy action • Trigger a policy based on the event count • Unique filenames for uploaded files • Upload files created by event scripts • XML schemata for Junos OS XML operational tag elements • jcs:open() extension function support for routing instances 	12.3R2
Configuration options to filter out interfaces from SNMP Get and GetNext operations	12.3R2
Enhanced SNMP support for logical switches and routing instances	12.3R2
Generating SNMP traps when MAC address table is full	12.3R2
Junos OS MIB support for VPLS	12.3R2
MIB support for VRF route entries	12.3R2
Proxy SNMP agent	12.3R2
SNMP MIB support for OSPFv3	12.3R2
SNMP poll and trap support for DHCP leases	12.3R2
SNMP support for the DHCP bindings table	12.3R2
SNMP support for the authd daemon and for radius-acc-server-mib and radius-auth-server-mib	12.3R2

Table 201: Network Management and Monitoring Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
SNMP support for spanning-tree protocols	12.3R2
Support for Internet draft draft-ietf-bfd-mib-02.txt—MIB for BFD liveness detection	12.3R2
Support for MIB objects in accounting profiles	12.3R2
Support for an enterprise-specific event MIB (mib-jnx-event.txt)	12.3R2
Support for sending traps over routing instances	12.3R2
Support for adding lists of clients to the SNMP community	12.3R2
Support for the enterprise-specific Packet Forwarding Engine MIB (mib-jnx-pfe.txt)	12.3R2
Support for the pimNeighborLoss trap	12.3R2
Support for trap spoofing	12.3R2
IEEE 802.3ah link fault management (LFM) for Ethernet OAM (also known as Ethernet in the First Mile, or EFM)	12.3R2
Port mirroring of Layer 2 VLAN and VPLS traffic	12.3R2
Fast update filters for dynamic profiles	12.3R2
Flow aggregation to multiple collectors	12.3R2
IPv6 flow aggregation templates	12.3R2
Inline flow monitoring	12.3R2

Table 202: Port Security Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Automatic recovery for port error disable conditions	10.1R1	9.6R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	10.0R1	N.S.
DHCP option 82	10.1R1	9.3R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	N.S.	9.4R1	12.3R2
DHCP snooping	10.1R1	9.0R2	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	10.3R1	N.S.

Table 202: Port Security Features on Switches by Junos OS Release (*continued*)

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Dynamic ARP inspection (DAI)	10.1R1	9.0R2	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	10.3R1	N.S.
IP source guard	10.1R1	9.2R1	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	10.3R1	N.S.
Layer 3 virtual private network (VPN) for IPv4 (RFC 2547 and 4364)	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	12.3R2
Layer 3 virtual private network (VPN) for IPv6 through IPv4 MPLS	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	N.S.	12.3R2
MAC limiting	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2
MAC address limit per port	10.1R1	9.0R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	10.3R1	12.3R2
MAC limiting per port and per VLAN (VLAN membership MAC limit)	12.3R1	12.3R1	12.3R1	13.2X50 - D10	12.3R1	12.3R1	12.3R1	12.3R1	12.3R2
MAC move limiting	10.1R1	9.0R2	11.3R1	13.2X50 - D10	N.S.	N.S.	11.3R2	N.S.	N.S.
Persistent MAC learning (sticky MAC)	11.4R1	11.4R1	12.3R1	13.2X50 - D10	11.4R1	12.2R1	11.4R1	11.4R1	N.S.
Persistent storage for DHCP snooping	10.1R1	9.4R1	11.3R1	13.2X50 - D10	12.1R1	12.2R1	11.3R2	10.3R1	N.S.
Self-signed digital certificates for enabling SSL services	11.1R1	11.1R1	N.S.	13.2X50 - D10	11.1R1	12.2R1	12.1R1	11.1R1	N.S.
Static ARP support	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	12.3R2

Table 203: Power over Ethernet (PoE) Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Link Layer Discovery Protocol (LLDP) with granular Power over Ethernet (PoE) management	12.2R1	12.2R1 (EX4200-24PX and EX4200-48PX models only)	12.2R1	13.2X50 - D10	N.S.	N.S.	12.2R1	12.2R1	N.S.
NOTE: EX4200 switches must be running PoE controller software firmware version 4.04 or later to support the Link Layer Discovery Protocol (LLDP) with granular Power over Ethernet (PoE) management feature. See show chassis firmware detail and request system firmware upgrade poe to check or upgrade this firmware.									N.S.
Power over Ethernet (PoE)	10.1R1	9.0R2	11.3R1	13.2X50 - D10	—	—	11.3R2	11.2R1	N.S.
Power over Ethernet Plus (PoE+)	10.3R1	11.2R1 (EX4200-24PX and EX4200-48PX models only)	11.3R1	13.2X50 - D10	—	—	11.3R2	11.2R1	N.S.
Power over Ethernet (PoE) power management mode	10.1R1	9.3R2	11.3R1	13.2X50 - D10	—	—	11.3R2	11.2R1	N.S.

Table 204: Routing Policy and Packet Filtering Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Dynamic allocation of TCAM memory to firewall filters	10.1R1	10.0R1	11.3R1	–	10.2R1	12.2R1	N.S.	10.3R1	See Table 203 for a list of EX9200 routing policy and firewall filter features.
Firewall filters and rate limiting	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
For a list of supported firewall filter match conditions and actions, see <i>Platform Support for Firewall Filter Match Conditions, Actions, and Action Modifiers on EX Series Switches</i> .									
Firewall filters on LAGs	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	N.S.	10.0R1	
Firewall filters on the loopback interface	10.1R1	9.2R1	11.3R1	13.2X50 - D10	11.1R1	12.2R1	12.1R1	9.6R1	
For a list of supported firewall filter match conditions and actions on a loopback interface, see <i>Support for Match Conditions and Actions for Loopback Firewall Filters on Switches</i> .									
Firewall filters on the management interface	11.3R1	10.4R1	N.S.	13.2X50 - D10	10.4R1	12.2R1	12.1R1	10.4R1	
Firewall filters on the virtual management interface	–	10.4R1 (EX4200 Virtual Chassis only)	N.S.	13.2X50 - D10	–	–	–	–	
Firewall filters with IPv6	11.3R1	10.1R1	12.3R1	13.2X50 - D10	12.1R1	12.2R1	12.1R1	10.3R1	
Policing	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Tricolor marking policers	11.2R1	11.2R1	11.4R8	N.S.	12.1R1	12.2R1	11.3R1	N.S.	

Table 205: Routing Policy and Firewall Filters on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Access and access-internal routes	12.3R2
Extension of numeric-range match conditions in firewall filters	12.3R2
Aggregate policer support for different family address types configured on the same interface	12.3R2

Table 205: Routing Policy and Firewall Filters on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Authentication for BFD (MD5/SHA1)	12.3R2
BGP multipath link-bandwidth attribute	12.3R2
DHCP state persistence for DHCP relay agent	12.3R2
Dynamic configuration support for routing policies	12.3R2
Extended DHCP relay agent	12.3R2
Filter-based forwarding to a specific outgoing interface or destination IP address	12.3R2
Firewall filters within logical systems	12.3R2
IEEE 802.1p priority match conditions for Layer 2 VPN firewall filters	12.3R2
Filter-based forwarding to a specific outgoing interface or destination IP address	12.3R2
Layer 2 Gigabit Ethernet logical interface extended policing support	12.3R2
Layer 2 support for firewall filter match conditions	12.3R2
Load balancing of VPLS traffic	12.3R2
Option 60 support for extended DHCP relay agents	12.3R2
Policers on physical interfaces	12.3R2
Firewall filters feature support	12.3R2
Support for policers that limit traffic on logical interfaces in ingress or egress directions	12.3R2
Support for policers that rate-limit based on a percentage of physical port speed on an interface	12.3R2
Support for the discard action for tricolor marking policers applied to firewall filters	12.3R2
Support for the prefix-list match condition for firewall filters for the VPLS protocol family	12.3R2
Support for enhanced policer statistics	12.3R2
Support for MAC address validation	12.3R2
Tricolor marking policers	12.3R2

Table 206: Spanning-Tree Protocols Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
BPDU protection for spanning-tree protocols	10.1R1	9.1R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	See Table 206 for a list of EX9200 spanning-tree protocols features.
BPDU filter	12.2R1	12.2R1	12.2R1	13.2X50 - D10	12.2R1	12.2R1	12.2R1	12.2R1	
Distributed periodic packet management (PPM) for Spanning Tree Protocols	N.S.	12.3R1	N.S.	N.S.	N.S.	N.S.	N.S.	12.3R1	
Loop protection for spanning-tree protocols	10.1R1	9.1R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Root protection for spanning-tree protocols	10.1R1	9.1R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Spanning tree: <ul style="list-style-type: none"> • RSTP and VSTP concurrent configuration 	N.S.	10.2R1	12.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	10.2R1	
Spanning tree: <ul style="list-style-type: none"> • Spanning Tree Protocol (STP) • Rapid Spanning Tree Protocol (RSTP) • Multiple Spanning Tree Protocol (MSTP) 	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Spanning tree: <ul style="list-style-type: none"> • VLAN Spanning Tree Protocol (VSTP) 	10.1R1	9.4R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.6R1	

Table 207: Spanning-Tree Protocols Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Spanning tree: <ul style="list-style-type: none">• Spanning Tree Protocol (STP)• Rapid Spanning Tree Protocol (RSTP)• Multiple Spanning Tree Protocol (MSTP)• VLAN Spanning Tree Protocol (VSTP)	12.3R2
Root protection for spanning-tree protocols	12.3R2
Loop protection for spanning-tree protocols	12.3R2
Distributed periodic packet management (PPM) for Spanning Tree Protocols	12.3R2
BPDU filter	12.3R2

Table 208: System Management Features on Switches by Junos OS Release

Feature	EX2200	EX3200, EX4200	EX3300	EX4300	EX4500	EX4550	EX6200	EX8200	EX9200
Autoinstallation of configuration files	10.1R1	9.4R1	11.3R1	N.S.	10.2R1	12.2R1	11.3R2	N.S.	See Table 209 for a list of EX9200 system management features.
Automatic software download	10.1R1	9.6R1	11.3R1	N.S.	10.2R1	12.2R1	11.3R2	9.6R1	
Automatic repair of corrupted partition when booting from alternate partition	12.3R1	12.3R1	12.3R1	13.2X50 - D10	12.3R1	12.3R1	12.3R1	12.3R1	
Configuration rollback	10.1R1	9.0R2	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R2	9.4R1	
Zero Touch Provisioning (EZ Touchless Provisioning using DHCP)	12.2R1	12.2R1	12.2R5	N.S.	12.2R1	12.2R1	N.S.	N.S.	
J-Web interface, for switch configuration and management	10.1R1 (12.1R1 for EX2200-C switches)	9.0R2	12.1R1	13.2X50 - D10	10.2R1	12.2R1	12.1R1	9.4R1	
Junos Space Service Now support	12.3R1	12.3R1	12.3R1	N.S.	12.3R1	12.3R1	12.3R1	12.3R1	
LCD panel management support	–	9.0R1	11.3R1	13.2X50 - D10	10.2R1	12.2R1	11.3R1	9.4R1	
Online insertion and removal (OIR) of uplink modules	–	10.0R1	–	13.2X50 - D10	11.1R1	12.2R1	–	–	

Table 209: System Management Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Configuration rollback	12.3R2

Table 210: User Interface and Configuration Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Device-initiated SSH connection (outbound SSH)	12.3R2
Dynamic IPv6 filters	12.3R2
Dynamic configuration of the switch advertisement protocol	12.3R2
Dynamic profiles support by extended DHCP local server and extended DHCP relay agent	12.3R2
Enhanced IPv6 statistics	12.3R2
Extended DHCP local server	12.3R2
IGMP dynamic profiles	12.3R2
Extended DHCP local server	12.3R2
Protection for device configuration	12.3R2
RADIUS MSCHAPv2 protocol support for administrator authentication, password aging, and update	12.3R2
Limit configuration command output	12.3R2
Remote tracing	12.3R2
Support for CLI edit mode wildcard range	12.3R2
Support for configuring ARP aging time for a logical interface	12.3R2
Support for configuring a proxy server for downloading licenses	12.3R2
Support for configuring time-based user access	12.3R2
Support for logical router system administrators	12.3R2

Table 211: VPN Features on EX9200 Switches by Junos OS Release

Feature	Junos OS Release
Aggregated Ethernet interfaces for VPLS routing instances	12.3R2
BGP autodiscovery for LDP VPLS (FEC 129)	12.3R2
Clearing MAC addresses for better convergence	12.3R2
Configurable label block sizes for VPLS	12.3R2
Disable TTL propagation behavior for the routes in a VRF routing instance	12.3R2

Table 211: VPN Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
EXP-based traffic classification for VPLS	12.3R2
Enhanced show interface command for Layer 3 VPN functionality	12.3R2
Expanded interface support for the vrf-table-label statement	12.3R2
Extranet next-generation MVPN GRE tunnels for Layer 3 VPNs	12.3R2
GRE tunnels for Layer 3 VPNs Ignore MTU mismatch on Layer 2 circuits Integrated routing and bridging support for inter-AS VPLS between BGP-signaled VPLS and LDP-signaled VPLS LDP-based VPLS Label allocation and substitution policy	12.3R2
Ignore MTU mismatch on Layer 2 circuits	12.3R2
Integrated routing and bridging support for inter-AS VPLS between BGP-signaled VPLS and LDP-signaled VPLS	12.3R2
LDP-based VPLS	12.3R2
Label allocation and substitution policy	12.3R2
Layer 2 VPN multihoming	12.3R2
Layer 3 VPN BGP routes and labels	12.3R2
Layer 3 VPN localization	12.3R2
Load balancing and IP header filtering for Layer 3 VPNs	12.3R2
Local switching support for the ignore-encapsulation-mismatch statement	12.3R2
Multipath load balancing for EBGp and IBGP VPNs	12.3R2
Multiple logical trunk interfaces per physical interface	12.3R2
Multiprotocol BGP-based multicast VPN	12.3R2
NTP support for IPv4 VRF and IPv6 VRF	12.3R2
Nonstop active routing support for Layer 3 VPNs	12.3R2
PIM source-specific multicast (PIM-SSM) provider tunnel support added to Multiprotocol BGP-based multicast VPNs	12.3R2
Point-to-multipoint LSP support for VPLS	12.3R2
Point-to-multipoint LSP support for multicast VPNs	12.3R2

Table 211: VPN Features on EX9200 Switches by Junos OS Release (*continued*)

Feature	Junos OS Release
Proxy BGP route target filtering	12.3R2
Static VPLS	12.3R2
Static route target filtering	12.3R2
Support for autorp, BSR, PIM dense mode and mtrace for next-generation multicast VPNs	12.3R2
VLAN range for Layer 2 VPN	12.3R2
VPLS automatic site ID	12.3R2
VPLS automatic site ID for nonstop active routing	12.3R2
VPLS ping	12.3R2
VPLS trunk interfaces	12.3R2
eBGP and iBGP load-balancing support for MVPN and PIM	12.3R2

**Related
Documentation**

- *EX Series Virtual Chassis Software Features Overview*
- *EX2200 Switches Hardware Overview*
- *EX3200 Switches Hardware Overview*
- *EX3300 Switches Hardware Overview*
- *EX4200 Switches Hardware Overview*
- <will add topic-ref to EX4300 HW overview topic>
- *EX4500 Switches Hardware Overview*
- *EX4550 Switches Hardware Overview*
- *EX6210 Switch Hardware Overview*
- *EX8208 Switch Hardware Overview*
- *EX8216 Switch Hardware Overview*
- *EX9204 Switch Hardware Overview*
- [EX9208 Switch Hardware Overview on page 115](#)
- *EX9214 Switch Hardware Overview*
- *Line Card Model and Version Compatibility in an EX6200 Switch*
- *Line Card Model and Version Compatibility in an EX8200 Switch*
- *Line Card Model and Version Compatibility in an EX9200 Switch*

- *XRE200 External Routing Engine Hardware Overview*
- *Layer 3 Protocols Supported on EX Series Switches*
- *Layer 3 Protocols Not Supported on EX Series Switches*

Configuration

- [Configuration Statements on page 2553](#)

Configuration Statements

- [\[edit protocols bgp\] Hierarchy Level on page 2553](#)

[edit protocols bgp] Hierarchy Level

Several statements in the **[edit protocols mpls]** hierarchy are valid at numerous locations within it. To make the complete hierarchy easier to read, the repeated statements are listed in “[Common BGP Family Options](#)” on page 369 and that section is referenced at the appropriate locations in “[Complete \[edit protocols bgp\] Hierarchy](#)” on page 369.

- [Common BGP Family Options on page 2553](#)
- [Complete \[edit protocols bgp\] Hierarchy on page 2554](#)

Common BGP Family Options

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit protocols bgp\] Hierarchy](#)” on page 369 instead of the statements being repeated.

- **[edit protocols bgp family inet (any | flow | labeled-unicast | multicast | unicast)]**
- **[edit protocols bgp family inet6 (any | labeled-unicast | multicast | unicast)]**
- **[edit protocols bgp family (inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) signaling]**
- **[edit protocols bgp family inet-vpn (any | flow | multicast | unicast)]**
- **[edit protocols bgp family inet6-vpn (any | multicast | unicast)]**
- **[edit protocols bgp family iso-vpn unicast]**

The common BGP family options are as follows:

```
accepted-prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
}
damping;
loops number;
prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
}
rib-group group-name;
topology name {
    community {
```

```
    target identifier;  
  }  
}
```

Complete [edit protocols bgp] Hierarchy

The statement hierarchy listed in this section can also be included at the [edit logical-systems *logical-system-name*] hierarchy level.

```
protocols {  
  bgp {  
    disable;  
    accept-remote-nexthop;  
    advertise-external <conditional>;  
    advertise-from-main-vpn-tables;  
    advertise-inactive;  
    (advertise-peer-as | no-advertise-peer-as);  
    authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);  
    authentication-key key;  
    authentication-key-chain key-chain;  
    bfd-liveness-detection {  
      authentication {  
        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |  
          meticulous-keyed-sha-1 | simple-password);  
        key-chain key-chain-name;  
        loose-check;  
      }  
      detection-time {  
        threshold milliseconds;  
      }  
      holddown-interval milliseconds;  
      minimum-interval milliseconds;  
      minimum-receive-interval milliseconds;  
      multiplier number;  
      no-adaptation;  
      session-mode (automatic | multihop | single-hop);  
      transmit-interval {  
        minimum-interval milliseconds;  
        threshold milliseconds;  
      }  
      version (1 | automatic);  
    }  
    cluster cluster-identifier;  
    damping;  
    description text-description;  
    export [ policy-names ];  
    family family-name {  
      ... the family subhierarchies appear after the main [edit protocols bgp] hierarchy ...  
    }  
    graceful-restart {  
      disable;  
      restart-time seconds;  
      stale-routes-time seconds;  
    }  
    group group-name {  
      ... the group subhierarchy appears after the main [edit protocols bgp] hierarchy ...  
    }  
  }  
}
```



```

}
hold-time seconds;
idle-after-switch-over (seconds | forever);
import [ policy-names ];
include-mp-next-hop;
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <loops number> <alias> <private>;
local-interface interface-name;
local-preference local-preference;
log-updown;
metric-out (metric | igp (delay-med-update | offset) | minimum-igp offset);
mtu-discovery;
multihop {
    no-nexthop-change;
    ttl ttl-value;
}
no-agggregator-id;
no-client-reflect;
out-delay seconds;
outbound-route-filter {
    bgp-orf-cisco-mode;
    prefix-based {
        accept {
            inet;
            inet6;
        }
    }
}
passive;
path-selection {
    always-compare-med;
    as-path-ignore;
    cisco-non-deterministic;
    external-router-id;
    med-plus-igp {
        igp-multiplier number;
        med-multiplier number;
    }
}
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
vpn-apply-export;
}

bgp {
    family inet {
        (any | multicast) {

```

```

    ... statements in Common BGP Family Options on page 369 ...
  }
  flow {
    ... statements in Common BGP Family Options on page 369 PLUS ...
    no-validate [ validation-procedure-names ];
  }
  labeled-unicast {
    ... statements in Common BGP Family Options on page 369 PLUS ...
    add-path {
      receive;
      send {
        path-count number;
        prefix-policy [ policy-names ];
      }
    }
    aggregate-label {
      community community-name;
    }
    aigp [disable];
    explicit-null connected-only;
    per-group-label;
    per-prefix-label;
    resolve-vpn;
    rib (inet.3 | inet6.3);
    traffic-statistics {
      file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
      interval seconds;
    }
  }
  unicast {
    ... statements in Common BGP Family Options on page 369 PLUS ...
    add-path {
      receive;
      send {
        path-count number;
        prefix-policy [ policy-names ];
      }
    }
    topology name {
      community target identifier;
    }
  }
}

bgp {
  family inet6 {
    (any | multicast) {
      ... statements in Common BGP Family Options on page 369 ...
    }
    labeled-unicast {
      ... statements in Common BGP Family Options on page 369 PLUS ...
      add-path {
        receive;
        send {

```

```

        path-count number;
        prefix-policy [ policy-names ];
    }
}
aggregate-label {
    community community-name;
}
aigp [disable];
explicit-null;
per-group-label;
traffic-statistics {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    interval seconds;
}
}
unicast {
    ... statements in Common BGP Family Options on page 369 PLUS ...
    topology name {
        community target identifier;
    }
}
}
}

bgp {
    family (inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) {
        signaling {
            ... statements in Common BGP Family Options on page 369 ...
        }
    }
}

bgp {
    family inet-vpn {
        (any | multicast | unicast) {
            ... statements in Common BGP Family Options on page 369 PLUS ...
            aggregate-label <community community-name>;
        }
        flow {
            ... statements in Common BGP Family Options on page 369 ...
        }
    }
}

bgp {
    family inet6-vpn {
        (any | multicast | unicast) {
            ... statements in Common BGP Family Options on page 369 PLUS ...
            aggregate-label <community community-name>;
        }
    }
}

bgp {
    family iso-vpn {

```

```

    unicast {
        ... statements in Common BGP Family Options on page 369 PLUS ...
        aggregate-label <community community-name>;
    }
}

bgp {
    family route-target {
        accepted-prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
        }
        advertise-default;
        external-paths number;
        prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
        }
        proxy-generate <route-target-policy route-target-policy-name>;
    }
}

bgp {
    group group-name {
        ... same statements as at the [edit protocols bgp] hierarchy level PLUS ...
        allow [ all ip-prefix</prefix-length> ];
        as-override;
        multipath <multiple-as>;
        neighbor address {
            ... the neighbor subhierarchy appears after the main [edit protocols bgp group
                group-name] hierarchy ...
        }
        type (external | internal);
        ... BUT NOT ...
        disable; # NOT valid at this level
        group group-name { ... } # NOT valid at this level
        path-selection { ... } # NOT valid at this level
    }

    group group-name {
        neighbor address {
            ... same statements as at the [edit protocols bgp] hierarchy level PLUS ...
            as-override;
            multipath <multiple-as>;
            ... BUT NOT ...
            disable; # NOT valid at this level
            group group-name { ... } # NOT valid at this level
            neighbor address { ... } # NOT valid at this level
            path-selection { ... } # NOT valid at this level
        }
    }
}


```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

accept-remote-nexthop

Syntax	accept-remote-nexthop;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Specify that a single-hop EBGP peer accepts a remote next hop with which it does not share a common subnet. Configure a separate import policy on the EBGP peer to specify the remote next hop. You cannot configure multihop and accept-remote-nexthop statements for the same EPBG peer.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Single-Hop EBGP Peers to Accept Remote Next Hops</i> • <i>Understanding Route Advertisement</i> • multipath on page 2617

advertise-external

Syntax	<code>advertise-external {conditional};</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>neighbor-address</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.3.</p> <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify BGP to advertise the best external route into an IBGP mesh group, a route reflector cluster, or an AS confederation even if the best route is an internal route. In order to configure the advertise-external statement on a route reflector, you must disable intracluster reflection with the no-client-reflect statement. The advertise-external statement is supported at both the group and neighbor level. If you configure the statement at the neighbor level, you must configure it for all neighbors in a group. Otherwise, the group is automatically split into different groups.</p>
	<div>  <p>NOTE: When configuring the advertise-external statement for an AS confederation, it is recommended that EBGp peers belonging to different autonomous systems are configured in a separate EBGp peer group. This ensures consistency while BGP sends the best external route to peers in the configured peer group.</p> </div>
Options	<p>conditional—(Optional) Advertise the best external path only if the route selection process reaches the point at which the multiple exit discriminator (MED) metric is evaluated. The conditional option restricts advertisement to when the best external path and the active path are equal until the MED step of the route selection process. This implies that external routes with a longer AS path length than the active path, for instance, are not advertised. The criteria used for selecting the best external path is the same whether or not the conditional option is configured.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring BGP Route Advertisement</i> • <i>Understanding Route Advertisement</i>

- [advertise-inactive on page 2561](#)

advertise-inactive

Syntax	advertise-inactive;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure the routing table to export to BGP the best route learned by BGP even if Junos OS did not select this route to be an active route.</p> <p>One way to achieve multivendor compatibility is to include the advertise-inactive statement in the external BGP (EBGP) configuration. By default, BGP stores the route information it receives from update messages in the Junos OS routing table, and the routing table exports only active routes into BGP, which BGP then advertises to its peers. The advertise-inactive statement causes Junos OS to advertise the best BGP route that is inactive because of IGP preference. When you use the advertise-inactive statement, the Junos OS device uses, for example, the OSPF route for forwarding, and the other vendor's device uses the EBGP route for forwarding. However, from the perspective of an EBGP peer in a neighboring AS, both vendors' devices appear to behave the same way.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring the Preference Value for BGP Routes</i> • <i>Example: Configuring BGP Route Preference (Administrative Distance)</i> • <i>Understanding Route Advertisement</i> • advertise-external on page 2560



advertise-peer-as

Syntax	advertise-peer-as;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Disable the default behavior of suppressing AS routes.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring BGP Route Advertisement</i>• <i>Understanding Route Advertisement</i>• <i>no-advertise-peer-as</i>


aggregate-label

Syntax	aggregate-label { community <i>community-name</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp family inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp family inet6 labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp family inet-vpn unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp family inet-vpn6 unicast], [edit protocols bgp family inet labeled-unicast], [edit protocols bgp family inet6 labeled-unicast], [edit protocols bgp family inet-vpn unicast], [edit protocols bgp family inet6-vpn unicast]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify matching criteria (in the form of a community) such that all routes which match are assigned the same VPN label, selected from one of the several routes in the set defined by this criteria. This reduces the number of VPN labels that the router must consider, and aggregates the received labels.
Options	community <i>community-name</i> —Specify the name of the community to which to apply the aggregate label.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Aggregate Labels for VPNs</i>

allow

Syntax	<code>allow (all [<i>network/mask-length</i>]);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Implicitly configure BGP peers, allowing peer connections from any of the specified networks or hosts. To configure multiple BGP peers, configure one or more networks and hosts within a single allow statement or include multiple allow statements.
<div>  <p>NOTE: You cannot define a BGP group with dynamic peers with BGP authentication enabled.</p> </div>	
Options	<p>all—Allow all addresses, which is equivalent to 0.0.0.0/0 (or ::/0).</p> <p><i>network/mask-length</i>—IPv6 or IPv4 network number of a single address or a range of allowable addresses for BGP peers, followed by the number of significant bits in the subnet mask.</p>
<div>  <p>NOTE: You cannot define a BGP group with dynamic peers with authentication enabled.</p> </div>	
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> neighbor on page 2618

as-override

Syntax	as-override;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Compare the AS path of an incoming advertised route with the AS number of the BGP peer under the group and replace all occurrences of the peer AS number in the AS path with its own AS number before advertising the route to the peer.
<div>  <p>NOTE: The as-override statement is specific to a particular BGP group. This statement does not affect peers from the same remote AS configured in different groups.</p> </div>	
<p>Enabling the AS override feature allows routes originating from an AS to be accepted by a router residing in the same AS. Without AS override enabled, the routing device refuses the route advertisement once the AS path shows that the route originated from its own AS. This is done by default to prevent route loops. The as-override statement overrides this default behavior.</p> <p>Note that enabling the AS override feature may result in routing loops. Use this feature only for specific applications that require this type of behavior, and in situations with strict network control. One application is the IGP protocol between the provider edge routing device and the customer edge routing device in a virtual private network.</p>	
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring a Layer 3 VPN with Route Reflection and AS Override</i> • <i>Understanding Route Advertisement</i> • <i>Junos OS VPNs Configuration Guide</i>

authentication-algorithm

Syntax	<code>authentication-algorithm <i>algorithm</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ldp session <i>session-address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp session <i>session-address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols ldp session <i>session-address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ldp session <i>session-address</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced for BGP in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure an authentication algorithm type.
Options	<p><i>algorithm</i>—Specify one of the following types of authentication algorithms:</p> <ul style="list-style-type: none"> • aes-128-cmac-96—Cipher-based message authentication code (AES128, 96 bits). • hmac-sha-1-96—Hash-based message authentication code (SHA1, 96 bits). • md5—Message digest 5.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding Route Authentication</i> • <i>Example: Configuring Route Authentication for BGP</i>

authentication-key (Protocols BGP)

Syntax	authentication-key <i>key</i> ;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure an MD5 authentication key (password). Neighboring routing devices use the same password to verify the authenticity of BGP packets sent from this system.
Options	key —Authentication password. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Route Authentication for BGP</i>

authentication-key-chain (Protocols BGP)

Syntax	<code>authentication-key-chain <i>key-chain</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Apply and enable an authentication keychain to the routing device. Note that the referenced key chain must be defined. When configuring the authentication key update mechanism for BGP, you cannot commit the 0.0.0.0/allow statement with authentication keys or key chains. The CLI issues a warning and fails to commit such configurations.
Options	key-chain —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Route Authentication for BGP</i> • Example: Configuring BFD Authentication for Static Routes on page 3210 • <i>Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</i>

autonomous-system

Syntax	<pre>autonomous-system <i>autonomous-system</i> <asdot-notation> <loops <i>number</i>> { independent-domain <no-attrset>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-options]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>asdot-notation option introduced in Junos OS Release 9.3.</p> <p>asdot-notation option introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>no-attrset option introduced in Junos OS Release 10.4.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Specify the routing device's AS number.</p> <p>An autonomous system (AS) is a set of routing devices that are under a single technical administration and that generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices. An AS appears to other ASs to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it. ASs are identified by a number that is assigned by the Network Information Center (NIC) in the United States (http://www.isi.edu).</p> <p>If you are using BGP on the routing device, you must configure an AS number.</p> <p>The AS path attribute is modified when a route is advertised to an EBGP peer. Each time a route is advertised to an EBGP peer, the local routing device prepends its AS number to the existing path attribute, and a value of 1 is added to the AS number.</p> <p>In Junos OS Release 9.1 and later, the numeric range is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, <i>BGP Support for Four-octet AS Number Space</i>. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.</p> <p>In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <i><16-bit high-order value in decimal>.<16-bit low-order value in decimal></i>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.</p>
Options	<p><i>autonomous-system</i>—AS number. Use a number assigned to you by the NIC.</p>

Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format for 4-byte AS numbers

In this example, the 4-byte AS number 65,546 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 65546;
}
```

Range: 0.0 through 65535.65535 in AS-dot notation format for 4-byte numbers

In this example, 1.10 is the AS-dot notation format for 65,546:

```
[edit]
routing-options {
  autonomous-system 1.10;
}
```

Range: 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)

In this example, the 2-byte AS number 60,000 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 60000;
}
```

asdot-notation—(Optional) Display the configured 4-byte autonomous system number in the AS-dot notation format.

Default: Even if a 4-byte AS number is configured in the AS-dot notation format, the default is to display the AS number in the plain-number format.

loops number—(Optional) Specify the number of times detection of the AS number in the AS_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.

Range: 1 through 10

Default: 1



.....

NOTE: When you specify the same AS number in more than one routing instance on the local routing device, you must configure the same number of loops for the AS number in each instance. For example, if you configure a value of 3 for the loops statement in a VRF routing instance that uses the same AS number as that of the master instance, you must also configure a value of 3 loops for the AS number in the master instance.

Use the **independent-domain** option if the loops statement must be enabled only on a subset of routing instances.

.....

The remaining statement is explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Examples: Configuring External BGP Peering</i>• <i>Examples: Configuring Internal BGP Peering</i>• <i>4-Byte Autonomous System Numbers Overview</i> in the <i>Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview</i>• <i>Juniper Networks Implementation of 4-Byte Autonomous System Numbers</i> in the <i>Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview</i>• <i>Configuring 4-Byte Autonomous System Numbers</i> in the <i>Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview</i>

bfd-liveness-detection (Protocols BGP)

Syntax `bfd-liveness-detection {
 authentication {
 algorithm algorithm-name;
 key-chain key-chain-name;
 loose-check;
 }
 detection-time {
 threshold milliseconds;
 }
 hold-down-interval milliseconds;
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 session-mode (automatic | multihop | single-hop);
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 version (1 | automatic);
 }
}`

Hierarchy Level `[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
bgp group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
bgp group group-name neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor
address]`

Release Information Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
detection-time threshold and **transmit-interval threshold** options introduced in Junos OS Release 8.2
Support for logical routers introduced in Junos OS Release 8.3.
Support for IBGP and multihop EBGP sessions introduced in Junos OS Release 8.3.
holddown-interval statement introduced in Junos OS Release 8.5. You can configure this statement only for EBGP peers at the **[edit protocols bgp group *group-name* neighbor *address*]** hierarchy level.
no-adaptation statement introduced in Junos OS Release 9.0.
Support for BFD authentication introduced in Junos OS Release 9.6.

Support for BFD on IPv6 interfaces with BGP introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure bidirectional failure detection (BFD) timers and authentication for BGP.

For IBGP and multihop EBGP support, configure the **bfd-liveness-detection** statement at the global **[edit bgp protocols]** hierarchy level. You can also configure IBGP and multihop support for a routing instance or a logical system.

Options **authentication algorithm** *algorithm-name* (Optional)—Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**.

authentication key-chain *key-chain-name* (Optional)—Associate a security key with the specified BFD session using the name of the security keychain. The keychain name must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

detection-time threshold *milliseconds* (Optional)—Configure a threshold. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

holddown-interval *milliseconds* (Optional)—Configure an interval specifying how long a BFD session must remain up before a state change notification is sent. When you configure the hold-down interval for the BFD protocol for EBGp, the BFD session is unaware of the BGP session during this time. In this case, if the BGP session goes down during the configured hold-down interval, BFD already assumes it is down and does not send a state change notification. The **holddown-interval** statement is supported only for EBGp peers at the **[edit protocols bgp group group-name neighbor address]** hierarchy level. If the BFD session goes down and then comes back up during the configured hold-down interval, the timer is restarted. You must configure the hold-down interval on both EBGp peers. If you configure the hold-down interval for a multihop EBGp session, you must also configure a local IP address by including the **local-address** statement at the **[edit protocols bgp group group-name]** hierarchy level.

Range: 0 through 255,000

Default: 0

minimum-interval *milliseconds* (Required)—Configure the minimum intervals at which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately (using the **minimum-receive-interval** and **transmit-interval** statements).

Range: 1 through 255,000

minimum-receive-interval *milliseconds* (Optional)—Configure only the minimum interval at which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session.

Range: 1 through 255,000

multiplier *number* (Optional)—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

no-adaptation (Optional)—Configure BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable to not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds* (Optional)—Configure a threshold. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

transmit-interval minimum-interval *milliseconds* (Optional)—Configure only the minimum interval at which the local routing device transmits hello packets to a neighbor with which it has established a BFD session.

Range: 1 through 255,000

version (Optional)—Configure the BFD version to detect.

Range: 1 or **automatic** (autodetect the BFD version)

Default: **automatic**

The remaining statements are explained separately.


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

Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BFD for Static Routes on page 3197 • Example: Configuring BFD Authentication for Static Routes on page 3210 • Example: Configuring BFD on Internal BGP Peer Sessions on page 3390 • Example: Configuring BFD Authentication for BGP • Understanding BFD for BGP
------------------------------	--

bgp

Syntax	<code>bgp { ... }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Enable BGP on the routing device or for a routing instance.
Default	BGP is disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>BGP Configuration Guide</i>

bgp-orf-cisco-mode

Syntax	<code>bgp-orf-cisco-mode;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options outbound-route-filter],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options outbound-route-filter],</p> <p>[edit protocols bgp outbound-route-filter],</p> <p>[edit protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> outbound-route-filter],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options outbound-route-filter],</p> <p>[edit routing-options outbound-route-filter]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.2.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	Enable interoperability with routing devices that use the vendor-specific outbound route filter compatibility code of 130 and code type of 128.
	<div>  <p>NOTE: To enable interoperability for all BGP peers configured on the routing device, include the statement at the [edit routing-options outbound-route-filter] hierarchy level.</p> </div>
Default	Disabled
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- *Example: Configuring BGP Prefix-Based Outbound Route Filtering*

cluster

Syntax	<code>cluster <i>cluster-identifier</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Specify the cluster identifier to be used by the route reflector cluster in an internal BGP group.



CAUTION:

If you configure both route reflection and VPNs on the same routing device, the following modifications to the route reflection configuration cause current BGP sessions to be reset:

- Adding a cluster ID—If a BGP session shares the same AS number with the group where you add the cluster ID, all BGP sessions are reset regardless of whether the BGP sessions are contained in the same group.
- Creating a new route reflector—If you have an IBGP group with an AS number and create a new route reflector group with the same AS number, all BGP sessions in the IBGP group and the new route reflector group are reset.



NOTE: If you change the address family specified in the [edit protocols bgp family] hierarchy level, all current BGP sessions on the routing device are dropped and then reestablished.

Options	<i>cluster-identifier</i> —4-byte number (such as an IPv4 address).
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring BGP Route Reflectors</i>• <i>Understanding External BGP Peering Sessions</i>• no-client-reflect on page 2622

damping (Protocols BGP)

Syntax	damping;
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp family <i>family</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp family <i>family</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>], [edit protocols bgp], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> family <i>family</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family <i>family</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>] [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>family</i>] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for flap damping at the address family level introduced in Junos OS Release 12.2.</p>
Description	<p>Enable route flap damping. BGP route flapping describes the situation in which BGP systems send an excessive number of update messages to advertise network reachability information. Flap damping reduces the number of update messages sent between BGP</p>

peers, thereby reducing the load on these peers, without adversely affecting the route convergence time for stable routes.

You typically apply flap damping to external BGP (EBGP) routes (that is, to routes in different ASs). You can also apply it within a confederation, between confederation member ASs. Because routing consistency within an AS is important, do not apply flap damping to internal BGP (IBGP) routes. (If you do, it is ignored.) The exception to this rule is when flap damping is applied at the address family level. When you apply flap damping at the address family level, it works for both IBGP and EBGP.

Default Flap damping is disabled on the routing device.

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

Related Documentation

- *Examples: Configuring BGP Flap Damping*
- *Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family*

description (Protocols BGP)

Syntax	<code>description text-description;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Provide a description of the global, group, or neighbor configuration. If the text includes one or more spaces, enclose it in quotation marks (" "). The text is displayed in the output of the show command and has no effect on the configuration.
Options	text-description —Text description of the configuration. It is limited to 255 characters.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>BGP Configuration Guide</i>

disable (Protocols BGP)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Disable BGP on the system.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>BGP Configuration Guide</i>

explicit-null (Protocols BGP)

Syntax	explicit-null;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>family</i> inet labeled-unicast],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>family</i> inet6 labeled-unicast],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> <i>family</i> inet labeled-unicast],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> <i>family</i> inet6 labeled-unicast],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet labeled-unicast],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet6 labeled-unicast],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ldap],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols bgp <i>family</i> inet labeled-unicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols bgp <i>family</i> inet6 labeled-unicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> inet labeled-unicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> inet6 labeled-unicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet labeled-unicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet6 labeled-unicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols ldap],</p> <p>[edit protocols mpls],</p> <p>[edit protocols bgp <i>family</i> inet labeled-unicast],</p> <p>[edit protocols bgp <i>family</i> inet6 labeled-unicast],</p> <p>[edit protocols bgp group <i>group-name</i> <i>family</i> inet labeled-unicast],</p> <p>[edit protocols bgp group <i>group-name</i> <i>family</i> inet6 labeled-unicast],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet labeled-unicast],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet6 labeled-unicast],</p> <p>[edit protocols ldap],</p> <p>[edit routing-instances <i>instance-name</i> protocols bgp <i>family</i> inet labeled-unicast],</p> <p>[edit routing-instances <i>instance-name</i> protocols bgp <i>family</i> inet6 labeled-unicast],</p> <p>[edit routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> inet labeled-unicast],</p> <p>[edit routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> inet6 labeled-unicast],</p> <p>[edit routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet labeled-unicast],</p> <p>[edit routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> inet6 labeled-unicast],</p> <p>[edit routing-instances <i>instance-name</i> protocols ldap]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Advertise label 0 to the egress routing device of an LSP.

Default	If you do not include the explicit-null statement in the configuration, label 3 (implicit null) is advertised.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Advertising Explicit Null Labels to BGP Peers</i>

export (Protocols BGP)

Syntax	<code>export [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name neighbor address</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp <i>group group-name</i>],</code> <code>[edit protocols bgp <i>group group-name neighbor address</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Apply one or more policies to routes being exported from the routing table into BGP.
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding Route Advertisement</i>• <i>Routing Policy Configuration Guide</i>• import on page 2598

family (Protocols BGP)

```
Syntax  family {
    (inet | inet6 | inet-vpn | inet6-vpn | iso-vpn) {
        (any | flow | labeled-unicast | multicast | unicast) {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage-threshold> idle-timeout (forever | minutes);
            }
            add-path {
                send {
                    path-count number;
                    prefix-policy [ policy-names ];
                }
                receive;
            }
            algp [disable];
            loops number;
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            protection;
            rib-group group-name;
            topology name {
                community {
                    target identifier;
                }
            }
        }
        flow {
            no-validate policy-name;
        }
        labeled-unicast {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            aggregate-label {
                community community-name;
            }
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            resolve-vpn;
            rib (inet.3 | inet6.3);
            rib-group group-name;
            traffic-statistics {
                file filename <world-readable | no-world-readable>;
                interval seconds;
            }
        }
    }
}
```

```
    }
  }
  route-target {
    accepted-prefix-limit {
      maximum number;
      proxy-generate <route-target-policy route-target-policy-name>;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
  }
  (inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) {
    signaling {
      accepted-prefix-limit {
        maximum number;
        teardown <percentage-threshold> idle-timeout (forever | minutes);
      }
      add-path {
        send {
          path-count number;
          prefix-policy [ policy-names ];
        }
        receive;
      }
      aigp [disable];
      damping;
      loops number;
      prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
      }
      rib-group group-name;
    }
  }
}
```

Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>inet-mvpn and inet6-mvpn statements introduced in Junos OS Release 8.4.</p> <p>inet-mdt statement introduced in Junos OS Release 9.4.</p> <p>Support for the loops statement introduced in Junos OS Release 9.6.</p>
Description	<p>Enable multiprotocol BGP (MP-BGP) by configuring BGP to carry network layer reachability information (NLRI) for address families other than unicast IPv4, to specify MP-BGP to carry NLRI for the IPv6 address family, or to carry NLRI for VPNs.</p>


- Options**
- any**—Configure the family type to be both unicast and multicast.
 - inet**—Configure NLRI parameters for IPv4.
 - inet6**—Configure NLRI parameters for IPv6.
 - inet-mdt**—Configure NLRI parameters for the multicast distribution tree (MDT) subaddress family identifier (SAFI) for IPv4 traffic in Layer 3 VPNs.
 - inet-mvpn**—Configure NLRI parameters for IPv4 for multicast VPNs.
 - inet6-mvpn**—Configure NLRI parameters for IPv6 for multicast VPNs.
 - inet-vpn**—Configure NLRI parameters for IPv4 for Layer 3 VPNs.
 - inet6-vpn**—Configure NLRI parameters for IPv6 for Layer 3 VPNs.
 - iso-vpn**—Configure NLRI parameters for IS-IS for Layer 3 VPNs.
 - l2vpn**—Configure NLRI parameters for IPv4 for MPLS-based Layer 2 VPNs and VPLS.
 - labeled-unicast**—Configure the family type to be labeled-unicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by labeled-unicast for resolving the labeled-unicast routes. This statement is supported only with **inet** and **inet6**.
 - multicast**—Configure the family type to be multicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by multicast for resolving the multicast routes.
 - unicast**—Configure the family type to be unicast. This means that the BGP peers only carry the unicast routes that are being used for unicast forwarding purposes. The default family type is **unicast**.

The remaining statements are explained separately.

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

- Related Documentation**
- [autonomous-system on page 2569](#)
 - [local-as on page 2606](#)
 - *Understanding Multiprotocol BGP*

graceful-restart (Protocols BGP)

Syntax	<pre> graceful-restart { disable; restart-time seconds; stale-routes-time seconds; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Enable graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. Graceful restart is disabled by default.</p> <p>To configure the duration of the BGP graceful restart period, include the restart-time statement at the [edit protocols bgp graceful-restart] hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the stale-routes-time statement at the [edit protocols bgp graceful-restart] hierarchy level.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> NOTE: If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.</p> </div> <p>Configure graceful restart globally at the [edit routing-options] or [edit routing-instances <i>instance-name</i> routing-options] hierarchy level to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Graceful Restart Options for BGP on page 2129 • <i>Configuring Graceful Restart for QFabric Systems</i> • <i>Junos OS High Availability Configuration Guide</i>

group (Protocols BGP)

```
Syntax  group group-name {
        advertise-inactive;
        allow [ network/mask-length ];
        authentication-key key;
        cluster cluster-identifier;
        damping;
        description text-description;
        export [ policy-names ];
        family {
            (inet | inet6 | inet-vpn | inet6-vpn | l2-vpn) {
                (any | multicast | unicast | signaling) {
                    accepted-prefix-limit {
                        maximum number;
                        teardown <percentage> <idle-timeout (forever | minutes)>;
                    }
                    add-path {
                        send {
                            path-count number;
                            prefix-policy [ policy-names ];
                        }
                        receive;
                    }
                    aigp [disable];
                    damping;
                    prefix-limit {
                        maximum number;
                        teardown <percentage> <idle-timeout (forever | minutes)>;
                    }
                    rib-group group-name;
                    topology name {
                        community {
                            target identifier;
                        }
                    }
                }
            }
            flow {
                no-validate policy-name;
            }
            labeled-unicast {
                accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                explicit-null {
                    connected-only;
                }
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                resolve-vpn;
                rib inet.3;
            }
        }
    }
```

```

        rib-group group-name;
    }
}
route-target {
    accepted-prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
}
}
hold-time seconds;
import [ policy-names ];
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
local-preference local-preference;
log-updown;
metric-out metric;
multihop <ttl-value>;
multipath {
    multiple-as;
}
mvpn-iana-rt-import;
no-agggregator-id;
no-client-reflect;
out-delay seconds;
passive;
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
type type;
neighbor address {
    ... peer-specific-options ...
}
}

```

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 bgp],
 [edit protocols bgp],
 [edit routing-instances *routing-instance-name* protocols bgp]

Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Define a BGP peer group. BGP peer groups share a common type, peer autonomous system (AS) number, and cluster ID, if present. To configure multiple BGP groups, include multiple group statements.</p> <p>By default, the group's options are identical to the global BGP options. To override the global options, include group-specific options within the group statement.</p> <p>The group statement is one of the statements you must include in the configuration to run BGP on the routing device.</p> <p>Each group must contain at least one peer.</p>
Options	<p>group-name—Name of the BGP group.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>BGP Configuration Guide</i>

hold-time (Protocols BGP)

Syntax	<code>hold-time seconds;</code>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp <i>group group-name neighbor address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the hold-time value to use when negotiating a connection with the peer. The hold-time value is advertised in open packets and indicates to the peer the length of time that it should consider the sender valid. If the peer does not receive a keepalive, update, or notification message within the specified hold time, the BGP connection to the peer is closed and routing devices through that peer become unavailable.</p> <p>The hold time is three times the interval at which keepalive messages are sent.</p> <p>BGP on the local routing device uses the smaller of either the local hold-time value or the peer's hold-time value received in the open message as the hold time for the BGP connection between the two peers.</p> <p>Starting in Junos OS Release 12.3, the BGP hold-time value can be zero (0). This implies that the speaker does not expect keepalive messages from its peer to maintain the BGP session. When negotiating between two peers, if one side requests a nonzero hold time and the other requests a zero hold time, the negotiation settles on the nonzero value and keepalive intervals are determined accordingly. Both sides must be set to zero for keepalive messages to stop being sent.</p>
Options	<p>seconds—Hold time.</p> <p>Range: 10 through 65,535 seconds (or 0 for infinite hold time)</p> <p>Default: 90 seconds</p>



TIP: When you set a hold-time value of 1 through 19 seconds, we recommend that you also configure the BGP `precision-timers` statement. The `precision-timers` statement ensures that if scheduler slip messages occur, the routing device continues to send keepalive messages. When the `precision-timers` statement is included, keepalive message generation is performed in a dedicated kernel thread, which helps to prevent BGP session flaps.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>BGP Messages Overview</i>• <i>precision-timers</i>

idle-after-switch-over

Syntax	<code>idle-after-switch-over (forever seconds);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure the routing device so that it does not automatically reestablish BGP peer sessions after a nonstop active routing (NSR) switchover. This feature is particularly useful if you are using dynamic routing policies because the dynamic database is not synchronized with the backup Routing Engine when NSR is enabled.
Options	<p>forever—Do not reestablish a BGP peer session after an non-stop routing switchover until the clear bgp neighbor command is issued.</p> <p>seconds—Do not reestablish a BGP peer session after an non-stop routing switchover until after the specified period.</p> <p>Range: 1 through 4,294,967,295 ($2^{32} - 1$)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Preventing Automatic Reestablishment of BGP Peer Sessions After NSR Switchovers</i> • <i>Routing Policy Configuration Guide</i> • <i>Junos OS High Availability Configuration Guide</i>

import (Protocols BGP)

Syntax	<code>import [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i></code> <code> <i>neighbor address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp <i>group group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp group <i>group-name</i> <i>neighbor address</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocols bgp <i>group group-name</i>],</code> <code>[edit protocols bgp group <i>group-name</i> <i>neighbor address</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i></code> <code> <i>neighbor address</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Apply one or more routing policies to routes being imported into the Junos OS routing table from BGP.
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring BGP Interactions with IGP</i>s• <i>Understanding Route Advertisement</i>• <i>Importing and Exporting Routes</i>• <i>Routing Policy Configuration Guide</i>• export on page 2586

include-mp-next-hop

Syntax	include-mp-next-hop;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Enable multiprotocol updates to contain next-hop reachability information.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Examples: Configuring Multiprotocol BGP</i>

keep

Syntax	keep (all none);
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Control whether or not Junos OS keeps in memory and hides certain routes.</p> <p>If the keep none statement is used, Junos OS does not retain in memory and hide routes that are rejected because of a BGP import policy. Nor does BGP keep in memory and hide routes that are declared unfeasible due to BGP sanity checks. The keep none statement causes Junos OS to discard from memory the routes that are rejected due to BGP-specific logic or BGP evaluation. When a route is rejected because of some non-BGP-specific reason, the keep none statement has no effect on this route. This rejected route is retained in memory and hidden even though keep none is configured. An example of this type of hidden route is a route for which the protocol nexthop is unresolved.</p> <p>The routing table can retain the route information learned from BGP in one of the following ways:</p> <ul style="list-style-type: none"> • Default (omit the keep statement)—Keep all route information that was learned from BGP, except for routes whose AS path is looped and whose loop includes the local AS. • keep all—Keep all route information that was learned from BGP. • keep none—Discard routes that were received from a peer and that were rejected by import policy or other sanity checking, such as AS path or next hop. When you configure keep none for the BGP session and the inbound policy changes, Junos OS forces readvertisement of the full set of routes advertised by the peer.

In an AS path healing situation, routes with looped paths theoretically could become usable during a soft reconfiguration when the AS path loop limit is changed. However, there is a significant memory usage difference between the default and **keep all**.

Consider the following scenarios:

- A peer readvertises routes back to the peer from which it learned them.

This can happen in the following cases:

- Another vendor's routing device advertises the routes back to the sending peer.
- The Junos OS peer's default behavior of not readvertising routes back to the sending peer is overridden by configuring **advertise-peer-as**.
- A provider edge (PE) routing device discards any VPN route that does not have any of the expected route targets.

When **keep all** is configured, the behavior of discarding routes received in the above scenarios is overridden.



CAUTION: When you configure **keep (all | none)**, the associated BGP sessions are restarted.

Default By default, BGP retains incoming rejected routes in memory and hides them. If you do not include the **keep** statement, most routes are retained in the routing table. BGP keeps all route information that was learned from BGP, except for routes whose AS path is looped and whose loop includes the local AS.

Options **all**—Retain all routes.

none—Discard routes that were received from a peer and that were rejected by import policy or other sanity checking. When **keep none** is configured for the BGP session and the inbound policy changes, Junos OS forces readvertisement of the full set of routes advertised by the peer.

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

Related Documentation

- *Understanding Route Advertisement*
- [out-delay on page 2625](#)

labeled-unicast (Protocols BGP)

Syntax	<pre> labeled-unicast { accepted-prefix-limit { maximum <i>number</i>; teardown <<i>percentage</i>> <idle-timeout (forever <i>minutes</i>)>; } aggregate-label { community <i>community-name</i>; } explicit-null { connected-only; } prefix-limit { maximum <i>number</i>; teardown <<i>percentage</i>> <idle-timeout (forever <i>minutes</i>)>; } resolve-vpn; rib (inet.3 inet6.3); rib-group <i>group-name</i>; } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols bgp <i>family</i> (inet inet6)], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> <i>family</i> (inet inet6)], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> (inet inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>family</i> (inet inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> (inet inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> (inet inet6)], [edit protocols bgp <i>family</i> (inet inet6)], [edit protocols bgp group <i>group-name</i> <i>family</i> (inet inet6)], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> (inet inet6)], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>family</i> (inet inet6)], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <i>family</i> (inet inet6)], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <i>family</i> (inet inet6)] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure the family type to be labeled-unicast.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

Related Documentation • *Examples: Configuring Multiprotocol BGP*

local-address (Protocols BGP)

Syntax	<code>local-address address;</code>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp <i>group group-name neighbor address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the address of the local end of a BGP session. This address is used to accept incoming connections to the peer and to establish connections to the remote peer. When none of the operational interfaces are configured with the specified local address, a session with a BGP peer is placed in the idle state.</p> <p>You generally configure a local address to explicitly configure the system's IP address from BGP's point of view. This IP address can be either an IPv6 or IPv4 address. Typically, an IP address is assigned to a loopback interface, and that IP address is configured here.</p> <p>For internal BGP (IBGP) peering sessions, generally the loopback interface (lo0) is used to establish connections between the IBGP peers. The loopback interface is always up as long as the device is operating. If there is a route to the loopback address, the IBGP peering session stays up. If a physical interface address is used instead and that interface goes up and down, the IBGP peering session also goes up and down. Thus, the loopback interface provides fault tolerance in case the physical interface or the link goes down, if the device has link redundancy.</p> <p>When a device peers with a remote device's loopback interface address, the local device expects BGP update messages to come from (be sourced by) the remote device's loopback interface address. The local-address statement enables you to specify the source information in BGP update messages. If you omit the local-address statement, the expected source of BGP update messages is based on the device's source address selection rules, which normally result in the egress interface address being the expected source of update messages. When this happens, the peering session is not established because a mismatch exists between the expected source address (the egress interface</p>

of the peer) and the actual source (the loopback interface of the peer). To ensure that the expected source address matches the actual source address, specify the loopback interface address in the **local-address** statement.



NOTE: Although a BGP session can be established when only one of the paired routing devices has **local-address** configured, we strongly recommend that you configure **local-address** on both paired routing devices for IBGP and multihop EBGP sessions. The **local-address** statement ensures that deterministic fixed addresses are used for the BGP session end-points.

If you include the **default-address-selection** statement in the configuration, the software chooses the system default address as the source for most locally generated IP packets. For protocols in which the local address is unconstrained by the protocol specification, for example IBGP and multihop EBGP, if you do not configure a specific local address when configuring the protocol, the local address is chosen using the same methods as other locally generated IP packets.

Default If you do not configure a local address, BGP uses the routing device's source address selection rules to set the local address.

Options **address**—IPv6 or IPv4 address of the local end of the connection.

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

Related Documentation

- [Example: Configuring Internal BGP Peering Sessions on Logical Systems on page 3365](#)
- [Example: Configuring Internal BGP Peer Sessions](#)
- [Understanding Internal BGP Peering Sessions](#)
- [router-id](#)

local-as

Syntax	<code>local-as <i>autonomous-system</i> <loops <i>number</i>> <private alias> <no-prepend-global-as>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <i>group group-name</i>],</p> <p>[edit protocols bgp <i>group group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><i>alias</i> option introduced in Junos OS Release 9.5.</p> <p><i>no-prepend-global-as</i> option introduced in Junos OS Release 9.6.</p>
Description	<p>Specify the local autonomous system (AS) number. An AS is a set of routing devices that are under a single technical administration and generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices.</p> <p>Internet service providers (ISPs) sometimes acquire networks that belong to a different AS. When this occur, there is no seamless method for moving the BGP peers of the acquired network to the AS of the acquiring ISP. The process of configuring the BGP peers with the new AS number can be time-consuming and cumbersome. In this case, it might not be desirable to modify peer arrangements or configuration. During this kind of transition period, it can be useful to configure BGP-enabled devices in the new AS to use the former AS number in BGP updates. This former AS number is called a <i>local</i> AS.</p>



NOTE: If you are using BGP on the routing device, you must configure an AS number before you specify the local as number.

In Junos OS Release 9.1 and later, the AS numeric range in plain-number format is extended to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*.

In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For

example, the 4-byte AS number of 65546 in plain-number format is represented as 1.10 in the AS-dot notation format.

Options **alias**—(Optional) Configure the local AS as an alias of the global AS number configured for the router at the **[edit routing-options]** hierarchy level. As a result, a BGP peer considers any local AS to which it is assigned as equivalent to the primary AS number configured for the routing device. When you use the **alias** option, only the AS (global or local) used to establish the BGP session is prepended in the AS path sent to the BGP neighbor.

autonomous-system—AS number.

Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format

Range: 0.0 through 65535.65535 in AS-dot notation format

loops number—(Optional) Specify the number of times detection of the AS number in the AS_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.



NOTE: If you configure the local AS values for any BGP group, the detection of routing loops is performed using both the AS and the local AS values for all BGP groups.

If the local AS for the EBGP or IBGP peer is the same as the current AS, do not use the **local-as** statement to specify the local AS number.

When you configure the local AS within a VRF, this impacts the AS path loop-detection mechanism. All of the **local-as** statements configured on the device are part of a single AS domain. The AS path loop-detection mechanism is based on looking for a matching AS present in the domain.

Range: 1 through 10

Default: 1

no-prepend-global-as—(Optional) Specify to strip the global AS and to prepend only the local AS in AS paths sent to external peers.

private—(Optional) Configure to use the local AS only during the establishment of the BGP session with a BGP neighbor but to hide it in the AS path sent to external BGP peers. Only the global AS is included in the AS path sent to external peers.



NOTE: The **private** and **alias** options are mutually exclusive. You cannot configure both options with the same **local-as** statement.

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

Related Documentation

- *Examples: Configuring BGP Local AS*
- *Example: Configuring a Local AS for EBGp Sessions*
- [autonomous-system on page 2569](#)
- [family on page 2587](#)

local-interface (IPv6)

Syntax local-interface *interface-name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp [group group-name neighbor ipv6-link-local-address](#)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* [neighbor ipv6-link-local-address](#)],
[edit protocols bgp group *group-name* [neighbor ipv6-link-local-address](#)],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* [neighbor ipv6-link-local-address](#)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Specify the interface name of the EBGp peer that uses IPv6 link-local addresses. This peer is link-local in scope.

Options *interface-name*—Interface name of the EBGp IPv6 peer.

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

Related Documentation

- [Example: Configuring Internal BGP Peering Sessions on Logical Systems on page 3365](#)
- *Example: Configuring Internal BGP Peer Sessions*
- [Example: Configuring External BGP on Logical Systems with IPv6 Interfaces on page 3375](#)
- *Understanding Internal BGP Peering Sessions*

local-preference

Syntax	<code>local-preference local-preference;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Modify the value of the LOCAL_PREF path attribute, which is a metric used by IBGP sessions to indicate the degree of preference for an external route. The route with the highest local preference value is preferred.</p> <p>The LOCAL_PREF path attribute always is advertised to internal BGP peers and to neighboring confederations. It is never advertised to external BGP peers.</p>
Default	If you omit this statement, the LOCAL_PREF path attribute, if present, is not modified.
Options	<p>local-preference—Preference to assign to routes learned from BGP or from the group or peer.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: If the LOCAL_PREF path attribute is present, do not modify its value. If a BGP route is received without a LOCAL_PREF attribute, the route is handled locally (it is stored in the routing table and advertised by BGP) as if it were received with a LOCAL_PREF value of 100. By default, non-BGP routes that are advertised by BGP are advertised with a LOCAL_PREF value of 100.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring the Local Preference Value for BGP Routes</i> • <i>Understanding Internal BGP Peering Sessions</i>

- [preference on page 2630](#)

log-updown (Protocols BGP)

Syntax	log-updown;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify to generate a log a message whenever a BGP peer makes a state transition. Messages are logged using the system logging mechanism located at the [edit system syslog] hierarchy level.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Example: Preventing BGP Session Resets</i>• <i>Junos OS System Basics Configuration Guide</i>• traceoptions on page 2637

metric-out (Protocols BGP)

Syntax	<code>metric-out (<i>metric</i> <i>minimum-igp offset</i> <i>igp</i> (<i>delay-med-update</i> <i>offset</i>);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <i>group group-name</i>],</p> <p>[edit protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Option delay-med-update introduced in Junos OS Release 9.0.</p>
Description	<p>Specify the metric for all routes sent using the multiple exit discriminator (MED, or MULTI_EXIT_DISC) path attribute in update messages. This path attribute is used to discriminate among multiple exit points to a neighboring AS. If all other factors are equal, the exit point with the lowest metric is preferred.</p> <p>You can specify a constant metric value by including the metric option. For configurations in which a BGP peer sends third-party next hops that require the local system to perform next-hop resolution—IBGP configurations, configurations within confederation peers, or EBGP configurations that include the multihop command—you can specify a variable metric by including the minimum-igp or igp option.</p> <p>You can increase or decrease the variable metric calculated from the IGP metric (either from the igp or minimum-igp statement) by specifying a value for offset. The metric is increased by specifying a positive value for offset, and decreased by specifying a negative value for offset.</p> <p>In Junos OS Release 9.0 and later, you can specify that a BGP group or peer not advertise updates for the MED path attributes used to calculate IGP costs for BGP next hops unless the MED is lower. You can also configure an interval to delay when MED updates are sent by including the med-igp-update-interval minutes statement at the [edit routing-options] hierarchy level.</p>
Options	<p>delay-med-update—Specify that a BGP group or peer configured with the metric-out igp statement not advertise MED updates unless the current MED value is lower than</p>

the previously advertised MED value, or another attribute associated with the route has changed, or the BGP peer is responding to a refresh route request.



NOTE: You cannot configure the `delay-med-update` statement at the global BGP level.

igp—Set the metric to the most recent metric value calculated in the IGP to get to the BGP next hop. Routes learned from an EBGP peer usually have a next hop on a directly connected interface and thus the IGP value is equal to zero. This is the value advertised.

metric—Primary metric on all routes sent to peers.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: No metric is sent.

minimum-igp—Set the metric to the minimum metric value calculated in the IGP to get to the BGP next hop. If a newly calculated metric is greater than the minimum metric value, the metric value remains unchanged. If a newly calculated metric is lower, the metric value is lowered to that value. When you change a neighbor's export policy from any configuration to a configuration that sets the minimum IGP offset on an exported route, the advertised MED is not updated if the value would increase as a result, even if the previous configuration does not use a minimum IGP-based MED value. This behavior helps to prevent unnecessary route flapping when an IGP cost changes, by not forcing a route update if the metric value increases past the previous lowest known value.

offset—Increases or decreases the metric by this value.

Range: -2^{31} through $2^{31} - 1$

Default: None

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


Related Documentation	<ul style="list-style-type: none">• <i>Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates</i>• <i>Examples: Configuring BGP MED</i>• <i>Example: Configuring the MED Attribute Directly</i>• <i>Understanding the MED Attribute</i>• <i>med-igp-update-interval</i>
------------------------------	---

mtu-discovery

Syntax	mtu-discovery;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure TCP path maximum transmission unit (MTU) discovery.</p> <p>TCP path MTU discovery enables BGP to automatically discover the best TCP path MTU for each BGP session. In Junos OS, TCP path MTU discovery is disabled by default for all BGP neighbor sessions.</p> <p>When MTU discovery is disabled, TCP sessions that are not directly connected transmit packets of 512-byte maximum segment size (MSS). These small packets minimize the chances of packet fragmentation at a device along the path to the destination. However, because most links use an MTU of at least 1500 bytes, 512-byte packets do not result in the most efficient use of link bandwidth. For directly connected EBGP sessions, MTU mismatches prevent the BGP session from being established. As a workaround, enable path MTU discovery within the EBGP group.</p> <p>Path MTU discovery dynamically determines the MTU size on the network path between the source and the destination, with the goal of avoiding IP fragmentation. Path MTU discovery works by setting the Don't Fragment (DF) bit in the IP headers of outgoing packets. When a device along the path has an MTU that is smaller than the packet, the device drops the packet. The device also sends back an ICMP Fragmentation Needed (Type 3, Code 4) message that contains the device's MTU, thus allowing the source to reduce its path MTU appropriately. The process repeats until the MTU is small enough to traverse the entire path without fragmentation.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- *Example: Limiting TCP Segment Size for BGP*
 - *Configuring the Junos OS for IPv6 Path MTU Discovery*
 - *Configuring the Junos OS for Path MTU Discovery on Outgoing GRE Tunnel Connections*

multihop

Syntax	<pre>multihop { no-nexthop-change; ttl <i>ttl-value</i>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure an EBGp multihop session.</p> <p>For Layer 3 VPNs, you configure the EBGp multihop session between the PE and CE routing devices. This allows you to configure one or more routing devices between the PE and CE routing devices.</p> <p>An external confederation peer is a special case that allows unconnected third-party next hops. You do not need to configure multihop sessions explicitly in this particular case because multihop behavior is implied.</p> <p>If you have external BGP confederation peer-to-loopback addresses, you still need the multihop configuration.</p>
	<div>  <p>NOTE: You cannot configure the <code>accept-remote-nexthop</code> statement at the same time.</p> </div>
Default	<p>If you omit this statement, all EBGp peers are assumed to be directly connected (that is, you are establishing a nonmultihop, or “regular,” BGP session), and the default time-to-live (TTL) value is 1.</p>

The remaining statements are explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring EBGp Multihop Sessions</i>• <i>Configuring EBGp Multihop Sessions Between PE and CE Routers in Layer 3 VPNs</i>• accept-remote-nextthop on page 2559• <i>no-nextthop-change</i>• <i>tth</i>

multipath (Protocols BGP)

Syntax	<pre> multipath { multiple-as; vpn-unequal-cost equal-external-internal; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Allow load sharing among multiple EBGP paths and multiple IBGP paths. A path is considered a BGP equal-cost path (and will be used for forwarding) if a tie-break is performed. The tie-break is performed after the BGP route path selection step that chooses the next-hop path that is resolved through the IGP route with the lowest metric. All paths with the same neighboring AS, learned by a multipath-enabled BGP neighbor, are considered.</p>
Options	<p>multiple-as—Disable the default check requiring that paths accepted by BGP multipath must have the same neighboring AS.</p> <p>vpn-unequal-cost equal-external-internal—Enable load-balancing in a Layer 3 VPN with unequal cost paths.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding BGP Path Selection</i> • <i>Example: Load Balancing BGP Traffic</i>

neighbor (Protocols BGP)

```
Syntax  neighbor address {
    accept-remote-nexthop;
    advertise-external <conditional>;
    advertise-inactive;
    (advertise-peer-as | no-advertise-peer-as);
    as-override;
    authentication-algorithm algorithm;
    authentication-key key;
    authentication-key-chain key-chain;
    cluster cluster-identifier;
    damping;
    description text-description;
    export [ policy-names ];
    family {
        (inet | inet6 | inet-mvpn | inet6-mpvn | inet-vpn | inet6-vpn | iso-vpn | l2-vpn) {
            (any | flow | multicast | unicast | signaling) {
                accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                damping;
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                rib-group group-name;
                topology name {
                    community {
                        target identifier;
                    }
                }
            }
        }
        flow {
            no-validate policy-name;
        }
        labeled-unicast {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            aggregate-label {
                community community-name;
            }
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            resolve-vpn;
            rib inet.3;
        }
    }
}
```



```

    rib-group group-name;
    topology name {
        community {
            target identifier;
        }
    }
}
route-target {
    advertise-default;
    external-paths number;
    accepted-prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
}
signaling {
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
}
}
graceful-restart {
    disable;
    restart-time seconds;
    stale-routes-time seconds;
}
hold-time seconds;
import [ policy-names ];
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
local-interface interface-name;
local-preference preference;
log-updown;
metric-out (metric | minimum-igp <offset> | igp <offset>);
mtu-discovery;
multihop <ttl-value>;
multipath {
    multiple-as;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
passive;
peer-as autonomous-system;
preference preference;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
}

```

```
    flag flag <flag-modifier> <disable>;  
  }  
  vpn-apply-export;  
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
bgp **group** *group-name*],
[edit protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Explicitly configure a neighbor (peer). To configure multiple BGP peers, include multiple **neighbor** statements.

By default, the peer's options are identical to those of the group. You can override these options by including peer-specific option statements within the **neighbor** statement.

The **neighbor** statement is one of the statements you can include in the configuration to define a minimal BGP configuration on the routing device. (You can include an **allow all** statement in place of a **neighbor** statement.)

Options **address**—IPv6 or IPv4 address of a single peer.

The remaining statements are explained separately.

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

Related Documentation

- *BGP Configuration Guide*

no-aggregator-id

Syntax	no-aggregator-id;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Prevent different routing devices within an AS from creating aggregate routes that contain different AS paths.</p> <p>Junos OS performs route aggregation, which is the process of combining the characteristics of different routes so that only a single route is advertised. Aggregation reduces the amount of information that BGP must store and exchange with other BGP systems. When aggregation occurs, the local routing device adds the local AS number and the router ID to the aggregator path attribute. The no-aggregator-id statement causes Junos OS to place a 0 in the router ID field and thus eliminate the possibility of having multiple aggregate advertisements in the network, each with different path information.</p>
Default	If you omit this statement, the router ID is included in the BGP aggregator path attribute.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>BGP Messages Overview</i>

no-client-reflect

Syntax	no-client-reflect;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Disable intracluster route redistribution by the system acting as the route reflector. Include this statement when the client cluster is fully meshed to prevent the sending of redundant route advertisements. Route reflection provides a way to decrease BGP control traffic and minimizing the number of update messages sent within the AS.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring BGP Route Reflectors</i>• cluster on page 2579

no-validate

Syntax	<code>no-validate <i>policy-name</i>;</code>
Hierarchy Level	<code>[edit protocols bgp group <i>group-name</i> family (inet inet flow)],</code> <code>[edit protocols bgp group <i>group-name</i> neighbor address <i>family</i> (inet inet flow)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family (inet inet flow)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor address <i>family</i> (inet inet flow)]</code>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>When BGP is carrying flow-specification network layer reachability information (NLRI) messages, the no-validate statement omits the flow route validation procedure after packets are accepted by a policy.</p> <p>The receiving BGP-enabled device accepts a flow route if it passes the following criteria:</p> <ul style="list-style-type: none"> • The originator of a flow route matches the originator of the best match unicast route for the destination address that is embedded in the route. • There are no more specific unicast routes, when compared to the destination address of the flow route, for which the active route has been received from a different next-hop autonomous system. <p>The first criterion ensures that the filter is being advertised by the next-hop used by unicast forwarding for the destination address embedded in the flow route. For example, if a flow route is given as 10.1.1.1, proto=6, port=80, the receiving BGP-enabled device selects the more specific unicast route in the unicast routing table that matches the destination prefix 10.1.1.1/32. On a unicast routing table containing 10.1/16 and 10.1.1/24, the latter is chosen as the unicast route to compare against. Only the active unicast route entry is considered. This follows the concept that a flow route is valid if advertised by the originator of the best unicast route.</p> <p>The second criterion addresses situations in which a given address block is allocated to different entities. Flows that resolve to a best-match unicast route that is an aggregate route are only accepted if they do not cover more specific routes that are being routed to different next-hop autonomous systems.</p> <p>You can bypass the validation process and use your own specific import policy. To disable the validation procedure and use an import policy instead, include the no-validate statement in the configuration.</p> <p>Flow routes configured for VPNs with family inet-vpn are not automatically validated, so the no-validate statement is not supported at the <code>[edit protocols bgp group <i>group-name</i> family inet-vpn]</code> hierarchy level. No validation is needed if the flow routes are configured locally between devices in a single AS.</p>

Options *policy-name*—Import policy to match NLRI messages.


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

Related • *Example: Configuring Flow Routes*
Documentation

out-delay

Syntax	<code>out-delay seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify how long a route must be present in the Junos OS routing table before it is exported to BGP. Use this time delay to help bundle routing updates.</p> <p>When configured, the out-delay value displays as Outbound Timer when using show bgp group or show bgp group neighbor commands.</p>
Default	If you omit this statement, routes are exported to BGP immediately after they have been added to the routing table.
Options	<p>seconds—Output delay time.</p> <p>Range: 0 through 65,535 seconds</p> <p>Default: 0 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Understanding Route Advertisement</i>

outbound-route-filter

Syntax	<pre> outbound-route-filter { bgp-orf-cisco-mode; prefix-based { accept { (inet inet6); } } } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>] </pre>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure a BGP peer to accept outbound route filters from a remote peer.
Options	<p>accept—Specify that outbound route filters from a BGP peer be accepted.</p> <p>inet—Specify that IPv4 prefix-based outbound route filters be accepted.</p> <p>inet6—Specify that IPv6 prefix-based outbound route filters be accepted.</p>
	<p> NOTE: You can specify that both IPv4 and IPv6 outbound route filters be accepted.</p>
	<p>prefix-based—Specify that prefix-based filters be accepted.</p> <p>The bgp-orf-cisco-mode statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

Related Documentation • *Example: Configuring BGP Prefix-Based Outbound Route Filtering*

passive (Protocols BGP)

Syntax	passive;
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>], [edit protocols bgp], [edit protocols bgp <i>group</i> <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure the routing device so that active open messages are not sent to the peer. Once you configure the routing device to be passive, the routing device will wait for the peer to issue an open request before a message is sent.
Default	If you omit this statement, all explicitly configured peers are active, and each peer periodically sends open requests until its peer responds.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	• <i>Example: Preventing BGP Session Flaps When VPN Families Are Configured</i>

peer-as (Protocols BGP)

Syntax	<code>peer-as <i>autonomous-system</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the neighbor (peer) autonomous system (AS) number.</p> <p>For EBGP, the peer is in another AS, so the AS number you specify in the peer-as statement must be different from the local router's AS number, which you specify in the autonomous-system statement. For IBGP, the peer is in the same AS, so the two AS numbers that you specify in the autonomous-system and peer-as statements must be the same.</p> <p>The AS numeric range in plain-number format has been extended in Junos OS Release 9.1 to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, <i>BGP Support for Four-octet AS Number Space</i>. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of the Junos OS support 2-byte AS numbers.</p> <p>In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <i><16-bit high-order value in decimal>.<16-bit low-order value in decimal></i>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.</p> <p>With the introduction of 4-byte AS numbers, you might have a combination of routers that support 4-byte AS numbers and 2-byte AS numbers. For more information about what happens when establishing BGP peer relationships between 4-byte and 2-byte capable routers, see the following topics:</p>

- *Establishing a Peer Relationship Between a 4-Byte Capable Router and a 2-Byte Capable Router Using a 2-Byte AS Number in the Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview.*
- *Establishing a Peer Relationship Between a 4-Byte Capable Router and a 2-Byte Capable Router Using a 4-Byte AS Number in the Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview.*

Options *autonomous-system*—AS number.

Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format for 4-byte AS numbers

Range: 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)

Range: 0.0 through 65535.65535 in AS-dot notation format for 4-byte AS numbers

Required Privilege routing—To view this statement in the configuration.

Level routing-control—To add this statement to the configuration.

Related Documentation

- *4-Byte Autonomous System Numbers Overview* in the [Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview](#)

- *Juniper Networks Implementation of 4-Byte Autonomous System Numbers* in the [Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview](#)

preference (Protocols BGP)

Syntax	<code>preference preference;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the preference for routes learned from BGP.</p> <p>At the BGP global level, the preference statement sets the preference for routes learned from BGP. You can override this preference in a BGP group or peer preference statement.</p> <p>At the group or peer level, the preference statement sets the preference for routes learned from the group or peer. Use this statement to override the preference set in the BGP global preference statement when you want to favor routes from one group or peer over those of another.</p>
Options	<p>preference—Preference to assign to routes learned from BGP or from the group or peer.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 170 for the primary preference</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • local-preference on page 2609 • <i>Example: Configuring the Preference Value for BGP Routes</i>

prefix-limit

Syntax	<pre>prefix-limit { maximum <i>number</i>; teardown <<i>percentage</i>> <idle-timeout (forever <i>minutes</i>)>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit protocols bgp family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit protocols bgp group <i>group-name</i> family (inet inet6) (any labeled-unicast multicast unicast)],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family (inet inet6) (any flow labeled-unicast multicast unicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6) (any flow labeled-unicast multicast unicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Limit the number of prefixes received on a BGP peer session and a rate-limit logging when injected prefixes exceed a set limit.
Options	<p>maximum <i>number</i>—When you set the maximum number of prefixes, a message is logged when that number is exceeded.</p> <p>Range: 1 through 4,294,967,295 ($2^{32} - 1$)</p> <p>teardown <<i>percentage</i>>—If you include the teardown statement, the session is torn down when the maximum number of prefixes is reached. If you specify a percentage, messages are logged when the number of prefixes exceeds that percentage. After the session is torn down, it is reestablished in a short time unless you include the idle-timeout statement. Then the session can be kept down for a specified amount of time, or forever. If you specify forever, the session is reestablished only after you issue a clear bgp neighbor command.</p> <p>Range: 1 through 100</p>

idle-timeout (**forever** | *timeout-in-minutes*)—(Optional) If you include the **idle-timeout** statement, the session is torn down for a specified amount of time, or forever. If you specify a period of time, the session is allowed to reestablish after this timeout period. If you specify **forever**, the session is reestablished only after you intervene with a **clear bgp neighbor** command.

Range: 1 through 2400

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

Related Documentation	<ul style="list-style-type: none">• <i>accepted-prefix-limit</i>• <i>Understanding Multiprotocol BGP</i>
------------------------------	---

remove-private

Syntax `remove-private all replace nearest;`

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
 [edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
 [edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp *group* *group-name* **neighbor** *address*],
 [edit protocols bgp],
 [edit protocols bgp **group** *group-name*],
 [edit protocols bgp *group* *group-name* **neighbor** *address*],
 [edit routing-instances *routing-instance-name* protocols bgp],
 [edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
 [edit routing-instances *routing-instance-name* protocols bgp *group* *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description When advertising AS paths to remote systems, have the local system strip private AS numbers from the AS path. The numbers are stripped from the AS path starting at the left end of the AS path (the end where AS paths have been most recently added). The routing device stops searching for private ASs when it finds the first nonprivate AS or a peer's private AS. If the AS path contains the AS number of the external BGP (EBGP) neighbor, BGP does not remove the private AS number.



NOTE: As of Junos OS 10.0R2 and higher, if there is a need to send prefixes to an EBGP peer that has an AS number that matches an AS number in the AS path, consider using the `as-override` statement instead of the `remove-private` statement.

The operation takes place after any confederation member ASs have already been removed from the AS path, if applicable.

The Junos OS recognizes the set of AS numbers that is considered private, a range that is defined in the Internet Assigned Numbers Authority (IANA) assigned numbers document.

The set of reserved AS numbers is in the range from 64,512 through 65,535.

Options **all**—Remove all private AS numbers from the original path. Do not stop the process of removing private AS numbers, even if a public AS number is encountered.

nearest—When you use the **all** and **replace** options, choose the last (right-most) public AS number encountered in the original AS path for the replacement value, as the AS path is processed from left to right. If no public AS number is encountered, the default replacement value is used. (See the **replace** option for information about the default replacement value.)

replace—When you use the **all** option, instead of removing private AS numbers, perform a replace operation. The default replacement value for the private AS number is the local AS number at the BGP group level for the BGP peer. If you are unsure about the replacement value, check the local AS value displayed in the output of the **show bgp group group-name** command.

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

Related Documentation	<ul style="list-style-type: none">• <i>Example: Removing Private AS Numbers from AS Paths</i>
------------------------------	---


rib-group (Protocols BGP)

Syntax	<code>rib-group group-name;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp family inet (labeled-unicast unicast multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family inet (labeled-unicast unicast multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet (labeled-unicast unicast multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family inet (labeled-unicast unicast multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family inet (labeled-unicast unicast multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet (labeled-unicast unicast multicast)],</p> <p>[edit protocols bgp family inet (labeled-unicast unicast multicast)],</p> <p>[edit protocols bgp group <i>group-name</i> family inet (labeled-unicast unicast multicast)],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet (labeled-unicast unicast multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp family inet (labeled-unicast unicast multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family inet (labeled-unicast unicast multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet (labeled-unicast unicast multicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Add unicast prefixes to unicast and multicast tables.
Options	group-name —Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens. You generally specify only one routing table group.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i> • <i>Example: Importing Direct and Static Routes Into a Routing Instance</i> • <i>Understanding Multiprotocol BGP</i>

tcp-mss (Protocols BGP)

Syntax	<code>tcp-mss <i>segment-size</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor</code> <code> <i>neighbor-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp group <i>group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> bgp group <i>group-name</i> neighbor <i>neighbor-name</i>],</code> <code>[edit protocols bgp],</code> <code>[edit protocol bgp group <i>group-name</i>],</code> <code>[edit protocols bgp group <i>group-name</i> neighbor <i>neighbor-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor</code> <code> <i>neighbor-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure the maximum segment size (MSS) for the TCP connection for BGP neighbors. The MSS is only valid in increments of 2 KB. The value used is based on the value set, but is rounded down to the nearest multiple of 2048.
Options	<i>segment-size</i> —MSS for the TCP connection. Range: 1 through 4096
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Limiting TCP Segment Size for BGP</i>

traceoptions (Protocols BGP)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp <i>group</i> <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>4byte-as statement introduced in Junos OS Release 9.2.</p> <p>4byte-as statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	Configure BGP protocol-level tracing options. To specify more than one tracing operation, include multiple flag statements.
<div>  <p>NOTE: The traceoptions statement is not supported on QFabric systems.</p> </div>	
Default	The default BGP protocol-level tracing options are inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level. The default group-level trace options are inherited from the BGP protocol-level traceoptions statement. The default peer-level trace options are inherited from the group-level traceoptions statement.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place BGP tracing output in the file bgp-log.</p>

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file.0***, then ***trace-file.1***, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

BGP Tracing Flags

- **4byte-as**—4-byte AS events.
- **bfd**—BFD protocol events.
- **damping**—Damping operations.
- **graceful-restart**—Graceful restart events.
- **keepalive**—BGP keepalive messages. If you enable the the BGP **update** flag only, received keepalive messages do not generate a trace message.
- **nsr-synchronization**—Nonstop routing synchronization events.
- **open**—Open packets. These packets are sent between peers when they are establishing a connection.
- **packets**—All BGP protocol packets.
- **refresh**—BGP refresh packets.
- **update**—Update packets. These packets provide routing updates to BGP systems. If you enable only this flag, received keepalive messages do not generate a trace message. Use the **keepalive** flag to generate a trace message for keepalive messages.

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **filter**—Provide filter trace information. Applies only to **route** and **damping** tracing flags.
- **receive**—Trace the packets being received.
- **send**—Trace the packets being transmitted.

no-world-readable—(Optional) Prevent any user from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • log-updown on page 2610 statement • <i>Understanding Trace Operations for BGP Protocol Traffic</i> • <i>Configuring OSPF Refresh and Flooding Reduction in Stable Topologies</i>

type (Protocols BGP)

Syntax	<code>type type;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Specify the type of BGP peer group.</p> <p>When configuring a BGP group, you can indicate whether the group is an IBGP group or an EBGP group. All peers in an IBGP group are in the same AS, while peers in an EBGP group are in different ASs and normally share a subnet.</p>
Options	<p>type—Type of group:</p> <ul style="list-style-type: none">• external—External group, which allows inter-AS BGP routing• internal—Internal group, which allows intra-AS BGP routing
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>BGP Configuration Guide</i>

Administration

- [Operational Commands on page 2640](#)

[Operational Commands](#)

clear bgp damping

Syntax	clear bgp damping <logical-system (all <i>logical-system-name</i>)> < <i>prefix</i> >
Syntax (EX Series Switch and QFX Series)	clear bgp damping < <i>prefix</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear BGP route flap damping information.
Options	<p>none—Clear all BGP route flap damping information.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>prefix</i>—(Optional) Clear route flap damping information for only the specified destination prefix.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show policy damping on page 2674 • <i>show route damping</i>
List of Sample Output	clear bgp damping on page 2641
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bgp damping

```
user@host> clear bgp damping
```

clear bgp neighbor

Syntax	<pre>clear bgp neighbor <as <i>as-number</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor> <soft soft-inbound> <soft-minimum-igp></pre>
Syntax (EX Series Switch and QFX Series)	<pre>clear bgp neighbor <as <i>as-number</i>> <instance <i>instance-name</i>> <neighbor> <soft soft-inbound> <soft-minimum-igp></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Perform one of the following tasks:</p> <ul style="list-style-type: none">• Change the state of one or more BGP neighbors to IDLE. For neighbors in the ESTABLISHED state, this command drops the TCP connection to the neighbors and then reestablishes the connection.• (soft or soft-inbound keyword only) Reapply export policies or import policies, respectively, and send refresh updates to one or more BGP neighbors without changing their state.
Options	<p>none—Change the state of all BGP neighbors to IDLE.</p> <p>as <i>as-number</i>—(Optional) Apply this command only to neighbors in the specified autonomous system (AS).</p> <p>instance <i>instance-name</i>—(Optional) Apply this command only to neighbors for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor—(Optional) IP address of a BGP peer. Apply this command only to the specified neighbor.</p> <p>soft—(Optional) Reapply any export policies and send refresh updates to neighbors without clearing the state.</p> <p>soft-inbound—(Optional) Reapply any import policies and send refresh updates to neighbors without clearing the state.</p>

soft-minimum-igp—(Optional) Provides soft refresh of the outbound state when the interior gateway protocol (IGP) metric is reset.

Required Privilege Level clear

Related Documentation • [show bgp neighbor on page 2655](#)

List of Sample Output [clear bgp neighbor on page 2643](#)

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

Sample Output

[clear bgp neighbor](#)

```
user@host> clear bgp neighbor
```

clear bgp table

Syntax	<code>clear bgp table <i>table-name</i></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switch and QFX Series)	<code>clear bgp table <i>table-name</i></code>
Release Information	Command introduced in Junos OS Release 9.0. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Request that BGP refresh routes in a specified routing table.
Options	<code>logical-system (all <i>logical-system-name</i>)</code> —(Optional) Perform this operation on all logical systems or on a particular logical system. <code>table-name</code> —Request that BGP refresh routes in the specified table.
Additional Information	In some cases, a prefix limit is associated with a routing table for a VPN instance. When this limit is exceeded (for example, because of a network misconfiguration), some routes might not be inserted in the table. Such routes need to be added to the table after the network issue is resolved. Use the clear bgp table command to request that BGP refresh routes in a VPN instance table.
Required Privilege Level	clear
List of Sample Output	clear bgp table private.inet.0 on page 2644 clear bgp table inet.6 logical-system all on page 2644 clear bgp table private.inet.6 logical-system ls1 on page 2644 clear bgp table logical-system all inet.0 on page 2644 clear bgp table logical-system ls2 private.inet.0 on page 2645
Output Fields	This command produces no output.

Sample Output

[clear bgp table private.inet.0](#)

```
user@host> clear bgp table private.inet.0
```

[clear bgp table inet.6 logical-system all](#)

```
user@host> clear bgp table inet.6 logical-system all
```

[clear bgp table private.inet.6 logical-system ls1](#)

```
user@host> clear bgp table private.inet.6 logical-system ls1
```

[clear bgp table logical-system all inet.0](#)

```
user@host> clear bgp table logical-system all inet.0
```

`clear bgp table logical-system ls2 private.inet.0`

`user@host> clear bgp table logical-system ls2 private.inet.0`

show bgp bmp

Syntax	show bgp bmp
Release Information	Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display information about the BGP Monitoring Protocol (BMP).
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show bgp bmp on page 2646
Output Fields	Table 212 on page 2646 lists the output fields for the show bgp bmp command. Output fields are listed in the approximate order in which they appear.

Table 212: show bgp bmp Output Fields

Field Name	Field Description
BMP station address/port	IP address and port number of the monitoring station to which BGP Monitoring Protocol (BMP) statistics are sent.
BMP session state	Status of the BMP session: UP or DOWN .
Memory consumed by BMP	Memory used by the active BMP session.
Statistics timeout	Amount of time, in seconds, between transmissions of BMP data to the monitoring station.
Memory limit	Threshold, in bytes, at which the routing device stops collecting BMP data.
Memory-connect retry timeout	Amount of time, in seconds, after which the routing device attempts to resume a BMP session that was ended after the configured memory threshold was exceeded.

Sample Output

show bgp bmp

```

user@host> show bgp bmp
  BMP station address/port: 172.24.24.157+5454
  BMP session state: DOWN
  Memory consumed by BMP: 0
  Statistics timeout: 15
  Memory limit: 10485760
  Memory connect retry timeout: 600

```


show bgp group

Syntax	<pre>show bgp group <brief detail summary> <group-name> <exact-instance instance-name> <instance instance-name> <logical-system (all logical-system-name)> <rtf></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bgp group <brief detail summary> <group-name> <exact-instance instance-name> <instance instance-name></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p>
Description	Display information about the configured BGP groups.
Options	<p>none—Display group information about all BGP groups.</p> <p>brief detail summary—(Optional) Display the specified level of output.</p> <p>group-name—(Optional) Display group information for the specified group.</p> <p>exact-instance instance-name—(Optional) Display information for the specified instance only.</p> <p>instance instance-name—(Optional) Display information about BGP groups for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp group instance cust1 command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>rtf—(Optional) Display BGP group route targeting information.</p>
Required Privilege Level	view
List of Sample Output	<p>show bgp group on page 2652</p> <p>show bgp group brief on page 2652</p> <p>show bgp group detail on page 2653</p> <p>show bgp group rtf detail on page 2654</p> <p>show bgp group summary on page 2654</p>

Output Fields Table 213 on page 2649 describes the output fields for the **show bgp group** command. Output fields are listed in the approximate order in which they appear.

Table 213: show bgp group Output Fields

Field Name	Field Description	Level of Output
Group Type or Group	Type of BGP group: Internal or External .	All levels
group-index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.	rtf detail
AS	AS number of the peer. For internal BGP (IBGP), this number is the same as Local AS .	brief detail none
Local AS	AS number of the local routing device.	brief detail none
Name	Name of a specific BGP group.	brief detail none
Index	Unique index number of a BGP group.	brief detail none
Flags	Flags associated with the BGP group. This field is used by Juniper Networks customer support.	brief detail none
Remove-private options	Options associated with the remove-private statement.	brief detail none
Holdtime	Maximum number of seconds allowed to elapse between successive keepalive or update messages that BGP receives from a peer in the BGP group, after which the connection to the peer is closed and routing devices through that peer become unavailable.	brief detail none
Export	Export policies configured for the BGP group with the export statement.	brief detail none
MED tracks IGP metric update delay	Time, in seconds, that updates to multiple exit discriminator (MED) are delayed. Also displays the time remaining before the interval is set to expire	All levels
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.	brief detail none
Total peers	Total number of peers in the group.	brief detail none
Established	Number of peers in the group that are in the established state.	All levels

Table 213: show bgp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active/Received/Accepted/Damped	<p>Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established in the main routing device or in a routing instance.</p> <ul style="list-style-type: none"> If a peer is not established, the field shows the state of the peer session: Active, Connect, or Idle. If a BGP session is established in the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the inet.0 (main) and inet.2 (multicast) routing tables. For example, 8/10/10/2 and 2/4/4/0 indicate the following: <ul style="list-style-type: none"> 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the inet.0 routing table. 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the inet.2 routing table. 	summary
ip-addresses	List of peers who are members of the group. The address is followed by the peer's port number.	All levels
Route Queue Timer	Number of seconds until queued routes are sent. If this time has already elapsed, this field displays the number of seconds by which the updates are delayed.	detail
Route Queue	Number of prefixes that are queued up for sending to the peers in the group.	detail
inet.number	<p>Number of active, received, accepted, and damped routes in the routing table. For example, inet.0: 7/10/9/0 indicates the following:</p> <ul style="list-style-type: none"> 7 active routes, 10 received routes, 9 accepted routes, and no damped routes from a BGP peer appear in the inet.0 routing table. 	none

Table 213: show bgp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
Table inet.number	Information about the routing table. <ul style="list-style-type: none"> • Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • Active prefixes—Number of prefixes received from the peer that are active in the routing table. • Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols. • Advertised prefixes—Number of prefixes advertised to a peer. • Received external prefixes—Total number of prefixes from the external BGP (EBGP) peers, both active and inactive, that are in the routing table. • Active external prefixes—Number of prefixes received from the EBGP peers that are active in the routing table. • Externals suppressed—Number of routes received from EBGP peers currently inactive because of damping or other reasons. • Received internal prefixes—Total number of prefixes from the IBGP peers, both active and inactive, that are in the routing table. • Active internal prefixes—Number of prefixes received from the IBGP peers that are active in the routing table. • Internals suppressed—Number of routes received from IBGP peers currently inactive because of damping or other reasons. • RIB State—Status of the graceful restart process for this routing table: BGP restart is complete, BGP restart in progress, VPN restart in progress, or VPN restart is complete. 	detail
Groups	Total number of groups.	All levels
Peers	Total number of peers.	All levels
External	Total number of external peers.	All levels
Internal	Total number of internal peers.	All levels
Down peers	Total number of unavailable peers.	All levels
Flaps	Total number of flaps that occurred.	All levels
Table	Name of a routing table.	brief , none
Tot Paths	Total number of routes.	brief , none
Act Paths	Number of active routes.	brief , none
Suppressed	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.	brief , none

Table 213: show bgp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
History	Number of withdrawn routes stored locally to keep track of damping history.	brief, none
Damp State	Number of active routes with a figure of merit greater than zero, but lower than the threshold at which suppression occurs.	brief, none
Pending	Routes being processed by the BGP import policy.	brief, none
Group	Group the peer belongs to in the BGP configuration.	detail
Receive mask	Mask of the received target included in the advertised route.	detail
Entries	Number of route entries received.	detail
Target	Route target that is to be passed by route-target filtering. If a route advertised from the provider edge (PE) routing device matches an entry in the route-target filter, the route is passed to the peer.	detail
Mask	Mask which specifies that the peer receive routes with the given route target.	detail

Sample Output

show bgp group

```

user@host> show bgp group
Groups: 2  Peers: 2   External: 0   Internal: 2   Down peers: 1   Flaps: 0
Table      Tot Paths  Act Paths  Suppressed   History Damp State   Pending

inet.0
          0         0         0           0         0         0

bgp.13vpn.0
          0         0         0           0         0         0

bgp.rtarget.0
          2         0         0           0         0         0

```

show bgp group brief

```

user@host> show bgp group brief
Groups: 2  Peers: 2   External: 0   Internal: 2   Down peers: 1   Flaps: 0
Table      Tot Paths  Act Paths  Suppressed   History Damp State   Pending

inet.0
          0         0         0           0         0         0

bgp.13vpn.0
          0         0         0           0         0         0

```

```

bgp.rtarget.0
                2          0          0          0          0          0

```

show bgp group detail

```

user@host> show bgp group detail
Group Type: Internal  AS: 1                      Local AS: 1
Name: ibgp           Index: 0                  Flags: <Export Eval>
Holdtime: 0
Total peers: 3       Established: 0
22.0.0.2
22.0.0.8
22.0.0.5

Groups: 1  Peers: 3  External: 0  Internal: 3  Down peers: 3  Flaps: 3
Table bgp.l3vpn.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table bgp.mdt.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table VPN-A.inet.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table VPN-A.mdt.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0

```

```
Externals suppressed:      0
Received internal prefixes: 0
Active internal prefixes:  0
Internals suppressed:      0
RIB State: BGP restart is complete
RIB State: VPN restart is complete
```

show bgp group rtf detail

```
user@host> show bgp group rtf detail
Group: internal (group-index: 0)
  Receive mask: 00000002
  Table: bgp.rtarget.0
    Target      Mask      Entries: 2
    100:100/64   00000002
    200:201/64   (Group)
Group: internal (group-index: 1)
  Table: bgp.rtarget.0
    Target      Mask      Entries: 1
    200:201/64   (Group)
```

show bgp group summary

```
user@host> show bgp group summary
Group      Type      Peers    Established    Active/Received/Accepted/Damped
ibgp       Internal  3         0
Groups: 1  Peers: 3    External: 0    Internal: 3    Down peers: 3    Flaps: 3
bgp.l3vpn.0 : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
bgp.mdt.0   : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
VPN-A.inet.0 : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
VPN-A.mdt.0 : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
```

show bgp neighbor

Syntax	<pre>show bgp neighbor <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor-address> <orf (detail <i>neighbor-address</i>)</pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bgp neighbor <instance <i>instance-name</i>> <exact-instance <i>instance-name</i>> <neighbor-address> <orf (<i>neighbor-address</i> detail)</pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>orf option introduced in Junos OS Release 9.2.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p>
Description	Display information about BGP peers.
Options	<p>none—Display information about all BGP peers.</p> <p>exact-instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>instance <i>instance-name</i>—(Optional) Display information about BGP peers for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp neighbor instance cust1 command).</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor-address—(Optional) Display information for only the BGP peer at the specified IP address.</p> <p>orf (detail <i>neighbor-address</i>)—(Optional) Display outbound route-filtering information for all BGP peers or only for the BGP peer at the specified IP address. The default is to display brief output. Use the detail option to display detailed output.</p>
Additional Information	For information about the local-address , nlri , hold-time , and preference statements, see the <i>Junos OS Routing Protocols Configuration Guide</i> .
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear bgp neighbor on page 2642

List of Sample Output

- [show bgp neighbor on page 2662](#)
- [show bgp neighbor \(CLNS\) on page 2663](#)
- [show bgp neighbor \(Layer 2 VPN\) on page 2663](#)
- [show bgp neighbor \(Layer 3 VPN\) on page 2665](#)
- [show bgp neighbor neighbor-address on page 2666](#)
- [show bgp neighbor neighbor-address on page 2667](#)
- [show bgp neighbor orf neighbor-address detail on page 2668](#)

Output Fields Table 214 on page 2656 describes the output fields for the **show bgp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 214: show bgp neighbor Output Fields

Field Name	Field Description
Peer	Address of the BGP neighbor. The address is followed by the neighbor port number.
AS	AS number of the peer.
Local	Address of the local routing device. The address is followed by the peer port number.
Type	Type of peer: Internal or External .
State	<p>Current state of the BGP session:</p> <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.
Flags	<p>Internal BGP flags:</p> <ul style="list-style-type: none"> • Aggregate Label—BGP has aggregated a set of incoming labels (labels received from the peer) into a single forwarding label. • CleanUp—The peer session is being shut down. • Delete—This peer has been deleted. • Idled—This peer has been permanently idled. • ImportEval—At the last commit operation, this peer was identified as needing to reevaluate all received routes. • Initializing—The peer session is initializing. • SendRtn—Messages are being sent to the peer. • Sync—This peer is synchronized with the rest of the peer group. • TryConnect—Another attempt is being made to connect to the peer. • Unconfigured—This peer is not configured. • WriteFailed—An attempt to write to this peer failed.

Table 214: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Last state	<p>Previous state of the BGP session:</p> <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.
Last event	<p>Last activity that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Closed—The BGP session closed. • ConnectRetry—The transport protocol connection failed, and BGP is trying again to connect. • HoldTime—The session ended because the hold timer expired. • KeepAlive—The local routing device sent a BGP keepalive message to the peer. • Open—The local routing device sent a BGP open message to the peer. • OpenFail—The local routing device did not receive an acknowledgment of a BGP open message from the peer. • RecvKeepAlive—The local routing device received a BGP keepalive message from the peer. • RecvNotify—The local routing device received a BGP notification message from the peer. • RecvOpen—The local routing device received a BGP open message from the peer. • RecvUpdate—The local routing device received a BGP update message from the peer. • Start—The peering session started. • Stop—The peering session stopped. • TransportError—A TCP error occurred.
Last error	<p>Last error that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Cease—An error occurred, such as a version mismatch, that caused the session to close. • Finite State Machine Error—In setting up the session, BGP received a message that it did not understand. • Hold Time Expired—The session's hold time expired. • Message Header Error—The header of a BGP message was malformed. • Open Message Error—A BGP open message contained an error. • None—No errors occurred in the BGP session. • Update Message Error—A BGP update message contained an error.
Export	Name of the export policy that is configured on the peer.
Import	Name of the import policy that is configured on the peer.

Table 214: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Options	Configured BGP options: <ul style="list-style-type: none"> • AddressFamily—Configured address family: inet or inet-vpn. • AuthKeyChain—Authentication key change is enabled. • DropPathAttributes—Certain path attributes are configured to be dropped from neighbor updates during inbound processing. • GracefulRestart—Graceful restart is configured. • HoldTime—Hold time configured with the hold-time statement. The hold time is three times the interval at which keepalive messages are sent. • IgnorePathAttributes—Certain path attributes are configured to be ignored in neighbor updates during inbound processing. • Local Address—Address configured with the local-address statement. • Multihop—Allow BGP connections to external peers that are not on a directly connected network. • NLRI—Configured MBGP state for the BGP group: multicast, unicast, or both if you have configured nlri any. • Peer AS—Configured peer autonomous system (AS). • Preference—Preference value configured with the preference statement. • Refresh—Configured to refresh automatically when the policy changes. • Rib-group—Configured routing table group.
Path-attributes dropped	Path attribute codes that are dropped from neighbor updates.
Path-attributes ignored	Path attribute codes that are ignored during neighbor updates.
Authentication key change	(appears only if the authentication-keychain statement has been configured) Name of the authentication keychain enabled.
Authentication algorithm	(appears only if the authentication-algorithm statement has been configured) Type of authentication algorithm enabled: hmac or md5 .
Address families configured	Names of configured address families for the VPN.
Local Address	Address of the local routing device.
Remove-private options	Options associated with the remove-private statement.
Holdtime	Hold time configured with the hold-time statement. The hold time is three times the interval at which keepalive messages are sent.
Flags for NLRI inet-label-unicast	Flags related to labeled-unicast: <ul style="list-style-type: none"> • TrafficStatistics—Collection of statistics for labeled-unicast traffic is enabled.

Table 214: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Traffic statistics	Information about labeled-unicast traffic statistics: <ul style="list-style-type: none"> • Options—Options configured for collecting statistics about labeled-unicast traffic. • File—Name and location of statistics log files. • size—Size of all the log files, in bytes. • files—Number of log files.
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.
Preference	Preference value configured with the preference statement.
Outbound Timer	Time for which the route is available in Junos OS routing table before it is exported to BGP. This field is displayed in the output only if the out-delay parameter is configured to a non-zero value.
Number of flaps	Number of times the BGP session has gone down and then come back up.
Peer ID	Router identifier of the peer.
Group index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.
Peer index	Index that is unique within the BGP group to which the peer belongs.
Local ID	Router identifier of the local routing device.
Local Interface	Name of the interface on the local routing device.
Active holdtime	Hold time that the local routing device negotiated with the peer.
Keepalive Interval	Keepalive interval, in seconds.
BFD	Status of BFD failure detection.
Local Address	Name of directly connected interface over which direct EBGP peering is established.
NLRI for restart configured on peer	Names of address families configured for restart.
NLRI advertised by peer	Address families supported by the peer: unicast or multicast .
NLRI for this session	Address families being used for this session.
Peer supports Refresh capability	Remote peer's ability to send and request full route table readvertisement (route refresh capability). For more information, see RFC 2918, <i>Route Refresh Capability for BGP-4</i> .
Restart time configured on peer	Configured time allowed for restart on the neighbor.

Table 214: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Stale routes from peer are kept for	When graceful restart is negotiated, the maximum time allowed to hold routes from neighbors after the BGP session has gone down.
Peer does not support Restarter functionality	Graceful restart restarter-mode is disabled on the peer.
Peer does not support Receiver functionality	Graceful restart helper-mode is disabled on the peer.
Restart time requested by this peer	Restart time requested by this neighbor during capability negotiation.
Restart flag received from the peer	When this field appears, the BGP speaker has restarted (Restarting), and this peer should not wait for the end-of-rib marker from the speaker before advertising routing information to the speaker.
NLRI that peer supports restart for	Neighbor supports graceful restart for this address family.
NLRI peer can save forwarding state	Neighbor supporting this address family saves all forwarding states.
NLRI that peer saved forwarding for	Neighbor saves all forwarding states for this address family.
NLRI that restart is negotiated for	Router supports graceful restart for this address family.
NLRI of received end-of-rib markers	Address families for which end-of-routing-table markers are received from the neighbor.
NLRI of all end-of-rib markers sent	Address families for which end-of-routing-table markers are sent to the neighbor.
Peer supports 4 byte AS extension (peer-as 1)	Peer understands 4-byte AS numbers in BGP messages. The peer is running Junos OS Release 9.1 or later.
NLRIs for which peer can receive multiple paths	Appears in the command output of the local router if the downstream peer is configured to receive multiple BGP routes to a single destination, instead of only receiving the active route. Possible value is inet-unicast .
NLRIs for which peer can send multiple paths: inet-unicast	Appears in the command output of the local router if the upstream peer is configured to send multiple BGP routes to a single destination, instead of only sending the active route. Possible value is inet-unicast .

Table 214: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Table inet.number	<p>Information about the routing table:</p> <ul style="list-style-type: none"> • RIB State—BGP is in the graceful restart process for this routing table: restart is complete or restart in progress. • Bit—Number that represents the entry in the routing table for this peer. • Send state—State of the BGP group: in sync, not in sync, or not advertising. • Active prefixes—Number of prefixes received from the peer that are active in the routing table. • Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • Accepted prefixes—Total number of prefixes from the peer that have been accepted by a routing policy. • Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
Last traffic (seconds)	Last time any traffic was received from the peer or sent to the peer, and the last time the local routing device checked.
Input messages	Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Output messages	Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Input dropped path attributes	<p>Information about dropped path attributes:</p> <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count.
Input ignored path attributes	<p>Information about ignored path attributes:</p> <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count.
Output queue	Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue 0 is for unicast NLRIs, and queue 1 is for multicast NLRIs.
Trace options	Configured tracing of BGP protocol packets and operations.
Trace file	Name of the file to receive the output of the tracing operation.
Filter Updates rcv	<p>(orf option only) Number of outbound-route filters received for each configured address family.</p> <p>NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.</p>

Table 214: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Immediate	(orf option only) Number of route updates received with the immediate flag set. The immediate flag indicates that the BGP peer should readvertise the updated routes. NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.
Filter	(orf option only) Type of prefix filter received: prefix-based or extended-community .
Received filter entries	(orf option only) List of received filters displayed.
seq	(orf option only) Numerical order assigned to this prefix entry among all the received outbound route filter prefix entries.
prefix	(orf option only) Address for the prefix entry that matches the filter.
minlength	(orf option only) Minimum prefix length, in bits, required to match this prefix.
maxlength	(orf option only) Maximum prefix length, in bits, required to match this prefix.
match	(orf option only) For this prefix match, whether to permit or deny route updates.

Sample Output

show bgp neighbor

```

user@host > show bgp neighbor
Peer: 10.255.7.250+179 AS 10   Local: 10.255.7.248+63740 AS 10
  Type: Internal   State: Established   Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ redist_static ]
  Options: <Preference LocalAddress PeerAS Refresh>
  Local Address: 10.255.7.248 Holdtime: 90 Preference: 170 Outbound Timer: 50
  Number of flaps: 0
  Peer ID: 10.255.7.250   Local ID: 10.255.7.248   Active Holdtime: 90
  Keepalive Interval: 30   Group index: 0   Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 10)
  Peer does not support Addpath
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:           1
    Received prefixes:         1

```

```

Accepted prefixes:          1
Suppressed due to damping:  0
Advertised prefixes:        1
Last traffic (seconds): Received 9    Sent 5    Checked 5
Input messages:  Total 36    Updates 2    Refreshes 0    Octets 718
Output messages: Total 37    Updates 1    Refreshes 0    Octets 796
Output Queue[0]: 0

Peer: 10.255.162.214+52193 AS 100 Local: 10.255.167.205+179 AS 100
Type: Internal    State: Established (route reflector client)Flags: <Sync>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
Options: <Preference LocalAddress Cluster AddressFamily Rib-group Refresh>
Address families configured: inet-unicast inet-vpn-unicast route-target
Local Address: 10.255.167.205 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.255.162.214    Local ID: 10.255.167.205    Active Holdtime: 90
Keepalive Interval: 30      Group index: 0    Peer index: 1

```

show bgp neighbor (CLNS)

```

user@host> show bgp neighbor
Peer: 10.245.245.1+179 AS 200 Local: 10.245.245.3+3770 AS 100
Type: External    State: Established    Flags: <ImportEval Sync>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS
Rib-group Refresh>
Address families configured: iso-vpn-unicast
Local Address: 10.245.245.3 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.245.245.1    Local ID: 10.245.245.3    Active Holdtime: 90
Keepalive Interval: 30      Peer index: 0
NLRI advertised by peer: iso-vpn-unicast
NLRI for this session: iso-vpn-unicast
Peer supports Refresh capability (2)
Table bgp.isovpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          3
  Received prefixes:        3
  Suppressed due to damping: 0
  Advertised prefixes:      3
Table aaa.iso.0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: not advertising
  Active prefixes:          3
  Received prefixes:        3
  Suppressed due to damping: 0
Last traffic (seconds): Received 6    Sent 5    Checked 5
Input messages:  Total 1736    Updates 4    Refreshes 0    Octets 33385
Output messages: Total 1738    Updates 3    Refreshes 0    Octets 33305
Output Queue[0]: 0
Output Queue[1]: 0

```

show bgp neighbor (Layer 2 VPN)

```

user@host> show bgp neighbor
Peer: 10.69.103.2    AS 65100 Local: 10.69.103.1    AS 65103
Type: External    State: Active    Flags: <ImportEval>

```

```

Last State: Idle           Last Event: Start
Last Error: None
Export: [ BGP-INET-import ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
Address families configured: inet-unicast
Local Address: 10.69.103.1 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer: 10.69.104.2          AS 65100 Local: 10.69.104.1          AS 65104
Type: External            State: Active              Flags: <ImportEval>
Last State: Idle          Last Event: Start
Last Error: None
Export: [ BGP-L-import ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
Address families configured: inet-labeled-unicast
Local Address: 10.69.104.1 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer: 10.255.14.182+179 AS 69    Local: 10.255.14.176+2131 AS 69
Type: Internal            State: Established        Flags: <ImportEval>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
Address families configured: inet-vpn-unicast 12vpn
Local Address: 10.255.14.176 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.255.14.182    Local ID: 10.255.14.176    Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast 12vpn
NLRI advertised by peer: inet-vpn-unicast 12vpn
NLRI for this session: inet-vpn-unicast 12vpn
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast 12vpn
NLRI peer can save forwarding state: inet-vpn-unicast 12vpn
NLRI that peer saved forwarding for: inet-vpn-unicast 12vpn
NLRI that restart is negotiated for: inet-vpn-unicast 12vpn
NLRI of received end-of-rib markers: inet-vpn-unicast 12vpn
Table bgp.13vpn.0 Bit: 10000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          10
  Received prefixes:        10
  Suppressed due to damping: 0
Table bgp.12vpn.0 Bit: 20000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2

```

```

    Suppressed due to damping: 0
Table BGP-L.inet.0 Bit: 40000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table LDP.inet.0 Bit: 50000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table L2VPN.l2vpn.0 Bit: 90000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Last traffic (seconds): Received 0    Sent 0    Checked 0
Input messages: Total 14    Updates 13    Refreshes 0    Octets 1053
Output messages: Total 3    Updates 0    Refreshes 0    Octets 105
Output Queue[0]: 0
Output Queue[1]: 0
Output Queue[2]: 0
Output Queue[3]: 0
Output Queue[4]: 0
Output Queue[5]: 0
Output Queue[6]: 0
Output Queue[7]: 0
Output Queue[8]: 0

```

show bgp neighbor (Layer 3 VPN)

```

user@host> show bgp neighbor
Peer: 4.4.4.4+179    AS 10045 Local: 5.5.5.5+1214    AS 10045
Type: Internal    State: Established    Flags: <ImportEval>

```

```

Last State: OpenConfirm   Last Event: RecvKeepAlive
Last Error: None
Export: [ match-all ] Import: [ match-all ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
        Rib-group Refresh>
Address families configured: inet-vpn-unicast
Local Address: 5.5.5.5 Holdtime: 90 Preference: 170
Flags for NLRI inet-labeled-unicast: TrafficStatistics
Traffic Statistics: Options: all File: /var/log/bstat.log
                               size 131072 files 10

Traffic Statistics Interval: 60
Number of flaps: 0
Peer ID: 192.168.1.110    Local ID: 192.168.1.111    Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast
NLRI advertised by peer: inet-vpn-unicast
NLRI for this session: inet-vpn-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast
NLRI peer can save forwarding state: inet-vpn-unicast
NLRI that peer saved forwarding for: inet-vpn-unicast
NLRI that restart is negotiated for: inet-vpn-unicast
NLRI of received end-of-rib markers: inet-vpn-unicast
NLRI of all end-of-rib markers sent: inet-vpn-unicast
Table bgp.13vpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Table vpn-green.inet.0 Bit: 20001
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
Last traffic (seconds): Received 15   Sent 20   Checked 20
Input messages: Total 40   Updates 2   Refreshes 0   Octets 856
Output messages: Total 44   Updates 2   Refreshes 0   Octets 1066
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: detail packets
Trace file: /var/log/bgpr.log size 131072 files 10

```

show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.1.111
Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
Type: Internal   State: Established (route reflector client)Flags: <Sync>
Last State: OpenConfirm   Last Event: RecvKeepAlive
Last Error: None
Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group
Refresh>
Address families configured: inet-vpn-unicast inet-labeled-unicast
Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
Flags for NLRI inet-vpn-unicast: AggregateLabel

```



```

Flags for NLRI inet-labeled-unicast: AggregateLabel
Number of flaps: 0
Peer ID: 10.255.245.12    Local ID: 10.255.245.13    Active Holdtime: 90
Keepalive Interval: 30
BFD: disabled
NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
NLRI for this session: inet-vpn-unicast inet-labeled-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 300
Stale routes from peer are kept for: 60
Restart time requested by this peer: 300
NLRI that peer supports restart for: inet-unicast inet6-unicast
NLRI that restart is negotiated for: inet-unicast inet6-unicast
NLRI of received end-of-rib markers: inet-unicast inet6-unicast
NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
Table inet.0 Bit: 10000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 4
  Received prefixes: 6
  Suppressed due to damping: 0
Table inet6.0 Bit: 20000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 0
  Received prefixes: 2
  Suppressed due to damping: 0
Last traffic (seconds): Received 3    Sent 3    Checked 3
Input messages: Total 9    Updates 6    Refreshes 0    Octets 403
Output messages: Total 7    Updates 3    Refreshes 0    Octets 365
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: detail packets
Trace file: /var/log/bgpr size 131072 files 10

```

show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.4.222
Peer: 192.168.4.222+4902 AS 65501 Local: 192.168.4.221+179 AS 65500
  Type: External    State: Established    Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: Cease
  Export: [ export-policy ] Import: [ import-policy ]
  Options: <Preference HoldTime AddressFamily PeerAS PrefixLimit Refresh>
  Address families configured: inet-unicast inet-multicast
  Holdtime: 60000 Preference: 170
  Number of flaps: 4
  Last flap event: RecvUpdate
  Error: 'Cease' Sent: 5 Recv: 0
  Peer ID: 10.255.245.6    Local ID: 10.255.245.5    Active Holdtime: 60000
  Keepalive Interval: 20000    Peer index: 0
  BFD: disabled, down
  Local Interface: fxp0.0
  NLRI advertised by peer: inet-unicast inet-multicast
  NLRI for this session: inet-unicast inet-multicast
  Peer supports Refresh capability (2)
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:      8
    Received prefixes:    10

```

```
Accepted prefixes:          10
Suppressed due to damping:  0
Advertised prefixes:        3
Table inet.2 Bit: 20000
RIB State: BGP restart is complete
Send state: in sync
Active prefixes:            0
Received prefixes:          0
Accepted prefixes:          0
Suppressed due to damping:  0
Advertised prefixes:        0
Last traffic (seconds): Received 357 Sent 357 Checked 357
Input messages: Total 4 Updates 2 Refreshes 0 Octets 211
Output messages: Total 4 Updates 1 Refreshes 0 Octets 147
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: all
Trace file: /var/log/bgp size 10485760 files 10
```

show bgp neighbor orf neighbor-address detail

```
user@host > show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56+179 Type: External
Group: ext1

inet-unicast
  Filter updates rcv:          1 Immediate:          1
  Filter: prefix-based receive
  Received filter entries:
    seq 1: prefix 2.2.2.2/32: minlen 32: maxlen 32: match deny:

inet6-unicast
  Filter updates rcv:          0 Immediate:          1
  Filter: prefix-based receive
  Received filter entries:
    *.*
```

show bgp summary

Syntax	<pre>show bgp summary <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bgp summary <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p>
Description	Display BGP summary information.
Options	<p>none—Display BGP summary information for all routing instances.</p> <p>exact-instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>instance <i>instance-name</i>—(Optional) Display information for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp summary instance cust1 command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	<p>show bgp summary (When a Peer Is Not Established) on page 2672</p> <p>show bgp summary (When a Peer Is Established) on page 2672</p> <p>show bgp summary (CLNS) on page 2672</p> <p>show bgp summary (Layer 2 VPN) on page 2672</p> <p>show bgp summary (Layer 3 VPN) on page 2673</p>
Output Fields	<p>Table 215 on page 2669 describes the output fields for the show bgp summary command. Output fields are listed in the approximate order in which they appear.</p>

Table 215: show bgp summary Output Fields

Field Name	Field Description
Groups	Number of BGP groups.
Peers	Number of BGP peers.

Table 215: show bgp summary Output Fields (*continued*)

Field Name	Field Description
Down peers	Number of down BGP peers.
Table	Name of routing table.
Tot Paths	Total number of paths.
Act Paths	Number of active routes.
Suppressed	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
History	Number of withdrawn routes stored locally to keep track of damping history.
Damp State	Number of routes with a figure of merit greater than zero, but still active because the value has not reached the threshold at which suppression occurs.
Pending	Routes in process by BGP import policy.
Peer	Address of each BGP peer. Each peer has one line of output.
AS	Peer's AS number.
InPkt	Number of packets received from the peer.
OutPkt	Number of packets sent to the peer.
OutQ	Number of BGP packets that are queued to be transmitted to a particular neighbor. It normally is 0 because the queue usually is emptied quickly.
Flaps	Number of times the BGP session has gone down and then come back up.
Last Up/Down	Last time since the neighbor transitioned to or from the established state.

Table 215: show bgp summary Output Fields (*continued*)

Field Name	Field Description
State #Active /Received/Accepted /Damped	<p>Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established on the main routing device or in a routing instance.</p> <ul style="list-style-type: none"> If a peer is not established, the field shows the state of the peer session: Active, Connect, or Idle. In general, the Idle state is the first stage of a connection. BGP is waiting for a Start event. A session can be idle for other reasons as well. The reason that a session is idle is sometimes displayed. For example: Idle (Removal in progress) or Idle (LicenseFailure). If a BGP session is established on the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the inet.0 (main) and inet.2 (multicast) routing tables. For example, 8/10/10/2 and 2/4/4/0 indicate the following: <ul style="list-style-type: none"> 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the inet.0 routing table. 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the inet.2 routing table. If a BGP session is established in a routing instance, the field indicates the established (Establ) state, identifies the specific routing table that receives BGP updates, and shows the number of active, received, and damped routes that are received from a neighbor. For example, Establ VPN-AB.inet.0: 2/4/0 indicates the following: <ul style="list-style-type: none"> The BGP session is established. Routes are received in the VPN-AB.inet.0 routing table. The local routing device has two active routes, four received routes, and no damped routes from a BGP peer. <p>When a BGP session is established, the peers are exchanging update messages.</p>

Sample Output

show bgp summary (When a Peer Is Not Established)

```

user@host> show bgp summary
Groups: 2 Peers: 4 Down peers: 1
Table      Tot Paths  Act Paths Suppressed  History  Damp State   Pending
inet.0      6          4          0          0        0      0        0
Peer        AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.3     65002      86       90       0        2      42:54 0/0/0

0/0/0
10.0.0.4     65002      90       91       0        1      42:54 0/2/0

0/0/0
10.0.0.6     65002      87       90       0        3          3 Active
10.1.12.1    65001      89       89       0        1      42:54 4/4/0

0/0/0

```

show bgp summary (When a Peer Is Established)

```

user@host> show bgp summary
Groups: 1 Peers: 3 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State   Pending
inet.0      6          4          0          0        0      0        0
Peer        AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.2     65002    88675    88652     0        2      42:38 2/4/0

0/0/0
10.0.0.3     65002    54528    54532     0        1     2w4d22h 0/0/0

0/0/0
10.0.0.4     65002    51597    51584     0        0     2w3d22h 2/2/0

0/0/0

```

show bgp summary (CLNS)

```

user@host> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Peer        AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.245.245.1 200     1735     1737     0        0    14:26:12 Establ
  bgp.isovpn.0: 3/3/0
  aaaa.iso.0: 3/3/0

```

show bgp summary (Layer 2 VPN)

```

user@host> show bgp summary
Groups: 1 Peers: 5 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State   Pending
bgp.l2vpn.0 1          1          0          0        0      0        0
inet.0      0          0          0          0        0      0        0
Peer        AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.255.245.35 65299     72       74       0        1     19:00 Establ
  bgp.l2vpn.0: 1/1/0
  frame-vpn.l2vpn.0: 1/1/0

```

```

10.255.245.36 65299      2164      2423      0        4        19:50 Establ
  bgp.12vpn.0: 0/0/0
  frame-vpn.12vpn.0: 0/0/0
10.255.245.37 65299      36         37         0        4        17:07 Establ
  inet.0: 0/0/0
10.255.245.39 65299      138        168         0        6        53:48 Establ
  bgp.12vpn.0: 0/0/0
  frame-vpn.12vpn.0: 0/0/0
10.255.245.69 65299      134        140         0        6        53:42 Establ
  inet.0: 0/0/0

```

show bgp summary (Layer 3 VPN)

```

user@host> show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths  Suppressed  History Damp State Pending
bgp.13vpn.0      2          2          0           0         0      0        0
Peer           AS      InPkt      OutPkt      OutQ      Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.39.1.5       2        21         22          0          0        6:26 Establ
  VPN-AB.inet.0: 1/1/0
10.255.71.15    1        19         21          0          0        6:17 Establ
  bgp.13vpn.0: 2/2/0
  VPN-A.inet.0: 1/1/0
  VPN-AB.inet.0: 2/2/0
  VPN-B.inet.0: 1/1/0

```

show policy damping

Syntax	show policy damping <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and QFX Series)	show policy damping
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display information about BGP route flap damping parameters.
Options	<p>none—Display information about BGP route flap damping parameters.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Additional Information	In the output from this command, figure-of-merit values correlate with the probability of future instability of a routing device. Routes with higher figure-of-merit values are suppressed for longer periods of time. The figure-of-merit value decays exponentially over time. A figure-of-merit value of zero is assigned to each new route. The value is increased each time the route is withdrawn or readvertised, or when one of its path attributes changes.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • “Configuring BGP Flap Damping Parameters” in the <i>Routing Policy Configuration Guide</i> • clear bgp damping on page 2641 • <i>show route damping</i>
List of Sample Output	show policy damping on page 2675
Output Fields	Table 216 on page 2674 describes the output fields for the show policy damping command. Output fields are listed in the approximate order in which they appear.

Table 216: show policy damping Output Fields

Field Name	Field Description
Halflife	Decay half-life, in minutes. The value represents the period during which the accumulated figure-of-merit value is reduced by half if the route remains stable. If a route has flapped, but then becomes stable, the figure-of-merit value for the route decays exponentially. For example, for a route with a figure-of-merit value of 1500, if no incidents occur, its figure-of-merit value is reduced to 750 after 15 minutes and to 375 after another 15 minutes.

Table 216: show policy damping Output Fields (*continued*)

Field Name	Field Description
Reuse merit	Figure-of-merit value below which a suppressed route can be used again. A suppressed route becomes reusable when its figure-of-merit value decays to a value below a reuse threshold, and the route once again is considered usable and can be installed in the forwarding table and exported from the routing table.
Suppress/cutoff merit	Figure-of-merit value above which a route is suppressed for use or inclusion in advertisements. When a route's figure-of-merit value reaches a particular level, called the cutoff or suppression threshold, the route is suppressed. When a route is suppressed, the routing table no longer installs the route into the forwarding table and no longer exports this route to any of the routing protocols.
Maximum suppress time	Maximum hold-down time, in minutes. The value represents the maximum time that a route can be suppressed no matter how unstable it has been before this period of stability.
Computed values	<ul style="list-style-type: none"> • Merit ceiling—Maximum merit that a flapping route can collect. • Maximum decay—Maximum decay half-life, in minutes.

Sample Output

show policy damping

```

user@host> show policy damping
Default damping information:
  Halflife: 15 minutes
  Reuse merit: 750 Suppress/cutoff merit: 3000
  Maximum suppress time: 60 minutes
  Computed values:
    Merit ceiling: 12110
    Maximum decay: 6193
Damping information for "standard-damping":
  Halflife: 10 minutes
  Reuse merit: 4000 Suppress/cutoff merit: 8000
  Maximum suppress time: 30 minutes
  Computed values:
    Merit ceiling: 32120
    Maximum decay: 12453

```

IS-IS

- [Configuration on page 2675](#)
- [Administration on page 2735](#)

Configuration

- [Configuration Statements on page 2675](#)

Configuration Statements

- [\[edit protocols isis\] Hierarchy Level on page 2676](#)

[edit protocols isis] Hierarchy Level

The following statement hierarchy can also be included at the [edit protocols isis] hierarchy level.

```
protocols {
  isis {
    disable;
    clns-routing;
    context-identifier ip-address </prefix> {
      level (1 | 2) <disable>;
    }
    export [ policy-names ];
    graceful-restart {
      disable;
      helper-disable;
      restart-duration seconds;
    }
    ignore-attached-bit;
    interface interface-name {
      ... the interface subhierarchy appears after the main [edit protocols isis] hierarchy ...
    }
    label-switched-path name level level metric metric;
    level (1 | 2) {
      disable;
      authentication-key key;
      authentication-type authentication;
      external-preference preference;
      no-csnp-authentication;
      no-hello-authentication;
      no-psnp-authentication;
      preference preference;
      prefix-export-limit number;
      wide-metrics-only;
    }
    loose-authentication-check;
    lsp-lifetime seconds;
    max-areas number;
    no-adjacency-holddown;
    no-authentication-check;
    no-ipv4-routing;
    no-ipv6-routing;
    overload {
      advertise-high-metrics;
      timeout seconds;
    }
    reference-bandwidth reference-bandwidth;
    rib-group {
      inet group-name;
      inet6 group-name;
    }
    spf-options {
      delay milliseconds;
      holddown milliseconds;
      rapid-runs number;
    }
  }
}
```

```

}
topologies {
  ipv4-multicast;
  ipv6-multicast;
  ipv6-unicast;
}
traceoptions {
  file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
traffic-engineering {
  disable;
  family inet {
    shortcuts {
      multicast-rpf-routes;
    }
  }
  family inet6 {
    shortcuts;
  }
}
ignore-lsp-metrics;
}

isis {
  interface interface-name {
    disable;
    bfd-liveness-detection {
      authentication {
        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
          meticulous-keyed-sha-1 | simple-password);
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      version (1 | automatic);
    }
  }
  checksum;
  csnp-interval (seconds | disable);
  hello-padding (adaptive | loose | strict);
  ldp-synchronization {
    disable;
    hold-time seconds;
  }
  level (1 | 2) {

```

```
    disable;  
    hello-authentication-key key;  
    hello-authentication-type authentication;  
    hello-interval seconds;  
    hold-time seconds;  
    ipv4-multicast-metric number;  
    ipv6-multicast-metric number;  
    ipv6-unicast-metric number;  
    metric metric;  
    passive;  
    priority number;  
    te-metric metric;  
  }  
  link-protection;  
  lsp-interval milliseconds;  
  mesh-group (value | blocked);  
  no-adjacency-down-notification;  
  no-eligible-backup;  
  no-ipv4-multicast;  
  no-ipv6-multicast;  
  no-ipv6-unicast;  
  no-unicast-topology;  
  node-link-protection;  
  passive;  
  point-to-point;  
}  
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

authentication-key (Protocols IS-IS)

Syntax	authentication-key <i>key</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level <i>level-number</i>], [edit protocols isis level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis level <i>level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Authentication key (password). Neighboring routing devices use the password to verify the authenticity of packets sent from this interface. For the key to work, you also must include the authentication-type statement.</p> <p>All routing devices must use the same password. If you are using the Junos OS IS-IS software with another implementation of IS-IS, the other implementation must be configured to use the same password for the domain, the area, and all interfaces adjacent to the Juniper Networks routing device.</p>
Default	If you do not include this statement and the authentication-type statement, IS-IS authentication is disabled.
Options	key —Authentication password. The password can be up to 1024 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").



CAUTION: A simple password for authentication is truncated if it exceeds 254 characters.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Hitless Authentication Key Rollover for IS-IS

authentication-type (Protocols IS-IS)

Syntax	<code>authentication-type <i>authentication</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Enable authentication and specify the authentication scheme for IS-IS. If you enable authentication, you must specify a password by including the authentication-key statement.
Default	If you do not include this statement and the authentication-key statement, IS-IS authentication is disabled.
Options	<i>authentication</i> —Authentication scheme: <ul style="list-style-type: none">• md5—Use HMAC authentication in combination with MD5. HMAC-MD5 authentication is defined in RFC 2104, <i>HMAC: Keyed-Hashing for Message Authentication</i>.• simple—Use a simple password for authentication. The password is included in the transmitted packet, making this method of authentication relatively insecure. We recommend that you <i>not</i> use this authentication method.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Hitless Authentication Key Rollover for IS-IS• authentication-key on page 2679• no-authentication-check on page 2712

bfd-liveness-detection (Protocols IS-IS)

Syntax	<pre> bfd-liveness-detection { authentication { algorithm <i>algorithm-name</i>; key-chain <i>key-chain-name</i>; loose-check; } detection-time { threshold <i>milliseconds</i>; } minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; multiplier <i>number</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } version (1 automatic); } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit protocols isis interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>detection-time threshold and transmit-interval threshold options added in Junos OS Release 8.2.</p> <p>Support for logical systems introduced in Junos OS Release 8.3.</p> <p>no-adaptation statement introduced in Junos OS Release 9.0.</p> <p>authentication algorithm, authentication key-chain, and authentication loose-check options introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Configure bidirectional failure detection timers and authentication.
Options	<p>authentication algorithm <i>algorithm-name</i> —Configure the algorithm used to authenticate the specified BFD session: simple-password, keyed-md5, keyed-sha-1, meticulous-keyed-md5, meticulous-keyed-sha-1.</p> <p>authentication key-chain <i>key-chain-name</i> —Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the authentication-key-chains key-chain statement at the [edit security] hierarchy level.</p> <p>authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication might not be configured at both ends of the BFD session.</p>

detection-time threshold *milliseconds*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

minimum-interval *milliseconds*—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.

Range: 1 through 255,000

minimum-receive-interval *milliseconds*—Configure the minimum interval after which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

Range: 1 through 255,000

multiplier *number*—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

no-adaptation—Specify that BFD sessions not adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds*—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

transmit-interval minimum-interval *milliseconds*—Configure a minimum interval after which the local routing device transmits hello packets to a neighbor. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

Range: 1 through 255,000

version—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version)

Default: automatic

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

Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring BFD for IS-IS</i>• <i>Example: Configuring BFD Authentication for IS-IS</i>
------------------------------	--

checksum (Protocols IS-IS)

Syntax	checksum;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Enable checksums for packets on this interface. Junos OS supports IS-IS checksums as documented in RFC 3358, <i>Optional Checksums in Intermediate System to Intermediate System (ISIS)</i> . The checksum cannot be enabled with MD5 hello authentication on the same interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Enabling Packet Checksums on IS-IS Interfaces</i>


csnp-interval

Syntax	<code>csnp-interval (seconds disable);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure the interval between complete sequence number PDUs (CSNPs) on a LAN interface.</p> <p>If the routing device is the designated router on a LAN, IS-IS sends CSN packets every 10 seconds. If the routing device is on a point-to-point interface, it sends CSN packets every 5 seconds. To protect against link-state PDU flooding, we recommend modifying the default interval.</p> <p>To modify the CSNP interval, include the csnp-interval statement.</p> <p>To configure the interface not to send any CSNPs, specify the disable option.</p>
Default	By default, IS-IS sends CSNPs periodically. If the routing device is the designated router on a LAN, IS-IS sends CSNPs every 10 seconds. If the routing device is on a point-to-point interface, it sends CSNPs every 5 seconds.
Options	<p>disable—Do not send CSNPs on this interface.</p> <p>seconds—Number of seconds between the sending of CSNPs.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: 10 seconds on LAN broadcast links. 5 seconds on point-to-point links.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces</i>

disable (Protocols IS-IS)

Syntax	disable;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> protocols isis traffic-engineering], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering], [edit protocols isis], [edit protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit protocols isis traffic-engineering], [edit routing-instances <i>routing-instance-name</i> protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Disable IS-IS on the routing device, on an interface, or on a level.</p> <p>At the [edit protocols isis traffic-engineering] hierarchy level, disable IS-IS support for traffic engineering.</p> <p>Enabling IS-IS on an interface (by including the interface statement at the [edit protocols isis] or the [edit routing-instances routing-instance-name protocols isis] hierarchy level), disabling it (by including the disable statement), and not actually having IS-IS run on an interface (by including the passive statement) are mutually exclusive states.</p>
Default	<p>IS-IS is enabled for Level 1 and Level 2 routers on all interfaces on which family iso is enabled.</p> <p>IS-IS support for traffic engineering is enabled.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Multi-Level IS-IS</i> • IS-IS Overview on page 3343

export (Protocols IS-IS)

Syntax	<code>export [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols isis],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>isis],</code> <code>[edit protocols isis],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols isis]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Apply one or more policies to routes being exported from the routing table into IS-IS.</p> <p>All routing protocols store the routes that they learn in the routing table. The routing table uses this collected route information to determine the active routes to destinations. The routing table then installs the active routes into its forwarding table and exports them into the routing protocols. It is these exported routes that the protocols advertise.</p> <p>For each protocol, you control which routes the protocol stores in the routing table and which routes the routing table exports into the protocol from the routing table by defining a <i>routing policy</i> for that protocol.</p> <div><p>NOTE: For IS-IS, you cannot apply routing policies that affect how routes are imported into the routing table; doing so with a link-state protocol can easily lead to an inconsistent topology database.</p></div>
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Redistributing OSPF Routes into IS-IS</i>• Example: Configuring an IS-IS Default Route Policy on Logical Systems on page 3356

external-preference (Protocols IS-IS)

Syntax	<code>external-preference <i>preference</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure the preference of external routes.
Options	<i>preference</i> —Preference value. Range: 0 through 4,294,967,295 ($2^{32} - 1$) Default: 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes)
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Route Preferences Overview</i> • <i>Example: Redistributing OSPF Routes into IS-IS</i> • <i>Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS</i> • preference on page 2725


graceful-restart (Protocols IS-IS)

Syntax	<pre>graceful-restart { disable; helper-disable; restart-duration <i>seconds</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit protocols isis]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Configure graceful restart parameters for IS-IS.</p> <p>Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled for all routing protocols at the [edit routing-options] hierarchy level. When graceful restart is enabled, the restarting routing device is not removed from the network topology during the restart period. The adjacencies are reestablished after restart is complete.</p> <p>On LAN interfaces where IS-IS is configured on a transit router that serves as the designated router (DR), a graceful restart causes:</p> <ul style="list-style-type: none">• The ingress router of the label-switched path (LSP), which passes through the DR, to break the LSP.• The ingress router to re-signal the LSP.
Options	<p>disable—Disable graceful restart for IS-IS.</p> <p>helper-disable—Disable graceful restart helper capability. Helper mode is enabled by default.</p> <p>restart-duration <i>seconds</i>—Time period for the restart to last, in seconds. Range: 30 through 300 seconds Default: 30 seconds</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Routing Protocols Graceful Restart on page 2128

hello-authentication-key

Syntax	<code>hello-authentication-key password;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Configure an authentication key (password) for hello packets. Neighboring routing devices use the password to verify the authenticity of packets sent from an interface. For the key to work, you also must include the hello-authentication-type statement.
Default	By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password.
Options	<p>password—Authentication password. The password can be up to 255 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • authentication-key on page 2679 • authentication-type on page 2680 • hello-authentication-type on page 2690

hello-authentication-type

Syntax	hello-authentication-type (md5 simple);
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Enable authentication on an interface for hello packets. If you enable authentication on hello packets, you must specify a password by including the hello-authentication-key statement.</p> <p>You can configure authentication for a given IS-IS level on an interface. On a point-to-point link, if you enable hello authentication for both IS-IS levels, the password configured for Level 1 is used for both levels.</p>
	<div>  <p>CAUTION: If no authentication is configured for Level 1 on a point-to-point link with both levels enabled, the hello packets are sent without any password, regardless of the Level 2 authentication configurations.</p> </div>
Default	By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password.
Options	<p>md5—Specifies Message Digest 5 as the packet verification type.</p> <p>simple—Specifies simple authentication as the packet verification type.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • authentication-key on page 2679 • authentication-type on page 2680 • hello-authentication-key on page 2689

hello-interval (Protocols IS-IS)

Syntax	<code>hello-interval seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Modify the frequency with which the routing device sends hello packets out of an interface, in seconds.</p> <p>Routing devices send hello packets at a fixed interval on all interfaces to establish and maintain neighbor relationships. This interval is advertised in the hello interval field in the hello packet.</p> <p>You can send out hello packets in subsecond intervals. To send out hello packets every 333 milliseconds, set the hello-interval value to 1.</p>
Options	<p>seconds—Frequency of transmission for hello packets.</p> <p>Range: 1 through 20,000 seconds</p> <p>Default: 3 seconds (for designated intermediate system [DIS] routers), 9 seconds (for non-DIS routers)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>hold-time</i>

hello-padding

Syntax	hello-padding (adaptive disable loose strict);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure padding on hello packets to accommodate asymmetrical maximum transfer units (MTUs) from different hosts.</p> <p>This helps to prevent a premature adjacency Up state when one routing device's MTU does not meet the requirements to establish the adjacency.</p> <p>As an OSI Layer 2 protocol, IS-IS does not support data fragmentation. Therefore, maximum packet sizes must be established and supported between two routers. During adjacency establishment, the IS-IS protocol makes sure that the link supports a packet size of 1492 bytes by padding outgoing hello packets up to the maximum packet size of 1492 bytes.</p> <p>This is the default behavior of the Junos OS IS-IS implementation. However, Junos OS provides an option to disable hello padding that can override this behavior.</p> <p>There are four types of hello padding:</p> <ul style="list-style-type: none">• Adaptive padding—On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state type, length, and value (TLV) tuple. If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface. Adaptive padding has more overhead than loose padding and is able to detect MTU asymmetry from one side of the connection. This one-sided detection can result in generation of extra link-state PDUs that are flooded throughout the network. Specify the adaptive option to configure enough padding to establish an adjacency to neighbors.• Disabled padding—Padding is disabled on all types of interfaces for all adjacency states. Specify the disable option to accommodate interfaces that support less than the default packet size of 1492 bytes.• Loose padding (the default)—The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state. Loose padding might not be able to detect certain situations such as asymmetrical MTUs between the routing devices. Specify the loose option to configure enough padding to initialize an adjacency to neighbors.

- **Strict padding**—Padding is done on all interface types and for all adjacency states, and is continuous. Strict padding has the most overhead. The advantage is that strict padding detects MTU issues on both sides of a link. Specify the **strict** option to configure padding to allow all adjacency states with neighbors.

Options **adaptive**—Configure padding until the neighbor adjacency is established and active.

disable—Disable padding on all types of interfaces for all adjacency states.

loose—Configure padding until the state of the adjacency is initialized.

strict—Configure padding for all adjacency states.

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

Related Documentation • *Example: Configuring IS-IS*

hold-time (Protocols IS-IS)

Syntax	<code>hold-time seconds;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code> <code>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Set the length of time a neighbor considers this router to be operative (up) after receiving a hello packet. If the neighbor does not receive another hello packet within the specified time, it marks this routing device as inoperative (down). The hold time itself is advertised in the hello packets.</p> <p>The hold time specifies how long a neighbor should consider this routing device to be operative without receiving another hello packet. If the neighbor does not receive a hello packet from this routing device within the hold time, it marks the routing device as being unavailable.</p>
Options	<p>seconds—Hold-time value, in seconds.</p> <p>Range: 3 through 65,535 seconds, or 1 to send out hello packets every 333 milliseconds</p> <p>Default: 9 seconds (for designated intermediate system [DIS] routers), 27 seconds (for non-DIS routers; three times the default hello interval)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring IS-IS</i>• hello-interval on page 2691

ignore-attached-bit


Syntax	ignore-attached-bit;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Ignore the attached bit on IS-IS Level 1 routers. Configuring this statement enables the routing device to ignore the attached bit on incoming Level 1 link-state PDUs. If the attached bit is ignored, no default route, which points to the routing device which has set the attached bit, is installed.</p> <p>There might be times, such as during a denial-of-service (DoS) attack, that you do not want a Level 1 router to be able to forward traffic based on a default route.</p> <p>To prevent a routing device from being able to reach interarea destinations, you can prevent the routing device from installing the default route without affecting the status of its IS-IS adjacencies. The ignore-attached-bit statement is used to tell the routing device to ignore the presence of the attached bit in Level 1 link-state PDUs, which blocks the installation of the IS-IS default route.</p>
Default	The ignore-attached-bit statement is disabled by default.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> •

interface (Protocols IS-IS)

```
Syntax interface (all | interface-name) {
    disable;
    bfd-liveness-detection {
        authentication {
            algorithm algorithm-name;
            key-chain key-chain-name;
            loose-check;
        }
        detection-time {
            threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        transmit-interval {
            threshold milliseconds;
            minimum-interval milliseconds;
        }
        multiplier number;
    }
    checksum;
    csnp-interval (seconds | disable);
    hello-padding (adaptive | loose | strict);
    ldp-synchronization {
        disable;
        hold-time seconds;
    }
    lsp-interval milliseconds;
    mesh-group (value | blocked);
    no-adjacency-holddown;
    no-ipv4-multicast;
    no-ipv6-multicast;
    no-ipv6-unicast;
    no-unicast-topology;
    passive;
    point-to-point;
    level level-number {
        disable;
        hello-authentication-key key;
        hello-authentication-key-chain key-chain-name;
        hello-authentication-type authentication;
        hello-interval seconds;
        hold-time seconds;
        ipv4-multicast-metric metric;
        ipv6-multicast-metric metric;
        ipv6-unicast-metric metric;
        metric metric;
        passive;
        priority number;
        te-metric metric;
    }
}
```

Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure interface-specific IS-IS properties. To configure more than one interface, include the interface statement multiple times.</p> <p>Enabling IS-IS on an interface (by including the interface statement at the [edit protocols isis] or the [edit routing-instances <i>routing-instance-name</i> protocols isis] hierarchy level), disabling it (by including the disable statement), and not actually having IS-IS run on an interface (by including the passive statement) are mutually exclusive states.</p>
Options	<p>all—Have Junos OS create IS-IS interfaces automatically. If you include this option, disable IS-IS on the management interface (fxp0).</p> <p>interface-name—Name of an interface. Specify the full interface name, including the physical and logical address components.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring IS-IS</i> • <i>Example: Configuring Multi-Level IS-IS</i>

ipv4-multicast

Syntax	ipv4-multicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis topologies], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis topologies], [edit protocols isis topologies], [edit routing-instances <i>routing-instance-name</i> protocols isis topologies]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure alternate IPv4 multicast topologies. <div> NOTE: The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics. You can also selectively disable interfaces from participating in the IPv6 topology while continuing to participate in the IPv4 topology. This lets you exercise control over the paths that unicast data takes through a network.</div>
Default	Multicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring IS-IS Multicast Topology</i>

ipv4-multicast-metric

Syntax	<code>ipv4-multicast-metric <i>metric</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify the multicast topology metric value for the level.
Options	<i>metric</i> —Metric value. Range: 0 through 16,777,215
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring IS-IS Multicast Topology</i>


ipv6-multicast

Syntax	<code>ipv6-multicast;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis topologies], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis topologies], [edit protocols isis topologies], [edit routing-instances <i>routing-instance-name</i> protocols isis topologies]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure alternate IPv6 multicast topologies.
Default	Multicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring IS-IS Multicast Topology</i>

ipv6-multicast-metric

Syntax	ipv6-multicast-metric <i>metric</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify the IPv6 alternate multicast topology metric value for the level.
Options	<i>metric</i> —Metric value. Range: 0 through 16,777,215
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring IS-IS Multicast Topology</i>

ipv6-unicast

Syntax	ipv6-unicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis topologies], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis topologies], [edit protocols isis topologies], [edit routing-instances <i>routing-instance-name</i> protocols isis topologies]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure alternate IPv6 unicast topologies. This statement causes IS-IS to calculate an alternate IPv6 unicast topology, in addition to the normal IPv4 unicast topology, and add the corresponding routes to inet6.0.
	<div>  <p>NOTE: The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics. You can also selectively disable interfaces from participating in the IPv6 topology while continuing to participate in the IPv4 topology. This lets you exercise control over the paths that unicast data takes through a network.</p> </div>
Default	IPv6 unicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies</i>

ipv6-unicast-metric

Syntax	ipv6-unicast-metric <i>metric</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit protocols isis interface <i>interface-name</i> level <i>level-number</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify the IPv6 unicast topology metric value for the level. The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics.
Options	<i>metric</i> —Metric value. Range: 0 through 16,777,215
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies</i>

isis

Syntax	isis { ... }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Enable IS-IS routing on the routing device or for a routing instance. The isis statement is the one statement you must include in the configuration to run IS-IS on the routing device or in a routing instance.
Default	IS-IS is disabled on the routing device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring IS-IS</i> • <i>Example: Configuring Multi-Level IS-IS</i>

level (Global IS-IS)

Syntax	<pre>level <i>level-number</i> { authentication-key <i>key</i>; authentication-key-chain (Protocols IS-IS) <i>key-chain-name</i>; authentication-type <i>type</i>; external-preference <i>preference</i>; no-csnp-authentication; no-hello-authentication; no-psnp-authentication; preference <i>preference</i>; wide-metrics-only; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>], [edit protocols <i>isis</i>], [edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure the global-level properties.</p> <p>You can administratively divide a single AS into smaller groups called areas. You configure each routing device interface to be in an area. Any interface can be in any area. The area address applies to the entire routing device. You cannot specify one interface to be in one area and another interface in a different area. To route between areas, you must have two adjacent Level 2 routers that communicate with each other.</p> <p>Level 1 routers can only route within their IS-IS area. To send traffic outside their area, Level 1 routers must send packets to the nearest intra-area Level 2 router. A routing device can be a Level 1 router, a Level 2 router, or both. You specify the router level on a per-interface basis, and a routing device becomes adjacent to other routing devices on the same level on that link only.</p> <p>You can configure one Level 1 routing process and one Level 2 routing process on each interface, and you can configure the two levels differently.</p>
Options	<p><i>level-number</i>—IS-IS level number.</p> <p>Values: 1 or 2</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring IS-IS</i>• <i>Example: Configuring Multi-Level IS-IS</i>

link-protection (Protocols IS-IS)

Syntax	link-protection;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Enable link protection on the specified IS-IS interface. Junos OS creates a backup loop-free alternate path to the primary next hop for all destination routes that traverse the protected interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Link and Node Protection for IS-IS Routes</i> • node-link-protection on page 2719

loose-authentication-check

Syntax	loose-authentication-check;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Allow the use of MD5 authentication without requiring network-wide deployment.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Hitless Authentication Key Rollover for IS-IS</i>

lsp-interval

Syntax	<code>lsp-interval milliseconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure the link-state PDU interval time.</p> <p>By default, the routing device sends one link-state PDU packet out an interface every 100 milliseconds. To disable the transmission of all link-state PDUs, set the interval to 0.</p> <p>Link-state PDU throttling by use of the lsp-interval statement controls the flooding pace to neighboring routing devices in order to not overload them.</p> <p>Also, consider that control traffic (such as link-state PDUs and related packets) might delay user traffic (information packets) because control traffic always has precedence in terms of scheduling on the routing device interface cards. Unfortunately, the control traffic transmission rate is not decreased on low-bandwidth interfaces, such as DS-0 or fractional T1 and E1 interface. Line control traffic stays the same. On a low-bandwidth circuit that is transmitting 30 full-MTU-sized packets, there is not much bandwidth left over for other types of packets.</p>
Default	By default, the routing device sends one link-state PDU out an interface every 100 milliseconds.
Options	<p>milliseconds—Number of milliseconds between the sending of link-state PDUs. Specifying a value of 0 blocks all link-state PDU transmission.</p> <p>Range: 0 through 1000 milliseconds</p> <p>Default: 100 milliseconds</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces</i>

lsp-lifetime

Syntax	<code>lsp-lifetime <i>seconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Specify how long a link-state PDU originating from the routing device should persist in the network. The routing device sends link-state PDUs often enough so that the link-state PDU lifetime never expires.</p> <p>Because link-state PDUs have a maximum lifetime, they need to be refreshed. Refreshing means that a routing device needs to re-originate its link-state PDUs periodically. The re-origination interval must be less than the link-state PDU's lifetime. For example, if the link-state PDU is valid for 1200 seconds, the routing device needs to refresh the link-state PDU in less than 1200 seconds to avoid removal of the link-state PDU from the link-state database by other routing devices. The recommended maximum link-state PDU origination interval is the lifetime minus 300 seconds. So, in a default environment this would be 900 seconds. In Junos OS, the refresh interval is derived from the lifetime and is equal to the lifetime minus 317 seconds. You can change the lifetime to a higher value to reduce the number of refreshes in the network. (You would rarely want to increase the number of refreshes.) Often these periodic link-state PDU refreshes are referred to as refresh noise, and network administrators want to reduce this noise as much as possible.</p> <p>The show isis overview command displays the link-state PDU lifetime.</p>
Default	By default, link-state PDUs are maintained in network databases for 1200 seconds (20 minutes) before being considered invalid. This length of time, called the <i>LSP lifetime</i> , normally is sufficient to guarantee that link-state PDUs never expire.
Options	<p><i>seconds</i>—link-state PDU lifetime, in seconds.</p> <p>Range: 350 through 65,535 seconds</p> <p>Default: 1200 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces</i> • http://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf

max-areas

Syntax	<code>max-areas <i>number</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code><i>isis</i>]</code> <code>[edit protocols <i>isis</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</code>
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Modify the maximum number of IS-IS areas advertised.</p> <p>This value is included in the Maximum Address Area field of the IS-IS common PDU header included in all outgoing PDUs.</p> <p>The maximum number of areas you can advertise is restricted to 36 to ensure that the IIH PDUs have enough space to include other type, length, and value (TLV) fields, such as the Authentication and IPv4 and IPv6 Interface Address TLVs.</p>
Options	<p><i>number</i>—Maximum number of areas to include in the IS-IS hello (IIH) PDUs and link-state PDUs.</p> <p>Range: 3 through 36</p> <p>Default: 3</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Multi-Level IS-IS</i>

mesh-group (Protocols IS-IS)

Syntax	mesh-group (blocked <i>value</i>);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure an interface to be part of a mesh group, which is a set of fully connected nodes.</p> <p>A <i>mesh group</i> is a set of routing devices that are fully connected. That is, they have a fully meshed topology. When link-state PDUs are being flooded throughout an area, each router within a mesh group receives only a single copy of a link-state PDU instead of receiving one copy from each neighbor, thus minimizing the overhead associated with the flooding of link-state PDUs.</p> <p>To create a mesh group and designate that an interface be part of the group, assign a mesh-group number to all the routing device interfaces in the group. To prevent an interface in the mesh group from flooding link-state PDUs, configure blocking on that interface.</p>
Options	<p>blocked—Configure the interface so that it does not flood link-state PDUs.</p> <p>value—Number that identifies the mesh group.</p> <p>Range: 1 through 4,294,967,295 ($2^{32} - 1$; 32 bits are allocated to identify a mesh group)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Mesh Groups of IS-IS Interfaces</i>

metric (Protocols IS-IS)

Syntax	<code>metric <i>metric</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Specify the metric value for the level.

All IS-IS routes have a cost, which is a routing metric that is used in the IS-IS link-state calculation. The cost is an arbitrary, dimensionless integer that can be from 1 through 63, or from 1 through 16,777,215 ($2^{24} - 1$) if you are using wide metrics.

Similar to other routing protocols, IS-IS provides a way of exporting routes from the routing table into the IS-IS network. When a route is exported into the IS-IS network without a specified metric, IS-IS uses default metric values for the route, depending on the protocol that was used to learn the route.

Table 217 on page 2710 depicts IS-IS route export metric default values.

Table 217: Default Metric Values for Routes Exported into IS-IS

Protocol Used for Learning the Route	Default Metric Value
Direct	10
Static	Same as reported by the protocol used for exporting the route
Aggregate	10
Generate	10
RIP	Same as reported by the protocol used for exporting the route
OSPF	Same as reported by the protocol used for exporting the route
BGP	10

The default metric values behavior can be customized by using routing policies.

Options	<i>metric</i> —Metric value.
	Range: 1 through 63, or 1 through 16,777,215 (if you have configured wide metrics)

Default: 10 (for all interfaces except lo0), 0 (for the lo0 interface)

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Enabling Wide IS-IS Metrics for Traffic Engineering</i> • <i>te-metric</i> • wide-metrics-only on page 2735

no-adjacency-holddown

Syntax	no-adjacency-holddown;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Disable the hold-down timer for IS-IS adjacencies.</p> <p>A hold-down timer delays the advertising of adjacencies by waiting until a time period has elapsed before labeling adjacencies in the up state. You can disable this hold-down timer, which labels adjacencies up faster. However, disabling the hold-down timer creates more frequent link-state PDU updates and SPF computation.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • hold-time on page 2694

no-authentication-check

Syntax	no-authentication-check;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Generate authenticated packets and check the authentication on received packets, but do not reject packets that cannot be authenticated.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">•• csnp-interval on page 2684• hello-authentication-type on page 2690

no-csnp-authentication

Syntax	no-csnp-authentication;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Suppress authentication check on complete sequence number PDU (CSNP) packets.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• csnp-interval on page 2684

no-eligible-backup (Protocols IS-IS)

Syntax	no-eligible-backup;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Exclude the specified interface as a backup interface for IS-IS interfaces on which link protection or node-link protection is enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Link and Node Protection for IS-IS Routes</i> • link-protection on page 2705 • node-link-protection on page 2719


no-hello-authentication

Syntax	no-hello-authentication;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Suppress authentication check on complete sequence number hello packets.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • hello-authentication-type on page 2690

no-ipv4-multicast

Syntax	no-ipv4-multicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Exclude an interface from IPv4 multicast topologies.
Default	Multicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring IS-IS Multicast Topology</i>

no-ipv4-routing

Syntax	no-ipv4-routing;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Disable IP version 4 (IPv4) routing.</p> <p>Disabling IPv4 routing has the following results:</p> <ul style="list-style-type: none"> • The routing device does not advertise the network layer protocol identifier (NLPID) for IPv4 in the Junos OS link-state PDU fragment zero. • The routing device does not advertise any IPv4 prefixes in Junos OS link-state PDUs. • The routing device does not advertise the NLPID for IPv4 in Junos OS hello packets. • The routing device does not advertise any IPv4 addresses in Junos OS hello packets. • The routing device does not calculate any IPv4 routes.
	<div>  <p>NOTE: Note: Even when <code>no-ipv4-routing</code> is configured, an IS-IS traceoptions log can list rejected IPv4 addresses. When a configuration is committed, IS-IS schedules a scan of the routing table to determine whether any routes need to be exported into the IS-IS link state database. The implicit default export policy action is to reject everything. IPv4 addresses from the routing table are examined for export, rejected by the default policy, and the rejections are logged.</p> </div>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies</i>

no-ipv6-multicast

Syntax	no-ipv6-multicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Exclude an interface from the IPv6 multicast topologies.
Default	Multicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring IS-IS Multicast Topology</i>

no-ipv6-routing

Syntax	no-ipv6-routing;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Disable IP version 6 (IPv6) routing. Disabling IPv6 routing has the following results: <ul style="list-style-type: none"> • The routing device does not advertise the network layer protocol identifier (NLPID) for IPv6 in the Junos OS link-state PDU fragment zero. • The routing device does not advertise any IPv6 prefixes in Junos OS link-state PDUs. • The routing device does not advertise the NLPID for IPv6 in Junos OS hello packets. • The routing device does not advertise any IPv6 addresses in Junos OS hello packets. • The routing device does not calculate any IPv6 routes.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies</i>

no-ipv6-unicast

Syntax	no-ipv6-unicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Exclude an interface from the IPv6 unicast topologies. This enables you to exercise control over the paths that unicast data takes through a network.
Default	IPv6 unicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies</i>

no-psnp-authentication

Syntax	no-psnp-authentication;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Suppress authentication check on partial sequence number PDU (PSNP) packets.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring IS-IS Authentication</i>

no-unicast-topology

Syntax	no-unicast-topology;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Exclude an interface from the IPv4 unicast topologies.
Default	IPv4 unicast topologies are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring IS-IS Multicast Topology</i>

node-link-protection (Protocols IS-IS)

Syntax	node-link-protection;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Enable node-link protection on the specified IS-IS interface. Junos OS creates an alternate loop-free path to the primary next hop for all destination routes that traverse a protected interface. This alternate path avoids the primary next-hop routing device altogether and establishes a path through a different routing device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Link and Node Protection for IS-IS Routes</i> • link-protection on page 2705

overload (Protocols IS-IS)

Syntax	<pre>overload { advertise-high-metrics; allow-route-leaking; timeout <i>seconds</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>], [edit protocols <i>isis</i>], [edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure the local routing device so that it appears to be overloaded. This statement causes the routing device to continue participating in IS-IS routing, but prevents it from being used for transit traffic. Traffic destined to immediately attached subnets continues to transit the routing device.</p> <p>You can also advertise maximum link metrics in network layer reachability information (NLRI) instead of setting the overload bit.</p> <p>You configure or disable overload mode in IS-IS with or without a timeout. Without a timeout, overload mode is set until it is explicitly deleted from the configuration. With a timeout, overload mode is set if the time elapsed since the IS-IS instance started is less than the specified timeout.</p> <p>A timer is started for the difference between the timeout and the time elapsed since the instance started. If the time elapsed after the IS-IS instance is enabled is less than the specified timeout, overload mode is set. When the timer expires, overload mode is cleared. In overload mode, the routing device IS-IS advertisements are originated with the overload bit set. This causes the transit traffic to take paths around the routing device. However, the overloaded routing device's own links are still accessible.</p> <p>The value of the overload bit depends on these three scenarios:</p> <ol style="list-style-type: none">1. When the overload bit has already been set to a given value and the routing process is restarted: Link-state PDUs are regenerated with the overload bit cleared.2. When the overload bit is reset to a lesser value while the routing process is running: Link-state PDUs are regenerated with the overload bit cleared.3. When the overload bit is reset to a greater value while the routing process is running: Link-state PDUs are regenerated with the overload bit set to the difference between the old and new value. <p>In overload mode, the routing device advertisement is originated with all the transit routing device links (except stub) set to a metric of 0xFFFF. The stub routing device links are</p>

advertised with the actual cost of the interfaces corresponding to the stub. This causes the transit traffic to avoid the overloaded routing device and take paths around the routing device.

To understand the reason for setting the overload bit, consider that BGP converges slowly. It is not very good at detecting that a neighbor is down because it has slow-paced keepalive timers. Once the BGP neighbor is determined to be down, it can take up to 2 minutes for a BGP router to declare the neighbor down. IS-IS is much quicker. IS-IS only takes 10-30 seconds to detect absent peers. It is the slowness of BGP, more precisely the slowness of internal BGP (IBGP), that necessitates the use of the overload bit. IS-IS and BGP routing are mutually dependent on each other. If both do not converge at the same time, traffic is dropped without notification (black holed).

You might want to configure the routing device so that it appears to be overloaded when you are restarting routing on the device. Setting the overload bit for a fixed amount of time right after a restart of the routing protocol process (rpd) ensures that the router does not receive transit traffic while the routing protocols (especially IBGP) are still converging.

Setting the overload bit is useful when performing hardware or software maintenance work on a routing device. After the maintenance work, clear the overload bit to carry on forwarding transit traffic. Manual clearing of the overload bit is not always possible. What is needed is an automated way of clearing the overload bit after some amount of time. Most networks use a time value of 300 seconds. This 5-minute value provides a good balance, allowing time to bring up even large internal IBGP meshes, while still relatively quick.

Another appropriate application for setting for the overload bit is on dedicated devices such as BGP route reflectors, which are intentionally not meant to carry any transit traffic. In this case, you would not use the timer.

You can verify that the overload bit is set by running the **show isis database** command.

Options **advertise-high-metrics**—Advertise maximum link metrics in NLRIs instead of setting the overload bit.

The **advertise-high-metric** setting is only valid while the routing device is in overload mode.

When **advertise-high-metric** is configured, IS-IS does not set the overload bit. Rather, it sets the metric to 63 or 16,777,214, depending whether wide metrics are enabled. This allows the overloaded routing device to be used for transit as a last resort.

An L1-L2 router in overload mode stops leaking route information between L1 and L2 levels and clears its attached bit. This is also true when **advertise-high-metrics** is configured.

allow-route-leaking—Enable leaking of route information into the network even if the overload bit is set.



NOTE: The **allow-route-leaking** option does not work if the routing device is in dynamic overload mode. Dynamic overload can occur if the device has exceeded its resource limits, such as the prefix limit.

timeout seconds—Number of seconds at which the overloading is reset.


Range: 60 through 1800 seconds

Default: 0 seconds

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

Related Documentation • *Example: Configuring IS-IS*

passive (Protocols IS-IS)

Syntax	<code>passive;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Advertise the direct interface addresses on an interface or into a level on the interface without actually running IS-IS on that interface or level.</p> <p>This statement effectively prevents IS-IS from running on the interface. To enable IS-IS on an interface, include the interface statement at the [edit protocols isis] or the [edit routing-instances <i>routing-instance-name</i> protocols isis] hierarchy level. To disable it, include the disable statement at those hierarchy levels. The three states—enabling, disabling, or not running IS-IS on an interface—are mutually exclusive.</p>
	<p> NOTE: Configuring IS-IS on a loopback interface automatically renders it as a passive interface, irrespective of whether the passive statement was used in the configuration of the interface.</p>
	<p>If neither passive mode nor the family iso option is configured on the IS-IS interface, then the routing device treats the interface as not being operational, and no direct IPv4/IPv6 routes are exported into IS-IS. (You configure the family iso option at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.)</p>
Default	By default, IS-IS must be configured on an interface or a level for direct interface addresses to be advertised into that level.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Multi-Level IS-IS</i> • <i>disable</i>

point-to-point

Syntax	point-to-point;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface interface-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface interface-name], [edit protocols isis interface interface-name], [edit routing-instances <i>routing-instance-name</i> protocols isis interface interface-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure an IS-IS interface to behave like a point-to-point connection.</p> <p>You can use the point-to-point statement to configure a LAN interface to act like a point-to-point interface for IS-IS. You do not need an unnumbered LAN interface, and it has no effect if configured on an interface that is already point-to-point.</p> <p>The point-to-point statement affects only IS-IS protocol procedures on that interface. All other protocols continue to treat the interface as a LAN interface. Only two IS-IS routing devices can be connected to the LAN interface, and both must be configured as point-to-point.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• IS-IS Overview on page 3343• <i>Understanding IS-IS Designated Routers</i>• <i>Example: Configuring Synchronization Between IS-IS and LDP</i>

preference (Protocols IS-IS)

Syntax	<code>preference <i>preference</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure the preference of internal routes.</p> <p>Route preferences (also known as administrative distances) are used to select which route is installed in the forwarding table when several protocols calculate routes to the same destination. The route with the lowest preference value is selected.</p> <p>To change the preference values, include the preference statement (for internal routes) or the external-preference statement.</p>
Options	<p><i>preference</i>—Preference value.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Route Preferences Overview</i> • <i>Example: Redistributing OSPF Routes into IS-IS</i> • <i>Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS</i> • external-preference on page 2687

prefix-export-limit (Protocols IS-IS)

Syntax	<code>prefix-export-limit <i>number</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols isis level level-number],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>isis level level-number],</code> <code>[edit protocols isis level level-number],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Configure a limit to the number of prefixes exported into IS-IS.</p> <p>By default, there is no limit to the number of prefixes that can be exported into IS-IS. To configure a limit to the number of prefixes that can be exported into IS-IS, include the prefix-export-limit statement. The prefix-export-limit statement protects the rest of the network from a malicious policy by applying a threshold filter for exported routes.</p> <p>The number of prefixes depends on the size of your network. Good design advice is to set it to double the total number of IS-IS Level 1 and Level 2 routing devices in your network.</p> <p>If the number of prefixes exported into IS-IS exceeds the configured limit, the overload bit is set and the overload state is reached. When other routers detect that this bit is set, they do not use this routing device for transit traffic, but they do use it for packets destined to the overloaded routing device's directly connected networks and IP prefixes. The overload state can be cleared by using the clear isis overload command.</p> <p>The show isis overview command displays the prefix export limit when it is configured.</p>
Options	<p><i>number</i>—Prefix limit.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: None</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS</i>• <i>Example: Redistributing OSPF Routes into IS-IS</i>

priority (Protocols IS-IS)

Syntax	<code>priority <i>number</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure the interface's priority for becoming the designated router. The interface with the highest priority value becomes that level's designated router.</p> <p>The priority value is meaningful only on a multiaccess network. It has no meaning on a point-to-point interface.</p> <p>A routing device advertises its priority to become a designated router in its hello packets. On all multiaccess networks, IS-IS uses the advertised priorities to elect a designated router for the network. This routing device is responsible for sending network link-state advertisements, which describe all the routing devices attached to the network. These advertisements are flooded throughout a single area.</p> <p>A routing device's priority for becoming the designated router is indicated by an arbitrary number from 0 through 127. Routing devices with a higher value are more likely to become the designated router.</p>
Options	<p><i>number</i>—Priority value.</p> <p>Range: 0 through 127</p> <p>Default: 64</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring IS-IS Designated Routers</i>

reference-bandwidth (Protocols IS-IS)

Syntax	<code>reference-bandwidth <i>reference-bandwidth</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>],</code> <code>[edit protocols <i>isis</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Optimize routing based on bandwidth by setting the reference bandwidth used in calculating the default interface cost.</p> <p>All IS-IS interfaces have a cost, which is a routing metric that is used in the IS-IS link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics. When there are several equal-cost routes to a destination, traffic is distributed equally among them.</p> <p>The cost of a route is described by a single dimensionless metric that is determined using the following formula:</p> $\text{cost} = \text{reference-bandwidth} / \text{bandwidth}$ <p>For example, if you set the reference bandwidth to 1 Gbps (that is, <i>reference-bandwidth</i> is set to 1,000,000,000), a 100-Mbps interface has a routing metric of 10.</p> <p>All IS-IS interfaces have a cost, which is a routing metric that is used in the IS-IS link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics.</p>
Options	<p><i>reference-bandwidth</i>—Reference bandwidth value in bits per second.</p> <p>Range: 9600 through 1,000,000,000,000 bps</p> <p>Default: None</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring IS-IS</i>• http://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf

rib-group (Protocols IS-IS)

Syntax	<pre> rib-group { inet <i>group-name</i>; inet6 <i>group-name</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols isis],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis],</p> <p>[edit protocols isis],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Install routes learned from IS-IS routing instances into routing tables in the IS-IS routing table group. You can install IPv4 routes or IPv6 routes.</p> <p>Support for IPv6 routing table groups in IS-IS enables IPv6 routes that are learned from IS-IS routing instances to be installed into other routing tables defined in an IS-IS routing table group.</p>
Options	<p><i>group-name</i>—Name of the routing table group.</p> <p>inet—Install IPv4 IS-IS routes.</p> <p>inet6—Install IPv6 IS-IS routes.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i> • <i>Example: Importing Direct and Static Routes Into a Routing Instance</i> • <i>Understanding Multiprotocol BGP</i>

spf-options (Protocols IS-IS)

Syntax	<pre>spf-options { delay <i>milliseconds</i>; holddown <i>milliseconds</i>; rapid-runs <i>number</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>], [edit protocols <i>isis</i>], [edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Configure options for running the shortest-path-first (SPF) algorithm.</p> <p>Running the SPF algorithm is usually the beginning of a series of larger system-wide events. For example, the SPF algorithm can lead to interior gateway protocol (IGP) prefix changes, which then lead to BGP nexthop resolution changes. Consider what happens if there are rapid link changes in the network. The local routing device can become overwhelmed. This is why it sometimes makes sense to throttle the scheduling of the SPF algorithm.</p> <p>You can configure the following SPF options:</p> <ul style="list-style-type: none">• The delay in the time between the detection of a topology change and when the SPF algorithm actually runs.• The maximum number of times that the SPF algorithm can run in succession before the hold-down timer begins.• The time to hold down, or wait, before running another SPF calculation after the SPF algorithm has run in succession the configured maximum number of times. <p>If the network stabilizes during the hold-down period and the SPF algorithm does not need to run again, the system reverts to the configured values for the delay and rapid-runs statements.</p>
Options	<p>delay <i>milliseconds</i>—Time interval between the detection of a topology change and when the SPF algorithm runs.</p> <p>Range: 50 through 1000 milliseconds</p> <p>Default: 200 milliseconds</p> <p>holddown <i>milliseconds</i>—Time interval to hold down, or wait before a subsequent SPF algorithm runs after the SPF algorithm has run the configured maximum number of times in succession.</p> <p>Range: 2000 through 10,000 milliseconds</p> <p>Default: 5000 milliseconds</p>

rapid-runs *number*—Maximum number of times the SPF algorithm can run in succession.
After the maximum is reached, the holddown interval begins.

Range: 1 through 5

Default: 3

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

Related Documentation

- *Example: Configuring Link and Node Protection for IS-IS Routes*

topologies (Protocols IS-IS)

Syntax

```
topologies {
  ipv4-multicast;
  ipv6-multicast;
  ipv6-unicast;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols *isis*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols *isis*],
[edit protocols *isis*],
[edit routing-instances *routing-instance-name* protocols *isis*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure alternate IS-IS topologies.

The remaining statements are explained separately.

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

Related Documentation

- *Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies*
- *Example: Configuring IS-IS Multicast Topology*

traceoptions (Protocols IS-IS)

Syntax	<pre>traceoptions { file <i>name</i> <size <i>size</i>> <files <i>number</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis], [edit protocols isis], [edit routing-instances <i>routing-instance-name</i> protocols isis]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure IS-IS protocol-level tracing options. To specify more than one tracing operation, include multiple flag statements.



NOTE: The **traceoptions** statement is not supported on QFabric systems.

Default	The default IS-IS protocol-level tracing options are those inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks (" "). All files are placed in the directory /var/log. We recommend that you place IS-IS tracing output in the file isis-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files</p> <p>Default: 10 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one flag, include multiple flag statements.</p> <p>IS-IS Protocol-Specific Tracing Flags</p>

- **csn**—Complete sequence number PDU (CSNP) packets
- **error**—Errored IS-IS packets
- **graceful-restart**—Graceful restart operation
- **hello**—Hello packets
- **ldp-synchronization**—Synchronization between IS-IS and LDP
- **lsp**—Link-state PDUs
- **lsp-generation**—Link-state PDU generation packets
- **packets**—All IS-IS protocol packets
- **psn**—Partial sequence number PDU (PSNP) packets
- **spf**—Shortest-path-first calculations

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations, including adjacency changes

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **receive**—Trace the packets being received.
- **send**—Trace the packets being transmitted.

no-world-readable—(Optional) Prevent any user from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level	routing and trace—To view this statement in the configuration.
	routing-control and trace-control—To add this statement to the configuration.
Related Documentation	• <i>Example: Configuring the Transmission Frequency for CSNPs on IS-IS Interfaces</i>
	• <i>Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces</i>
	• <i>Example: Enabling Packet Checksums on IS-IS Interfaces</i>

wide-metrics-only

Syntax	wide-metrics-only;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis level level-number], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis level level-number], [edit protocols isis level level-number], [edit routing-instances <i>routing-instance-name</i> protocols isis level level-number]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Configure IS-IS to generate metric values greater than 63 on a per IS-IS level basis. Normally, IS-IS metrics can have values up to 63, and IS-IS generates two type, length, and value (TLV) tuples, one for an IS-IS adjacency and the second for an IP prefix. To allow IS-IS to support traffic engineering, a second pair of TLVs has been added to IS-IS, one for IP prefixes and the second for IS-IS adjacency and traffic engineering information. With these TLVs, IS-IS metrics can have values up to 16,777,215 ($2^{24} - 1$). To configure IS-IS to generate only the new pair of TLVs and thus to allow the wider range of metric values, include the wide-metrics-only statement.
Default	By default, Junos OS supports the sending and receiving of wide metrics. Junos OS allows a maximum metric value of 63 and generates both pairs of TLVs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Enabling Wide IS-IS Metrics for Traffic Engineering</i> • <i>te-metric</i>

Administration

- [Operational Commands on page 2735](#)

Operational Commands

clear isis adjacency

Syntax	clear isis adjacency <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> <neighbor>
Syntax (EX Series Switches and QFX Series)	clear isis adjacency <instance <i>instance-name</i> > <interface <i>interface-name</i> > <neighbor>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Remove entries from the IS-IS adjacency database.
Options	none —Remove all entries from the adjacency database. instance <i>instance-name</i> —(Optional) Clear all adjacencies for the specified routing instance only. interface <i>interface-name</i> —(Optional) Clear all adjacencies for the specified interface only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. neighbor —(Optional) Clear adjacencies for the specified neighbor only.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show isis adjacency on page 2744
List of Sample Output	clear isis adjacency on page 2736
Output Fields	See show isis adjacency for an explanation of output fields.

Sample Output

clear isis adjacency

The following sample output displays IS-IS adjacency database information before and after the **clear isis adjacency** command is entered:

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface      System          L State      Hold (secs) SNPA
so-1/0/0.0     karaku1         3 Up          26
so-1/1/3.0     1921.6800.5080 3 Up          23
```

```
so-5/0/0.0    1921.6800.5080 3 Up                19
```


```
user@host> clear isis adjacency karakul
```

```
user@host> show isis adjacency
```

```
IS-IS adjacency database:
```

Interface	System	L State	Hold (secs)	SNPA
so-1/0/0.0	karakul	3 Initializing	26	
so-1/1/3.0	1921.6800.5080	3 Up	24	
so-5/0/0.0	1921.6800.5080	3 Up	21	

clear isis database

Syntax	clear isis database <entries> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <purge>
Syntax (EX Series Switches and QFX Series)	clear isis database <entries> <instance <i>instance-name</i> > <purge>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. purge option introduced in Junos OS Release 9.0. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Remove the entries from the IS-IS link-state database, which contains prefixes and topology information. You can also use purge with any of the options to initiate a network-wide purge of link-state PDUs rather than the local deletion of entries from the IS-IS link-state database.
<div> CAUTION: In a production network, the purge command option might cause short-term network-wide traffic disruptions.</div>	
Options	none —Remove all entries from the IS-IS link-state database for all routing instances. entries —(Optional) Name of the database entry. instance <i>instance-name</i> —(Optional) Clear all entries for the specified routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. purge —(Optional) Discard all entries in the IS-IS link-state database.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show isis database on page 2758
List of Sample Output	clear isis database on page 2739
Output Fields	See show isis database for an explanation of output fields.

Sample Output

clear isis database

The following sample output displays IS-IS link-state database information before and after the **clear isis database** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)
crater.00-00          0x12   0x84dd             1139
  1 LSPs
IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)
crater.00-00          0x19   0xe92c             1134
badlands.00-00        0x16   0x1454             985
carlsbad.00-00        0x33   0x220b            1015
ranier.00-00          0x2e   0xfc31             1007
1921.6800.5066.00-00  0x11   0x7313             566
1921.6800.5067.00-00  0x14   0xd9d4             939
  6 LSPs
```

```
user@host> clear isis database
```

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)

IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)
```

clear isis overload

Syntax	clear isis overload <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	clear isis overload <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Reset the IS-IS dynamic overload bit. This command can appear to not work, continuing to display overload after execution. The bit is reset only if the root cause is corrected by configuration remotely or locally.</p> <p>When other routers detect that the overload bit is set, they do not use this routing device for transit traffic, but they do use it for packets destined to the overloaded routing device's directly connected networks and IP prefixes.</p>
Options	<p>none—Reset the IS-IS dynamic overload bit.</p> <p>instance <i>instance-name</i>—(Optional) Reset the IS-IS dynamic overload bit for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show isis database on page 2758
List of Sample Output	clear isis overload on page 2740
Output Fields	See show isis database for an explanation of output fields.

Sample Output

clear isis overload

The following sample output displays IS-IS database information before and after the **clear isis overload** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
pro3-c.00-00          0x4    0x10db    1185 L1 L2 Overload

  1 LSPs
```

IS-IS level 2 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
pro3-c.00-00	0x5	0x429f	1185	L1 L2 Overload

pro2-a.00-00	0x91e	0x2589	874	L1 L2
pro2-a.02-00	0x1	0xcbc	874	L1 L2

3 LSPs

user@host> clear isis overload

user@host> show isis database

IS-IS level 1 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
pro3-c.00-00	0xa	0x429e	1183	L1 L2

1 LSPs

IS-IS level 2 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
pro3-c.00-00	0xc	0x9c39	1183	L1 L2
pro2-a.00-00	0x91e	0x2589	783	L1 L2
pro2-a.02-00	0x1	0xcbc	783	L1 L2

3 LSPs

clear isis statistics

Syntax	clear isis statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	clear isis statistics <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Set statistics about IS-IS traffic to zero.
Options	<p>none—Set IS-IS traffic statistics to zero for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Set IS-IS traffic statistics to zero for the specified routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show isis statistics on page 2782
List of Sample Output	clear isis statistics on page 2742
Output Fields	See show isis statistics for an explanation of output fields.

Sample Output

clear isis statistics

The following sample output displays IS-IS statistics before and after the **clear isis statistics** command is entered:

```
user@host> show isis statistics
IS-IS statistics for merino:
```

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	12793	12793	0	8666	719
IIH	116751	116751	0	118834	0
CSNP	203956	203956	0	204080	0
PSNP	7356	7350	6	8635	0
Unknown	0	0	0	0	0
Totals	340856	340850	6	340215	719

Total packets received: 340856 Sent: 340934

SNP queue length: 0 Drops: 0

LSP queue length: 0 Drops: 0

SPF runs: 1064
Fragments rebuilt: 1087
LSP regenerations: 436
Purges initiated: 0

user@host> clear isis statistics

user@host> show isis statistics
IS-IS statistics for merino:

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	0	0	0	0	0
IIH	3	3	0	3	0
CSNP	2	2	0	4	0
PSNP	0	0	0	0	0
Unknown	0	0	0	0	0
Totals	5	5	0	7	0

Total packets received: 5 Sent: 7

SNP queue length: 0 Drops: 0
LSP queue length: 0 Drops: 0

SPF runs: 0
Fragments rebuilt: 0
LSP regenerations: 0
Purges initiated: 0

show isis adjacency

Syntax	show isis adjacency <system-id> <brief detail extensive> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>	
Syntax (EX Series Switches and QFX Series)	show isis adjacency <system-id> <brief detail extensive> <instance <i>instance-name</i> >	
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.	
Description	Display information about IS-IS neighbors.	
Options	<p>none—Display standard information about IS-IS neighbors for all routing instances.</p> <p>system id—(Optional) Display information about IS-IS neighbors for the specified intermediate system.</p> <p>brief detail extensive—(Optional) Display standard information about IS-IS neighbors with the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display information about IS-IS neighbors for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Display information about IS-IS neighbors for all logical systems or for a particular logical system.</p>	
Required Privilege Level	view	
Related Documentation	<ul style="list-style-type: none"> • clear isis adjacency on page 2736 	
List of Sample Output	show isis adjacency on page 2746 show isis adjacency brief on page 2746 show isis adjacency detail on page 2747 show isis adjacency extensive on page 2747	
Output Fields	Table 218 on page 2744 describes the output fields for the show isis adjacency command. Output fields are listed in the approximate order in which they appear.	

Table 218: show isis adjacency Output Fields

Field Name	Field Description	Level of Output
Interface	Interface through which the neighbor is reachable.	All levels

Table 218: show isis adjacency Output Fields (*continued*)

Field Name	Field Description	Level of Output
System	System identifier (sysid), displayed as a name, if possible.	brief
L or Level	Level: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 An exclamation point (!) preceding the level number indicates that the adjacency is missing an IP address.	All levels
State	State of the adjacency: Up , Down , New , One-way , Initializing , or Rejected .	All levels
Hold (secs)	Remaining hold time of the adjacency.	brief
SNPA	Subnetwork point of attachment (MAC address of the next hop).	brief
Expires in	How long until the adjacency expires, in seconds.	detail
Priority	Priority to become the designated intermediate system.	detail extensive
Up/Down transitions	Count of adjacency status changes from Up to Down or from Down to Up .	detail
Last transition	Time of the last Up/Down transition.	detail
Circuit type	Bit mask of levels on this interface: 1=Level 1 router; 2=Level 2 router; 3=both Level 1 and Level 2 router.	detail
Speaks	Protocols supported by this neighbor.	detail extensive
MAC address	MAC address of the interface.	detail extensive
Topologies	Supported topologies.	detail extensive
Restart capable	Whether a neighbor is capable of graceful restart: Yes or No .	detail extensive
Adjacency advertisement: Advertise	This routing device has signaled to advertise this interface to its neighbors in their link-state PDUs.	detail extensive
Adjacency advertisement: Suppress	This neighbor has signaled not to advertise the interface in the routing device's outbound link-state PDUs.	detail extensive
IP addresses	IP address of this neighbor.	detail extensive

Table 218: show isis adjacency Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transition log	<p>List of recent transitions, including:</p> <ul style="list-style-type: none"> • When—Time at which an IS-IS adjacency transition occurred. • State—Current state of the IS-IS adjacency (up, down, or rejected). <ul style="list-style-type: none"> • Up—Adjacency is up and operational. • Down—Adjacency is down and not available. • Rejected—Adjacency has been rejected. • Event—Type of transition that occurred. <ul style="list-style-type: none"> • Seenself—Possible routing loop has been detected. • Interface down—IS-IS interface has gone down and is no longer available. • Error—Adjacency error. • Down reason—Reason that an IS-IS adjacency is down: <ul style="list-style-type: none"> • 3-Way Handshake Failed—Connection establishment failed. • Address Mismatch—Address mismatch caused link failure. • Aged Out—Link expired. • ISO Area Mismatch—IS-IS area mismatch caused link failure. • Bad Hello—Unacceptable hello message caused link failure. • BFD Session Down—Bidirectional failure detection caused link failure. • Interface Disabled—IS-IS interface is disabled. • Interface Down—IS-IS interface is unavailable. • Interface Level Disabled—IS-IS level is disabled. • Level Changed—IS-IS level has changed on the adjacency. • Level Mismatch—Levels on adjacency are not compatible. • MPLS LSP Down—Label-switched path (LSP) is unavailable. • MT Topology Changed—IS-IS topology has changed. • MT Topology Mismatch—IS-IS topology is mismatched. • Remote System ID Changed—Adjacency peer system ID changed. • Protocol Shutdown—IS-IS protocol is disabled. • CLI Command—Adjacency brought down by user. • Unknown—Unknown. 	extensive

Sample Output

show isis adjacency

```

user@host> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
at-2/3/0.0         ranier      3 Up          23

```

show isis adjacency brief

The output for the **show isis adjacency brief** command is identical to that for the **show isis adjacency** command. For sample output, see [show isis adjacency on page 2746](#).

show isis adjacency detail

```
user@host> show isis adjacency detail
ranier
  Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 21 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 00:01:09 ago
  Circuit type: 3, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 11.1.1.2
```

show isis adjacency extensive

```
user@host> show isis adjacency extensive
ranier
  Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 22 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 00:01:16 ago
  Circuit type: 3, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 11.1.1.2
  Transition log:
  When           State      Event      Down reason
  Wed Nov  8 21:24:25  Up        Seenself
```

show isis authentication

Syntax	show isis authentication <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis authentication <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 7.5. Command introduced in Junos OS Release 9.0 for EX Series switches. Support for hitless authentication key rollover introduced in Junos OS Release 11.2. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display information about IS-IS authentication.
Options	<p>none—Display information about IS-IS authentication.</p> <p>instance <i>instance-name</i>—(Optional) Display IS-IS authentication for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show isis authentication on page 2749 show isis authentication (With Hitless Authentication Key Rollover Configured) on page 2749
Output Fields	Table 219 on page 2748 describes the output fields for the show isis authentication command. Output fields are listed in the approximate order in which they appear.

Table 219: show isis authentication Output Fields

Field Name	Field Description
Interface	Interface name.
Level	IS-IS level.
IIH Auth	IS-IS Hello (IIH) packet authentication type. Displays the name of the active keychain if hitless authentication key rollover is configured.
CSN Auth	Complete sequence number authentication type.
PSN Auth	Partial sequence number authentication type.

Table 219: show isis authentication Output Fields (*continued*)

Field Name	Field Description
L1 LSP Authentication	Layer 1 link-state PDU authentication type.
L2 LSP Authentication	Layer 2 link-state PDU authentication type.

Sample Output

show isis authentication

```

user@host> show isis authentication
Interface          Level IIH Auth  CSN Auth  PSN Auth
at-2/3/0.0         1      Simple   Simple    Simple
                   2      MD5      MD5       MD5

L1 LSP Authentication: Simple
L2 LSP Authentication: MD5

```

show isis authentication (With Hitless Authentication Key Rollover Configured)

```

user@host> show isis authentication
Interface          Level IIH Auth  CSN Auth  PSN Auth
so-0/1/3.0         2      hakrhello MD5       MD5

L2 LSP Authentication: MD5

```

show isis backup coverage

Syntax	show isis backup coverage <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis backup coverage <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display information about the level of backup coverage available.
Options	<p>none—Display information about the level of backup coverage available for all the nodes and prefixes in the network.</p> <p>instance <i>instance-name</i>—(Optional) Display information about the level of backup coverage for a specific IS-IS routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Link and Node Protection for IS-IS Routes</i> • show isis backup label-switched-path on page 2752
List of Sample Output	show isis backup coverage on page 2751
Output Fields	Table 220 on page 2750 lists the output fields for the show isis backup coverage command. Output fields are listed in the approximate order in which they appear.

Table 220: show isis backup coverage Output Fields

Field Name	Field Description
Topology	Type of topology or address family: IPV4 Unicast or IPV6 Unicast .
Level	IS-IS level: <ul style="list-style-type: none"> • 1—Level 1 • 2—Level 2
Node	By topology, the percentage of all routes configured on the node that are protected through backup coverage.

Table 220: show isis backup coverage Output Fields (*continued*)

Field Name	Field Description
IPv4	Percentage of IPv4 unicast routes that are protected through backup coverage.
IPv6	Percentage of IPv6 unicast routes that are protected through backup coverage.
CLNS	Percentage of Connectionless Network Service (CLNS) routes that are protected through backup coverage.

Sample Output

show isis backup coverage

```
user@host> show isis backup coverage
Backup Coverage:
  Topology  Level  Node   IPv4   IPv6   CLNS
  IPv4 Unicast    2  28.57% 22.22% 0.00% 0.00%
  IPv6 Unicast    2   0.00% 0.00% 0.00% 0.00%
```

show isis backup label-switched-path

Syntax	show isis backup label-switched-path <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis backup label-switched-path
Release Information	Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display information about MPLS label-switched-paths (LSPs) designated as backup routes for IS-IS routes.
Options	none —Display information about MPLS LSPs designated as backup routes for IS-IS routes. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Link and Node Protection for IS-IS Routes</i> • show isis backup coverage on page 2750
List of Sample Output	show isis backup label-switched-path on page 2753
Output Fields	Table 221 on page 2752 lists the output fields for the show isis backup label-switched-path command. Output fields are listed in the approximate order in which they appear.

Table 221: show isis backup label-switched-path Output Fields

Field Name	Field Description
Backup MPLS LSPs	List of MPLS LSPs designated as backup paths for IS-IS routes.
Egress	IP address of the egress routing device for the LSP.
Status	State of the LSP: <ul style="list-style-type: none"> • Up—The routing device can detect RSVP hello messages from the neighbor. • Down—The routing device has received one of the following indications: <ul style="list-style-type: none"> • Communication failure from the neighbor. • Communication from IGP that the neighbor is unavailable. • Change in the sequence numbers in the RSVP hello messages sent by the neighbor. • Deleted—LSP is no longer available as a backup path.

Table 221: show isis backup label-switched-path Output Fields (*continued*)

Field Name	Field Description
Last change	Time elapsed since the neighbor state changed either from up to down or from down to up. The format is <i>hh:mm:ss</i> .
TE-metric	Configured traffic engineering metric.
Metric	Configured metric.

Sample Output

show isis backup label-switched-path

```
user@host> show isis backup label-switched-path
Backup MPLS LSPs:
f-to-g, Egress: 192.168.1.4, Status: up, Last change: 06:12:03
TE-metric: 9, Metric: 0
```

show isis backup spf results

Syntax	<code>show isis backup spf results</code> <code><instance <i>instance-name</i>></code> <code><level (1 2)></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><no-coverage></code> <code><topology (ipv4-unicast ipv6-multicast ipv6-unicast unicast)></code>
Syntax (EX Series Switches)	<code>show isis backup spf results</code> <code><instance <i>instance-name</i>></code> <code><level (1 2)></code> <code><no-coverage></code> <code><topology (ipv4-unicast unicast)></code>
Release Information	Command introduced in Junos OS Release 9.5.
Description	Display information about IS-IS shortest-path-first (SPF) calculations for backup paths.
Options	<p>none—Display information about IS-IS SPF calculations for all backup paths for all destination nodes.</p> <p>instance <i>instance-name</i>—(Optional) Display SPF calculations for backup paths for the specified routing instance.</p> <p>level (1 2)—(Optional) Display SPF calculations for the backup paths for the specified IS-IS level.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Display SPF calculations for the backup paths for all logical systems or on a particular logical system.</p> <p>no-coverage—(Optional) Display SPF calculations only for destinations that do not have backup coverage.</p> <p>topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)—(Optional) Display SPF calculations for backup paths for the specified topology only.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Link and Node Protection for IS-IS Routes</i>• show isis backup coverage on page 2750
List of Sample Output	show isis backup spf results on page 2755 show isis backup spf results no-coverage on page 2756
Output Fields	Table 222 on page 2755 lists the output fields for the show isis backup spf results command. Output fields are listed in the approximate order in which they appear.

Table 222: show isis backup spf results Output Fields

Field Name	Field Description
<i>node-name</i>	Name of the destination node.
Address	Address of the destination node.
Primary next-hop	Interface and name of the node of the primary next hop to reach the destination.
Root	Name of the next-hop neighbor.
Metric	Metric to the node.
Eligible	Indicates that the next-hop neighbor has been designated as a backup path to the destination node.
Backup next-hop	Name of the interface of the backup next hop.
SNPA	Subnetwork point of attachment (MAC address of the next hop).
LSP	Name of the MPLS label-switched path (LSP) designated as a backup path.
Not eligible	Indicates that the next-hop neighbor cannot function as a backup path to the destination.
Reason	Describes why the next-hop neighbor is designated as Not eligible as a backup path.

Sample Output

show isis backup spf results

```

user@host> show isis backup spf results

IS-IS level 1 SPF results:
  0 nodes

IS-IS level 2 SPF results:
banff.00
  Primary next-hop: so-6/0/0.0, IPV4, olympic
  Primary next-hop: ae0.0, IPV4, camaro, SNPA: 0:90:69:f:67:f0
  Primary next-hop: so-6/0/0.0, IPV6, olympic
  Primary next-hop: ae0.0, IPV6, camaro, SNPA: 0:90:69:f:67:f0
  Root: camaro, Root Metric: 10, Metric: 10
  Not eligible, Reason: Primary next-hop multipath
  Root: olympic, Root Metric: 10, Metric: 10
  Not eligible, Reason: Primary next-hop multipath
  Root: glacier, Root Metric: 10, Metric: 25
  Not eligible, Reason: Primary next-hop multipath
crater.00
  Primary next-hop: so-6/0/0.0, IPV4, olympic
  Primary next-hop: so-6/0/0.0, IPV6, olympic

```

```

Root: olympic, Root Metric: 10, Metric: 10
  Not eligible, Reason: Primary next-hop link fate sharing
Root: glacier, Root Metric: 10, Metric: 15
  Eligible, Backup next-hop: as0.0, IPV4, glacier
  Eligible, Backup next-hop: as0.0, IPV6, glacier
Root: camaro, Root Metric: 10, Metric: 20
  Not eligible, Reason: Interface is already covered
olympic.00
Primary next-hop: so-6/0/0.0, IPV4, olympic
Primary next-hop: so-6/0/0.0, IPV6, olympic
Root: olympic, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing
Root: camaro, Root Metric: 10, Metric: 20
  track-item: olympic.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
Root: glacier, Root Metric: 10, Metric: 20
  track-item: olympic.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
camaro.00
Primary next-hop: ae0.0, IPV4, camaro, SNPA: 0:90:69:f:67:f0
Primary next-hop: ae0.0, IPV6, camaro, SNPA: 0:90:69:f:67:f0
Root: camaro, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing
Root: glacier, Root Metric: 10, Metric: 20
  track-item: camaro.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
Root: olympic, Root Metric: 10, Metric: 20
  track-item: camaro.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
glacier.00
Primary next-hop: as0.0, IPV4, glacier
Primary next-hop: as0.0, IPV6, glacier
Root: glacier, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing
Root: camaro, Root Metric: 10, Metric: 20
  track-item: glacier.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
Root: olympic, Root Metric: 10, Metric: 20
  track-item: glacier.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
5 nodes

```

show isis backup spf results no-coverage

```
user@host> show isis backup spf results no-coverage
```

```
IS-IS level 1 SPF results:
0 nodes
```

```
IS-IS level 2 SPF results:
olympic.00
Primary next-hop: so-6/0/0.0, IPV4, olympic
Primary next-hop: so-6/0/0.0, IPV6, olympic
Root: olympic, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing

```

```
Root: camaro, Root Metric: 10, Metric: 20
  track-item: olympic.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
Root: glacier, Root Metric: 10, Metric: 20
  track-item: olympic.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
camaro.00
  Primary next-hop: ae0.0, IPV4, camaro, SNPA: 0:90:69:f:67:f0
  Primary next-hop: ae0.0, IPV6, camaro, SNPA: 0:90:69:f:67:f0
  Root: camaro, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing
  Root: glacier, Root Metric: 10, Metric: 20
  track-item: camaro.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
  Root: olympic, Root Metric: 10, Metric: 20
  track-item: camaro.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
glacier.00
  Primary next-hop: as0.0, IPV4, glacier
  Primary next-hop: as0.0, IPV6, glacier
  Root: glacier, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next-hop link fate sharing
  Root: camaro, Root Metric: 10, Metric: 20
  track-item: glacier.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
  Root: olympic, Root Metric: 10, Metric: 20
  track-item: glacier.00-00
  track-item: kobuk.00-00
  Not eligible, Reason: Path loops
3 nodes
```

show isis database

Syntax	<code>show isis database</code> <code><system-id></code> <code><brief detail extensive></code> <code><instance <i>instance-name</i>></code> <code><level (1 2)></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches and QFX Series)	<code>show isis database</code> <code><system-id></code> <code><brief detail extensive></code> <code><level (1 2)></code> <code><instance <i>instance-name</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display the entries in the IS-IS link-state database, which contains data about PDU packets.
Options	<p>none—Display standard information about IS-IS link-state database entries for all routing instances.</p> <p><i>system id</i>—(Optional) Display IS-IS link-state database entries for the specified intermediate system.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display IS-IS link-state database entries for the specified routing instance.</p> <p>level (1 2)—(Optional) Display IS-IS link-state database entries for the specified IS-IS level.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Display standard information about IS-IS link-state database entries for all logical systems or for a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear isis database on page 2738
List of Sample Output	show isis database on page 2760 show isis database brief on page 2761 show isis database detail on page 2761 show isis database extensive on page 2761

Output Fields Table 223 on page 2759 describes the output fields for the **show isis database** command. Output fields are listed in the approximate order in which they appear. Fields that contain internal IS-IS information useful only in troubleshooting obscure problems are not described in the table. For more details about these fields, contact your customer support representative.

Table 223: show isis database Output Fields

Field Name	Field Description	Level of Output
Interface name	Name of the interface on which the link-state PDU has been received; always IS-IS for this command.	All levels
level	Level of intermediate system: <ul style="list-style-type: none"> • 1—Intermediate system routes within an area; when the destination is outside an area, it routes toward a Level 2 system. • 2—Intermediate system routes between areas and toward other ASs. 	All levels
LSP ID	Link-state PDU identifier.	All levels
Sequence	Sequence number of the link-state PDU.	All levels
Checksum	Checksum value of the link-state PDU.	All levels
Lifetime (secs)	Remaining lifetime of the link-state PDU, in seconds.	All levels
Attributes	Attributes of the specified database: L1 , L2 , Overload , or Attached (L1 only).	none brief
# LSPs	Total number of link-state PDUs in the specified link-state database.	none brief
IP prefix	Prefix advertised by this link-state PDU.	detail extensive
IS neighbor	IS-IS neighbor of the advertising system.	detail extensive
ES neighbor	(J Series routers only) An ES-IS neighbor of the advertising system.	detail extensive
IP prefix	IPv4 prefix advertised by this link-state PDU.	detail extensive
V6 prefix	IPv6 prefix advertised by this link-state PDU.	detail extensive
Metric	Metric of the prefix or neighbor.	detail extensive
Header	<ul style="list-style-type: none"> • LSP ID—Link state PDU identifier of the header. • Length—Header length. • Allocated Length—Amount of length available for the header. • Router ID—Address of the local routing device. • Remaining Lifetime—Remaining lifetime of the link-state PDU, in seconds. 	extensive

Table 223: show isis database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Packet	<ul style="list-style-type: none"> • LSP ID—The identifier for the link-state PDU. • Length—Packet length. • Lifetime—Remaining lifetime, in seconds. • Checksum—The checksum of the link-state PDU. • Sequence—The sequence number of the link-state PDU. Every time the link-state PDU is updated, this number increments. • Attributes—Packet attributes. • NLPID—Network layer protocol identifier. • Fixed length—Specifies the set length for the packet. 	extensive
TLVs	<ul style="list-style-type: none"> • Area Address—Area addresses that the routing device can reach. • Speaks—Supported routing protocols. • IP router id—ID of the routing device (usually the IP address). • IP address—IPv4 address. • Hostname—Assigned name of the routing device. • IP prefix—IP prefix of the routing device. • Metric—IS-IS metric that measures the cost of the adjacency between the originating routing device and the advertised routing device. • IP extended prefix—Extended IP prefix of the routing device. • IS neighbor—Directly attached neighbor's name and metric. • IS extended neighbor—Directly attached neighbor's name, metric, and IP address. 	extensive

Sample Output

show isis database

```

user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
kobuk.00-00           0x3    0x3167    1057 L1 L2
camaro.00-00          0x5    0x770e    1091 L1 L2
ranier.00-00          0x4    0xaa95    1091 L1 L2
glacier.00-00         0x4    0x206f    1089 L1 L2
glacier.02-00         0x1    0xd141    1089 L1 L2
badlands.00-00        0x3    0x87a2    1093 L1 L2
    6 LSPs

IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
kobuk.00-00           0x6    0x8d6b    1096 L1 L2
camaro.00-00          0x9    0x877b    1101 L1 L2
ranier.00-00          0x8    0x855d    1103 L1 L2
glacier.00-00         0x7    0xf892    1098 L1 L2
glacier.02-00         0x1    0xd141    1089 L1 L2
badlands.00-00        0x6    0x562     1105 L1 L2
    6 LSPs

```

show isis database brief

The output for the **show isis database brief** command is identical to that for the **show isis database** command. For sample output, see [show isis database on page 2760](#).

show isis database detail

```
user@host> show isis database logical-system CE3 sisira.00-00 detail
```

IS-IS level 1 link-state database:

```
sisira.00-00 Sequence: 0x11, Checksum: 0x10fc, Lifetime: 975 secs
  IS neighbor: hemantha-CE3.02           Metric:      10
  ES neighbor: 0015.0015.0015           Metric:      10 Down
  ES neighbor: 0025.0025.0025           Metric:      10 Down
  ES neighbor: 0030.0030.0030           Metric:      10 Down
  ES neighbor: 0040.0040.0040           Metric:      10 Down
  ES neighbor: sisira                     Metric:       0
  IP prefix: 1.0.0.0/24                  Metric:      10 External Down
  IP prefix: 3.0.0.0/24                  Metric:      10 External Down
  IP prefix: 4.0.0.0/24                  Metric:      10 External Down
  IP prefix: 5.0.0.0/24                  Metric:      10 Internal Up
  IP prefix: 15.15.15.15/32              Metric:      10 External Down
  IP prefix: 25.25.25.25/32              Metric:      10 External Down
  IP prefix: 30.30.30.30/32              Metric:      10 External Down
  IP prefix: 40.40.40.40/32              Metric:      10 External Down
  IP prefix: 60.60.60.60/32              Metric:       0 Internal Up
```

IS-IS level 2 link-state database:

```
sisira.00-00 Sequence: 0x13, Checksum: 0x69ac, Lifetime: 993 secs
  IS neighbor: hemantha-CE3.02           Metric:      10
  IP prefix: 1.0.0.0/24                  Metric:      10 External Down
  IP prefix: 3.0.0.0/24                  Metric:      10 External Down
  IP prefix: 4.0.0.0/24                  Metric:      10 External Down
  IP prefix: 5.0.0.0/24                  Metric:      10 Internal Up
  IP prefix: 15.15.15.15/32              Metric:      10 External Down
  IP prefix: 25.25.25.25/32              Metric:      10 External Down
  IP prefix: 30.30.30.30/32              Metric:      10 External Down
  IP prefix: 40.40.40.40/32              Metric:      10 External Down
  IP prefix: 50.50.50.50/32              Metric:      10 Internal Up
  IP prefix: 60.60.60.60/32              Metric:       0 Internal Up
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
                                          Metric:       0 Internal Up
```

show isis database extensive

```
user@host> show isis database logical-system CE3 sisira.00-00 extensive
```

IS-IS level 1 link-state database:

```
sisira.00-00 Sequence: 0x11, Checksum: 0x10fc, Lifetime: 970 secs
```

```

IS neighbor: hemantha-CE3.02           Metric:      10
Two-way fragment: hemantha-CE3.02-00, Two-way first fragment:
hemantha-CE3.02-00
ES neighbor: 0015.0015.0015           Metric:      10 Down
ES neighbor: 0025.0025.0025           Metric:      10 Down
ES neighbor: 0030.0030.0030           Metric:      10 Down
ES neighbor: 0040.0040.0040           Metric:      10 Down
ES neighbor: sisira                    Metric:      0
IP prefix: 1.0.0.0/24                  Metric:      10 External Down
IP prefix: 3.0.0.0/24                  Metric:      10 External Down
IP prefix: 4.0.0.0/24                  Metric:      10 External Down
IP prefix: 5.0.0.0/24                  Metric:      10 Internal Up
IP prefix: 15.15.15.15/32              Metric:      10 External Down
IP prefix: 25.25.25.25/32              Metric:      10 External Down
IP prefix: 30.30.30.30/32              Metric:      10 External Down
IP prefix: 40.40.40.40/32              Metric:      10 External Down
IP prefix: 60.60.60.60/32              Metric:      0 Internal Up

```

```

Header: LSP ID: sisira.00-00, Length: 336 bytes
Allocated length: 336 bytes, Router ID: 0.0.0.0
Remaining lifetime: 970 secs, Level: 1, Interface: 333
Estimated free bytes: 144, Actual free bytes: 0
Aging timer expires in: 970 secs
Protocols: IP, IPv6, CLNS

```

```

Packet: LSP ID: sisira.00-00, Length: 336 bytes, Lifetime : 1198 secs
Checksum: 0x10fc, Sequence: 0x11, Attributes: 0xb L1 L2 Attached
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
Packet type: 18, Packet version: 1, Max area: 0

```

TLVs:

```

Area address: 60.0006.80ff.f800.0000.0108.0001 (13)
Speaks: IP
Speaks: IPV6
Speaks: CLNP
Hostname: sisira
ES neighbor TLV: Internal, Metric: default 0, Up
  ES: sisira
IS neighbor: hemantha-CE3.02, Internal, Metric: default 10
IS extended neighbor: hemantha-CE3.02, Metric: default 10
ES neighbor TLV: External, Metric: default 10, Down
  ES: 0040.0040.0040
ES neighbor TLV: External, Metric: default 10, Down
  ES: 0025.0025.0025
ES neighbor TLV: External, Metric: default 10, Down
  ES: 0015.0015.0015
ES neighbor TLV: External, Metric: default 10, Down
  ES: 0030.0030.0030
IP external prefix: 3.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 40.40.40.40/32, Internal, Metric: default 10, Down
IP external prefix: 4.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 25.25.25.25/32, Internal, Metric: default 10, Down
IP external prefix: 15.15.15.15/32, Internal, Metric: default 10, Down
IP external prefix: 1.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 30.30.30.30/32, Internal, Metric: default 10, Down
IP extended prefix: 3.0.0.0/24 metric 10 down
IP extended prefix: 40.40.40.40/32 metric 10 down
IP extended prefix: 4.0.0.0/24 metric 10 down
IP extended prefix: 25.25.25.25/32 metric 10 down
IP extended prefix: 15.15.15.15/32 metric 10 down
IP extended prefix: 1.0.0.0/24 metric 10 down

```



```

IP extended prefix: 30.30.30.30/32 metric 10 down
IP prefix: 60.60.60.60/32, Internal, Metric: default 0, Up
IP prefix: 5.0.0.0/24, Internal, Metric: default 10, Up
IP extended prefix: 60.60.60.60/32 metric 0 up
IP extended prefix: 5.0.0.0/24 metric 10 up
No queued transmissions

```

IS-IS level 2 link-state database:

```

sisira.00-00 Sequence: 0x13, Checksum: 0x69ac, Lifetime: 988 secs
IS neighbor: hemantha-CE3.02 Metric: 10
Two-way fragment: hemantha-CE3.02-00, Two-way first fragment:
hemantha-CE3.02-00
IP prefix: 1.0.0.0/24 Metric: 10 External Down
IP prefix: 3.0.0.0/24 Metric: 10 External Down
IP prefix: 4.0.0.0/24 Metric: 10 External Down
IP prefix: 5.0.0.0/24 Metric: 10 Internal Up
IP prefix: 15.15.15.15/32 Metric: 10 External Down
IP prefix: 25.25.25.25/32 Metric: 10 External Down
IP prefix: 30.30.30.30/32 Metric: 10 External Down
IP prefix: 40.40.40.40/32 Metric: 10 External Down
IP prefix: 50.50.50.50/32 Metric: 10 Internal Up
IP prefix: 60.60.60.60/32 Metric: 0 Internal Up
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
Metric: 10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
Metric: 10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
Metric: 10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
Metric: 10 External Down
ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
Metric: 0 Internal Up

```

```

Header: LSP ID: sisira.00-00, Length: 427 bytes
Allocated length: 427 bytes, Router ID: 0.0.0.0
Remaining lifetime: 988 secs, Level: 2, Interface: 333
Estimated free bytes: 130, Actual free bytes: 0
Aging timer expires in: 988 secs
Protocols: IP, IPv6, CLNS

```

```

Packet: LSP ID: sisira.00-00, Length: 427 bytes, Lifetime : 1198 secs
Checksum: 0x69ac, Sequence: 0x13, Attributes: 0x3 L1 L2
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
Packet type: 20, Packet version: 1, Max area: 0

```

TLVs:

```

Area address: 60.0006.80ff.f800.0000.0108.0001 (13)
Speaks: IP
Speaks: IPV6
Speaks: CLNP
Hostname: sisira
IS neighbor: hemantha-CE3.02, Internal, Metric: default 10
IS extended neighbor: hemantha-CE3.02, Metric: default 10
IP external prefix: 3.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 40.40.40.40/32, Internal, Metric: default 10, Down
IP external prefix: 4.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 25.25.25.25/32, Internal, Metric: default 10, Down
IP external prefix: 15.15.15.15/32, Internal, Metric: default 10, Down
IP external prefix: 1.0.0.0/24, Internal, Metric: default 10, Down
IP external prefix: 30.30.30.30/32, Internal, Metric: default 10, Down

```

```
IP extended prefix: 3.0.0.0/24 metric 10 down
IP extended prefix: 40.40.40.40/32 metric 10 down
IP extended prefix: 4.0.0.0/24 metric 10 down
IP extended prefix: 25.25.25.25/32 metric 10 down
IP extended prefix: 15.15.15.15/32 metric 10 down
IP extended prefix: 1.0.0.0/24 metric 10 down
IP extended prefix: 30.30.30.30/32 metric 10 down
ISO prefix-neighbor TLV: Internal, Metric: default 0, Up
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
ISO prefix-neighbor TLV: External, Metric: default 10, Down
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
ISO prefix-neighbor TLV: External, Metric: default 10, Down
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
ISO prefix-neighbor TLV: External, Metric: default 10, Down
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
ISO prefix-neighbor TLV: External, Metric: default 10, Down
  Prefix : 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
IP prefix: 60.60.60.60/32, Internal, Metric: default 0, Up
IP prefix: 5.0.0.0/24, Internal, Metric: default 10, Up
IP prefix: 50.50.50.50/32, Internal, Metric: default 10, Up
IP extended prefix: 60.60.60.60/32 metric 0 up
IP extended prefix: 5.0.0.0/24 metric 10 up
IP extended prefix: 50.50.50.50/32 metric 10 up
No queued transmissions
```

show isis hostname

Syntax	show isis hostname <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis hostname
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display IS-IS hostname database information.
Options	none —Display IS-IS hostname database information. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show isis hostname on page 2765
Output Fields	Table 224 on page 2765 describes the output fields for the show isis hostname command. Output fields are listed in the approximate order in which they appear.

Table 224: show isis hostname Output Fields

Field Name	Field Description
System Id	System identifier mapped to the hostname.
Hostname	Hostname mapped to the system identifier.
Type	Type of mapping between system identifier and hostname. <ul style="list-style-type: none">• Dynamic—Hostname mapping determined as described in RFC 2763, <i>Dynamic Hostname Exchange Mechanism for IS-IS</i>.• Static—Hostname mapping configured by user.

Sample Output

show isis hostname

user@host> show isis hostname		
IS-IS hostname database:		
System Id	Hostname	Type
1921.6800.4201	isis1	Dynamic
1921.6800.4202	isis2	Static
1921.6800.4203	isis3	Dynamic

show isis interface

Syntax	<pre>show isis interface <brief detail extensive> <interface-name> <logical-system (all logical-system-name)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis interface <brief detail extensive> <interface-name></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Display status information about IS-IS-enabled interfaces.</p> <div data-bbox="474 823 542 892" data-label="Image"> </div> <p>NOTE: If the configured metric for an IS-IS level is above 63, and the wide-metrics-only statement is not configured, the show isis interface detail command and the show isis interface extensive command display 63 as the metric value for that level. Configure the wide-metrics-only statement to generate metric values greater than 63 on a per IS-IS level basis.</p> <p>The show isis interface command displays the configured metric value for an IS-IS level irrespective of whether is configured or not.</p>
Options	<p>none—Display standard information about all IS-IS-enabled interfaces.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>interface-name—(Optional) Display information about the specified interface only.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> <i>Example: Enabling Wide IS-IS Metrics for Traffic Engineering</i>
List of Sample Output	<p>show isis interface on page 2768</p> <p>show isis interface brief on page 2768</p> <p>show isis interface detail on page 2769</p> <p>show isis interface extensive on page 2769</p> <p>show isis interface extensive (With LDP) on page 2769</p>
Output Fields	<p>Table 225 on page 2767 describes the output fields for the show isis interface command. Output fields are listed in the approximate order in which they appear.</p>

Table 225: show isis interface Output Fields

Field Name	Field Description	Level of Output
<i>interface-name</i>	Name of the interface.	detail
Designated router	Routing device selected by other routers that is responsible for sending link-state advertisements that describe the network. Used only on broadcast networks.	detail
Index	Interface index assigned by the Junos OS kernel.	detail
State	Internal implementation information.	detail
Circuit id	Circuit identifier.	detail
Circuit type	Circuit type: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 	detail
LSP interval	Interval between link-state PDUs sent from the interface.	detail
CSNP interval	Interval between complete sequence number PDUs sent from the interface.	detail extensive
Sysid	System identifier.	detail
Interface	Interface through which the adjacency is made.	none brief
L or Level	Level: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 	All levels
CirID	Circuit identifier.	none brief
Level 1 DR	Level 1 designated intermediate system.	none brief
Level 2 DR	Level 2 designated intermediate system.	none brief
L1/L2 Metric	Interface's metric for Level 1 and Level 2. If there is no information, the metric is 0.	none brief
Adjacency advertisement: Advertise	This routing device has signaled to advertise this interface to its neighbors in their label-switched paths (LSPs).	detail extensive
Adjacency advertisement: Suppress	This neighbor has signaled not to advertise this interface in the routing device's outbound LSPs.	detail extensive
Adjacencies	Number of adjacencies established on this interface.	detail

Table 225: show isis interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Priority	Priority value for this interface.	detail
Metric	Metric value for this interface.	detail
Hello(s) / Hello Interval	Interface's hello interval.	detail extensive
Hold(s) / Hold Time	Interface's hold time.	detail extensive
Designated Router	Router responsible for sending network link-state advertisements, which describe all the routing devices attached to the network.	detail
Hello padding	Type of hello padding: <ul style="list-style-type: none"> • Adaptive—On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state TLV. If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface. • Loose—(Default) The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state. • Strict—Padding is performed on all interface types and for all adjacency states, and is continuous. 	extensive
LDP sync state	Current LDP synchronization state: in sync , in holddown , or not supported .	extensive
reason	Reason for being in the LDP sync state.	extensive
config holdtime	Configured value of the hold timer.	extensive
remaining	If the state is not in sync and the hold time is not infinity, then this field displays the remaining hold time in seconds.	extensive

Sample Output

show isis interface

```

user@host> show isis interface
IS-IS interface database:
Interface          L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
at-2/3/0.0         3   0x1 Point to Point    Point to Point    10/10
1o0.0              0   0x1 Passive           Passive           0/0

```

show isis interface brief

The output for the **show isis interface brief** command is identical to that for the **show isis interface** command. For sample output, see [show isis interface on page 2768](#).

show isis interface detail

```

user@host> show isis interface detail
IS-IS interface database:
at-2/3/0.0
  Index: 66, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 5 s
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1           1           64      10      9.000      27
    2           1           64      10      9.000      27
lo0.0
  Index: 64, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1           0           64       0      0 Passive
    2           0           64       0      0 Passive

```

show isis interface extensive

```

user@host> show isis interface extensive
IS-IS interface database:
at-2/3/0.0
  Index: 66, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 5 s, Loose Hello padding
  Level 1
    Adjacencies: 1, Priority: 64, Metric: 10
    Hello Interval: 9.000 s, Hold Time: 27 s
  Level 2
    Adjacencies: 1, Priority: 64, Metric: 10
    Hello Interval: 9.000 s, Hold Time: 27 s
lo0.0
  Index: 64, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled, Loose Hello padding
  Level 1
    Adjacencies: 0, Priority: 64, Metric: 0
    Passive
  Level 2
    Adjacencies: 0, Priority: 64, Metric: 0
    Passive

```

show isis interface extensive (With LDP)

```

user@host> show isis interface extensive
IS-IS interface database:
so-1/1/2.0
  Index: 114, State: 0x6, Circuit id: 0x1, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 20 s, Loose Hello padding
  Adjacency advertisement: Advertise
  LDP sync state: in sync, for: 00:01:28, reason: LDP up during config
  config holdtime: 20 seconds
  Level 2
    Adjacencies: 1, Priority: 64, Metric: 11
    Hello Interval: 9.000 s, Hold Time: 27 s
    IPV4 MulticastMetric: 10
    IPV6 UnicastMetric: 10

```

show isis overview

Syntax	show isis overview <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis overview <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 8.5. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display IS-IS overview information.
Options	none —Display standard overview information about IS-IS for all routing instances. instance <i>instance-name</i> —(Optional) Display overview information for the specified routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show isis overview on page 2771
Output Fields	Table 226 on page 2770 lists the output fields for the show isis overview command. Output fields are listed in the approximate order in which they appear.

Table 226: show isis overview Output Fields

Field Name	Field Description
Instance	IS-IS routing instance.
Router ID	Router ID of the routing device.
Adjacency holddown	Adjacency holddown capability: enabled or disabled .
Maximum Areas	Maximum number of IS-IS areas advertised by the routing device.
LSP life time	Lifetime of the link-state PDU, in seconds.
Attached bit evaluation	Attached bit capability: enabled or disabled .
SPF delay	Delay before performing consecutive shortest-path-first (SPF) calculations.
SPF holddown	Delay before performing additional SPF calculations after the maximum number of consecutive SPF calculations is reached.

Table 226: show isis overview Output Fields (*continued*)

Field Name	Field Description
SPF rapid runs	Maximum number of SPF calculations that can be performed in succession before the holddown timer begins.
Overload bit at startup is set	Overload bit capability is enabled.
Overload high metrics	Overload high metrics capability: enabled or disabled .
Overload timeout	Time period after which overload is reset and the time that remains before the timer is set to expire.
Traffic engineering	Traffic engineering capability: enabled or disabled .
Restart	Graceful restart capability: enabled or disabled .
Restart duration	Time period for complete reacquisition of IS-IS neighbors.
Helper mode	Graceful restart helper capability: enabled or disabled .
Level	IS-IS level: <ul style="list-style-type: none"> • 1—Level 1 information • 2—Level 2 information
IPv4 is enabled	IP Protocol version 4 capability is enabled.
IPv6 is enabled	IP Protocol version 6 capability is enabled.
CLNS is enabled	(J Series routers only) OSI CLNP capability is enabled.
Internal route preference	Preference value of internal routes.
External route preference	Preference value of external routes.
Wide area metrics are enabled	Wide area metrics capability is enabled.
Narrow metrics are enabled	Narrow metrics capability is enabled.

Sample Output

show isis overview

```

user@host> show isis overview
Instance: master
Router ID: 192.168.1.220
Adjacency holddown: enabled

```

Maximum Areas: 3
LSP life time: 65535
Attached bit evaluation: enabled
SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
Overload bit at startup is set
 Overload high metrics: disabled
 Overload timeout: 300 sec, expires in 295 seconds
IPv4 is enabled, IPv6 is enabled
Traffic engineering: enabled
Restart: Enabled
 Restart duration: 210 sec
 Helper mode: Enabled
Level 1
 Internal route preference: 15
 External route preference: 160
 Wide metrics are enabled, Narrow metrics are enabled
Level 2
 Internal route preference: 18
 External route preference: 165
 Wide metrics are enabled

show isis route

Syntax	<pre>show isis route <destination> <inet inet6> <instance instance-name> <logical-system (all logical-system-name)> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis route <destination> <inet inet6> <instance instance-name> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Display the routes in the IS-IS routing table.
Options	<p>none—Display all routes in the IS-IS routing table for all supported address families for all routing instances.</p> <p>destination—(Optional) Destination address for the route.</p> <p>inet inet6—(Optional) Display inet (IPv4) or inet6 (IPv6) routes, respectively.</p> <p>instance instance-name—(Optional) Display routes for the specified routing instance only.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)—(Optional) Display routes for the specified topology only, or use unicast to display information, if available, for both IPv4 and IPv6 unicast topologies.</p>
Required Privilege Level	view
List of Sample Output	<p>show isis route logical-system on page 2774</p> <p>show isis route (CLNS) on page 2774</p> <p>show isis route on page 2775</p>
Output Fields	<p>Table 227 on page 2773 describes the output fields for the show isis route command. Output fields are listed in the approximate order in which they appear.</p>

Table 227: show isis route Output Fields

Field Name	Field Description
Current version	Number of the current version of the IS-IS routing table.

Table 227: show isis route Output Fields (*continued*)

Field Name	Field Description
L1	Version of Level 1 SPF that was run.
L2	Version of Level 2 SPF that was run.
Prefix	Destination of the route.
L	IS-IS level: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2
Version	Version of SPF that generated the route.
Metric	Metric value associated with the route.
Type	Metric type: int (internal) or ext (external).
Interface	Interface to the next hop.
Via	System identifier of the next hop, displayed as a name if possible.
ISO Routes	ISO routing table entries.
snpa	MAC address.

Sample Output

show isis route logical-system

```

user@host> show isis route logical-system ls1
IS-IS routing table                      Current version: L1: 8 L2: 11
Prefix      L Version Metric Type Interface  Via
10.9.7.0/30  2      11    20 int  gr-0/2/0.0  h
10.9.201.1/32 2      11    60 int  gr-0/2/0.0  h
IPv6 Unicast IS-IS routing table        Current version: L1: 9 L2: 11
Prefix      L Version Metric Type Interface  Via
8009:3::a09:3200/126 2      11    20 int  gr-0/2/0.0  h

```

show isis route (CLNS)

```

user@host> show isis route
IS-IS routing table                      Current version: L1: 10 L2: 8
IPv4/IPv6 Routes
Prefix      L Version Metric Type Interface  Via
0.0.0.0/0   1      10    10 int  fe-0/0/1.0  ISIS.0
ISO Routes
Prefix L   Version Metric Type Interface  Via  snpa
0/0    1      10    10 int  fe-0/0/1.0  isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001/104

```

```

1          10          0 int
47.0005.80ff.f800.0000.0108.0001.1921.6800.4001/152
1          10          10 int fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001.1921.6800.4002/152
1          10          20 int fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0002/104
1          10          0 int
47.0005.80ff.f800.0000.0108.0002.1921.6800.4001/152
1          10          10 int fe-0/0/1.0 isis.0 0:12:0:34:0:56

```

show isis route

```
user@host> show isis route
```

```

IS-IS routing table          Current version: L1: 4 L2: 13
IPv4/IPv6 Routes
-----
Prefix                      L   Version  Metric Type Interface      NH   Via
10.255.71.52/32             2    13        10   int  ae0.0                 IPV4 camaro
10.255.71.238/32           2    13        20   int  so-6/0/0.0           IPV4 olympic
                           as0.0                 IPV4 glacier
10.255.71.239/32           2    13        20   int  so-6/0/0.0           IPV4 olympic
                           ae0.0                 IPV4 camaro
10.255.71.242/32           2    13        10   int  as0.0                 IPV4 glacier
10.255.71.243/32           2    13        10   int  so-6/0/0.0           IPV4 olympic
12.13.0.0/30                2    13        20   int  so-6/0/0.0           IPV4 olympic
12.15.0.0/30                2    13        20   int  so-6/0/0.0           IPV4 olympic
13.15.0.0/30                2    13        30   int  ae0.0                 IPV4 camaro
                           so-6/0/0.0           IPV4 olympic
                           as0.0                 IPV4 glacier
13.16.0.0/30                2    13        25   int  as0.0                 IPV4 glacier
14.15.0.0/30                2    13        20   int  ae0.0                 IPV4 camaro
192.2.1.0/30                2    13        30   int  so-6/0/0.0           IPV4 olympic
                           as0.0                 IPV4 glacier
1eee::/64                   2    13        30   int  so-6/0/0.0           IPV6 olympic
                           as0.0                 IPV6 glacier
abcd::10:255:71:52/128      2    13        10   int  ae0.0                 IPV6 camaro
abcd::10:255:71:238/128     2    13        20   int  so-6/0/0.0           IPV6 olympic
                           as0.0                 IPV6 glacier
abcd::10:255:71:239/128     2    13        20   int  so-6/0/0.0           IPV6 olympic

```

					ae0.0	IPv6 camaro
abcd::10:255:71:242/128	2	13	10	int	as0.0	IPv6 glacier
abcd::10:255:71:243/128	2	13	10	int	so-6/0/0.0	IPv6 olympic

show isis spf

Syntax	show isis spf (brief log results) <instance <i>instance-name</i> > <level (1 2)> <logical-system (all <i>logical-system-name</i>)> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)>
Syntax (EX Series Switches)	show isis spf (brief log results) <instance <i>instance-name</i> > <level (1 2)> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display information about IS-IS shortest-path-first (SPF) calculations.
Options	<p>brief—Display an overview of SPF calculations.</p> <p>instance <i>instance instance-name</i>—(Optional) Display SPF calculations for the specified routing instance.</p> <p>level (1 2)—(Optional) Display SPF calculations for the specified IS-IS level.</p> <p>log—Display the log of SPF calculations.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>results—Display the results of SPF calculations.</p> <p>topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)—(Optional) Display SPF calculations for the specified topology only.</p>
Required Privilege Level	view
List of Sample Output	show isis spf log on page 2778 show isis spf results logical-system on page 2779 show isis spf results (CLNS) on page 2780
Output Fields	Table 228 on page 2777 describes the output fields for the show isis spf command. Output fields are listed in the approximate order in which they appear.

Table 228: show isis spf Output Fields

Field Name	Field Description
Node	System ID of a node.
Metric	Metric to the node.

Table 228: show isis spf Output Fields (*continued*)

Field Name	Field Description
Interface	Interface of the next hop.
Via	System ID of the next hop.
SNPA	Subnetwork point of attachment (MAC address of the next hop).
Start time	(log option only) Time that the SPF computation started.
Elapsed (secs)	(log option only) Length of time, in seconds, required to complete the SPF computation.
Count	(log option only) Number of times the SPF was triggered.
Reason	(log option only) Reason that the SPF computation was completed.

Sample Output

show isis spf log

```

user@host> show isis spf log logical-system lsl
IS-IS level 1 SPF log:
Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000069    1 Reconfig
Fri Oct 31 12:41:18   0.000107    3 Updated LSP fix.00-00
Fri Oct 31 12:41:18   0.000050    3 Address change on so-1/2/2.0
Fri Oct 31 12:41:23   0.000033    1 Updated LSP fix.00-00
Fri Oct 31 12:41:28   0.000178    5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59   0.000060    1 Updated LSP fix.00-00
Fri Oct 31 12:42:30   0.000161    2 Multi area attachment change
Fri Oct 31 12:56:58   0.000198    1 Periodic SPF
Fri Oct 31 13:10:29   0.000209    1 Periodic SPF
IS-IS level 2 SPF log:
Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000035    1 Reconfig
Fri Oct 31 12:41:18   0.000047    2 Updated LSP fix.00-00
Fri Oct 31 12:41:18   0.000043    5 Address change on gr-0/2/0.0
Fri Oct 31 12:41:23   0.000022    1 Updated LSP fix.00-00
Fri Oct 31 12:41:59   0.000144    3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30   0.000257    3 New LSP skag.00-00
Fri Oct 31 12:54:37   0.000195    1 Periodic SPF
Fri Oct 31 12:55:50   0.000178    1 Updated LSP fix.00-00
Fri Oct 31 12:55:55   0.000174    1 Updated LSP h.00-00
Fri Oct 31 12:55:58   0.000176    1 Updated LSP skag.00-00
Fri Oct 31 13:08:14   0.000198    1 Periodic SPF
IPv6 Unicast IS-IS level 1 SPF log:
Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000028    1 Reconfig
Fri Oct 31 12:41:18   0.000043    3 Updated LSP fix.00-00
Fri Oct 31 12:41:18   0.000112    4 Updated LSP fix.00-00
Fri Oct 31 12:41:23   0.000059    1 Updated LSP fix.00-00
Fri Oct 31 12:41:25   0.000041    1 Updated LSP fix.00-00

```



```

Fri Oct 31 12:41:28      0.000103    5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59      0.000040    1 Updated LSP fix.00-00
Fri Oct 31 12:42:30      0.000118    2 Multi area attachment change
Fri Oct 31 12:56:08      0.000289    1 Periodic SPF
Fri Oct 31 13:11:07      0.000214    1 Periodic SPF
IPv6 Unicast IS-IS level 2 SPF log:

```

```

Start time      Elapsed (secs) Count Reason
Fri Oct 31 12:41:18 0.000027     1 Reconfig
Fri Oct 31 12:41:18 0.000039     2 Updated LSP fix.00-00
Fri Oct 31 12:41:18 0.000049     6 Updated LSP fix.00-00
Fri Oct 31 12:41:23 0.000025     1 Updated LSP fix.00-00
Fri Oct 31 12:41:25 0.000023     1 Updated LSP fix.00-00
Fri Oct 31 12:41:59 0.000087     3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30 0.000123     3 New LSP skag.00-00
Fri Oct 31 12:55:50 0.000121     1 Updated LSP fix.00-00
Fri Oct 31 12:55:55 0.000121     1 Updated LSP h.00-00
Fri Oct 31 12:55:58 0.000121     1 Updated LSP skag.00-00
Fri Oct 31 13:09:46 0.000201     1 Periodic SPF
...

```

show isis spf results logical-system

```
user@host> show isis spf results logical-system ls1
```

```
IS-IS level 1 SPF results:
```

Node	Metric	Interface	Via	SNPA
scat.00	10	ge-1/1/0.0	scat	0:90:69:a6:48:9d
	20	10.9.1.0/30		
fix.02	10			
fix.00	0			
	10	10.9.1.0/30		
	10	10.9.5.0/30		
	10	10.9.6.0/30		
	20	10.9.7.0/30		
	60	10.9.201.1/32		

```
3 nodes
```

```
IS-IS level 2 SPF results:
```

Node	Metric	Interface	Via	SNPA
skag.00	20	gr-0/2/0.0	h	
	30	10.9.7.0/30		
skag.02	20	gr-0/2/0.0	h	
h.00	10	gr-0/2/0.0	h	
	20	10.9.6.0/30		
	20	10.9.7.0/30		
	60	10.9.201.1/32		
fix.00	0			
	10	10.9.1.0/30		
	10	10.9.5.0/30		
	10	10.9.6.0/30		

```
4 nodes
```

```
IPv6 Unicast IS-IS level 1 SPF results:
```

Node	Metric	Interface	Via	SNPA
scat.00	10	ge-1/1/0.0	scat	0:90:69:a6:48:9d
		ge-1/1/0.0	scat	0:90:69:a6:48:9d
	20	8009:1::a09:1400/126		
fix.02	10			
fix.00	0			
	10	8009:1::a09:1400/126		
	10	8009:2::a09:1e00/126		

```

                20      8009:3::a09:3200/126
                10      8009:4::a09:2800/126
    3 nodes

IPv6 Unicast IS-IS level 2 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00    20      gr-0/2/0.0    h
           30      8009:3::a09:3200/126
skag.02    20      gr-0/2/0.0    h
           10      gr-0/2/0.0    h
h.00       10      gr-0/2/0.0    h
           20      8009:3::a09:3200/126
           20      8009:4::a09:2800/126
fix.00     0
           10      8009:1::a09:1400/126
           10      8009:2::a09:1e00/126
           10      8009:4::a09:2800/126
    4 nodes

Multicast IS-IS level 1 SPF results:
Node      Metric      Interface      Via      SNPA
scat.00    10      ge-1/1/0.0    scat    0:90:69:a6:48:9d
fix.02     10
fix.00     0
    3 nodes

Multicast IS-IS level 2 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00    20      gr-0/2/0.0    h
skag.02    20      gr-0/2/0.0    h
h.00       10      gr-0/2/0.0    h
fix.00     0
    4 nodes
...

```

show isis spf results (CLNS)

```

user@host> show isis spf results
IS-IS level 1 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00 10      fe-0/0/1.0     toothache 0:12:0:34:0:56
           10      fe-0/0/1.0     toothache 0:12:0:34:0:56
           20      192.168.37.64/29
           10      1921.6800.4001
           20      1921.6800.4002
pro1-a.02 10
pro1-a.00 0
           0      10.255.245.1/32
           10      192.168.37.64/29
           0      1921.6800.4211
    3 nodes

IS-IS level 2 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00 10      fe-0/0/1.0     toothache 0:12:0:34:0:56
           10      fe-0/0/1.0     toothache 0:12:0:34:0:56
           20      10.255.245.1/32
           20      192.168.37.64/29
           20      47.0005.80ff.f800.0000.0109.0010/104

```

pro1-a.02	10	
pro1-a.00	0	
	0	10.255.245.1/32
	10	192.168.37.64/29
3 nodes		

show isis statistics

Syntax	show isis statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis statistics <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display statistics about IS-IS traffic.
Options	none —Display IS-IS traffic statistics for all routing instances. instance <i>instance-name</i> —(Optional) Display statistics for the specified routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear isis statistics on page 2742
List of Sample Output	show isis statistics on page 2784
Output Fields	Table 229 on page 2783 describes the output fields for the show isis statistics command. Output fields are listed in the approximate order in which they appear.

Table 229: show isis statistics Output Fields

Field Name	Field Description
PDU type	<p>PDU type:</p> <ul style="list-style-type: none"> • CSNP—Complete sequence number PDUs contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases. The designated router multicasts CSNPs on broadcast links in place of sending explicit acknowledgments for each link-state PDU. • IIH—IS-IS hello packets are broadcast to discover the identity of neighboring IS-IS systems and to determine whether the neighbors are Level 1 or Level 2 intermediate systems. • LSP—Link-state PDUs contain information about the state of adjacencies to neighboring IS-IS systems. Link-state PDUs are flooded periodically throughout an area. • PSNP—Partial sequence number PDUs are sent multicast by a receiver when it detects that it is missing a link-state PDU (when its link-state PDU database is out of date). The receiver sends a PSNP to the system that transmitted the CSNP, effectively requesting that the missing link-state PDU be transmitted. That routing device, in turn, forwards the missing link-state PDU to the requesting routing device. • Unknown—The PDU type is unknown.
Received	Number of PDUs received since IS-IS started or since the statistics were set to zero.
Processed	Number of PDUs received less the number dropped.
Drops	Number of PDUs dropped.
Sent	Number of PDUs transmitted since IS-IS started or since the statistics were set to zero.
Rexmit	Number of PDUs retransmitted since IS-IS started or since the statistics were set to zero.
Total packets received/sent	Total number of PDUs received and transmitted since IS-IS started or since the statistics were set to zero.
SNP queue length	Number of CSPN and PSNP packets currently waiting in the queue for processing. This value is almost always 0.
LSP queue length	Number of link-state PDUs waiting in the queue for processing. This value is almost always 0.
SPF runs	Number of shortest-path-first (SPF) calculations that have been performed. If this number is incrementing rapidly, it indicates that the network is unstable.
Fragments rebuilt	Number of link-state PDU fragments that the local system has computed.
LSP regenerations	Number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed.
Purges initiated	Number of purges that the system initiated. A purge is initiated if the software decides that a link-state PDU must be removed from the network.

Sample Output

show isis statistics

```
user@host> show isis statistics
IS-IS statistics for merino:

PDU type      Received  Processed  Drops      Sent      Rexmit
LSP           12227    12227     0          8184     683
IIH           113808   113808     0        115817     0
CSNP          198868   198868     0        198934     0
PSNP           6985     6979      6         8274     0
Unknown        0         0         0          0        0
Totals       331888   331882     6        331209   683

Total packets received: 331888 Sent: 331892

SNP queue length:      0 Drops:      0
LSP queue length:      0 Drops:      0

SPF runs:              1014
Fragments rebuilt:     1038
LSP regenerations:     425
Purges initiated:      0
```

OSPF

- [Configuration on page 2784](#)
- [Administration on page 2855](#)

Configuration

- [Configuration Statements on page 2784](#)

Configuration Statements

- [\[edit protocols ospf\] Hierarchy Level on page 2784](#)
- [\[edit protocols ospf3\] Hierarchy Level on page 2788](#)

[edit protocols ospf] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```
protocols {
  ospf {
    disable;
    area area-id {
      ... the area subhierarchy appears after the main [edit protocols ospf] hierarchy ...
    }
    backup-spf-options {
      disable;
      downstream-paths-only;
      no-install;
    }
    database-protection {
```

```

ignore-count number;
ignore-time seconds;
maximum-lsa number;
reset-time seconds;
warning-only;
warning-threshold percent;
}
export [ policy-names ];
external-preference preference;
graceful-restart {
    disable;
    helper-disable <both | restart-signaling | standard>;
    no-strict-lsa-checking;
    notify-duration seconds;
    restart-duration seconds;
}
import [ policy-names ];
lsa-refresh-interval;
no-nssa-abr;
no-rfc-1583;
overload <timeout seconds>;
preference preference;
prefix-export-limit number;
reference-bandwidth reference-bandwidth;
rib-group group-name;
spf-options {
    delay milliseconds;
    holddown milliseconds;
    rapid-runs number;
}
topology (default | ipv4-multicast | name) {
    backup-spf-options {
        disable;
        downstream-paths-only;
        no-install;
    }
    overload;
    prefix-export-limit number;
    spf-options {
        delay milliseconds;
        holddown milliseconds;
        rapid-runs number;
    }
    topology-id number;
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
traffic-engineering {
    advertise-unnumbered-interfaces;
    credibility-protocol-preference;
    ignore-lsp-metrics;
    multicast-rpf-routes;
    no-topology;
}

```

```

    shortcuts <lsp-metric-into-summary>;
  }
}

ospf {
  area area-id {
    area-range ip-prefix </prefix-length> <exact> <override-metric metric> <restrict>;
    context-identifier identifier
    interface interface-name {
      ... the interface subhierarchy appears after the main [edit ospf area area-id] hierarchy
        level ...
    }
    label-switched-path name {
      disable;
      metric metric;
      topology (name | default | ipv4-multicast) {
        disable;
        metric metric;
      }
    }
    network-summary-export [ policy-names ];
    network-summary-import [ policy-names ];
    nssa {
      area-range ip-prefix </prefix-length> <exact> <override-metric metric> <restrict>;
      default-lsa {
        default-metric metric;
        metric-type type;
        type-7;
      }
      (summaries | no-summaries);
    }
    peer-interface interface-name {
      disable;
      authentication {
        md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
        simple-password key-string;
      }
      dead-interval seconds;
      demand-circuit;
      flood-reduction;
      hello-interval seconds;
      no-neighbor-down-notification;
      retransmit-interval seconds;
      transit-delay seconds;
    }
    stub <default-metric metric> <summaries | no-summaries>;
    virtual-link neighbor-id router-id transit-area area-id {
      disable;
      authentication {
        md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
        simple-password key-string;
      }
      dead-interval seconds;
      demand-circuit;
      flood-reduction;
      hello-interval seconds;
    }
  }
}

```



```

ipsec-sa sa-name;
no-neighbor-down-notification;
retransmit-interval seconds;
topology (name | default | ipv4-multicast) {
    disable;
    metric metric;
}
transit-delay seconds;
}
}

area area-id {
    interface interface-name {
        disable;
        authentication {
            md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
            simple-password key-string;
        }
        bandwidth-based-metrics {
            bandwidth value metric number;
        }
        bfd-liveness-detection {
            authentication {
                algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            full-neighbors-only;
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (1 | automatic);
        }
        dead-interval seconds;
        demand-circuit;
        dynamic-neighbors;
        flood-reduction;
        hello-interval seconds;
        interface-type (nbma | p2mp | p2p);
        ipsec-sa sa-name;
        ldp-synchronization {
            disable;
            hold-time seconds;
        }
        (link-protection | node-link-protection);
        metric metric;
        neighbor address <eligible>;
    }
}

```

```
no-eligible-backup;
no-interface-state-traps;
no-neighbor-down-notification;
passive {
    traffic-engineering {
        remote-node-id address;
    }
}
poll-interval seconds;
priority number;
retransmit-interval seconds;
secondary;
te-metric metric;
topology (name | default | ipv4-multicast) {
    disable;
    bandwidth-based-metrics {
        bandwidth value;
        metric number;
    }
    metric metric;
}
transit-delay seconds;
}
}
}
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

[edit protocols ospf3] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```
protocols {
    ospf3 {
        disable;
        area area-id {
            ... the area subhierarchy appears after the main [edit protocols ospf3] hierarchy ...
        }
        backup-spf-options {
            disable;
            downstream-paths-only;
            no-install;
        }
        database-protection {
            ignore-count number;
            ignore-time seconds;
            maximum-lsa number;
            reset-time seconds;
            warning-only;
            warning-threshold percent;
        }
        export [ policy-names ];
    }
}
```

```

external-preference preference;
graceful-restart {
    disable;
    helper-disable;
    no-strict-lsa-checking;
    notify-duration seconds;
    restart-duration seconds;
}
import [ policy-names ];
lsa-refresh-interval;
no-nssa-abr;
no-rfc-1583;
overload <timeout seconds>;
preference preference;
prefix-export-limit number;
realm (ipv4-multicast | ipv4-unicast | ipv6-multicast | ipv6-unicast) {
    ... the realm subhierarchies appear after the main [edit protocols ospf3] hierarchy ...
}
reference-bandwidth reference-bandwidth;
rib-group group-name;
spf-options {
    delay milliseconds;
    holddown milliseconds;
    no-ignore-our-externals;
    rapid-runs number;
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
traffic-engineering {
    ignore-lsp-metrics;
    shortcuts <lsp-metric-into-summary>;
}
}

ospf3 {
    area area-id {
        area-range ip-prefix</prefix-length> <exact> <override-metric metric> <restrict>;
        inter-area-prefix-export [ policy-names ];
        inter-area-prefix-import [ policy-names ];
        interface interface-name {
            ... the interface subhierarchy appears after the main [edit ospf3 area area-id]
                hierarchy level ...
        }
        nssa {
            area-range ip-prefix</prefix-length> <exact> <override-metric metric> <restrict>;
            default-lsa {
                default-metric metric;
                metric-type type;
                type-7;
            }
            (summaries | no-summaries);
        }
        stub <default-metric metric> <summaries | no-summaries>;
    }
}

```

```
virtual-link neighbor-id router-id transit-area area-id {
  disable;
  dead-interval seconds;
  demand-circuit;
  flood-reduction;
  hello-interval seconds;
  ipsec-sa sa-name;
  retransmit-interval seconds;
  transit-delay seconds;
}
}

area area-id {
  interface interface-name {
    disable;
    bandwidth-based-metrics {
      bandwidth value metric number;
    }
    bfd-liveness-detection {
      authentication {
        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
          meticulous-keyed-sha-1 | simple-password);
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
    }
    full-neighbors-only;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    version (1 | automatic);
  }
  dead-interval seconds;
  demand-circuit;
  flood-reduction;
  hello-interval seconds;
  interface-type (p2mp-over-lan | p2p);
  ipsec-sa sa-name;
  (link-protection | node-link-protection);
  metric metric;
  no-eligible-backup;
  own-router-lsa;
  passive {
    traffic-engineering {
      remote-node-id address;
    }
  }
  priority number;
  retransmit-interval seconds;
```

```

        transit-delay seconds;
    }
}

ospf3 {
    realm (ipv4-multicast| ipv6-multicast) {
        ... same statements as at the [edit protocols ospf3] hierarchy level, EXCEPT FOR ...
        area area-id {
            interface interface-name {
                no-eligible-backup; # NOT valid at this level
            }
            virtual-link { ... } # NOT valid at this level
        }
        backup-spf-options { ... } # NOT valid at this level
        realm realm-identifier { ... } # NOT valid at this level
        traffic-engineering { ... } # NOT valid at this level
    }
}

ospf3 {
    realm ipv4-unicast {
        ... same statements as at the [edit protocols ospf3] hierarchy level, PLUS ...
        area area-id {
            interface interface-name {
                ldp-synchronization {
                    disable;
                    hold-time seconds;
                }
            }
        }

        ... BUT NOT ...
        area area-id {
            virtual-link { ... } # NOT valid at this level
        }
        realm realm-identifier { ... } # NOT valid at this level
        traffic-engineering { ... } # NOT valid at this level
    }
}

ospf3 {
    realm ipv6-unicast {
        disable;
        backup-spf-options {
            disable;
            downstream-paths-only;
            no-install;
        }
    }
}

```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

area

Syntax	<pre> area <i>area-id</i> { interface <i>interface-name</i> { passive; topology (ipv4-multicast <i>name</i>) { disable; } } virtual-link neighbor-id <i>router-id</i> transit-area <i>area-id</i> { topology (ipv4-multicast <i>name</i>) { disable; } } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Specify the area identifier for this routing device to use when participating in OSPF routing. All routing devices in an area must use the same area identifier to establish adjacencies.</p> <p>Specify multiple area statements to configure the routing device as an area border router. An area border router does not automatically summarize routes between areas. Use the area-range statement to configure route summarization. By definition, an area border router must be connected to the backbone area either through a physical link or through a virtual link. To create a virtual link, include the virtual-link statement.</p> <p>To specify that the routing device is directly connected to the OSPF backbone, include the area 0.0.0.0 statement.</p> <p>All routing devices on the backbone must be contiguous. If they are not, use the virtual-link statement to create the appearance of connectivity to the backbone.</p>

You can also configure any interface that belongs to one or more topologies to advertise the direct interface addresses without actually running OSPF on that interface. By default, OSPF must be configured on an interface in order for direct interface addresses to be advertised as interior routes.



NOTE: If you configure an interface with the **passive** statement, it applies to all the topologies to which the interface belongs. You cannot configure an interface as passive for only one specific topology and have it remain active for any other topologies to which it belongs.

Options	area-id —Area identifier. The identifier can be up to 32 bits. It is common to specify the area number as a simple integer or an IP address. Area number 0.0.0.0 is reserved for the OSPF backbone area.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>OSPF Areas and Router Functionality Overview</i>• <i>Understanding Multiple Address Families for OSPFv3</i>• virtual-link on page 2855

area-range

Syntax	area-range <i>network/mask-length</i> <exact> <override-metric <i>metric</i> > <restrict>;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit routing-instances <i>routing-instance-name</i> realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>(Area border routers only) For an area, summarize a range of IP addresses when sending summary link advertisements (within an area). To summarize multiple ranges, include multiple area-range statements.</p> <p>For a not-so-stubby area (NSSA), summarize a range of IP addresses when sending NSSA link-state advertisements. The specified prefixes are used to aggregate external routes learned within the area when the routes are advertised to other areas. To specify multiple prefixes, include multiple area-range statements. All external routes learned within the area that do not fall into one of the prefixes are advertised individually to other areas.</p>
Default	By default, area border routing devices do not summarize routes being sent from one area to other areas, but rather send all routes explicitly.
Options	<p>exact—(Optional) Summarization of a route is advertised only when an exact match is made with the configured summary range.</p> <p>mask-length—Number of significant bits in the network mask.</p> <p>network—IP address. You can specify one or more IP addresses.</p>

override-metric *metric*—(Optional) Override the metric for the IP address range and configure a specific metric value.

restrict—(Optional) Do not advertise the configured summary. This hides all routes that are contained within the summary, effectively creating a route filter.

Range: 1 through 16,777,215

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

Related Documentation	<ul style="list-style-type: none">• <i>Example: Summarizing Ranges of Routes in OSPF Link-State Advertisements</i>
------------------------------	--

bandwidth-based-metrics

Syntax	<pre>bandwidth-based-metrics { bandwidth <i>value</i>; metric <i>number</i>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology <i>topology-name</i>], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology <i>topology-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instances</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology <i>topology-name</i>], [edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology <i>topology-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</pre>
Release Information	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 9.5 for EX Series switches.</p>
Description	<p>Specify a set of bandwidth threshold values and associated metric values for an OSPF interface or for a topology on an OSPF interface. When the bandwidth of an interface changes, Junos OS automatically sets the interface metric to the value associated with the appropriate bandwidth threshold value.</p>
Options	<p>bandwidth <i>value</i>—Specify the bandwidth threshold in bits per second.</p> <p>Range: 9600 through 1,000,000,000,000,000</p> <p>metric <i>number</i>—Specify a metric value to associate with a specific bandwidth value.</p> <p>Range: 1 through 65,535</p>



NOTE: You must also configure a static metric value for the OSPF interface or topology with the metric statement. Junos OS uses this value to calculate the cost of a route from the OSPF interface or topology if the bandwidth for the interface is higher than of any bandwidth threshold values configured for bandwidth-based metrics.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth</i>• metric on page 2824• <i>Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth</i>

bfd-liveness-detection (Protocols OSPF)

Syntax `bfd-liveness-detection {`
 `authentication {`
 `algorithm algorithm-name;`
 `key-chain key-chain-name;`
 `loose-check;`
 `}`
 `detection-time {`
 `threshold milliseconds;`
 `}`
 `full-neighbors-only`
 `minimum-interval milliseconds;`
 `minimum-receive-interval milliseconds;`
 `multiplier number;`
 `no-adaptation;`
 `transmit-interval {`
 `minimum-interval milliseconds;`
 `threshold milliseconds;`
 `}`
 `version (1 | automatic);`
 `}`

Hierarchy Level `[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id interface interface-name],`
 `[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast |`
 `ipv4-multicast | ipv6-multicast) area area-id interface interface-name],`
 `[edit logical-systems logical-system-name routing-instances routing-instance-name protocols`
 `(ospf | ospf3) area area-id interface interface-name],`
 `[edit logical-systems logical-system-name routing-instances routing-instance-name protocols`
 `ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id interface`
 `interface-name],`
 `[edit protocols (ospf | ospf3) area area-id interface interface-name],`
 `[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id`
 `interface interface-name],`
 `[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id interface`
 `interface-name],`
 `[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast |`
 `ipv4-multicast | ipv6-multicast) area area-id interface interface-name]`

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 detection-time threshold and **transmit-interval threshold** options added in Junos OS Release 8.2.
 Support for logical systems introduced in Junos OS Release 8.3.
 no-adaptation option introduced in Junos OS Release 9.0.
 no-adaptation option introduced in Junos OS Release 9.0 for EX Series switches.
 Support for OSPFv3 introduced in Junos OS Release 9.3.
 Support for OSPFv3 introduced in Junos OS Release 9.3 for EX Series switches.
 full-neighbors-only option introduced in Junos OS Release 9.5.
 full-neighbors-only option introduced in Junos OS Release 9.5 for EX Series switches.

authentication algorithm, **authentication key-chain**, and **authentication loose-check** options introduced in Junos OS Release 9.6.

Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure bidirectional failure detection timers and authentication for OSPF.

The remaining statements are explained separately.

Options **authentication algorithm** *algorithm-name*—Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.

authentication key-chain *key-chain-name*—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

detection-time threshold *milliseconds*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

full-neighbors-only—Establish BFD sessions only for OSPF neighbors in the full state. The default behavior is to establish BFD sessions for all OSPF neighbors.

minimum-interval *milliseconds*—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements.

Range: 1 through 255,000 milliseconds

minimum-receive-interval *milliseconds*—Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

Range: 1 through 255,000 milliseconds

multiplier *number*—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

no-adaptation—Specify that BFD sessions should not adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds*—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

transmit-interval minimum-interval *milliseconds*—Configure the minimum interval at which the routing device transmits hello packets to a neighbor with which it has established

a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

Range: 1 through 255,000

version—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

Default: **automatic**

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

Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring BFD for OSPF</i>• <i>Example: Configuring BFD Authentication for OSPF</i>
------------------------------	--

dead-interval

Syntax	<code>dead-interval seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	Specify how long OSPF waits before declaring that a neighboring routing device is unavailable. This is an interval during which the routing device receives no hello packets from the neighbor.
Options	<p>seconds—Interval to wait.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: Four times the hello interval—40 seconds (broadcast and point-to-point networks); 120 seconds (nonbroadcast multiple access (NBMA) networks)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring OSPF Timers</i> • <i>Configuring RSVP and OSPF for LMP Peer Interfaces</i>

- [hello-interval on page 2814](#)

default-lsa

Syntax	<pre>default-lsa { default-metric metric; metric-type type; type-7; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>On area border routers only, for a not-so-stubby area (NSSA), inject a default link-state advertisement (LSA) with a specified metric value into the area. The default route matches any destination that is not explicitly reachable from within the area.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • OSPF Areas and Router Functionality Overview • Example: Configuring OSPF Not-So-Stubby Areas • nssa on page 2829 • stub on page 2846

default-metric

Syntax	<code>default-metric <i>metric</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area area-id nssa default-lsa],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area area-id stub],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa default-lsa],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> stub],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area area-id nssa default-lsa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area area-id stub],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa default-lsa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> stub],</p> <p>[edit protocols (ospf ospf3) area area-id nssa default-lsa],</p> <p>[edit protocols (ospf ospf3) area area-id stub],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa default-lsa],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> stub],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area area-id nssa default-lsa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area area-id stub],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa default-lsa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> stub]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	On area border routing devices only, for a stub area, inject a default route with a specified metric value into the area. The default route matches any destination that is not explicitly reachable from within the area.
Options	<p><i>metric</i>—Metric value.</p> <p>Range: 1 through 16,777,215</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • OSPF Areas and Router Functionality Overview • nssa on page 2829

- [stub on page 2846](#)

disable (OSPF)

Syntax	disable;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instances</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) virtual-link],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> virtual-link neighbor-id <i>router-id</i> transit-area <i>area-id</i>],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Disable OSPF, an OSPF interface, or an OSPF virtual link.</p> <p>By default, control packets sent to the remote end of a virtual link must be forwarded using the default topology. In addition, the transit area path consists only of links that</p>

are in the default topology. You can disable a virtual link for a configured topology, but not for a default topology. Include the **disable** statement at the **[edit protocols ospf area *area-id* virtual-link neighbor-id router-id transit-area *area-id* topology *name*]** hierarchy level.



NOTE: If you disable the virtual link by including the **disable** statement at the **[edit protocols ospf area *area-id* virtual-link neighbor-id router-id transit-area *area-id*]** hierarchy level, you disable the virtual link for all topologies, including the default topology. You cannot disable the virtual link only in the default topology.

Default	The configured object is enabled (operational) unless explicitly disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>OSPF Configuration Overview</i> • <i>Configuring RSVP and OSPF for LMP Peer Interfaces</i>

domain-id

Syntax	<code>domain-id <i>domain-id</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<i>ospf</i> <i>ospf3</i>)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols (<i>ospf</i> <i>ospf3</i>)]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify a domain ID for a route. The domain ID identifies the OSPF domain from which the route originated.
Options	<p><i>domain-id</i>—You can specify either an IP address or an IP address and a local identifier using the following format: <i>ip-address:local-identifier</i>. If you do not specify a local identifier with the IP address, the identifier is assumed to have a value of 0.</p> <p>Default: If the router ID is not configured in the routing instance, the router ID is derived from an interface address belonging to the routing instance.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Routing Between PE and CE Routers in Layer 3 VPNs</i>

domain-vpn-tag

Syntax	<code>domain-vpn-tag <i>number</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Set a virtual private network (VPN) tag for OSPFv2 external routes generated by the provider edge (PE) routing device.
Options	<i>number</i> —VPN tag.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Routing Between PE and CE Routers in Layer 3 VPNs</i>

export (Protocols OSPF)

Syntax	<code>export [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast </code> <code> ipv4-multicast ipv6-multicast)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> (ospf ospf3)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</code> <code>[edit protocols (ospf ospf3)],</code> <code>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast </code> <code> ipv4-multicast ipv6-multicast)]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for the realm statement introduced in Junos OS Release 9.2. Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Apply one or more policies to routes being exported from the routing table into OSPF.
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding OSPF Routing Policy on page 3320• <i>Import and Export Policies for Network Summaries Overview</i>• import on page 2816• <i>Routing Policy Configuration Guide</i>

external-preference (Protocols OSPF)

Syntax	<code>external-preference <i>preference</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Set the route preference for OSPF external routes.
Options	<p><i>preference</i>—Preference value.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 150</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Controlling OSPF Route Preferences • preference on page 2835

graceful-restart (Protocols OSPF)

Syntax	<pre> graceful-restart { disable; helper-disable (standard restart-signaling both); no-strict-lsa-checking; notify-duration <i>seconds</i>; restart-duration <i>seconds</i>; }</pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for the no-strict-lsa-checking statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the helper mode standard, restart-signaling, and both options introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	<p>Configure graceful restart for OSPF.</p> <p>Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled for all routing protocols at the [edit routing-options] hierarchy level.</p>
Options	<p>disable—Disable graceful restart for OSPF.</p> <p>helper-disable (standard restart-signaling both)—Disable helper mode for graceful restart. When helper mode is disabled, a device cannot help a neighboring device that is attempting to restart. Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. The standard, restart-signaling, and both options are only supported for OSPFv2. Specify standard to disable helper mode for standard graceful restart (based on RFC 3623). Specify restart-signaling to disable helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and RFC 4813). Specify both to disable helper mode for both standard and restart signaling-based graceful restart. The last committed statement takes precedence over the previously configured statement.</p> <p>Default: Helper mode is enabled by default. For OSPFv2, both standard and restart-signaling based helper modes are enabled by default.</p> <p>no-strict-lsa-checking—Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router. LSA checking is enabled by default.</p>



NOTE: The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both

statements at the same time, the routing device displays a warning message when you enter the `show protocols (ospf | ospf3)` command.

.....
notify-duration seconds—Estimated time needed to send out purged grace LSAs over all the interfaces.

Range: 1 through 3600 seconds

Default: 30 seconds

restart-duration seconds—Estimated time needed to reacquire a full OSPF neighbor from each area.

Range: 1 through 3600 seconds

Default: 180 seconds

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

Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Graceful Restart for OSPF</i> • <i>Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart</i> • <i>Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart</i> • <i>Example: Disabling Strict LSA Checking for OSPF Graceful Restart</i> • <i>Configuring Graceful Restart for QFabric Systems</i> • <i>Junos OS High Availability Configuration Guide</i>
------------------------------	--

hello-interval (Protocols OSPF)

Syntax	<code>hello-interval seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	Specify how often the routing device sends hello packets out the interface. The hello interval must be the same for all routing devices on a shared logical IP network.
Options	<p>seconds—Time between hello packets, in seconds.</p> <p>Range: 1 through 255 seconds</p> <p>Default: 10 seconds (broadcast and point-to-point networks); 30 seconds (nonbroadcast multiple access [NBMA] networks)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring OSPF Timers</i> • <i>Configuring RSVP and OSPF for LMP Peer Interfaces</i> • dead-interval on page 2803

ignore-lsp-metrics

Syntax	ignore-lsp-metrics;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf traffic-engineering shortcuts], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf traffic-engineering shortcuts], [edit protocols ospf traffic-engineering], [edit routing-instances <i>routing-instance-name</i> protocols ospf traffic-engineering shortcuts]
Release Information	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for (OSPFv3) introduced in Junos OS Release 9.4. Support for (OSPFv3) introduced in Junos OS Release 9.4 for EX Series switches.
Description	Ignore RSVP LSP metrics in OSPF traffic engineering shortcut calculations.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Enabling OSPF Traffic Engineering Support</i>

import (Protocols OSPF)

Syntax	<code>import [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast </code> <code> ipv4-multicast ipv6-multicast)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> (ospf ospf3)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</code> <code>[edit protocols (ospf ospf3)],</code> <code>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast </code> <code> ipv4-multicast ipv6-multicast)]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for the realm statement introduced in Junos OS Release 9.2. Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Filter OSPF routes from being added to the routing table.
Options	<i>policy-names</i> —Name of one or more policies.
Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding OSPF Routing Policy on page 3320• <i>Import and Export Policies for Network Summaries Overview</i>• export on page 2810• <i>Routing Policy Configuration Guide</i>

inter-area-prefix-export

Syntax	<code>inter-area-prefix-export [<i>policy-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ip4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>],</p> <p>[edit protocols ospf3 area <i>area-id</i>],</p> <p>[edit protocols ospf3 realm (ip4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 area <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ip4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Statement introduced in Junos OS Release 9.1 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	Apply an export policy for OSPFv3 to specify which interarea prefix link-state advertisements (LSAs) are flooded into an area.
Options	<i>policy-name</i> —Name of a policy configured at the [edit policy-options policy-statement <i>policy-name</i> term <i>term-name</i>] hierarchy level.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Import and Export Policies for Network Summaries Overview</i> • inter-area-prefix-import on page 2818 • <i>Routing Policy Configuration Guide</i>

inter-area-prefix-import

Syntax	<code>inter-area-prefix-import [<i>policy-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <i>area area-id</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ospf3 <i>realm</i> (ipv4-unicast </code> <code> ipv4-multicast ipv6-multicast) <i>area area-id</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> ospf3 <i>area area-id</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> ospf3 <i>realm</i> (ipv4-unicast ipv4-multicast ipv6-multicast) <i>area area-id</i>],</code> <code>[edit protocols ospf3 <i>area area-id</i>],</code> <code>[edit protocols ospf3 <i>realm</i> (ip4-unicast ipv4-multicast ipv6-multicast)], <i>area area-id</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <i>area area-id</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 <i>realm</i> (ipv4-unicast </code> <code> ipv4-multicast ipv6-multicast) <i>area area-id</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 9.1 for EX Series switches. Support for the realm statement introduced in Junos OS Release 9.2. Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.
Description	Apply an import policy for OSPFv3 to specify which routes learned from an area are used to generate interarea prefixes into other areas.
Options	<i>policy-name</i> —Name of a policy configured at the <code>[edit policy-options policy-statement <i>policy-name</i> term <i>term-name</i>]</code> hierarchy level.
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Import and Export Policies for Network Summaries Overview</i>• inter-area-prefix-export on page 2817• <i>Routing Policy Configuration Guide</i>

interface (Protocols OSPF)

Syntax interface *interface-name* {
 disable;
 authentication *key* <*key-id identifier*>;
 bfd-liveness-detection {
 authentication {
 algorithm *algorithm-name*;
 key-chain *key-chain-name*;
 loose-check;
 }
 detection-time {
 threshold *milliseconds*;
 }
 minimum-interval *milliseconds*;
 minimum-receive-interval *milliseconds*;
 transmit-interval {
 threshold *milliseconds*;
 minimum-interval *milliseconds*;
 }
 multiplier *number*;
 }
 dead-interval *seconds*;
 demand-circuit;
 hello-interval *seconds*;
 ipsec-sa *name*;
 interface-type *type*;
 ldp-synchronization {
 disable;
 hold-time *seconds*;
 }
 metric *metric*;
 neighbor *address* <*eligible*>;
 no-interface-state-traps;
 passive;
 poll-interval *seconds*;
 priority *number*;
 retransmit-interval *seconds*;
 te-metric *metric*;
 topology (ipv4-multicast | *name*) {
 metric *metric*;
 }
 transit-delay *seconds*;
 transmit-interval *seconds*;
}

Hierarchy Level [edit logical-systems *logical-system-name* protocols (ospf | ospf3) *area area-id*],
 [edit logical-systems *logical-system-name* protocols ospf3 *realm* (ipv4-unicast |
 ipv4-multicast | ipv6-multicast) *area area-id*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 (ospf | ospf3) *area area-id*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols
 ospf3 *realm* (ipv4-unicast | ipv4-multicast | ipv6-multicast) *area area-id*],
 [edit protocols (ospf | ospf3) *area area-id*],

```
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id],  
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id],  
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast |  
  ipv4-multicast | ipv6-multicast) area area-id]
```

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for the **topology** statement introduced in Junos OS Release 9.0.
Support for the **topology** statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for the **realm** statement introduced in Junos OS Release 9.2.
Support for the **realm** statement introduced in Junos OS Release 9.2 for EX Series switches.
Support for the **no-interface-state-traps** statement introduced in Junos OS Release 10.3.
This statement is supported only for OSPFv2.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Enable OSPF routing on a routing device interface.

You must include at least one **interface** statement in the configuration to enable OSPF on the routing device.

Options **interface-name**—Name of the interface. Specify the interface by IP address or interface name for OSPFv2, or only the interface name for OSPFv3. Using both the interface name and IP address of the same interface produces an invalid configuration. To configure all interfaces, you can specify **all**. Specifying a particular interface and **all** produces an invalid configuration.



.....
NOTE: For nonbroadcast interfaces, specify the IP address of the nonbroadcast interface as *interface-name*.
.....

The remaining statements are explained separately.



.....
NOTE: You cannot run both OSPF and ethernet-tcc encapsulation between two Juniper Networks routing devices.
.....

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

- Related Documentation**
- *OSPF Configuration Overview*
 - *Example: Configuring Multitopology Routing Based on Applications*
 - *Example: Configuring Multitopology Routing Based on a Multicast Source*
 - *Example: Configuring Multiple Address Families for OSPFv3*
 - *neighbor*

interface-type (Protocols OSPF)

Syntax	<code>interface-type (nbma p2mp p2p);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-multicast ipv4-unicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-multicast ipv4-unicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf3 realm (ipv4-multicast ipv4-unicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-multicast ipv4-unicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for OSPFv3 for interface type p2p only introduced in Junos OS Release 9.4. You cannot configure other interface types for OSPFv3.</p> <p>Support for OSPFv3 for interface type p2p only introduced in Junos OS Release 9.4 for EX Series switches.</p>
Description	<p>Specify the type of interface.</p> <p>By default, the software chooses the correct interface type based on the type of physical interface. Therefore, you should never have to set the interface type. The exception to this is for NBMA interfaces, which default to an interface type of point-to-multipoint. To have these interfaces explicitly run in Nonbroadcast multiaccess (NBMA) mode, configure the nbma interface type, using the IP address of the local ATM interface.</p> <p>In Junos OS Release 9.3 and later, a point-to-point interface can be an Ethernet interface without a subnet.</p>
Default	The software chooses the correct interface type based on the type of physical interface.
Options	<p>nbma (OSPFv2 only)—Nonbroadcast multiaccess (NBMA) interface.</p> <p>p2mp (OSPFv2 only)—Point-to-multipoint interface.</p> <p>p2p—Point-to-point interface.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- *About OSPF Interfaces*
 - *Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network*

lsp-metric-into-summary

Syntax	<code>lsp-metric-into-summary;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) traffic-engineering shortcuts],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) traffic-engineering shortcuts],</p> <p>[edit protocols (ospf ospf3) traffic-engineering shortcuts],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) traffic-engineering shortcuts]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4.</p> <p>Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4 for EX Series switches.</p>
Description	Advertise the LSP metric in summary LSAs.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>OSPF Support for Traffic Engineering</i> • <i>Example: Enabling OSPF Traffic Engineering Support</i>

metric (Protocols OSPF Interface)

Syntax	<code>metric <i>metric</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (ipv4-multicast <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> sham-link-remote],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (ipv4-multicast <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (ipv4-multicast <i>name</i>)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> sham-link-remote],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i> topology (ipv4-multicast <i>name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Specify the cost of an OSPF interface. The cost is a routing metric that is used in the link-state calculation.</p> <p>To set the cost of routes exported into OSPF, configure the appropriate routing policy.</p>
Options	<p>metric—Cost of the route.</p> <p>Range: 1 through 65,535</p> <p>Default: By default, the cost of an OSPF route is calculated by dividing the reference-bandwidth value by the bandwidth of the physical interface. Any specific value you configure for the metric overrides the default behavior of using the reference-bandwidth value to calculate the cost of the route for that interface.</p>

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Controlling the Cost of Individual OSPF Network Segments</i>• <i>Example: Configuring OSPFv2 Sham Links</i>• <i>Example: Configuring Multitopology Routing Based on Applications</i>• <i>Example: Configuring Multitopology Routing Based on a Multicast Source</i>• bandwidth-based-metrics on page 2797• reference-bandwidth on page 2839

metric-type

Syntax	<code>metric-type type;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)) area <i>area-id</i> nssadefault-lsa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)) area <i>area-id</i> nssa default-lsa],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)) area <i>area-id</i> nssa default-lsa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa],</p> <p>[edit routing-instances <i>routing-instances</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)) area <i>area-id</i> nssa default-lsa]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Specify the external metric type for the default LSA.</p> <p>The configured metric determines the method used to compute the cost to a destination:</p> <ul style="list-style-type: none"> • The Type 1 external metric is equivalent to the link-state metric. The path cost uses the advertised external path cost and the path cost to the AS boundary router (the route is equal to the sum of all internal costs and the external cost). • The Type 2 external metric uses the cost assigned by the AS boundary router (the route is equal to the external cost alone). By default, OSPF uses the Type 2 external metric.
Options	type —Metric type: 1 or 2
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>OSPF Areas and Router Functionality Overview</i> • <i>Example: Configuring OSPF Not-So-Stubby Areas</i>

no-nssa-abr

Syntax	no-nssa-abr;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</p>
Release Information	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Disable exporting Type 7 link-state advertisements into not-so-stubby-areas (NSSAs) for an autonomous system boundary router (ASBR) or an area border router (ABR).
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Example: Configuring OSPF Not-So-Stubby Areas</i>

no-rfc-1583

Syntax	no-rfc-1583;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit protocols (ospf ospf3)], [edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for the realm statement introduced in Junos OS Release 9.2. Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Disable compatibility with RFC 1583, <i>OSPF Version 2</i> . If the same external destination is advertised by AS boundary routers that belong to different OSPF areas, disabling compatibility with RFC 1583 can prevent routing loops.
Default	Compatibility with RFC 1583 is enabled by default.
Required Privilege Level	routing—To view this statement in the configuration. routing-control-level—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Disabling OSPFv2 Compatibility with RFC 1583</i>

nssa

Syntax	<pre>nssa { area-range network/mask-length <restrict> <exact> <override-metric metric>; default-lsa { default-metric metric; metric-type type; type-7; } (no-summaries summaries); }</pre>
Hierarchy Level	<pre>[edit logical-systems logical-system-name protocols (ospf ospf3) area area-id], [edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit logical-systems logical-system-name routing-instances routing-instance-name protocols (ospf ospf3) area area-id], [edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit protocols (ospf ospf3) area area-id], [edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit routing-instances routing-instance-name protocols (ospf ospf3) area area-id], [edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Configure a not-so-stubby area (NSSA). An NSSA allows external routes to be flooded within the area. These routes are then leaked into other areas.</p> <p>You cannot configure an area as being both a stub area and an NSSA.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>OSPF Areas and Router Functionality Overview</i> • <i>Example: Configuring OSPF Not-So-Stubby Areas</i> • stub on page 2846


ospf

Syntax	ospf { ... }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Enable OSPF routing on the routing device. You must include the ospf statement to enable OSPF on the routing device.
Default	OSPF is disabled on the routing device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>OSPF Configuration Overview</i>• [edit protocols ospf] Hierarchy Level on page 391

ospf3

Syntax	ospf3 { ... }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Enable OSPFv3 routing on the routing device. You must include the ospf3 statement to enable OSPFv3.
Default	OSPFv3 is disabled on the routing device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>OSPF Configuration Overview</i> • [edit protocols ospf3] Hierarchy Level on page 395

overload (Protocols OSPF)

Syntax	<pre>overload { timeout <i>seconds</i>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols ospf topology (default ipv4-multicast <i>name</i>)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure the local routing device so that it appears to be overloaded. You might do this when you want the routing device to participate in OSPF routing, but do not want it to be used for transit traffic.</p>
	<div>  <p>NOTE: Traffic destined to directly attached interfaces continues to reach the routing device.</p> </div>
Options	<p>timeout <i>seconds</i>—(Optional) Number of seconds at which the overloading is reset. If no timeout interval is specified, the routing device remains in overload state until the overload statement is deleted or a timeout is set.</p> <p>Range: 60 through 1800 seconds</p> <p>Default: 0 seconds</p>



NOTE: Multitopology Routing does not support the timeout option.

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

Related Documentation

- *Example: Configuring OSPF to Make Routing Devices Appear Overloaded*
- *Example: Configuring Multitopology Routing Based on Applications*
- *Example: Configuring Multitopology Routing Based on a Multicast Source*

passive (Protocols OSPF)

Syntax	<pre> passive { traffic-engineering { remote-node-id address; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>traffic-engineering and remote-node-id address statements introduced in Junos OS Release 8.0.</p> <p>traffic-engineering and remote-node-id address statements introduced in Junos OS Release 8.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Advertise the direct interface addresses on an interface without actually running OSPF on that interface. A passive interface is one for which the address information is advertised as an internal route in OSPF, but on which the protocol does not run.</p> <p>To configure an interface in OSPF passive traffic engineering mode, include the traffic-engineering statement. Configuring OSPF passive traffic engineering mode enables the dynamic discovery of OSPF AS boundary routers.</p> <p>Enable OSPF on an interface by including the interface statement at the [edit protocols (ospf ospf3) area <i>area-id</i>] or the [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i>] hierarchy levels. Disable it by including the disable statement. To prevent OSPF from running on an interface, include the passive statement. These three states are mutually exclusive.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- *Example: Configuring a Passive OSPF Interface*
 - *Example: Configuring OSPF Passive Traffic Engineering Mode*
 - [disable on page 2807](#)

preference (Protocols OSPF)

Syntax	<code>preference <i>preference</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Set the route preference for OSPF internal routes.
Options	<p><i>preference</i>—Preference value.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: 10</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Controlling OSPF Route Preferences</i> • external-preference on page 2811

prefix-export-limit (Protocols OSPF)

Syntax	<code>prefix-export-limit <i>number</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols ospf topology (default ipv4-multicast <i>name</i>)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Configure a limit to the number of prefixes exported into OSPF.
Options	<p><i>number</i>—Prefix limit.</p> <p>Range: 0 through 4,294,967,295 ($2^{32} - 1$)</p> <p>Default: None</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Limiting the Number of Prefixes Exported to OSPF</i> • <i>Example: Configuring Multitopology Routing Based on Applications</i> • <i>Example: Configuring Multitopology Routing Based on a Multicast Source</i>


priority (Protocols OSPF)

Syntax	<code>priority number;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)] area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)] area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)] area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)] area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	Specify the routing device's priority for becoming the designated routing device. The routing device that has the highest priority value on the logical IP network or subnet becomes the network's designated router. You must configure at least one routing device on each logical IP network or subnet to be the designated router. You also should specify a routing device's priority for becoming the designated router on point-to-point interfaces.
Options	<p>number—Routing device's priority for becoming the designated router. A priority value of 0 means that the routing device never becomes the designated router. A value of 1 means that the routing device has the least chance of becoming a designated router.</p> <p>Range: 0 through 255</p> <p>Default: 128</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>OSPF Designated Router Overview</i> • <i>Example: Controlling OSPF Designated Router Election</i>

realm

Syntax	<pre>realm (ipv4-unicast ipv4-multicast ipv6-unicast) { area area-id { interface interface-name; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf3], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3], [edit protocols ospf3], [edit routing-instances <i>routing-instance-name</i> protocols ospf3]
Release Information	Statement introduced in Junos OS Release 9.2. Statement introduced in Junos OS Release 9.2 for EX Series switches.
Description	Configure OSPFv3 to advertise address families other than unicast IPv6. Junos OS maps each address family you configure to a separate realm with its own set of neighbors and link-state database.
Options	ipv4-unicast —Configure a realm for IPv4 unicast routes. ipv4-multicast —Configure a realm for IPv4 multicast routes. ipv6-multicast —Configure a realm for IPv6 multicast routes. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Multiple Address Families for OSPFv3</i>

reference-bandwidth (Protocols OSPF)

Syntax	<code>reference-bandwidth <i>reference-bandwidth</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Set the reference bandwidth used in calculating the default interface cost. The cost is calculated using the following formula:</p> $\text{cost} = \text{ref-bandwidth} / \text{bandwidth}$
Options	<p><i>reference-bandwidth</i>—Reference bandwidth, in bits per second.</p> <p>Range: 9600 through 1,000,000,000,000 bits</p> <p>Default: 100 Mbps (100,000,000 bits)</p>
<div>  <p>NOTE: The default behavior is to use the reference-bandwidth value to calculate the cost of OSPF interfaces. You can override this behavior for any OSPF interface by configuring a specific cost with the metric statement.</p> </div>	
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Controlling the Cost of Individual OSPF Network Segments • metric on page 2824

retransmit-interval (OSPF)

Syntax	<code>retransmit-interval seconds;</code>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>], [edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols (ospf ospf3) area <i>area-id</i> virtual-link], [edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	Specify how long the routing device waits to receive a link-state acknowledgment packet before retransmitting link-state advertisements (LSAs) to an interface's neighbors.
Options	<p>seconds—Interval to wait.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: 5 seconds</p>



NOTE: You must configure LSA retransmit intervals to be equal to or greater than 3 seconds to avoid triggering a retransmit trap, because Junos OS delays LSA acknowledgments by up to 2 seconds.

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

Related Documentation

- *Example: Configuring OSPF Timers*
- *Configuring RSVP and OSPF for LMP Peer Interfaces*

rib-group (Protocols OSPF)

Syntax	<code>rib-group group-name;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Install routes learned from OSPF routing instances into routing tables in the OSPF routing table group.
Options	group-name —Name of the routing table group.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i> • <i>Example: Importing Direct and Static Routes Into a Routing Instance</i> • <i>Understanding Multiprotocol BGP</i> • <i>interface-routes</i> • <i>rib-group</i>

route-type-community

Syntax	<code>route-type-community (iana vendor);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Specify an extended community value to encode the OSPF route type. Each extended community is coded as an eight-octet value. This statement sets the most significant bit to either an IANA or vendor-specific route type.
Options	iana —Encode a route type with the value 0x0306 . This is the default value. vendor —Encode the route type with the value 0x8000 .
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Routing Between PE and CE Routers in Layer 3 VPNs</i>

shortcuts (Protocols OSPF)

Syntax	shortcuts { lsp-metric-into-summary; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) traffic-engineering], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) traffic-engineering], [edit protocols (ospf ospf3) traffic-engineering], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) traffic-engineering]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4. Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4 for EX Series switches.
Description	Configure OSPF to use MPLS label-switched paths (LSPs) as shortcut next hops. By default, shortcut routes calculated through OSPFv2 are installed in the inet.3 routing table, and shortcut routes calculated through OSPFv3 are installed in the inet6.3 routing table.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> <i>Example: Enabling OSPF Traffic Engineering Support</i>

spf-options (Protocols OSPF)

Syntax	<pre>spf-options { delay <i>milliseconds</i>; holddown <i>milliseconds</i>; rapid-runs <i>number</i>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit protocols (ospf ospf3)], [edit protocols ospf topology (default ipv4-multicast <i>name</i>)], [edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf topology (default ipv4-multicast <i>name</i>)], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</pre>
Release Information	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure options for running the shortest-path-first (SPF) algorithm. You can configure the following:</p> <ul style="list-style-type: none">• A delay for when to run the SPF algorithm after a network topology change is detected.• The maximum number of times the SPF algorithm can run in succession.• A hold-down interval after the SPF algorithm runs the maximum number of times. <p>Running the SPF algorithm is usually the beginning of a series of larger system-wide events. For example, the SPF algorithm can lead to interior gateway protocol (IGP) prefix changes, which then lead to BGP nexthop resolution changes. Consider what happens if there are rapid link changes in the network. The local routing device can become overwhelmed. This is why it sometimes makes sense to throttle the scheduling of the SPF algorithm.</p>

Options	delay <i>milliseconds</i> —Time interval between the detection of a topology change and when the SPF algorithm runs. Range: 50 through 8000 milliseconds Default: 200 milliseconds
	holddown <i>milliseconds</i> —Time interval to hold down, or to wait before a subsequent SPF algorithm runs after the SPF algorithm has run the configured maximum number of times in succession. Range: 2000 through 20,000 milliseconds Default: 5000 milliseconds
	rapid-runs <i>number</i> —Maximum number of times the SPF algorithm can run in succession. After the maximum is reached, the hold down interval begins. Range: 1 through 10 Default: 3
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring SPF Algorithm Options for OSPF</i>• <i>Example: Configuring Multitopology Routing Based on Applications</i>• <i>Example: Configuring Multitopology Routing Based on a Multicast Source</i>

stub

Syntax	<code>stub <default-metric metric> <(no-summaries summaries)>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area area-id],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area area-id],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols (ospf ospf3) area area-id],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area area-id],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Specify that this area not be flooded with AS external link-state advertisements (LSAs). You must include the stub statement when configuring all routing devices that are in the stub area.</p> <p>The backbone cannot be configured as a stub area.</p> <p>You cannot configure an area to be both a stub area and a not-so-stubby area (NSSA).</p>
Options	<p>no-summaries—(Optional) Do not advertise routes into the stub area. If you include the default-metric option, only the default route is advertised.</p> <p>summaries—(Optional) Flood summary LSAs into the stub area.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>OSPF Areas and Router Functionality Overview</i> • <i>Example: Configuring OSPF Stub and Totally Stubby Areas</i> • nssa on page 2829

summaries

Syntax	(summaries no-summaries);
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area area-id nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area area-id nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa],</p> <p>[edit protocols (ospf ospf3) area area-id nssa],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)] area <i>area-id</i> nssa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area area-id nssa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Configure whether or not area border routers advertise summary routes into an not-so-stubby area (NSSA):</p> <ul style="list-style-type: none"> • summaries—Flood summary link-state advertisements (LSAs) into the NSSA. • no-summaries—Prevent area border routers from advertising summaries into an NSSA. If default-metric is configured for an NSSA, a Type 3 LSA is injected into the area by default.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>OSPF Areas and Router Functionality Overview</i> • <i>Example: Configuring OSPF Not-So-Stubby Areas</i> • nssa on page 2829 • stub on page 2846

traceoptions (Protocols OSPF)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols (<i>ospf</i> <i>ospf3</i>)], [edit logical-systems <i>logical-system-name</i> protocols ospf3 <i>realm</i> (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<i>ospf</i> <i>ospf3</i>)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 <i>realm</i> (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit protocols (<i>ospf</i> <i>ospf3</i>)], [edit protocols ospf3 <i>realm</i> (ipv4-unicast ipv4-multicast ipv6-multicast)], [edit routing-instances <i>routing-instance-name</i> protocols (<i>ospf</i> <i>ospf3</i>)], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 <i>realm</i> (ipv4-unicast ipv4-multicast ipv6-multicast)]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Configure OSPF protocol-level tracing options.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>



NOTE: The **traceoptions** statement is not supported on QFabric systems.

Default	The default OSPF protocol-level tracing options are those inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place OSPF tracing output in the file ospf-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p>

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

OSPF Tracing Flags

- **database-description**—Database description packets, which are used in synchronizing the OSPF and OSPFv3 topological database.
- **error**—OSPF and OSPFv3 error packets.
- **event**—OSPF and OSPFv3 state transitions.
- **flooding**—Link-state flooding packets.
- **graceful-restart**—Graceful-restart events.
- **hello**—Hello packets, which are used to establish neighbor adjacencies and to determine whether neighbors are reachable.
- **ldp-synchronization**—Synchronization events between OSPF and LDP.
- **lsa-ack**—Link-state acknowledgment packets, which are used in synchronizing the OSPF topological database.
- **lsa-analysis**—Link-state analysis. Specific to the Juniper Networks implementation of OSPF, Junos OS performs LSA analysis before running the shortest-path-first (SPF) algorithm. LSA analysis helps to speed the calculations performed by the SPF algorithm.
- **lsa-request**—Link-state request packets, which are used in synchronizing the OSPF topological database.
- **lsa-update**—Link-state updates packets, which are used in synchronizing the OSPF topological database.
- **nsr-synchronization**—Nonstop routing synchronization events.
- **on-demand**—Trace demand circuit extensions.
- **packet-dump**—Content of selected packet types.
- **packets**—All OSPF packets.
- **restart-signaling**—(OSPFv2 only) Restart-signaling graceful restart events.
- **spf**—Shortest-path-first (SPF) calculations.

Global Tracing Flags

- **all**—All tracing operations.
- **general**—A combination of the **normal** and **route** trace operations.
- **normal**—All normal operations. If you do not specify this option, only unusual or abnormal operations are traced.
- **policy**—Policy operations and actions.
- **route**—Routing table changes.
- **state**—State transitions.
- **task**—Routing protocol task processing.
- **timer**—Routing protocol timer processing.

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

no-world-readable—(Optional) Prevent any user from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level	routing and trace—To view this statement in the configuration.
	routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Tracing OSPF Protocol Traffic</i>

traffic-engineering (OSPF)

Syntax	<pre> traffic-engineering { <advertise-unnumbered-interfaces>; <credibility-protocol-preference>; ignore-lsp-metrics; multicast-rpf-routes; no-topology; shortcuts { lsp-metric-into-summary; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>multicast-rpf-routes option introduced in Junos OS Release 7.5.</p> <p>advertise-unnumbered-interfaces option introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4.</p> <p>Support for OSPFv3 (ospf3) introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>credibility-protocol-preference statement introduced in Junos OS Release 9.4.</p> <p>credibility-protocol-preference statement introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Enable the OSPF traffic engineering features.
Default	Traffic engineering support is disabled.
Options	<p>advertise-unnumbered-interfaces—(Optional) (OSPFv2 only) Include the link-local identifier in the link-local traffic-engineering link-state advertisement. This statement must be included on both ends of an unnumbered link to allow an ingress LER to update the link in its traffic engineering database and use it for CSPF calculations. The link-local identifier is then used by RSVP to signal unnumbered interfaces as defined in RFC 3477.</p> <p>credibility-protocol-preference—(Optional) (OSPFv2 only) Use the configured preference value for OSPF routes to calculate the traffic engineering database credibility value used to select IGP routes. Use this statement to override the default behavior, in which the traffic engineering database prefers IS-IS routes even if OSPF routes are configured with a lower, that is, preferred, preference value. For example, OSPF routes have a default preference value of 10, whereas IS-IS Level 1 routes have a default preference value of 15. When protocol preference is enabled, the credibility value is determined by deducting the protocol preference value from a base value of 512. Using default protocol preference values, OSPF has a credibility value of 502,</p>

whereas IS-IS has a credibility value of 497. Because the traffic engineering database prefers IGP routes with the highest credibility value, OSPF routes are now preferred.

multicast-rpf-routes—(Optional) (OSPFv2 only) Install routes for multicast RPF checks into the **inet.2** routing table. The **inet.2** routing table consists of unicast routes used for multicast RPF lookup. RPF is an antispoofing mechanism used to check whether the packet is coming in on an interface that is also sending data back to the packet source.



NOTE: You must enable OSPF traffic engineering shortcuts to use the **multicast-rpf-routes** statement. You must not allow LSP advertisements into OSPF when configuring the **multicast-rpf-routes** statement.

no-topology—(Optional) (OSPFv2 only) Disable the dissemination of the link-state topology information.

The remaining statements are explained separately.

Required Privilege Level	routing—To view this statement in the configuration.
	routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Enabling OSPF Traffic Engineering Support</i>

transit-delay (OSPF)

Syntax	<code>transit-delay seconds;</code>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> virtual-link], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>], [edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>], [edit protocols (ospf ospf3) area <i>area-id</i> virtual-link], [edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)] area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> virtual-link], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Set the estimated time required to transmit a link-state update on the interface. When calculating this time, make sure to account for transmission and propagation delays.</p> <p>You should never have to modify the transit delay time.</p>
Options	<p>seconds—Estimated time, in seconds.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: 1 second</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring OSPF Timers</i> • <i>Configuring RSVP and OSPF for LMP Peer Interfaces</i>

type-7

Syntax	type-7;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa default-lsa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa default-lsa],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa default-lsa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> nssa default-lsa],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> nssa default-lsa]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Flood Type 7 default link-state advertisements (LSAs) if the no-summaries statement is configured.</p> <p>By default, when the no-summaries statement is configured, a Type 3 LSA is injected into not-so-stubby areas (NSSAs) for Junos OS Release 5.0 and later. To support backward compatibility with earlier Junos OS releases, include the type-7 statement. This statement enables NSSA ABRs to advertise a Type 7 default LSA into the NSSA if you have also included the no-summaries statement in the configuration.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>OSPF Areas and Router Functionality Overview</i> • <i>Example: Configuring OSPF Not-So-Stubby Areas</i> • no-summaries on page 2847

virtual-link

Syntax	<pre>virtual-link neighbor-id <i>router-id</i> transit-area <i>area-id</i> { disable; authentication key <key-id identifier>; dead-interval <i>seconds</i>; hello-interval <i>seconds</i>; ipsec-sa <i>name</i>; retransmit-interval <i>seconds</i>; transit-delay <i>seconds</i>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) <i>area area-id</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf <i>area area-id</i>], [edit protocols (ospf ospf3) <i>area area-id</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf <i>area area-id</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	<p>For backbone areas only, create a virtual link to use in place of an actual physical link. All area border routers and other routing devices on the backbone must be contiguous. If this is not possible and there is a break in OSPF connectivity, use virtual links to create connectivity to the OSPF backbone. When configuring virtual links, you must configure links on the two routing devices that form the end points of the link, and both of these routing devices must be area border routers. You cannot configure links through stub areas.</p>
Options	<p>neighbor-id <i>router-id</i>—IP address of the routing device at the remote end of the virtual link.</p> <p>transit-area <i>area-id</i>—Area identifier of the area through which the virtual link transits. Virtual links are not allowed to transit the backbone area.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>OSPF Areas and Router Functionality Overview</i> • <i>Example: Configuring OSPF Virtual Links</i>

Administration

- [Operational Commands on page 2855](#)

Operational Commands

clear (ospf | ospf3) database

Syntax clear (ospf | ospf3) database
<advertising-router (*router-id* | self)>
<area *area-id*>
<asbrsummary>
<external>
<instance *instance-name*>
<inter-area-prefix>
<inter-area-router>
<intra-area-prefix>
<link-local>
<logical-system (all | *logical-system-name*)>
<lsa-id *lsa-id*>
<netsummary>
<network>
<nssa>
<opaque-area>
<purge>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
<router>

Syntax (EX Series Switch and QFX Series) clear (ospf | ospf3) database
<advertising-router (*router-id* | self)>
<area *area-id*>
<asbrsummary>
<external>
<instance *instance-name*>
<inter-area-prefix>
<inter-area-router>
<intra-area-prefix>
<link-local>
<lsa-id *lsa-id*>
<netsummary>
<network>
<nssa>
<opaque-area>
<purge>
<router>

Release Information Command introduced before Junos OS Release 7.4.
advertising-router *router-id*, **area** *area-id*, **asbrsummary**, **external**, **inter-area-prefix**, **inter-area-router**, **intra-area-prefix**, **link-local**, **lsa-id** *lsa-id*, **netsummary**, **network**, **nssa**, **opaque-area**, and **router** options added in Junos OS Release 8.3. You must use the **purge** command with these options.
Command introduced in Junos OS Release 9.0 for EX Series switches.
realm option added in Junos OS Release 9.2.
advertising-router (*router-id* | **self**) option added in Junos OS Release 9.5.
advertising-router (*router-id* | **self**) option introduced in Junos OS Release 9.5 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.

Description With the master Routing Engine, delete entries in the Open Shortest Path First (OSPF) link-state advertisement (LSA) database. With the backup Routing Engine, delete the OSPF LSA database and sync the new database with the master Routing Engine. You can also use the **purge** command with any of the options to discard rather than delete the specified LSA entries.



CAUTION: This command is useful only for testing. Use it with care, because it causes significant network disruption.

Options

- none**—Delete all LSAs other than the system's own LSAs, which are regenerated. To resynchronize the database, the system destroys all adjacent neighbors that are in the state **EXSTART** or higher. The neighbors are then reacquired and the databases are synchronized.
- advertising-router** (*router-id* | **self**)—(Optional) Discard entries for the LSA entries advertised by the specified routing device or by this routing device.
- area** *area-id*—(Optional) Discard entries for the LSAs in the specified area.
- asbrsummary**—(Optional) Discard summary AS boundary router LSA entries.
- external**—(Optional) Discard external LSAs.
- instance** *instance-name*—(Optional) Delete or discard entries for the specified routing instance only.
- inter-area-prefix**—(OSPFv3 only) (Optional) Discard interarea prefix LSAs.
- inter-area-router**—(OSPFv3 only) (Optional) Discard interarea router LSAs.
- intra-area-prefix**—(OSPFv3 only) (Optional) Discard intra-area prefix LSAs.
- logical-system** (**all** | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.
- link-local**—(Optional) Delete link-local LSAs.
- lsa-id** *lsa-id*—(Optional) Discard the LSA entries with the specified LSA identifier.
- netsummary**—(Optional) Discard summary network LSAs.
- network**—(Optional) Discard network LSAs.
- nssa**—(Optional) Discard not-so-stubby area (NSSA) LSAs.
- opaque-area**—(Optional) Discard opaque area-scope LSAs.
- realm** (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(OSPFv3 only) (Optional) Delete the entries for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

router—(Optional) Discard router LSAs.

purge—(Optional) Discard all entries in the link-state advertisement database. All link-state advertisements are set to **MAXAGE** and are flooded. The database is repopulated when the originators of the link-state advertisements receive the **MAXAGE** link-state advertisements and reissue them.

Required Privilege Level

clear

Related Documentation

- [show ospf database on page 2897](#)
- [show ospf3 database on page 2905](#)

List of Sample Output [clear ospf database on page 2858](#)

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

Sample Output

[clear ospf database](#)

```
user@host> clear ospf database
```


clear (ospf | ospf3) io-statistics

Syntax	clear (ospf ospf3) io-statistics <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and QFX Series)	clear (ospf ospf3) io-statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear Open Shortest Path First (OSPF) input and output statistics.
Options	none —Clear OSPF input and output statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
List of Sample Output	clear ospf io-statistics on page 2859
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf io-statistics

```
user@host> clear ospf io-statistics
```

clear (ospf | ospf3) neighbor

Syntax	clear (ospf ospf3) neighbor <area <i>area-id</i> > <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> <neighbor> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)>
Syntax (EX Series Switch and QFX Series)	clear (ospf ospf3) neighbor <area <i>area-id</i> > <instance <i>instance-name</i> > <interface <i>interface-name</i> > <neighbor>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Tear down Open Shortest Path First (OSPF) neighbor connections.
Options	none —Tear down OSPF connections with all neighbors for all routing instances. area <i>area-id</i> —(Optional) Tear down neighbor connections for the specified area only. instance <i>instance-name</i> —(Optional) Tear down neighbor connections for the specified routing instance only. interface <i>interface-name</i> —(Optional) Tear down neighbor connections for the specified interface only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. neighbor —(Optional) Clear the state of the specified neighbor only. realm (ipv4-multicast ipv4-unicast ipv6-multicast) —(Optional) (OSPFv3 only) Clear the state of the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show (ospf ospf3) neighbor on page 2876
List of Sample Output	clear ospf neighbor on page 2861
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf neighbor

```
user@host> clear ospf neighbor
```

clear (ospf | ospf3) statistics

Syntax	clear (ospf ospf3) statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)>
Syntax (EX Series Switch and QFX Series)	clear (ospf ospf3) statistics <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear Open Shortest Path First (OSPF) statistics.
Options	<p>none—Clear OSPF statistics.</p> <p>instance <i>instance-name</i>—(Optional) Clear statistics for the specified routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(Optional) (OSPFv3 only) Clear statistics for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> show (ospf ospf3) statistics on page 2893
List of Sample Output	clear ospf statistics on page 2862
Output Fields	See show (ospf ospf3) statistics for an explanation of output fields.

Sample Output

clear ospf statistics

The following sample output displays OSPF statistics before and after the **clear ospf statistics** command is entered:

```
user@host> show ospf statistics
```

Packet type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Hello	3254	2268	3	1
DbD	41	46	0	0

```

LSReq          8          7          0          0
LSUpdate       212       154        0          0
LSAck          65        98          0          0

DBDs retransmitted :          3, last 5 seconds :          0
LSAs flooded       :          12, last 5 seconds :          0
LSAs flooded high-prio :          0, last 5 seconds :          0
LSAs retransmitted :          0, last 5 seconds :          0
LSAs transmitted to nbr:          3, last 5 seconds :          0
LSAs requested     :          5, last 5 seconds :          0
LSAs acknowledged :          19, last 5 seconds :          0

Flood queue depth :          0
Total rexmit entries :          0
db summaries      :          0
lsreq entries     :          0

Receive errors:
  626 subnet mismatches

```

```
user@host> clear ospf statistics
```

```
user@host> show ospf statistics
```

```

Packet type      Total
                  Sent   Received
Hello            3       1
  DbD             0       0
  LSReq           0       0
LSUpdate         0       0
LSAck            0       0

                  Last 5 seconds
                  Sent   Received
Hello            3       1
  DbD             0       0
  LSReq           0       0
LSUpdate         0       0
LSAck            0       0

DBDs retransmitted :          0, last 5 seconds :          0
LSAs flooded       :          0, last 5 seconds :          0
LSAs flooded high-prio :          0, last 5 seconds :          0
LSAs retransmitted :          0, last 5 seconds :          0
LSAs transmitted to nbr:          0, last 5 seconds :          0
LSAs requested     :          0, last 5 seconds :          0
LSAs acknowledged :          0, last 5 seconds :          0

Flood queue depth :          0
Total rexmit entries :          0
db summaries      :          0
lsreq entries     :          0

Receive errors:
  None

```

clear (ospf | ospf3) overload

Syntax	clear (ospf ospf3) overload <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	clear (ospf ospf3) overload <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear the Open Shortest Path First (OSPF) overload bit and rebuild link-state advertisements (LSAs).
Options	none —Clear the overload bit and rebuild LSAs for all routing instances. instance <i>instance-name</i> —(Optional) Clear the overload bit and rebuild LSAs for the specified routing instance only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
List of Sample Output	clear ospf overload on page 2864
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf overload

```
user@host> clear ospf overload
```

show (ospf | ospf3) interface

Syntax	<pre>show (ospf ospf3) interface <brief detail extensive> <area <i>area-id</i>> <<i>interface-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show (ospf ospf3) interface <brief detail extensive> <area <i>area-id</i>> <<i>interface-name</i>> <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>area option introduced in Junos OS Release 9.2.</p> <p>area option introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>realm option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display the status of OSPF interfaces.
Options	<p>none—Display standard information about the status of all OSPF interfaces for all routing instances</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>area <i>area-id</i>—(Optional) Display information about the interfaces that belong to the specified area.</p> <p><i>interface-name</i>—(Optional) Display information for the specified interface.</p> <p>instance <i>instance-name</i>—(Optional) Display all OSPF interfaces under the named routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(OSPFv3 only) (Optional) Display information about the interfaces for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	view
List of Sample Output	<p>show ospf interface brief on page 2868</p> <p>show ospf interface detail on page 2868</p> <p>show ospf3 interface detail on page 2868</p>

[show ospf interface detail \(When Multiarea Adjacency Is Configured\) on page 2868](#)
[show ospf interface area area-id on page 2869](#)
[show ospf interface extensive \(When Flooding Reduction Is Enabled\) on page 2870](#)
[show ospf interface extensive \(When LDP Synchronization Is Configured\) on page 2870](#)

Output Fields Table 230 on page 2866 lists the output fields for the **show (ospf | ospf3) interface** command. Output fields are listed in the approximate order in which they appear.

Table 230: show (ospf | ospf3) interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface running OSPF version 2 or OSPF version 3.	All levels
State	State of the interface: BDR , Down , DR , DRother , Loop , PtToPt , or Waiting .	All levels
Area	Number of the area that the interface is in.	All levels
DR ID	Address of the area's designated router.	All levels
BDR ID	Backup designated router for a particular subnet.	All levels
Nbrs	Number of neighbors on this interface.	All levels
Type	Type of interface: LAN , NBMA , P2MP , P2P , or Virtual .	detail extensive
Address	IP address of the neighbor.	detail extensive
Mask	Netmask of the neighbor.	detail extensive
Prefix-length	(OSPFv3) IPv6 prefix length, in bits.	detail extensive
OSPF3-Intf-Index	(OSPFv3) OSPF version 3 interface index.	detail extensive
MTU	Interface maximum transmission unit (MTU).	detail extensive
Cost	Interface cost (metric).	detail extensive
DR addr	Address of the designated router.	detail extensive
BDR addr	Address of the backup designated router.	detail extensive
Adj count	Number of adjacent neighbors.	detail extensive
Secondary	Indicates that this interface is configured as a secondary interface for this area. This interface can belong to more than one area, but can be designated as a primary interface for only one area.	detail extensive
Flood Reduction	Indicates that this interface is configured with flooding reduction. All self-originated LSAs from this interface are initially sent with the DoNotAge bit set. As a result, LSAs are refreshed only when a change occurs.	extensive

Table 230: show (ospf | ospf3) interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Priority	Router priority used in designated router (DR) election on this interface.	detail extensive
Flood list	List of link-state advertisements (LSAs) that might be about to flood this interface.	extensive
Ack list	Acknowledgment list. List of pending acknowledgments on this interface.	extensive
Descriptor list	List of packet descriptors.	extensive
Hello	Configured value for the hello timer.	detail extensive
Dead	Configured value for the dead timer.	detail extensive
Auth type	(OSPFv2) Authentication mechanism for sending and receiving OSPF protocol packets: <ul style="list-style-type: none"> • MD5—The MD5 mechanism is configured in accordance with RFC 2328. • None—No authentication method is configured. • Password—A simple password (RFC 2328) is configured. 	detail extensive
Topology	(Multiarea adjacency) Name of topology: default or <i>name</i> .	
LDP sync state	(OSPFv2 and LDP synchronization) Current state of LDP synchronization: in sync , in holddown , and not supported .	extensive
reason	(OSPFv2 and LDP synchronization) Reason for the current state of LDP synchronization. The LDP session might be up or down, or adjacency might be up or down.	extensive
config holdtime	(OSPFv2 and LDP synchronization) Configured value of the hold timer. If the state is not synchronized, and the hold time is not infinity, the remaining field displays the number of seconds that remain until the configured hold timer expires.	extensive
IPSec SA name	(OSPFv2) Name of the IPSec security association name.	detail extensive
Active key ID	(OSPFv2 and MD5) Number from 0 to 255 that uniquely identifies an MD5 key.	detail extensive
Start time	(OSPFv2 and MD5) Time at which the routing device starts using an MD5 key to authenticate OSPF packets transmitted on the interface on which this key is configured. To authenticate received OSPF protocol packets, the key becomes effective immediately after the configuration is committed. If the start time option is not configured, the key is effective immediately for send and receive and is displayed as Start time 1970 Jan 01 00:00:00 PST .	detail extensive
ReXmit	Configured value for the Retransmit timer.	detail extensive
Stub, Not Stub, or Stub NSSA	Type of area.	detail extensive

Sample Output

show ospf interface brief

```
user@host> show ospf interface brief
```

Intf	State	Area	DR ID	BDR ID	Nbrs
at-5/1/0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
ge-2/3/0.0	DR	0.0.0.0	192.168.4.16	192.168.4.15	1
lo0.0	DR	0.0.0.0	192.168.4.16	0.0.0.0	0
so-0/0/0.0	Down	0.0.0.0	0.0.0.0	0.0.0.0	0
so-6/0/1.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-6/0/2.0	Down	0.0.0.0	0.0.0.0	0.0.0.0	0
so-6/0/3.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

show ospf interface detail

```
user@host> show ospf interface detail
```

Interface	State	Area	DR ID	BDR ID	Nbrs
fe-0/0/1.0	BDR	0.0.0.0	192.168.37.12	10.255.245.215	1

Type LAN, address 192.168.37.11, Mask 255.255.255.248, MTU 4460, Cost 40
 DR addr 192.168.37.12, BDR addr 192.168.37.11, Adj count 1, Priority 128
 Hello 10, Dead 40, ReXmit 5, Not Stub

Interface	State	Area	DR ID	BDR ID	Nbrs
tl-0/2/1.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	0

Type P2P, Address 0.0.0.0, Mask 0.0.0.0, MTU 1500, Cost 2604
 Adj count 0
 Hello 10, Dead 40, ReXmit 5, Not Stub
 Auth type: MD5, Active key ID 3, Start time 2002 Nov 19 10:00:00 PST
 IPsec SA Name: sa

show ospf3 interface detail

```
user@host> show ospf3 interface so-0/0/3.0 detail
```

Interface	State	Area	DR-ID	BDR-ID	Nbrs
so-0/0/3.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

Address fe80::2a0:a5ff:fe28:1dfc, Prefix-length 64
 OSPF3-Intf-index 1, Type P2P, MTU 4470, Cost 12, Adj-count 1
 Hello 10, Dead 40, ReXmit 5, Not Stub

show ospf interface detail (When Multiarea Adjacency Is Configured)

```
user@host> show ospf interface detail
```

```
regress@router> show ospf interface detail
```

Interface	State	Area	DR ID	BDR ID	Nbrs
lo0.0	DR	0.0.0.0	10.255.245.2	0.0.0.0	0

Type: LAN, Address: 127.0.0.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0
 DR addr: 127.0.0.1, Adj count: 0, Priority: 128
 Hello: 10, Dead: 40, ReXmit: 5, Not Stub
 Auth type: None
 Topology default (ID 0) -> Cost: 0

Interface	State	Area	DR ID	BDR ID	Nbrs
lo0.0	DR	0.0.0.0	10.255.245.2	0.0.0.0	0

Type: LAN, Address: 10.255.245.2, Mask: 255.255.255.255, MTU: 65535, Cost: 0
 DR addr: 10.255.245.2, Adj count: 0, Priority: 128
 Hello: 10, Dead: 40, ReXmit: 5, Not Stub
 Auth type: None
 Topology default (ID 0) -> Cost: 0

Interface	State	Area	DR ID	BDR ID	Nbrs
so-0/0/0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1

```

Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-0/0/0.0      PtToPt  0.0.0.0      0.0.0.0      0.0.0.0      0

Type: P2P, Address: 192.168.37.46, Mask: 255.255.255.254, MTU: 4470, Cost: 1
Adj count: 0, , Passive
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Passive, Cost: 1
so-1/0/0.0      PtToPt  0.0.0.0      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  0.0.0.0      0.0.0.0      0.0.0.0      0

Type: P2P, Address: 192.168.37.54, Mask: 255.255.255.254, MTU: 4470, Cost: 1
Adj count: 0, , Passive
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Passive, Cost: 1
so-0/0/0.0      PtToPt  1.1.1.1      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  1.1.1.1      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-0/0/0.0      PtToPt  2.2.2.2      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  2.2.2.2      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1

```

show ospf interface area area-id

```
user@host> show ospf interface area 1.1.1.1
```

Interface	State	Area	DR ID	BDR ID	Nbrs
so-0/0/0.0	PtToPt	1.1.1.1	0.0.0.0	0.0.0.0	1
so-1/0/0.0	PtToPt	1.1.1.1	0.0.0.0	0.0.0.0	1

show ospf interface extensive
(When Flooding Reduction Is Enabled)

```
user@host> show ospf interface extensive
```

Interface	State	Area	DR ID	BDR ID	Nbrs
fe-0/0/0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	0

Type: P2P, Address: 10.10.10.1, Mask: 255.255.255.0, MTU: 1500, Cost: 1
Adj count: 0
Secondary, Flood Reduction
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1

show ospf interface extensive
(When LDP Synchronization Is Configured)

```
user@host> show ospf interface extensive
```

Interface	State	Area	DR ID	BDR ID
so-1/0/3.0	Down	0.0.0.0	0.0.0.0	0.0.0.0

Nbrs
0

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 65535
Adj count: 0
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
LDP sync state: in holddown, for: 00:00:08, reason: LDP down during config
config holddtime: 10 seconds, remaining: 1

show (ospf | ospf3) io-statistics

Syntax	show (ospf ospf3) io-statistics <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and QFX Series)	show (ospf ospf3) io-statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display Open Shortest Path First (OSPF) input and output statistics.
Options	<p>none—Display OSPF input and output statistics.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear (ospf ospf3) statistics on page 2862
List of Sample Output	show ospf io-statistics on page 2871
Output Fields	Table 231 on page 2871 lists the output fields for the show ospf io-statistics command. Output fields are listed in the approximate order in which they appear.

Table 231: show (ospf | ospf3) io-statistics Output Fields

Field Name	Field Description
Packets read	Number of OSPF packets read since the last time the routing protocol was started.
average per run	Total number of packets divided by the total number of times the OSPF read operation is scheduled to run.
max run	Maximum number of packets for a given run among all scheduled runs.
Receive errors	Number of faulty packets received with errors.

Sample Output

show ospf io-statistics

```
user@host> show ospf io-statistics
```

```
Packets read: 7361, average per run: 1.00, max run: 1
```

Receive errors:
None

show (ospf | ospf3) log

Syntax	show (ospf ospf3) log <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> <topology <i>topology-name</i> >
Syntax (EX Series Switch and QFX Series)	show (ospf ospf3) log <instance <i>instance-name</i> > <topology <i>topology-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. topology option introduced in Junos OS Release 9.0. topology option introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display the entries in the Open Shortest Path First (OSPF) log of SPF calculations.
Options	<p>none—Display entries in the OSPF log of SPF calculations for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display entries for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>topology <i>topology-name</i>—(Optional) (OSPFv2 only) Display entries for the specified topology.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(OSPFv3 only) (Optional) Display entries for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	view
List of Sample Output	show ospf log on page 2874 show ospf log topology voice on page 2874
Output Fields	Table 232 on page 2873 lists the output fields for the show (ospf ospf3) log command. Output fields are listed in the approximate order in which they appear.

Table 232: show (ospf | ospf3) log Output Fields

Field Name	Field Description
When	Time, in weeks (w) and days (d), since the SPF calculation was made.

Table 232: show (ospf | ospf3) log Output Fields (*continued*)

Field Name	Field Description
Type	Type of calculation: Cleanup, External, Interarea, NSSA, Redist, SPF, Stub, Total, or Virtuallink.
Elapsed	Amount of time, in seconds, that elapsed during the operation, or the time required to complete the SPF calculation. The start time is the time displayed in the When field.

Sample Output

show ospf log

```

user@host> show ospf log
When          Type          Elapsed
1w4d 17:25:58 Stub          0.000017
1w4d 17:25:58 SPF            0.000070
1w4d 17:25:58 Stub          0.000019
1w4d 17:25:58 Interarea      0.000054
1w4d 17:25:58 External       0.000005
1w4d 17:25:58 Cleanup        0.000203
1w4d 17:25:58 Total         0.000537
1w4d 17:24:48 SPF            0.000125
1w4d 17:24:48 Stub          0.000017
1w4d 17:24:48 SPF            0.000100
1w4d 17:24:48 Stub          0.000016
1w4d 17:24:48 Interarea      0.000056
1w4d 17:24:48 External       0.000005
1w4d 17:24:48 Cleanup        0.000238
1w4d 17:24:48 Total         0.000600
...

```

show ospf log topology voice

```

user@host> show ospf log topology voice
Topology voice SPF log:

    Last instance of each event type
When          Type          Elapsed
00:06:11      SPF            0.000116
00:06:11      Stub            0.000114
00:06:11      Interarea        0.000126
00:06:11      External         0.000067
00:06:11      NSSA             0.000037
00:06:11      Cleanup          0.000186

    Maximum length of each event type
When          Type          Elapsed
00:13:43      SPF            0.000140
00:13:33      Stub            0.000116
00:13:43      Interarea        0.000128
00:13:33      External         0.000075
00:13:38      NSSA             0.000039
00:13:53      Cleanup          0.000657

    Last 100 events

```


When	Type	Elapsed
00:13:53	SPF	0.000090
00:13:53	Stub	0.000041
00:13:53	Interarea	0.000123
00:13:53	External	0.000040
00:13:53	NSSA	0.000038
00:13:53	Cleanup	0.000657
00:13:53	Total	0.001252
.		
.		
00:06:11	SPF	0.000116
00:06:11	Stub	0.000114
00:06:11	Interarea	0.000126
00:06:11	External	0.000067
00:06:11	NSSA	0.000037
00:06:11	Cleanup	0.000186
00:06:11	Total	0.000818

show (ospf | ospf3) neighbor

Syntax	<pre>show (ospf ospf3) neighbor <brief detail extensive> <area <i>area-id</i>> <instance (all <i>instance-name</i>)> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show (ospf ospf3) neighbor <brief detail extensive> <area <i>area-id</i>> <instance (all <i>instance-name</i>)> <interface <i>interface-name</i>> <neighbor></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>instance all option introduced in Junos OS Release 9.1.</p> <p>instance all option introduced in Junos OS Release 9.1 for EX Series switches.</p> <p>area, interface, and realm options introduced in Junos OS Release 9.2.</p> <p>area and interface options introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Display information about OSPF neighbors.</p> <p>CPU utilization might increase while the device learns its OSPF neighbors. We recommend that you use the show (ospf ospf3) neighbor command after the device learns and establishes OSPF neighbor adjacencies. Depending on the size of your network, this might take several minutes. If you receive a “timeout communicating with routing daemon” error when using the show (ospf ospf3) neighbor command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the CLI.</p>
Options	<p>none—Display standard information about all OSPF neighbors for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>area <i>area-id</i>—(Optional) Display information about the OSPF neighbors for the specified area.</p> <p>instance (all <i>instance-name</i>)—(Optional) Display all OSPF interfaces for all routing instances or under the named routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Display information about OSPF neighbors for the specified logical interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>

neighbor—(Optional) Display information about the specified OSPF neighbor.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(OSPFv3 only) (Optional) Display information about the OSPF neighbors for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear (ospf ospf3) neighbor on page 2860
List of Sample Output	show ospf neighbor brief on page 2879 show ospf neighbor detail on page 2879 show ospf neighbor extensive on page 2880 show ospf3 neighbor detail on page 2881 show ospf neighbor area area-id on page 2881 show ospf neighbor interface interface-name on page 2881 show ospf3 neighbor instance all (OSPFv3 Multiple Family Address Support Enabled) on page 2881
Output Fields	Table 233 on page 2877 lists the output fields for the show (ospf ospf3) neighbor command. Output fields are listed in the approximate order in which they appear.

Table 233: show (ospf | ospf3) neighbor Output Fields

Field Name	Field Description	Level of Output
Address	Address of the neighbor.	All levels
Interface	Interface through which the neighbor is reachable.	All levels

Table 233: show (ospf | ospf3) neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
State	<p>State of the neighbor:</p> <ul style="list-style-type: none"> • Attempt—Valid only for neighbors attached to nonbroadcast networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort must be made to contact the neighbor. • Down—Initial state of a neighbor conversation. It indicates that no recent information has been received from the neighbor. Hello packets might continue to be sent to neighbors in the Down state, although at a reduced frequency. • Exchange—Routing device is describing its entire link-state database by sending database description packets to the neighbor. Each packet has a sequence number and is explicitly acknowledged. • ExStart—First step in creating an adjacency between the two neighboring routing devices. The goal of this step is to determine which routing device is the master, and to determine the initial sequence number. • Full—Neighboring routing devices are fully adjacent. These adjacencies appear in router link and network link advertisements. • Init—A hello packet has recently been sent by the neighbor. However, bidirectional communication has not yet been established with the neighbor. This state might occur, for example, because the routing device itself did not appear in the neighbor's hello packet. • Loading—Link-state request packets are sent to the neighbor to acquire more recent advertisements that have been discovered (but not yet received) in the Exchange state. • 2Way—Communication between the two routing devices is bidirectional. This state has been ensured by the operation of the Hello Protocol. This is the most advanced state short of beginning adjacency establishment. The (backup) designated router is selected from the set of neighbors in state 2Way or greater. 	All levels
ID	Router ID of the neighbor.	All levels
Pri	Priority of the neighbor to become the designated router.	All levels
Dead	Number of seconds until the neighbor becomes unreachable.	All levels
Link state acknowledgment list	Number of link-state acknowledgments received.	extensive
Link state retransmission list	<p>Total number of link-state advertisements retransmitted. For extensive output only, the following information is also displayed:</p> <ul style="list-style-type: none"> • Type—Type of link advertisement: ASBR, Sum, Extern, Network, NSSA, OpaqArea, Router, or Summary. • LSA ID—LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device. • Adv rtr—Address of the routing device that sent the advertisement. • Seq—Link sequence number of the advertisement. 	detail extensive

Table 233: show (ospf | ospf3) neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Neighbor-address	(OSPFv3 only) If the neighbor uses virtual links, the Neighbor-address is the site-local, local, or global address. If the neighbor uses a physical interface, the Neighbor-address is an IPv6 link-local address.	detail extensive
area	Area that the neighbor is in.	detail extensive
OSPF3-Intf-Index	(OSPFv3 only) Displays the OSPFv3 interface index.	detail extensive
opt	Option bits received in the hello packets from the neighbor.	detail extensive
DR or DR-ID	Address of the designated router.	detail extensive
BDR or BDR-ID	Address of the backup designated router.	detail extensive
Up	Length of time since the neighbor came up.	detail extensive
adjacent	Length of time since the adjacency with the neighbor was established.	detail extensive

Sample Output

show ospf neighbor brief

```

user@host> show ospf neighbor brief
  Address      Intf      State      ID          Pri  Dead
192.168.254.225 fxp3.0    2Way       10.250.240.32 128  36
192.168.254.230 fxp3.0    Full       10.250.240.8  128  38
192.168.254.229 fxp3.0    Full       10.250.240.35 128  33
10.1.1.129      fxp2.0    Full       10.250.240.12 128  37
10.1.1.131      fxp2.0    Full       10.250.240.11 128  38
10.1.2.1        fxp1.0    Full       10.250.240.9  128  32
10.1.2.81       fxp0.0    Full       10.250.240.10 128  33

```

show ospf neighbor detail

```

user@host> show ospf neighbor detail
  Address      Interface      State      ID          Pri  Dead
10.5.1.2      ge-1/2/0.1    Full       10.5.1.2    128  37
area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
Up 06:09:28, adjacent 05:17:36
Link state acknowledgment list: 3 entries

Link state retransmission list: 9 entries

10.5.10.2      ge-1/2/0.10    ExStart    10.5.1.38   128  34
area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
Up 06:09:28
master, seq 0xac1530f8, rexmit DBD in 3 sec
rexmit LSREQ in 0 sec
10.5.11.2      ge-1/2/0.11    Full       10.5.1.42   128  38
area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
Up 06:09:28, adjacent 05:26:46
Link state retransmission list: 1 entries

```

```

10.5.12.2      ge-1/2/0.12      ExStart  10.5.1.46      128    33
area 0.0.0.1, opt 0x42, DR 10.5.12.2, BDR 10.5.12.1
Up 06:09:28
master, seq 0xac188a68, rexmit DBD in 2 sec
rexmit LSREQ in 0 sec

```

show ospf neighbor extensive

```

user@host> show ospf neighbor extensive
Address      Interface      State      ID      Pri  Dead
10.5.1.2      ge-1/2/0.1     Full      10.5.1.2  128   33
area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
Up 06:09:42, adjacent 05:17:50
Link state retransmission list:

  Type      LSA ID      Adv rtr      Seq
Summary 10.8.56.0    172.25.27.82 0x8000004d
Router  10.5.1.94    10.5.1.94    0x8000005c
Network 10.5.24.2    10.5.1.94    0x80000036
Summary 10.8.57.0    172.25.27.82 0x80000024
Extern  1.10.90.0    10.8.1.2     0x80000041
Extern  1.4.109.0     10.6.1.2     0x80000041
Router  10.5.1.190    10.5.1.190   0x8000005f
Network 10.5.48.2    10.5.1.190   0x8000003d
Summary 10.8.58.0    172.25.27.82 0x8000004d
Extern  1.10.91.0    10.8.1.2     0x80000041
Extern  1.4.110.0     10.6.1.2     0x80000041
Router  10.5.1.18     10.5.1.18    0x8000005f
Network 10.5.5.2     10.5.1.18    0x80000033
Summary 10.8.59.0    172.25.27.82 0x8000003a
Summary 10.8.62.0    172.25.27.82 0x80000025

10.5.10.2     ge-1/2/0.10     ExStart  10.5.1.38      128    38
area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
Up 06:09:42
master, seq 0xac1530f8, rexmit DBD in 2 sec
rexmit LSREQ in 0 sec
10.5.11.2     ge-1/2/0.11     Full      10.5.1.42      128    33
area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
Up 06:09:42, adjacent 05:27:00
Link state retransmission list:

  Type      LSA ID      Adv rtr      Seq
Summary 10.8.58.0    172.25.27.82 0x8000004d

```

Extern	1.10.91.0	10.8.1.2	0x80000041
Extern	1.1.247.0	10.5.1.2	0x8000003f
Extern	1.4.110.0	10.6.1.2	0x80000041
Router	10.5.1.18	10.5.1.18	0x8000005f
Network	10.5.5.2	10.5.1.18	0x80000033
Summary	10.8.59.0	172.25.27.82	0x8000003a

show ospf3 neighbor detail

```
user@host> show ospf3 neighbor detail
ID          Interface          State    Pri    Dead
10.255.71.13 fe-0/0/2.0          Full     128    30
Neighbor-address fe80::290:69ff:fe9b:e002
Area 0.0.0.0, opt 0x13, OSPF3-Intf-Index 2
DR-ID 10.255.71.13, BDR-ID 10.255.71.12
Up 02:51:43, adjacent 02:51:43
```

show ospf neighbor area area-id

```
user@host >show ospf neighbor area 1.1.1.1
Address      Interface          State    ID          Pri    Dead
192.168.37.47 so-0/0/0.0          Full     10.255.245.4 128    33
Area 1.1.1.1
192.168.37.55 so-1/0/0.0          Full     10.255.245.5 128    37
Area 1.1.1.1
```

show ospf neighbor interface interface-name

```
user@host >show ospf neighbor interface so-0/0/0.0
Address      Interface          State    ID          Pri    Dead
192.168.37.47 so-0/0/0.0          Full     10.255.245.4 128    37
Area 0.0.0.0
192.168.37.47 so-0/0/0.0          Full     10.255.245.4 128    33
Area 1.1.1.1
192.168.37.47 so-0/0/0.0          Full     10.255.245.4 128    32
Area 2.2.2.2
```

show ospf3 neighbor instance all (OSPFv3 Multiple Family Address Support Enabled)

```
user @host > show ospf3 neighbor instance all
Instance: ina
Realm: ipv6-unicast
ID          Interface          State    Pri    Dead
100.1.1.1    fe-0/0/2.0          Full     128    37
Neighbor-address fe80::217:cb00:c87c:8c03
Instance: inb
Realm: ipv4-unicast
ID          Interface          State    Pri    Dead
100.1.2.1    fe-0/0/2.1          Full     128    33
Neighbor-address fe80::217:cb00:c97c:8c03
```

show (ospf | ospf3) overview

Syntax	show (ospf ospf3) overview <brief extensive> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)>
Syntax (EX Series Switch and QFX Series)	show (ospf ospf3) overview <brief extensive> <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Database protection introduced in Junos 10.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display Open Shortest Path First (OSPF) overview information.
Options	<p>none—Display standard information about all OSPF neighbors for all routing instances.</p> <p>brief extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display all OSPF interfaces under the named routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(Optional) (OSPFv3 only) Display information about the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	view
List of Sample Output	show ospf overview on page 2884 show ospf overview (With Database Protection) on page 2885 show ospf3 overview (With Database Protection) on page 2885 show ospf overview extensive on page 2885
Output Fields	Table 234 on page 2882 lists the output fields for the show ospf overview command. Output fields are listed in the approximate order in which they appear.

Table 234: show ospf overview Output Fields

Field name	Field Description	Level of Output
Instance	OSPF routing instance.	All levels

Table 234: show ospf overview Output Fields (*continued*)

Field name	Field Description	Level of Output
Router ID	Router ID of the routing device.	All levels
Route table index	Route table index.	All levels
Configured overload	Overload capability is enabled. If the overload timer is also configured, display the time that remains before it is set to expire. This field is not displayed after the timer expires.	All levels
Topology	Topology identifier.	All levels
Prefix export count	Number of prefixes exported into OSPF.	All levels
Full SPF runs	Number of complete Shortest Path First calculations.	All levels
SPF delay	Delay before performing consecutive Shortest Path First calculations.	All levels
SPF holddown	Delay before performing additional Shortest Path First (SPF) calculations after the maximum number of consecutive SPF calculations is reached.	All levels
SPF rapid runs	Maximum number of Shortest Path First calculations that can be performed in succession before the hold-down timer begins.	All levels
LSA refresh time	Refresh period for link-state advertisement (in minutes).	All levels
Database protection state	Current state of database protection.	All levels
Warning threshold	Threshold at which a warning message is logged (percentage of maximum LSA count).	All levels
Non self-generated LSAs	Number of LSAs whose router ID is not equal to the local router ID: Current , Warning (threshold), and Allowed .	All levels
Ignore time	How long the database has been in the ignore state.	All levels
Reset time	How long the database must stay out of the ignore or isolated state before it returns to normal operations.	All levels
Ignore count	Number of times the database has been in the ignore state: Current and Allowed .	All levels
Restart	Graceful restart capability: enabled or disabled .	All levels
Restart duration	Time period for complete reacquisition of OSPF neighbors.	All levels
Restart grace period	Time period for which the neighbors should consider the restarting routing device as part of the topology.	All levels

Table 234: show ospf overview Output Fields (*continued*)

Field name	Field Description	Level of Output
Graceful restart helper mode	(OSPFv2) Standard graceful restart helper capability (based on RFC 3623): enabled or disabled .	All levels
Restart-signaling helper mode	(OSPFv2) Restart signaling-based graceful restart helper capability (based on RFC 4811, RFC 4812, and RFC 4813): enabled or disabled .	All levels
Helper mode	(OSPFv3) Graceful restart helper capability: enabled or disabled .	All levels
Trace options	OSPF-specific trace options.	extensive
Trace file	Name of the file to receive the output of the tracing operation.	extensive
Area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Stub type	Stub type of area: Normal Stub , Not Stub , or Not so Stubby Stub .	All levels
Authentication Type	Type of authentication: None , Password , or MD5 . NOTE: The Authentication Type field refers to the authentication configured at the [edit protocols ospf area area-id] level. Any authentication configured for an interface in this area will not affect the value of this field.	All levels
Area border routers	Number of area border routers.	All levels
Neighbors	Number of autonomous system boundary routers.	All levels

Sample Output

show ospf overview

```

user@host> show ospf overview
Instance: master
  Router ID: 10.255.245.6
  Route table index: 0
  Configured overload, expires in 118 seconds
  LSA refresh time: 50 minutes
  Restart: Enabled
    Restart duration: 20 sec
    Restart grace period: 40 sec
    Helper mode: enabled
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 0
  Topology: default (ID 0)
  Prefix export count: 0
  Full SPF runs: 1

```

SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3

show ospf overview (With Database Protection)

```
user@host> show ospf overview
Instance: master
Router ID: 10.255.112.218
Route table index: 0
LSA refresh time: 50 minutes
Traffic engineering
Restart: Enabled
  Restart duration: 180 sec
  Restart grace period: 210 sec
  Graceful restart helper mode: Enabled
  Restart-signaling helper mode: Enabled
Database protection state: Normal
Warning threshold: 70 percent
Non self-generated LSAs: Current 582, Warning 700, Allowed 1000
Ignore time: 30, Reset time: 60
Ignore count: Current 0, Allowed 1
Area: 0.0.0.0
  Stub type: Not Stub
  Authentication Type: None
  Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 160
Topology: default (ID 0)
Prefix export count: 0
Full SPF runs: 70
SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
Backup SPF: Not Needed
```

show ospf3 overview (With Database Protection)

```
user@host> show ospf3 overview
Instance: master
Router ID: 10.255.112.128
Route table index: 0
LSA refresh time: 50 minutes
Database protection state: Normal
Warning threshold: 80 percent
Non self-generated LSAs: Current 3, Warning 8, Allowed 10
Ignore time: 30, Reset time: 60
Ignore count: Current 0, Allowed 2
Area: 0.0.0.0
  Stub type: Not Stub
  Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
Topology: default (ID 0)
Prefix export count: 0
Full SPF runs: 7
SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
Backup SPF: Not Needed
```

show ospf overview extensive

```
user@host> show ospf overview extensive
Instance: master
Router ID: 1.1.1.103
Route table index: 0
```

```
Full SPF runs: 13, SPF delay: 0.200000 sec
LSA refresh time: 50 minutes
Restart: Disabled
Trace options: lsa
Trace file: /var/log/ospf size 131072 files 10
Area: 0.0.0.0
  Stub type: Not Stub
  Authentication Type: None
  Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
```

show (ospf | ospf3) route

Syntax	<pre>show (ospf ospf3) route <brief detail extensive> <abr asbr extern inter intra> <destination> <instance (default ipv4-multicast <i>instance-name</i>)> <logical-system (default ipv4-multicast <i>logical-system-name</i>)> <network> <no-backup-coverage> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> <router> <topology (default ipv4-multicast <i>topology-name</i>)> <transit></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show (ospf ospf3) route <brief detail extensive> <abr asbr extern inter intra> <destination> <instance <i>instance-name</i> <network> <no-backup-coverage> <router> <topology (default ipv4-multicast <i>topology-name</i>)> <transit></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>topology option introduced in Junos OS Release 9.0.</p> <p>realm option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display the entries in the Open Shortest Path First (OSPF) routing table.
Options	<p>none—Display standard information about all entries in the OSPF routing table for all routing instances and all topologies.</p> <p>destination—Display routes to the specified IP address (with optional destination prefix length).</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>abr—(Optional) Display routes to area border routers.</p> <p>asbr—(Optional) Display routes to autonomous system border routers.</p> <p>extern—(Optional) Display external routes.</p> <p>inter—(Optional) Display interarea routes.</p> <p>intra—(Optional) Display intra-area routes.</p>

instance (**default** | **ipv4-multicast** | *instance-name*)—(Optional) Display entries for the default routing instance, the IPv4 multicast routing instance, or for the specified routing instance.

logical-system (**default** | **ipv4-multicast** | *logical-system-name*)—(Optional) Perform this operation on the default logical system, the IPv4 multicast logical system, or on a particular logical system.

network—(Optional) Display routes to networks.

no-backup-coverage—(Optional) Display routes with no backup coverage.

realm (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(OSPFv3 only) (Optional) Display entries in the routing table for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

router—(Optional) Display routes to all routers.

topology (**default** | **ipv4-multicast** | *topology-name*)—(OSPFv2 only) (Optional) Display routes for the default OSPF topology, IPv4 multicast topology, or for a particular topology.

transit—(Optional) (OSPFv3 only) Display OSPFv3 routes to pseudonodes.

Required Privilege Level view

List of Sample Output [show ospf route on page 2890](#)
[show ospf route detail on page 2890](#)
[show ospf3 route on page 2890](#)
[show ospf3 route detail on page 2891](#)
[show ospf route topology voice on page 2891](#)

Output Fields [Table 235 on page 2888](#) list the output fields for the **show (ospf | ospf3) route** command. Output fields are listed in the approximate order in which they appear.

Table 235: show (ospf | ospf3) route Output Fields

Field Name	Field Description	Output Level
Topology	Name of the topology.	All levels
Prefix	Destination of the route.	All levels
Path type	How the route was learned: <ul style="list-style-type: none"> Inter—Interarea route Ext1—External type 1 route Ext2—External type 2 route Intra—Intra-area route 	All levels

Table 235: show (ospf | ospf3) route Output Fields (*continued*)

Field Name	Field Description	Output Level
Route type	The type of routing device from which the route was learned: <ul style="list-style-type: none"> • AS BR—Route to AS border router. • Area BR—Route to area border router. • Area/AS BR—Route to router that is both an Area BR and AS BR. • Network—Network router. • Router—Route to a router that is neither an Area BR nor an AS BR. • Transit—(OSPFv3 only) Route to a pseudonode representing a transit network, LAN, or nonbroadcast multiaccess (NBMA) link. • Discard—Route to a summary discard. 	All levels
NH Type	Next-hop type: LSP or IP .	All levels
Metric	Route's metric value.	All levels
NH-interface	(OSPFv3 only) Interface through which the route's next hop is reachable.	All levels
NH-addr	(OSPFv3 only) IPv6 address of the next hop.	All levels
NextHop Interface	(OSPFv2 only) Interface through which the route's next hop is reachable.	All levels
Nexthop addr/label	(OSPFv2 only) If the NH Type is IP , then it is the address of the next hop. If the NH Type is LSP , then it is the name of the label-switched path.	All levels
Area	Area ID of the route.	detail
Origin	Router from which the route was learned.	detail
Type 7	Route was learned through a not-so-stubby area (NSSA) link-state advertisement (LSA).	detail
P-bit	Route was learned through NSSA LSA and the propagate bit was set.	detail
Fwd NZ	Forwarding address is nonzero. Fwd NZ is only displayed if the route is learned through an NSSA LSA.	detail
optional-capability	Optional capabilities propagated in the router LSA. This field is in the output for intra-area router routes only (when Route Type is Area BR , AS BR , Area/AS BR , or Router), not for interarea router routes or network routes. Three bits in this field are defined as follows: <ul style="list-style-type: none"> • 0x4 (V)—Routing device is at the end of a virtual active link. • 0x2 (E)—Routing device is an autonomous system boundary router. • 0x1 (B)—Routing device is an area border router. 	detail

Table 235: show (ospf | ospf3) route Output Fields (*continued*)

Field Name	Field Description	Output Level
priority	The priority assigned to the prefix: <ul style="list-style-type: none"> • high • medium • low <p>NOTE: The priority field applies only to routes of type Network.</p>	detail

Sample Output

show ospf route

```
user@host> show ospf route
Prefix          Path    Route    NH    Metric  NextHop    Nexthop
                Type    Type      Type                Interface  addr/label
10.255.71.12     Intra  Router   IP     1       fe-0/0/2.0 192.16.22.86
10.255.71.13/32  Intra  Network  IP     0       lo0.0
192.168.222.84/30 Intra  Network  LSP    1       fe-0/0/2.0 1sp-ab
```

show ospf route detail

```
user@host> show ospf route detail
Topology default Route Table:

Prefix          Path    Route    NH    Metric  NextHop    Nexthop
                Type    Type      Type                Interface  addr/label
10.255.14.174     Inter  AS BR     IP     210     t1-3/0/1.0
  area 0.0.0.2, origin 10.255.14.185
10.255.14.178     Intra  Router   IP     200     t3-3/1/3.0
  area 0.0.0.2, origin 10.255.14.178, optional-capability 0x0
10.210.1.0/30     Intra  Network  IP     10      t3-3/1/2.0
  area 0.0.0.2, origin 10.255.14.172, priority medium
100.1.1.1/32      Inter  Network  IP     210     t1-3/0/1.0
  area 0.0.0.2, origin 10.255.14.185, priority low
112.3.1.0/24      Ext2   Network  IP     0       t1-3/0/1.0
  area 0.0.0.0, origin 10.255.14.174, priority high
200.3.3.0/30      Inter  Network  IP     220     t1-3/0/1.0
  area 0.0.0.2, origin 10.255.14.185, priority high
```

show ospf3 route

```
user@host> show ospf3 route
Prefix          Path    Route    NH    Metric  NextHop    Nexthop
                Type    Type      Type                Interface  addr/label
10.255.71.13     Intra  Router   IP     1
  NH-interface fe-0/0/2.0, NH-addr fe80::290:69ff:fe9b:e002
10.255.71.13;0.0.0.2
10.255.245.1      Intra  Router   IP     40     fxp1.1      192.168.36.17
  area 0.0.0.0, origin 10.255.245.1 optional-capability 0x0,
10.255.245.3      Intra  AS BR     IP     1      fxp2.3      192.168.36.34
  area 0.0.0.0, origin 10.255.245.3 optional-capability 0x0,
10.255.245.1/32   Intra  Network  IP     40     fxp1.1      192.168.36.17
```



```

    area 0.0.0.0, origin 10.255.245.1, priority high
10.255.245.2/32      Intra Network   IP      0  lo0.0
    area 0.0.0.0, origin 10.255.245.2, priority medium
10.255.245.3/32      Intra Network   IP      1  fxp2.3      192.168.36.34

    area 0.0.0.0, origin 10.255.245.3, priority low
        Intra Transit   IP      1
    NH-interface fe-0/0/2.0
192::168:222:84/126 Intra Network   IP      1
    NH-interface fe-0/0/2.0
abcd::71:12/128     Intra Network   IP      0
    NH-interface lo0.0
abcd::71:13/128     Intra Network   LSP     1
    NH-interface fe-0/0/2.0, NH-addr lsp-cd

```

show ospf3 route detail

```

user@host> show ospf3 route detail
Prefix                                Path   Route   NH   Metric
                                type  type   type
10.255.14.174                        Intra  Area/AS BR IP    110
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Optional-capability 0x3
10.255.14.178                        Intra  Router  IP    200
    NH-interface t3-3/1/3.0
    Area 0.0.0.0, Origin 10.255.14.178, Optional-capability 0x0
10.255.14.185;0.0.0.2                Intra  Transit IP    200
    NH-interface t1-3/0/1.0
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.185
1000:1:1::1/128                     Inter  Network IP    110
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Priority low
1001:2:1::/48                       Ext1   Network IP    110
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority medium
1002:1:7::/48                       Ext2   Network IP    0
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority low
1002:3:4::/48                       Ext2   Network IP    0
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority high
abcd::10:255:14:172/128             Intra  Network IP    0
    NH-interface lo0.0
    Area 0.0.0.0, Origin 10.255.14.172, Priority low

```

show ospf route topology voice

```

user@host> show ospf route topology voice
Topology voice Route Table:
Prefix          Path   Route   NH   Metric  NextHop      Nexthop
                Type  Type   Type
10.255.8.2      Intra  Router  IP    1    so-0/2/0.0
10.255.8.3      Intra  Router  IP    2    so-0/2/0.0
10.255.8.1/32   Intra  Network IP    0    lo0.0
10.255.8.2/32   Intra  Network IP    1    so-0/2/0.0
10.255.8.3/32   Intra  Network IP    2    so-0/2/0.0
192.168.8.0/29  Intra  Network IP    2    so-0/2/0.0
192.168.8.44/30 Intra  Network IP    2    so-0/2/0.0
192.168.8.46/32 Intra  Network IP    1    so-0/2/0.0

```

192.168.8.48/30	Intra	Network	IP	1	so-0/2/1.0
192.168.8.52/30	Intra	Network	IP	2	so-0/2/0.0
192.168.9.44/30	Intra	Network	IP	1	so-0/2/0.0
192.168.9.45/32	Intra	Network	IP	2	so-0/2/0.0

show (ospf | ospf3) statistics

Syntax	show (ospf ospf3) statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)>
Syntax (EX Series Switch and QFX Series)	show (ospf ospf3) statistics <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display OSPF statistics.
Options	<p>none—Display OSPF statistics for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display all statistics for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(Optional) (OSPFv3 only) Display all statistics for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear (ospf ospf3) statistics on page 2862
List of Sample Output	show ospf statistics on page 2895 show ospf statistics logical-system all on page 2895 show ospf3 statistics on page 2896
Output Fields	Table 236 on page 2893 lists the output fields for the show (ospf ospf3) statistics command. Output fields are listed in the approximate order in which they appear.

Table 236: show (ospf | ospf3) statistics Output Fields

Field Name	Field Description
Packet type	Type of OSPF packet.
Total Sent/Total Received	Total number of packets sent and received.
Last 5 seconds Sent/Last 5 seconds Received	Total number of packets sent and received in the last 5 seconds.

Table 236: show (ospf | ospf3) statistics Output Fields (*continued*)

Field Name	Field Description
DBDs retransmitted	Total number of database description packets retransmitted, and number retransmitted in the last 5 seconds.
LSAs flooded	Total number of link-state advertisements flooded, and number flooded in the last 5 seconds.
LSAs flooded high-prio	<p>Total number of high priority link-state advertisements flooded, and number flooded in the last 5 seconds.</p> <p>A link-state advertisement is deemed a high priority if it has changed since it was last sent.</p>
LSAs retransmitted	Total number of link-state advertisements retransmitted, and number retransmitted in the last 5 seconds.
LSAs transmitted to nbr	Total number of link-state advertisements transmitted to a neighbor, and number transmitted in the last 5 seconds.
LSAs requested	Total number of link-state advertisements requested by neighboring devices, and number requested in the last 5 seconds.
LSAs acknowledged	Total number of link-state advertisements acknowledged, and number acknowledged in the last 5 seconds.
Flood queue depth	Total number of entries in the extended queue.
Total rexmit entries	Total number of retransmission entries waiting to be sent from the OSPF routing instance.
db summaries	Total number of database description summaries waiting to be sent from the OSPF routing instance.
lsreq entries	Total number of link-state request entries waiting to be sent from the OSPF routing instance.
Receive errors	<p>Number and type of receive errors. Some sample receive errors include:</p> <ul style="list-style-type: none"> • mtu mismatches • no interface found • no virtual link found • nssa mismatches • stub area mismatches • subnet mismatches <p>If there are no receive errors, the output displays none.</p>

Sample Output

show ospf statistics

```

user@host> show ospf statistics
Packet type          Total
                   Sent      Received
Hello                31         14
  DbD                 9         10
  LSReq               2          2
LSUpdate             8         16
  LSAck              9          9
                   Last 5 seconds
                   Sent      Received
Hello                2          2
  DbD                 0          0
  LSReq               0          0
LSUpdate             0          0
  LSAck              0          0

DBDs retransmitted   :          3, last 5 seconds :          0
LSAs flooded         :         12, last 5 seconds :          0
LSAs flooded high-prio :          0, last 5 seconds :          0
LSAs retransmitted   :          0, last 5 seconds :          0
LSAs transmitted to nbr:          3, last 5 seconds :          0
LSAs requested       :          5, last 5 seconds :          0
LSAs acknowledged    :         19, last 5 seconds :          0

Flood queue depth    :          0
Total rexmit entries :          0
db summaries         :          0
lsreq entries        :          0

Receive errors:
  862 no interface found
  115923 no virtual link found

```

show ospf statistics logical-system all

```

user@host> show ospf statistics logical-system all
logical-system: C
OSPF instance is not running
-----

logical-system: B
Packet type          Total
                   Sent      Received
Hello              313740      313653
  DbD                3          2
  LSReq              1          1
LSUpdate           2752       1825
  LSAck            1821       2747
                   Last 5 seconds
                   Sent      Received
Hello                1          0
  DbD                 0          0
  LSReq               0          0
LSUpdate            0          0
  LSAck              0          0

DBDs retransmitted   :          0, last 5 seconds :          0
LSAs flooded         :        2741, last 5 seconds :          0
LSAs flooded high-prio :         10, last 5 seconds :          0
LSAs retransmitted   :          0, last 5 seconds :          0
LSAs transmitted to nbr:          2, last 5 seconds :          0
LSAs requested       :          1, last 5 seconds :          0
LSAs acknowledged    :       1831, last 5 seconds :          0

Flood queue depth    :          0
Total rexmit entries :          0
db summaries         :          0
lsreq entries        :          0

Receive errors:

```

```

None
-----

logical-system: A

Packet type          Total          Last 5 seconds
                   Sent      Received      Sent      Received
Hello                313698      313695         0         0
  DbD                  2         3         0         0
  LSReq                1         1         0         0
LSUpdate             1825      2752         0         0
LSAck                2747      1821         0         0

DBDs retransmitted   :                0, last 5 seconds :      0
LSAs flooded         :             1825, last 5 seconds :      0
LSAs flooded high-prio :             10, last 5 seconds :      0
LSAs retransmitted   :                0, last 5 seconds :      0
LSAs transmitted to nbr:             1, last 5 seconds :      0
LSAs requested       :                2, last 5 seconds :      0
LSAs acknowledged   :             2748, last 5 seconds :      0

Flood queue depth    :                0
Total rexmit entries :                0
db summaries         :                0
lsreq entries        :                0

Receive errors:
None
-----

```

show ospf3 statistics

```

user@host> show ospf3 statistics

Packet type          Total          Last 5 seconds
                   Sent      Received      Sent      Received
Hello                0         0         0         0
  DbD                  0         0         0         0
  LSReq                0         0         0         0
LSUpdate             0         0         0         0
LSAck                0         0         0         0

DBDs retransmitted   :                0, last 5 seconds :      0
LSAs flooded         :                0, last 5 seconds :      0
LSAs flooded high-prio :                0, last 5 seconds :      0
LSAs retransmitted   :                0, last 5 seconds :      0
LSAs transmitted to nbr:                0, last 5 seconds :      0
LSAs requested       :                0, last 5 seconds :      0
LSAs acknowledged   :                0, last 5 seconds :      0

Flood queue depth    :                0
Total rexmit entries :                0
db summaries         :                0
lsreq entries        :                0

Receive errors:
None

```

show ospf database

Syntax	<pre>show ospf database <brief detail extensive summary> <advertising-router (address self)> <area area-id> <asbrsummary> <external> <instance instance-name> <link-local> <logical-system (all logical-system-name)> <lsa-id lsa-id> <netsummary> <network> <nssa> <opaque-area> <router></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show ospf database <brief detail extensive summary> <advertising-router (address self)> <area area-id> <asbrsummary> <external> <instance instance-name> <link-local> <lsa-id lsa-id> <netsummary> <network> <nssa> <opaque-area> <router></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>advertising-router self (address self) option introduced in Junos OS Release 9.5.</p> <p>advertising-router self (address self) option introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display the entries in the OSPF version 2 (OSPFv2) link-state database, which contains data about link-state advertisement (LSA) packets.
Options	<p>none—Display standard information about entries in the OSPFv2 link-state database for all routing instances.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>advertising-router (address self)—(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.</p> <p>area area-id—(Optional) Display the LSAs in a particular area.</p>

asbrsummary—(Optional) Display summary AS boundary router LSA entries.

external—(Optional) Display external LSAs.

instance *instance-name*—(Optional) Display all OSPF database information under the named routing instance.

link-local—(Optional) Display information about link-local LSAs.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsa-id *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.

netsummary—(Optional) Display summary network LSAs.

network—(Optional) Display information about network LSAs.

nssa—(Optional) Display information about not-so-stubby area (NSSA) LSAs.

opaque-area—(Optional) Display opaque area-scope LSAs.

router—(Optional) Display information about router LSAs.

Required Privilege Level

view

Related Documentation

- [clear \(ospf | ospf3\) database on page 2856](#)

List of Sample Output

[show ospf database on page 2900](#)
[show ospf database brief on page 2900](#)
[show ospf database detail on page 2900](#)
[show ospf database extensive on page 2902](#)
[show ospf database summary on page 2904](#)

Output Fields

[Table 237 on page 2898](#) describes the output fields for the **show ospf database** command. Output fields are listed in the approximate order in which they appear.

Table 237: show ospf database Output Fields

Field Name	Field Description	Level of Output
area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Type	Type of link advertisement: ASBRSum , Extern , Network , NSSA , OpaqArea , Router , or Summary .	All levels
ID	LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device.	All levels
Adv Rtr	Address of the routing device that sent the advertisement.	All levels
Seq	Link sequence number of the advertisement.	All levels

Table 237: show ospf database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Age	Time elapsed since the LSA was originated, in seconds.	All levels
Opt	Optional OSPF capabilities associated with the LSA.	All levels
Cksum	Checksum value of the LSA.	All levels
Len	Length of the advertisement, in bytes.	All levels
Router	Router link-state advertisement information: <ul style="list-style-type: none"> • bits—Flags describing the routing device that generated the LSP. • link count—Number of links in the advertisement. • id—ID of a routing device or subnet on the link. • data—For stub networks, the subnet mask. Otherwise, the IP address of the routing device that generated the LSP. • type—Type of link. It can be PointToPoint, Transit, Stub, or Virtual. • TOS count—Number of type-of-service (ToS) entries in the advertisement. • TOS 0 metric—Metric for ToS 0. • TOS—Type-of-service (ToS) value. • metric—Metric for the ToS. 	detail extensive
Network	Network link-state advertisement information: <ul style="list-style-type: none"> • mask—Network mask. • attached router—ID of the attached neighbor. 	detail extensive
Summary	Summary link-state advertisement information: <ul style="list-style-type: none"> • mask—Network mask. • TOS—Type-of-service (ToS) value. • metric—Metric for the ToS. 	detail extensive
Gen timer	How long until the LSA is regenerated.	extensive
Aging timer	How long until the LSA expires.	extensive
Installed <i>hh:mm:ss</i> ago	How long ago the route was installed.	extensive
expires in <i>hh:mm:ss</i>	How long until the route expires.	extensive
sent <i>hh:mm:ss</i> ago	How long ago the LSA was sent.	extensive
Last changed <i>hh:mm:ss</i> ago	How long ago the route was changed.	extensive
Change count	Number of times the route has changed.	extensive

Table 237: show ospf database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Ours	Indicates that this is a local advertisement.	extensive
Router LSAs	Number of router link-state advertisements in the link-state database.	summary
Network LSAs	Number of network link-state advertisements in the link-state database.	summary
Summary LSAs	Number of summary link-state advertisements in the link-state database.	summary
NSSA LSAs	Number of not-so-stubby area link-state advertisements in the link-state database.	summary

Sample Output

show ospf database

```

user@host> show ospf database
OSPF link state database, Area 0.0.0.1
  Type      ID            Adv Rtr      Seq          Age  Opt  Cksum  Len
Router     10.255.70.103   10.255.70.103 0x80000002   215  0x20 0x4112  48
Router     *10.255.71.242  10.255.71.242 0x80000002   214  0x20 0x11b1  48
Summary    *23.1.1.0       10.255.71.242 0x80000002   172  0x20 0x6d72  28
Summary    *24.1.1.0       10.255.71.242 0x80000002   177  0x20 0x607e  28
NSSA       *33.1.1.1       10.255.71.242 0x80000002   217  0x28 0x73bd  36

      OSPF link state database, Area 0.0.0.2
  Type      ID            Adv Rtr      Seq          Age  Opt  Cksum  Len
Router     10.255.71.52   10.255.71.52  0x80000004   174  0x20 0xd021  36
Router     *10.255.71.242  10.255.71.242 0x80000003   173  0x20 0xe191  36
Network    *23.1.1.1       10.255.71.242 0x80000002   173  0x20 0x9c76  32
Summary    *12.1.1.0       10.255.71.242 0x80000001   217  0x20 0xfeec  28
Summary    *24.1.1.0       10.255.71.242 0x80000002   177  0x20 0x607e  28
NSSA       *33.1.1.1       10.255.71.242 0x80000001   222  0x28 0xe047  36

      OSPF link state database, Area 0.0.0.3
  Type      ID            Adv Rtr      Seq          Age  Opt  Cksum  Len
Router     10.255.71.238   10.255.71.238 0x80000003   179  0x20 0x3942  36
Router     *10.255.71.242  10.255.71.242 0x80000003   177  0x20 0xf37d  36
Network    *24.1.1.1       10.255.71.242 0x80000002   177  0x20 0xc591  32
Summary    *12.1.1.0       10.255.71.242 0x80000001   217  0x20 0xfeec  28
Summary    *23.1.1.0       10.255.71.242 0x80000002   172  0x20 0x6d72  28
NSSA       *33.1.1.1       10.255.71.242 0x80000001   222  0x28 0xeb3b  36

```

show ospf database brief

The output for the **show ospf database brief** command is identical to that for the **show ospf database** command. For sample output, see [show ospf database on page 2900](#).

show ospf database detail

```

user@host> show ospf database detail
OSPF link state database, Area 0.0.0.1
  Type      ID            Adv Rtr      Seq          Age  Opt  Cksum  Len
Router     10.255.70.103   10.255.70.103 0x80000002   261  0x20 0x4112  48

```

```

bits 0x0, link count 2
id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
TOS count 0, TOS 0 metric 1
id 12.1.1.0, data 255.255.255.0, Type Stub (3)
TOS count 0, TOS 0 metric 1
Router *10.255.71.242 10.255.71.242 0x80000002 260 0x20 0x11b1 48
bits 0x3, link count 2
id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
TOS count 0, TOS 0 metric 1
id 12.1.1.0, data 255.255.255.0, Type Stub (3)
TOS count 0, TOS 0 metric 1
Summary *23.1.1.0 10.255.71.242 0x80000002 218 0x20 0x6d72 28
mask 255.255.255.0
TOS 0x0, metric 1
Summary *24.1.1.0 10.255.71.242 0x80000002 223 0x20 0x607e 28
mask 255.255.255.0
TOS 0x0, metric 1
NSSA *33.1.1.1 10.255.71.242 0x80000002 263 0x28 0x73bd 36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0

```

OSPF link state database, Area 0.0.0.2

Type	ID	Adv Rtr	Seq	Age	Opt	Cksum	Len
Router	10.255.71.52	10.255.71.52	0x80000004	220	0x20	0xd021	36
bits 0x0, link count 1							
id 23.1.1.1, data 23.1.1.2, Type Transit (2)							
TOS count 0, TOS 0 metric 1							
Router	*10.255.71.242	10.255.71.242	0x80000003	219	0x20	0xe191	36
bits 0x3, link count 1							
id 23.1.1.1, data 23.1.1.1, Type Transit (2)							
TOS count 0, TOS 0 metric 1							
Network	*23.1.1.1	10.255.71.242	0x80000002	219	0x20	0x9c76	32
mask 255.255.255.0							
attached router 10.255.71.242							
attached router 10.255.71.52							
Summary	*12.1.1.0	10.255.71.242	0x80000001	263	0x20	0xfeec	28
mask 255.255.255.0							
TOS 0x0, metric 1							
Summary	*24.1.1.0	10.255.71.242	0x80000002	223	0x20	0x607e	28
mask 255.255.255.0							
TOS 0x0, metric 1							
NSSA	*33.1.1.1	10.255.71.242	0x80000001	268	0x28	0xe047	36
mask 255.255.255.255							
Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0							

OSPF link state database, Area 0.0.0.3

Type	ID	Adv Rtr	Seq	Age	Opt	Cksum	Len
Router	10.255.71.238	10.255.71.238	0x80000003	225	0x20	0x3942	36
bits 0x0, link count 1							
id 24.1.1.1, data 24.1.1.2, Type Transit (2)							
TOS count 0, TOS 0 metric 1							
Router	*10.255.71.242	10.255.71.242	0x80000003	223	0x20	0xf37d	36
bits 0x3, link count 1							
id 24.1.1.1, data 24.1.1.1, Type Transit (2)							
TOS count 0, TOS 0 metric 1							
Network	*24.1.1.1	10.255.71.242	0x80000002	223	0x20	0xc591	32
mask 255.255.255.0							
attached router 10.255.71.242							
attached router 10.255.71.238							
Summary	*12.1.1.0	10.255.71.242	0x80000001	263	0x20	0xfeec	28
mask 255.255.255.0							

```

TOS 0x0, metric 1
Summary *23.1.1.0      10.255.71.242    0x80000002    218  0x20 0x6d72  28
mask 255.255.255.0
TOS 0x0, metric 1
NSSA  *33.1.1.1      10.255.71.242    0x80000001    268  0x28 0xeb3b  36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0

```

show ospf database extensive

```

user@host> show ospf database extensive
  OSPF link state database, Area 0.0.0.1
  Type      ID          Adv Rtr          Seq      Age  Opt  Cksum  Len
Router 10.255.70.103  10.255.70.103    0x80000002    286  0x20 0x4112  48
  bits 0x0, link count 2
  id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
  Aging timer 00:55:14
  Installed 00:04:43 ago, expires in 00:55:14
  Last changed 00:04:43 ago, Change count: 2
Router *10.255.71.242  10.255.71.242    0x80000002    285  0x20 0x11b1  48
  bits 0x3, link count 2
  id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
  Gen timer 00:45:15
  Aging timer 00:55:15
  Installed 00:04:45 ago, expires in 00:55:15, sent 00:04:43 ago
  Last changed 00:04:45 ago, Change count: 2, Ours
Summary *23.1.1.0      10.255.71.242    0x80000002    243  0x20 0x6d72  28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:57
Aging timer 00:55:57
Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0      10.255.71.242    0x80000002    248  0x20 0x607e  28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
Last changed 00:04:48 ago, Change count: 1, Ours
NSSA  *33.1.1.1      10.255.71.242    0x80000002    288  0x28 0x73bd  36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0
Gen timer 00:45:12
Aging timer 00:55:12
Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:48 ago
Last changed 00:04:48 ago, Change count: 2, Ours

  OSPF link state database, Area 0.0.0.2
  Type      ID          Adv Rtr          Seq      Age  Opt  Cksum  Len
Router 10.255.71.52   10.255.71.52     0x80000004    245  0x20 0xd021  36
  bits 0x0, link count 1
  id 23.1.1.1, data 23.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
  Aging timer 00:55:55

```

```

    Installed 00:04:02 ago, expires in 00:55:55
    Last changed 00:04:02 ago, Change count: 2
Router *10.255.71.242    10.255.71.242    0x80000003    244    0x20    0xe191    36
    bits 0x3, link count 1
    id 23.1.1.1, data 23.1.1.1, Type Transit (2)
    TOS count 0, TOS 0 metric 1
    Gen timer 00:45:56
    Aging timer 00:55:56
    Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
    Last changed 00:04:04 ago, Change count: 2, Ours
Network *23.1.1.1        10.255.71.242    0x80000002    244    0x20    0x9c76    32
    mask 255.255.255.0
    attached router 10.255.71.242
    attached router 10.255.71.52
    Gen timer 00:45:56
    Aging timer 00:55:56
    Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
    Last changed 00:04:04 ago, Change count: 1, Ours
Summary *12.1.1.0        10.255.71.242    0x80000001    288    0x20    0xfeec    28
    mask 255.255.255.0
    TOS 0x0, metric 1
    Gen timer 00:45:12
    Aging timer 00:55:12
    Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:04 ago
    Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0        10.255.71.242    0x80000002    248    0x20    0x607e    28
    mask 255.255.255.0
    TOS 0x0, metric 1
    Gen timer 00:45:52
    Aging timer 00:55:52
    Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:04 ago
    Last changed 00:04:48 ago, Change count: 1, Ours
NSSA *33.1.1.1          10.255.71.242    0x80000001    293    0x28    0xe047    36
    mask 255.255.255.255
    Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0
    Gen timer 00:45:07
    Aging timer 00:55:07
    Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:04 ago
    Last changed 00:04:53 ago, Change count: 1, Ours

```

OSPF link state database, Area 0.0.0.3

Type	ID	Adv Rtr	Seq	Age	Opt	Cksum	Len
Router	10.255.71.238	10.255.71.238	0x80000003	250	0x20	0x3942	36
bits 0x0, link count 1 id 24.1.1.1, data 24.1.1.2, Type Transit (2) TOS count 0, TOS 0 metric 1 Aging timer 00:55:50 Installed 00:04:07 ago, expires in 00:55:50 Last changed 00:04:07 ago, Change count: 2							
Router	*10.255.71.242	10.255.71.242	0x80000003	248	0x20	0xf37d	36
bits 0x3, link count 1 id 24.1.1.1, data 24.1.1.1, Type Transit (2) TOS count 0, TOS 0 metric 1 Gen timer 00:45:52 Aging timer 00:55:52 Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago Last changed 00:04:08 ago, Change count: 2, Ours							
Network	*24.1.1.1	10.255.71.242	0x80000002	248	0x20	0xc591	32
mask 255.255.255.0 attached router 10.255.71.242 attached router 10.255.71.238							

```
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
Last changed 00:04:08 ago, Change count: 1, Ours
Summary *12.1.1.0      10.255.71.242    0x80000001    288  0x20 0xfeec  28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:12
Aging timer 00:55:12
Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:13 ago
Last changed 00:04:48 ago, Change count: 1, Ours
Summary *23.1.1.0      10.255.71.242    0x80000002    243  0x20 0x6d72  28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:57
Aging timer 00:55:57
Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
Last changed 00:04:48 ago, Change count: 1, Ours
NSSA  *33.1.1.1        10.255.71.242    0x80000001    293  0x28 0xeb3b  36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0
Gen timer 00:45:07
Aging timer 00:55:07
Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:13 ago
Last changed 00:04:53 ago, Change count: 1, Ours
```

show ospf database summary

```
user@host> show ospf database summary
Area 0.0.0.1:
  2 Router LSAs
  2 Summary LSAs
  1 NSSA LSAs
Area 0.0.0.2:
  2 Router LSAs
  1 Network LSAs
  2 Summary LSAs
  1 NSSA LSAs
Area 0.0.0.3:
  2 Router LSAs
  1 Network LSAs
  2 Summary LSAs
  1 NSSA LSAs
Externals:
Interface fe-2/2/1.0:
Interface ge-0/3/2.0:
Interface so-0/1/2.0:
Interface so-0/1/2.0:
```

show ospf3 database

Syntax	<pre>show ospf3 database <brief detail extensive summary> <advertising-router (address self)> <area area-id> <external> <instance instance-name> <inter-area-prefix> <inter-area-router> <intra-area-prefix> <link> <link-local> <logical-system (all logical-system-name)> <lsa-id lsa-id> <network> <nssa> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> <router></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show ospf3 database <brief detail extensive summary> <advertising-router (address self)> <area area-id> <external> <instance instance-name> <inter-area-prefix> <inter-area-router> <intra-area-prefix> <link> <link-local> <lsa-id lsa-id> <network> <nssa> <router></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>realm option introduced in Junos OS Release 9.2.</p> <p>advertising-router (address self) option introduced in Junos Release 9.5.</p> <p>advertising-router (address self) option introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display the entries in the OSPF version 3 (OSPFv3) link-state database, which contains data about link-state advertisement (LSA) packets.
Options	<p>none—Display standard information about all entries in the OSPFv3 link-state database.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>advertising-router (address self)—(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.</p>

area *area-id*—(Optional) Display the LSAs in a particular area.

external—(Optional) Display external LSAs.

instance *instance-name*—(Optional) Display all OSPF database information under the named routing instance.

inter-area-prefix—(Optional) Display information about interarea-prefix LSAs.

inter-area-router—(Optional) Display information about interarea-router LSAs.

intra-area-prefix—(Optional) Display information about intra-area-prefix LSAs.

link—(Optional) Display information about link LSAs.

link-local—(Optional) Display information about link-local LSAs.

logical-system (**all** | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsa-id *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.

network—(Optional) Display information about network LSAs.

nssa—(Optional) Display information about not-so-stubby area (NSSA) LSAs.

realm (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(Optional) Display information about the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family other than IPv6 unicast, which is the default.

router—(Optional) Display information about router LSAs.

Required Privilege Level view

Related Documentation • [clear \(ospf | ospf3\) database on page 2856](#)

List of Sample Output [show ospf3 database brief on page 2911](#)
[show ospf3 database extensive on page 2911](#)
[show ospf3 database summary on page 2914](#)

Output Fields [Table 238 on page 2906](#) lists the output fields for the **show ospf3 database** command. Output fields are listed in the approximate order in which they appear.

Table 238: show ospf3 database Output Fields

Field Name	Field Description	Level of Output
OSPF link state database, area <i>area-number</i>	Entries in the link-state database for this area.	brief detail extensive
OSPF AS SCOPE link state database	Entries in the AS scope link-state database.	brief detail extensive

Table 238: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
OSPF Link-Local link state database, interface <i>interface-name</i>	Entries in the link-local link-state database for this interface.	brief detail extensive
area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Type	Type of link advertisement: Extern , InterArPfx , InterArRtr , IntraArPrx , Link , Network , NSSA , or Router .	brief detail extensive
ID	Link identifier included in the advertisement. An asterisk (*) preceding the identifier marks database entries that originated from the local routing device.	brief detail extensive
Adv Rtr	Address of the routing device that sent the advertisement.	brief detail extensive
Seq	Link sequence number of the advertisement.	brief detail extensive
Age	Time elapsed since the LSA was originated, in seconds.	brief detail extensive
Cksum	Checksum value of the LSA.	brief detail extensive
Len	Length of the advertisement, in bytes.	brief detail extensive
Router (Router Link-State Advertisements)		
bits	Flags describing the routing device that generated the LSP.	detail extensive
Options	Option bits carried in the router LSA.	detail extensive
For Each Router Link		
Type	Type of interface. The value of all other output fields describing a routing device interface depends on the interface's type: <ul style="list-style-type: none"> • PointToPoint (1)—Point-to-point connection to another routing device. • Transit (2)—Connection to a transit network. • Virtual (4)—Virtual link. 	detail extensive
Loc-if-id	Local interface ID assigned to the interface that uniquely identifies the interface with the routing device.	detail extensive
Nbr-if-id	Interface ID of the neighbor's interface for this routing device link.	detail extensive
Nbr-rtr-id	Router ID of the neighbor routing device (for type 2 interfaces, the attached link's designated router).	detail extensive
Metric	Cost of the router link.	detail extensive
Gen timer	How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .	extensive

Table 238: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive
Installed <i>nn:nn:nn</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>nn:nn:nn</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>nn:nn:nn</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
Ours	Indicates that this is a local advertisement.	extensive
Network (Network Link-State Advertisements)		
Options	Option bits carried in the network LSA.	detail extensive
Attached Router	Router IDs of each of the routing devices attached to the link. Only routing devices that are fully adjacent to the designated router are listed. The designated router includes itself in this list.	detail extensive
InterArPfx (Interarea-Prefix Link-State Advertisements)		
Prefix	IPv6 address prefix.	detail extensive
Prefix-options	Option bit associated with the prefix.	detail extensive
Metric	Cost of this route. Expressed in the same units as the interface costs in the router LSAs. When the interarea-prefix LSA is describing a route to a range of addresses, the cost is set to the maximum cost to any reachable component of the address range.	detail extensive
Gen timer	How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .	extensive
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive
Installed <i>nn:nn:nn</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>nn:nn:nn</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>nn:nn:nn</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
Ours	Indicates that this is a local advertisement.	extensive
InterArRtr (Interarea-Router Link-State Advertisements)		
Dest-router-id	Router ID of the routing device described by the LSA.	detail extensive
options	Optional capabilities supported by the routing device.	detail extensive

Table 238: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Metric	Cost of this route. Expressed in the same units as the interface costs in the router LSAs. When the interarea-prefix LSA is describing a route to a range of addresses, the cost is set to the maximum cost to any reachable component of the address range.	detail extensive
Prefix	IPv6 address prefix.	extensive
Prefix-options	Option bit associated with the prefix.	extensive
Extern (External Link-State Advertisements)		
Prefix	IPv6 address prefix.	detail extensive
Prefix-options	Option bit associated with the prefix.	detail extensive
Metric	Cost of the route, which depends on the value of Type .	detail extensive
Type <i>n</i>	Type of external metric: Type 1 or Type 2 .	detail extensive
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive
Installed <i>nn:nn:nn</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>nn:nn:nn</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>nn:nn:nn</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
Link (Link-State Advertisements)		
IPv6-Address	IPv6 link-local address on the link for which this link LSA originated.	detail extensive
Options	Option bits carried in the link LSA.	detail extensive
priority	Router priority of the interface attaching the originating routing device to the link.	detail extensive
Prefix-count	Number of IPv6 address prefixes contained in the LSA. The rest of the link LSA contains a list of IPv6 prefixes to be associated with the link.	detail extensive
Prefix	IPv6 address prefix.	detail extensive
Prefix-options	Option bit associated with the prefix.	detail extensive
Gen timer	How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .	extensive
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive

Table 238: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Installed <i>nn:nn:nn</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>nn:nn:nn</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>nn:nn:nn</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
Ours	Indicates that this is a local advertisement.	extensive
IntraArPfx (Intra-Area-Prefix Link-State Advertisements)		
Ref-lsa-type	LSA type of the referenced LSA. <ul style="list-style-type: none"> Router—Address prefixes are associated with a router LSA. Network—Address prefixes are associated with a network LSA. 	detail extensive
Ref-lsa-id	Link-state ID of the referenced LSA.	detail extensive
Ref-router-id	Advertising router ID of the referenced LSA.	detail extensive
Prefix-count	Number of IPv6 address prefixes contained in the LSA. The rest of the link LSA contains a list of IPv6 prefixes to be associated with the link.	detail extensive
Prefix	IPv6 address prefix.	detail extensive
Prefix-options	Option bit associated with the prefix.	detail extensive
Metric	Cost of this prefix. Expressed in the same units as the interface costs in the router LSAs.	detail extensive
Gen timer	How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .	extensive
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive
Installed <i>hh:mm:ss</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>hh:mm:ss</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>hh:mm:ss</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
<i>n</i> Router LSAs	Number of router LSAs in the link-state database.	summary
<i>n</i> Network LSAs	Number of network LSAs in the link-state database.	summary
<i>n</i> InterArPfx LSAs	Number of interarea-prefix LSAs in the link-state database.	summary

Table 238: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
<i>n</i> InterArRtr LSAs	Number of interarea-router LSAs in the link-state database.	summary
<i>n</i> IntraArPfx LSAs	Number of intra-area-prefix LSAs in the link-state database.	summary
Externals	Display of the external LSA database.	summary
<i>n</i> Extern LSAs	Number of external LSAs in the link-state database.	summary
Interface <i>interface-name</i>	Name of the interface for which link-local LSA information is displayed.	summary
<i>n</i> Link LSAs	Number of link LSAs in the link-state database.	summary

Sample Output

show ospf3 database brief

```

user@host> show ospf3 database brief
      OSPF3 link state database, area 0.0.0.0
      Type      ID          Adv Rtr      Seq          Age    Cksum  Len
      Router    0.0.0.1      10.255.4.85  0x80000003   885    0xa697  40
      Router    *0.0.0.1     10.255.4.93  0x80000002   953    0xc677  40
      InterArPfx *0.0.0.2     10.255.4.93  0x80000001   910    0xb96f  44
      InterArRtr *0.0.0.1     10.255.4.93  0x80000001   910    0xe159  32
      IntraArPfx *0.0.0.1     10.255.4.93  0x80000002   432    0x788f  72

      OSPF3 link state database, area 0.0.0.1
      Type      ID          Adv Rtr      Seq          Age    Cksum  Len
      Router    *0.0.0.1     10.255.4.93  0x80000003   916    0xea40  40
      Router    0.0.0.1     10.255.4.97  0x80000006   851    0xc95b  40
      Network    0.0.0.2     10.255.4.97  0x80000002   916    0x4598  32
      InterArPfx *0.0.0.1     10.255.4.93  0x80000002   117    0xa980  44
      InterArPfx *0.0.0.2     10.255.4.93  0x80000002    62    0xd47e  44
      NSSA       0.0.0.1     10.255.4.97  0x80000002   362    0x45ee  44
      IntraArPfx 0.0.0.1     10.255.4.97  0x80000006   851    0x2f77  52

      OSPF3 AS SCOPE link state database
      Type      ID          Adv Rtr      Seq          Age    Cksum  Len
      Extern     0.0.0.1     10.255.4.85  0x80000002    63    0x9b86  44
      Extern     *0.0.0.1     10.255.4.93  0x80000001   910    0x59c9  44

      OSPF3 Link-Local link state database, interface ge-1/3/0.0
      Type      ID          Adv Rtr      Seq          Age    Cksum  Len
      Link       *0.0.0.2     10.255.4.93  0x80000003   916    0x4dab  64

```

show ospf3 database extensive

```

user@host> show ospf3 database extensive
      OSPF3 link state database, area 0.0.0.0
      Type      ID          Adv Rtr      Seq          Age    Cksum  Len
      Router    0.0.0.1     10.255.4.85  0x80000003  1028    0xa697  40
      bits 0x2, Options 0x13
      Type PointToPoint (1), Metric 10

```

```

    Loc-If-Id 2, Nbr-If-Id 3, Nbr-Rtr-Id 10.255.4.93
    Aging timer 00:42:51
    Installed 00:17:05 ago, expires in 00:42:52, sent 02:37:54 ago
Router    *0.0.0.1          10.255.4.93      0x80000002  1096  0xc677  40
    bits 0x3, Options 0x13
    Type PointToPoint (1), Metric 10
    Loc-If-Id 3, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.85
    Gen timer 00:00:40
    Aging timer 00:41:44
    Installed 00:18:16 ago, expires in 00:41:44, sent 00:18:14 ago
    Ours
InterArPfx *0.0.0.2          10.255.4.93      0x80000001  1053  0xb96f  44
    Prefix feee::10:10:2:0/126
    Prefix-options 0x0, Metric 10
    Gen timer 00:17:02
    Aging timer 00:42:26
    Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
    Ours
InterArPfx *0.0.0.3          10.255.4.93      0x80000001  1053  0x71d3  44
    Prefix feee::10:255:4:97/128
    Prefix-options 0x0, Metric 10
    Gen timer 00:21:07
    Aging timer 00:42:26
    Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
    Ours
InterArRtr *0.0.0.1          10.255.4.93      0x80000001  1053  0xe159  32
    Dest-router-id 10.255.4.97, Options 0x19, Metric 10
    Gen timer 00:29:18
    Aging timer 00:42:26
    Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
    Ours
IntraArPfx 0.0.0.1          10.255.4.85      0x80000002  1028  0x2403  72
    Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.85
    Prefix-count 2
    Prefix feee::10:255:4:85/128
    Prefix-options 0x2, Metric 0
    Prefix feee::10:10:1:0/126
    Prefix-options 0x0, Metric 10
    Aging timer 00:42:51
    Installed 00:17:05 ago, expires in 00:42:52, sent 02:37:54 ago
IntraArPfx *0.0.0.1          10.255.4.93      0x80000002  575  0x788f  72
    Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.93
    Prefix-count 2
    Prefix feee::10:255:4:93/128
    Prefix-options 0x2, Metric 0
    Prefix feee::10:10:1:0/126
    Prefix-options 0x0, Metric 10
    Gen timer 00:33:23
    Aging timer 00:50:24
    Installed 00:09:35 ago, expires in 00:50:25, sent 00:09:33 ago
    OSPF3 link state database, area 0.0.0.1
Type      ID              Adv Rtr          Seq            Age  Cksum  Len
Router    *0.0.0.1          10.255.4.93      0x80000003  1059  0xea40  40
    bits 0x3, Options 0x19
    Type Transit (2), Metric 10
    Loc-If-Id 2, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.97
    Gen timer 00:08:51
    Aging timer 00:42:20
    Installed 00:17:39 ago, expires in 00:42:21, sent 00:17:37 ago
Router     0.0.0.1          10.255.4.97      0x80000006  994  0xc95b  40
    bits 0x2, Options 0x19

```

```

Type Transit (2), Metric 10
  Loc-If-Id 2, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.97
Aging timer 00:43:25
  Installed 00:16:31 ago, expires in 00:43:26, sent 02:37:54 ago
Network    0.0.0.2          10.255.4.97      0x80000002  1059  0x4598  32
Options 0x11
  Attached router 10.255.4.97
  Attached router 10.255.4.93
Aging timer 00:42:20
  Installed 00:17:36 ago, expires in 00:42:21, sent 02:37:54 ago
InterArPfx *0.0.0.1      10.255.4.93      0x80000002   260  0xa980  44
  Prefix feee::10:10:1:0/126
  Prefix-options 0x0, Metric 10
  Gen timer 00:45:39
  Aging timer 00:55:39
  Installed 00:04:20 ago, expires in 00:55:40, sent 00:04:18 ago
  Ours
InterArPfx *0.0.0.2      10.255.4.93      0x80000002   205  0xd47e  44
  Prefix feee::10:255:4:93/128
  Prefix-options 0x0, Metric 0
  Gen timer 00:46:35
  Aging timer 00:56:35
  Installed 00:03:25 ago, expires in 00:56:35, sent 00:03:23 ago
  Ours
InterArPfx *0.0.0.3      10.255.4.93      0x80000001  1089  0x9bbb  44
  Prefix feee::10:255:4:85/128
  Prefix-options 0x0, Metric 10
  Gen timer 00:04:46
  Aging timer 00:41:51
  Installed 00:18:09 ago, expires in 00:41:51, sent 00:17:43 ago
  Ours
NSSA      0.0.0.1          10.255.4.97      0x80000002   505  0x45ee  44
  Prefix feee::200:200:1:0/124
  Prefix-options 0x8, Metric 10, Type 2,
  Aging timer 00:51:35
  Installed 00:08:22 ago, expires in 00:51:35, sent 02:37:54 ago
IntraArPfx 0.0.0.1      10.255.4.97      0x80000006   994  0x2f77  52
  Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.97
  Prefix-count 1
  Prefix feee::10:255:4:97/128
  Prefix-options 0x2, Metric 0
  Aging timer 00:43:25
  Installed 00:16:31 ago, expires in 00:43:26, sent 02:37:54 ago
IntraArPfx 0.0.0.3      10.255.4.97      0x80000002  1059  0x4446  52
  Ref-lsa-type Network, Ref-lsa-id 0.0.0.2, Ref-router-id 10.255.4.97
  Prefix-count 1
  Prefix feee::10:10:2:0/126
  Prefix-options 0x0, Metric 0
  Aging timer 00:42:20
  Installed 00:17:36 ago, expires in 00:42:21, sent 02:37:54 ago
  OSPF3 AS SCOPE link state database
Type      ID          Adv Rtr          Seq          Age  Cksum  Len
Extern    0.0.0.1          10.255.4.85      0x80000002   206  0x9b86  44
  Prefix feee::100:100:1:0/124
  Prefix-options 0x0, Metric 20, Type 2,
  Aging timer 00:56:34
  Installed 00:03:23 ago, expires in 00:56:34, sent 02:37:54 ago
Extern    *0.0.0.1          10.255.4.93      0x80000001  1053  0x59c9  44
  Prefix feee::200:200:1:0/124
  Prefix-options 0x0, Metric 10, Type 2,
  Gen timer 00:25:12

```

```

Aging timer 00:42:26
Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago

    OSPF3 Link-Local link state database, interface ge-1/3/0.0
    Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Link         *0.0.0.2      10.255.4.93  0x80000003   1059 0x4dab 64
fe80::290:69ff:fe39:1cdb
Options 0x11, priority 128
Prefix-count 1
Prefix feee::10:10:2:0/126 Prefix-options 0x0
Gen timer 00:12:56
Aging timer 00:42:20
Installed 00:17:39 ago, expires in 00:42:21, sent 00:17:37 ago
Link         0.0.0.2      10.255.4.97  0x80000003   205  0xa87d 64
fe80::290:69ff:fe38:883e
Options 0x11, priority 128
Prefix-count 1
Prefix feee::10:10:2:0/126 Prefix-options 0x0
Aging timer 00:56:35
Installed 00:03:22 ago, expires in 00:56:35, sent 02:37:54 ago

    OSPF3 Link-Local link state database, interface so-2/2/0.0
    Type      ID          Adv Rtr      Seq          Age  Cksum  Len
Link         0.0.0.2      10.255.4.85  0x80000002   506  0x42bb 64
fe80::280:42ff:fe10:f169
Options 0x13, priority 128
Prefix-count 1
Prefix feee::10:10:1:0/126 Prefix-options 0x0
Aging timer 00:51:34
Installed 00:08:23 ago, expires in 00:51:34, sent 02:37:54 ago
Link         *0.0.0.3      10.255.4.93  0x80000002   505  0x6b7a 64
fe80::280:42ff:fe10:f177
Options 0x13, priority 128
Prefix-count 1
Prefix feee::10:10:1:0/126 Prefix-options 0x0
Gen timer 00:37:28
Aging timer 00:51:35
Installed 00:08:25 ago, expires in 00:51:35, sent 00:08:23 ago
Ours

```

show ospf3 database summary

```

user@host> show ospf3 database summary
Area 0.0.0.0:
  2 Router LSAs
  1 InterArPfx LSAs
  1 InterArRtr LSAs
  1 IntraArPfx LSAs
Area 0.0.0.1:
  2 Router LSAs
  1 Network LSAs
  2 InterArPfx LSAs
  1 NSSA LSAs
  1 IntraArPfx LSAs
Externals:
  2 Extern LSAs
Interface ge-1/3/0.0:
  1 Link LSAs
Interface lo0.0:

```


Interface so-2/2/0.0:
1 Link LSAs

RIP and RIPng

- [Overview on page 2915](#)
- [Configuration on page 2921](#)
- [Administration on page 3030](#)

Overview

- [RIP on page 2915](#)

RIP

- [RIP Overview on page 2915](#)
- [RIP Configuration Overview on page 2920](#)
- [Supported RIP and RIPng Standards on page 2920](#)

RIP Overview

RIP is an interior gateway protocol (IGP) that uses a distance-vector algorithm to determine the best route to a destination, using the hop count as the metric.

In a RIP network, each router's forwarding table is distributed among the nodes through the flooding of routing table information. Because topology changes are flooded throughout the network, every node maintains the same list of destinations. Packets are then routed to these destinations based on path-cost calculations done at each node in the network.




NOTE: In general, the term *RIP* refers to RIP version 1 and RIP version 2.

This topic contains the following sections:

- [Distance-Vector Routing Protocols on page 2915](#)
- [RIP Protocol Overview on page 2916](#)
- [RIP Packets on page 2917](#)
- [Maximizing Hop Count on page 2918](#)
- [Split Horizon and Poison Reverse Efficiency Techniques on page 2918](#)
- [Limitations of Unidirectional Connectivity on page 2919](#)

Distance-Vector Routing Protocols

Distance-vector routing protocols transmit routing information that includes a distance vector, typically expressed as the number of hops to the destination. This information is flooded out all protocol-enabled interfaces at regular intervals (every 30 seconds in the case of RIP) to create a network map that is stored in each node's local topology database. [Figure 29 on page 2916](#) shows how distance-vector routing works.

Figure 29: Distance-Vector Protocol

g015005

In [Figure 29 on page 2916](#), Routers A and B have RIP enabled on adjacent interfaces. Router A has known RIP neighbors Routers C, D, and E, which are 1, 2, and 3 hops away, respectively. Router B has known RIP neighbors Routers X, Y, and Z, which are 1, 2, and 3 hops away, respectively. Every 30 seconds, each router floods its entire routing table information out all RIP-enabled interfaces. In this case, flooding exchanges routing table information across the RIP link.

When Router A receives routing information from Router B, it adds 1 to the hop count to determine the new hop count. For example, Router X has a hop count of 1, but when Router A imports the route to X, the new hop count is 2. The imported route also includes information about where the route was learned, so that the original route is imported as a route to Router X through Router B with a hop count of 2.

When multiple routes to the same host are received, RIP uses the distance-vector algorithm to determine which path to import into the forwarding table. The route with the smallest hop count is imported. If there are multiple routes with the same hop count, all are imported into the forwarding table, and traffic is sent along the paths in round-robin fashion.

RIP Protocol Overview

The RIP IGP uses the Bellman-Ford, or *distance-vector*, algorithm to determine the best route to a destination. RIP uses the hop count as the metric. RIP enables hosts and routers to exchange information for computing routes through an IP-based network. RIP is intended to be used as an IGP in reasonably homogeneous networks of moderate size.

The Junos® operating system (Junos OS) supports RIP versions 1 and 2.



NOTE: RIP is not supported for multipoint interfaces.

RIP version 1 packets contain the minimal information necessary to route packets through a network. However, this version of RIP does not support authentication or subnetting.

RIP uses User Datagram Protocol (UDP) port 520.

RIP has the following architectural limitations:

- The longest network path cannot exceed 15 hops (assuming that each network, or hop, has a cost of 1).
- RIP depends on counting to infinity to resolve certain unusual situations—When the network consists of several hundred routers, and when a routing loop has formed, the amount of time and network bandwidth required to resolve a next hop might be great.
- RIP uses only a fixed metric to select a route. Other IGPs use additional parameters, such as measured delay, reliability, and load.

RIP Packets

RIP packets contain the following fields:

- **Command**—Indicates whether the packet is a request or response message. Request messages seek information for the router's routing table. Response messages are sent periodically and also when a request message is received. Periodic response messages are called *update messages*. Update messages contain the command and version fields and 25 destinations (by default), each of which includes the destination IP address and the metric to reach that destination.



NOTE: Beginning with Junos OS Release 11.1, three additional command field types are available to support RIP demand circuits. When you configure an interface for RIP demand circuits, the command field indicates whether the packet is an update request, update response, or update acknowledge message. Neighbor interfaces send updates on demand, not periodically. These command field types are only valid on interfaces configured for RIP demand circuits. For more detailed information, see [“RIP Demand Circuits Overview” on page 2984](#).

- **Version number**—Version of RIP that the originating router is running.
- **Address family identifier**—Address family used by the originating router. The family is always IP.
- **Address**—IP address included in the packet.
- **Metric**—Value of the metric advertised for the address.
- **Mask**—Mask associated with the IP address (RIP version 2 only).
- **Next hop**—IP address of the next-hop router (RIP version 2 only).

Routing information is exchanged in a RIP network by RIP request and RIP response packets. A router that has just booted can broadcast a RIP request on all RIP-enabled interfaces. Any routers running RIP on those links receive the request and respond by sending a RIP response packet immediately to the router. The response packet contains the routing table information required to build the local copy of the network topology map.

In the absence of RIP request packets, all RIP routers broadcast a RIP response packet every 30 seconds on all RIP-enabled interfaces. The RIP broadcast is the primary way in which topology information is flooded throughout the network.

Once a router learns about a particular destination through RIP, it starts a timer. Every time it receives a new response packet with information about the destination, the router resets the timer to zero. However, if the router receives no updates about a particular destination for 180 seconds, it removes the destination from its RIP routing table.

In addition to the regular transmission of RIP packets every 30 seconds, if a router detects a new neighbor or detects that an interface is unavailable, it generates a triggered update. The new routing information is immediately broadcast out all RIP-enabled interfaces, and the change is reflected in all subsequent RIP response packets.

Maximizing Hop Count

The successful routing of traffic across a RIP network requires that every node in the network maintain the same view of the topology. Topology information is broadcast between RIP neighbors every 30 seconds. If Router A is many hops away from a new host, Router B, the route to B might take significant time to propagate through the network and be imported into Router A's routing table. If the two routers are 5 hops away from each other, Router A cannot import the route to Router B until 2.5 minutes after Router B is online (30 seconds per hop). For large numbers of hops, the delay becomes prohibitive. To help prevent this delay from growing arbitrarily large, RIP enforces a maximum hop count of 15 hops. Any prefix that is more than 15 hops away is treated as unreachable and assigned a hop count equal to infinity. This maximum hop count is called the *network diameter*.


Split Horizon and Poison Reverse Efficiency Techniques

Because RIP functions by periodically flooding the entire routing table out to the network, it generates a lot of traffic. The split horizon and poison reverse techniques can help reduce the amount of network traffic originated by RIP hosts and make the transmission of routing information more efficient.

If a router receives a set of route advertisements on a particular interface, RIP determines that those advertisements do not need to be retransmitted out the same interface. This technique, known as *split horizon*, helps limit the amount of RIP routing traffic by eliminating information that other neighbors on that interface have already learned.

[Figure 30 on page 2918](#) shows an example of the split horizon technique.

Figure 30: Split Horizon Example




In [Figure 30 on page 2918](#), Router A advertises routes to Routers C, D, and E to Router B. In this example, Router A can reach Router C in 2 hops. When Router A advertises the route to Router B, Router B imports it as a route to Router C through Router A in 3 hops. If Router

B then readvertised this route to Router A, Router A would import it as a route to Router C through Router B in 4 hops. However, the advertisement from Router B to Router A is unnecessary, because Router A can already reach the route in 2 hops. The split horizon technique helps reduce extra traffic by eliminating this type of route advertisement.

Similarly, the poison reverse technique helps to optimize the transmission of routing information and improve the time to reach network convergence. If Router A learns about unreachable routes through one of its interfaces, it advertises those routes as unreachable (hop count of 16) out the same interface. [Figure 31 on page 2919](#) shows an example of the poison reverse technique.

Figure 31: Poison Reverse Example




In [Figure 31 on page 2919](#), Router A learns through one of its interfaces that routes to Routers C, D, and E are unreachable. Router A readvertises those routes out the same interface as unreachable. The advertisement informs Router B that Routers C, D, and E are definitely not reachable through Router A.

Limitations of Unidirectional Connectivity

Because RIP processes routing information based solely on the receipt of routing table updates, it cannot ensure bidirectional connectivity. As [Figure 32 on page 2919](#) shows, RIP networks are limited by their unidirectional connectivity.

Figure 32: Limitations of Unidirectional Connectivity



In [Figure 32 on page 2919](#), Routers A and D flood their routing table information to Router B. Because the path to Router E has the fewest hops when routed through Router A, that route is imported into Router B's forwarding table. However, suppose that Router A can

transmit traffic but is not receiving traffic from Router B because of an unavailable link or invalid routing policy. If the only route to Router E is through Router A, any traffic destined for Router A is lost, because bidirectional connectivity was never established.

OSPF establishes bidirectional connectivity with a three-way handshake.

**Related
Documentation**

- *Junos OS Feature Support Reference for SRX Series and J Series Devices*
- [RIP Configuration Overview on page 2920](#)
- [Example: Configuring RIP on page 2921](#)

RIP Configuration Overview

To achieve basic connectivity between all RIP hosts in a RIP network, you enable RIP on every interface that is expected to transmit and receive RIP traffic, as described in the steps that follow.

To configure a RIP network:

1. Configure network interfaces. See the *Junos OS Interfaces Configuration Guide for Security Devices*.
2. Define RIP groups, which are logical groupings of interfaces, and add interfaces to the groups. Then, configure a routing policy to export directly connected routes and routes learned through RIP routing exchanges. See [“Example: Configuring a Basic RIP Network” on page 2922](#).
3. (Optional) Configure metrics to control traffic through the RIP network. See [“Example: Controlling Traffic in a RIP Network with an Incoming Metric” on page 2960](#) and [“Example: Controlling Traffic in a RIP Network with an Outgoing Metric” on page 2962](#).
4. (Optional) Configure authentication to ensure that only trusted routers participate in the autonomous system’s routing. See [“Enabling Authentication with Plain-Text Passwords \(CLI Procedure\)” on page 2933](#) and [“Enabling Authentication with MD5 Authentication \(CLI Procedure\)” on page 2933](#).

**Related
Documentation**

- *Junos OS Feature Support Reference for SRX Series and J Series Devices*
- [RIP Overview on page 2915](#)
- [Verifying a RIP Configuration on page 2994](#)

Supported RIP and RIPng Standards

Junos OS substantially supports the following RFCs, which define standards for RIP (for IP version 4 [IPv4]) and RIP next generation (RIPng, for IP version 6 [IPv6]).

Junos OS supports authentication for all RIP protocol exchanges (MD5 or simple authentication).

- RFC 1058, *Routing Information Protocol*
- RFC 2080, *RIPng for IPv6*

- RFC 2082, *RIP-2 MD5 Authentication*

Multiple keys using distinct key IDs are not supported.

- RFC 2453, *RIP Version 2*

The following RFC does not define a standard, but provides information about RIPvng. The IETF classifies it as “Informational.”

- RFC 2081, *RIPvng Protocol Applicability Statement*

Related Documentation

- *Supported IPv4, TCP, and UDP Standards*
- *Supported IPv6 Standards*
- *Accessing Standards Documents on the Internet*

Configuration

- [Configuration Examples on page 2921](#)
- [Configuration Statements on page 2996](#)

Configuration Examples

- [Example: Configuring RIP on page 2921](#)
- [Example: Configuring Authentication for RIP Routes on page 2928](#)
- [Example: Configuring BFD for RIP on page 2934](#)
- [Example: Configuring BFD Authentication for RIP on page 2939](#)
- [Example: Configuring Point-to-Multipoint RIP Networks on page 2947](#)
- [Example: Applying Policies to RIP Routes Imported from Neighbors on page 2954](#)
- [Examples: Controlling Traffic with Metrics in a RIP Network on page 2959](#)
- [Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 2967](#)
- [Example: Redistributing Routes Among RIP Instances on page 2971](#)
- [Example: Configuring RIP Timers on page 2977](#)
- [Example: Configuring RIP Demand Circuits on page 2984](#)
- [Example: Tracing RIP Protocol Traffic on page 2989](#)
- [Verifying a RIP Configuration on page 2994](#)

Example: Configuring RIP

- [Understanding Basic RIP Routing on page 2921](#)
- [Example: Configuring a Basic RIP Network on page 2922](#)

Understanding Basic RIP Routing

RIP is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). By default, RIP does not advertise the subnets that are directly connected through the device's interfaces. For traffic to pass through a RIP network, you must create a routing policy to export these routes. Advertising only the direct routes propagates the

routes to the immediately adjacent RIP-enabled router only. To propagate all routes through the entire RIP network, you must configure the routing policy to export the routes learned through RIP.

Example: Configuring a Basic RIP Network

This example shows how to configure a basic RIP network.

- [Requirements on page 2922](#)
- [Overview on page 2922](#)
- [Configuration on page 2922](#)
- [Verification on page 2925](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview


In this example, you configure a basic RIP network, create a RIP group called **rip-group**, and add the directly connected interfaces to the RIP group. Then you configure a routing policy to advertise direct routes using policy statement **advertise-routes-through-rip**.

By default, Junos OS does not advertise RIP routes, not even routes that are learned through RIP. To advertise RIP routes, you must configure and apply an export routing policy that advertises RIP-learned and direct routes.

In Junos OS, you do not need to configure the RIP version. RIP version 2 is used by default.

To use RIP on the device, you must configure RIP on all of the RIP interfaces within the network. [Figure 33 on page 2922](#) shows the topology used in this example.

Figure 33: Sample RIP Network Topology



“CLI Quick Configuration” on [page 2922](#) shows the configuration for all of the devices in [Figure 33 on page 2922](#). The section “Step-by-Step Procedure” on [page 2923](#) describes the steps on Device R1.

Configuration

CLI Quick Configuration

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

Device R1

```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

Device R2

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

Device R3

```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure a basic RIP network:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

```

```

user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32

```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 172.16.0.1/32;
      address 192.168.1.1/32;
    }
  }
}
```

```
user@R1# show protocols
rip {
  group rip-group {
    export advertise-routes-through-rip;
    neighbor fe-1/2/0.1;
  }
}
```

```
user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Checking the Routing Table on page 2925](#)
- [Looking at the Routes That Device R1 Is Advertising to Device R2 on page 2925](#)
- [Looking at the Routes That Device R1 Is Receiving from Device R2 on page 2926](#)
- [Verifying the RIP-Enabled Interfaces on page 2926](#)
- [Verifying the Exchange of RIP Messages on page 2926](#)
- [Verifying Reachability of All Hosts in the RIP Network on page 2927](#)

Checking the Routing Table

Purpose Verify that the routing table is populated with the expected routes..

Action From operational mode, enter the **show route protocol rip** command.

```
user@R1> show route protocol rip
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30      *[RIP/100] 00:59:15, metric 2, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
172.16.2.2/32   *[RIP/100] 02:52:48, metric 2, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
172.16.3.3/32   *[RIP/100] 00:45:05, metric 3, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32  *[RIP/100] 02:52:48, metric 2, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
192.168.3.3/32  *[RIP/100] 00:45:05, metric 3, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
224.0.0.9/32    *[RIP/100] 00:45:09, metric 1
                 MultiRecv
```

Meaning The output shows that the routes have been learned from Device R2 and Device R3.

If you were to delete the **from protocol rip** condition in the routing policy on Device R2, the remote routes from Device R3 would not be learned on Device R1.

Looking at the Routes That Device R1 Is Advertising to Device R2

Purpose Verify that Device R1 is sending the expected routes.

Action From operational mode, enter the **show route advertising-protocol rip** command.

```
user@R1> show route advertising-protocol rip 10.0.0.1
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.1/32   *[Direct/0] 05:18:26
                 > via lo0.1
192.168.1.1/32  *[Direct/0] 05:18:25
                 > via lo0.1
```

Meaning Device R1 is sending routes to its directly connected networks.

Looking at the Routes That Device R1 Is Receiving from Device R2

Purpose Verify that Device R1 is receiving the expected routes.

Action From operational mode, enter the **show route receive-protocol rip** command.

```
user@R1> show route receive-protocol rip 10.0.0.2
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30          * [RIP/100] 02:31:22, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
172.16.2.2/32       * [RIP/100] 04:24:55, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
172.16.3.3/32       * [RIP/100] 02:17:12, metric 3, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32      * [RIP/100] 04:24:55, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
192.168.3.3/32      * [RIP/100] 02:17:12, metric 3, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
```

Meaning Device R1 is receiving from Device R2 all of Device R2's directly connected networks. Device R1 is also receiving from Device R2 all of Device R3's directly connected networks, which Device R2 learned from Device R3 through RIP.

Verifying the RIP-Enabled Interfaces

Purpose Verify that all RIP-enabled Interfaces are available and active.

Action From operational mode, enter the **show rip neighbor** command.

```
user@R1> show rip neighbor
```

Neighbor	Local State	Source Address	Destination Address	Send Mode	Receive Mode	In Met
fe-1/2/0.1	Up	10.0.0.1	224.0.0.9	mcast	both	1

Meaning The output shows that the RIP-enabled interface on Device R1 is operational.

In general for this command, the output shows a list of the RIP neighbors that are configured on the device. Verify the following information:

- Each configured interface is present. Interfaces are listed in alphabetical order.
- Each configured interface is up. The state of the interface is listed in the **Local State** column. A state of **Up** indicates that the link is passing RIP traffic. A state of **Dn** indicates that the link is not passing RIP traffic. In a point-to-point link, this state generally means that either the end point is not configured for RIP or the link is unavailable.

Verifying the Exchange of RIP Messages

Purpose Verify that RIP messages are being sent and received on all RIP-enabled interfaces.

Action From operational mode, enter the **show rip statistics** command.

```

user@R1> show rip statistics
RIPv2 info: port 520; holddown 120s.
      rts learned  rts held down  rqsts dropped  resps dropped
              5              0              0              0

fe-1/2/0.1: 5 routes learned; 2 routes advertised; timeout 180s; update interval
30s
Counter              Total    Last 5 min  Last minute
-----
Updates Sent          2669         10         2
Triggered Updates Sent      2         0         0
Responses Sent          0         0         0
Bad Messages           0         0         0
RIPv1 Updates Received     0         0         0
RIPv1 Bad Route Entries    0         0         0
RIPv1 Updates Ignored      0         0         0
RIPv2 Updates Received    2675        11         2
RIPv2 Bad Route Entries    0         0         0
RIPv2 Updates Ignored      0         0         0
Authentication Failures    0         0         0
RIP Requests Received     0         0         0
RIP Requests Ignored       0         0         0
none                     0         0         0

```

Meaning The output shows the number of RIP routes learned. It also shows the number of RIP updates sent and received on the RIP-enabled interfaces. Verify the following information:

- The number of RIP routes learned matches the number of expected routes learned. Subnets learned by direct connectivity through an outgoing interface are not listed as RIP routes.
- RIP updates are being sent on each RIP-enabled interface. If no updates are being sent, the routing policy might not be configured to export routes.
- RIP updates are being received on each RIP-enabled interface. If no updates are being received, the routing policy might not be configured to export routes on the host connected to that subnet. The lack of updates might also indicate an authentication error.

Verifying Reachability of All Hosts in the RIP Network

Purpose Use the **traceroute** command on each loopback address in the network to verify that all hosts in the RIP network are reachable from each Juniper Networks device.

Action From operational mode, enter the **traceroute** command.

```

user@R1> traceroute 192.168.3.3
traceroute to 192.168.3.3 (192.168.3.3), 30 hops max, 40 byte packets
 1  10.0.0.2 (10.0.0.2)  1.094 ms  1.028 ms  0.957 ms
 2  192.168.3.3 (192.168.3.3)  1.344 ms  2.245 ms  2.125 ms

```

Meaning Each numbered row in the output indicates a routing hop in the path to the host. The three-time increments indicate the round-trip time (RTT) between the device and the hop for each traceroute packet.

To ensure that the RIP network is healthy, verify the following information:

- The final hop in the list is the host you want to reach.
- The number of expected hops to the host matches the number of hops in the traceroute output. The appearance of more hops than expected in the output indicates that a network segment is probably unreachable. It might also indicate that the incoming or outgoing metric on one or more hosts has been set unexpectedly.

**Related
Documentation**

- [Example: Configuring Point-to-Multipoint RIP Networks on page 2947](#)

Example: Configuring Authentication for RIP Routes

- [Understanding RIP Authentication on page 2928](#)
- [Example: Configuring Route Authentication for RIP on page 2928](#)
- [Enabling Authentication with Plain-Text Passwords \(CLI Procedure\) on page 2933](#)
- [Enabling Authentication with MD5 Authentication \(CLI Procedure\) on page 2933](#)

Understanding RIP Authentication

RIPv2 provides authentication support so that RIP links can require authentication keys (passwords) before they become active. Authentication provides an additional layer of security on the network beyond the other security features. By default, this authentication is disabled.

Authentication keys can be specified in either plain-text or MD5 form. Authentication requires all routers within the RIP network or subnetwork to have the same authentication type and key (password) configured.

This type of authentication is not supported on RIPv1 networks.

Example: Configuring Route Authentication for RIP

This example shows how to configure authentication for a RIP network.

- [Requirements on page 2928](#)
- [Overview on page 2928](#)
- [Configuration on page 2929](#)
- [Verification on page 2932](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

You can configure the router to authenticate RIP route queries. By default, authentication is disabled. You can use one of the following authentication methods:


- Simple authentication—Uses a text password that is included in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet.

- MD5 authentication—Creates an encoded checksum that is included in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet's MD5 checksum.

This example shows MD5 authentication.

Figure 34 on page 2929 shows the topology used in this example.

Figure 34: RIP Authentication Network Topology



"CLI Quick Configuration" on page 2929 shows the configuration for all of the devices in Figure 34 on page 2929. The section "Step-by-Step Procedure" on page 2930 describes the steps on Device R1.

Configuration

CLI Quick Configuration

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

Device R1

```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set protocols rip authentication-type md5
set protocols rip authentication-key "$9$ONLRBhreK87dsM8i.5FAtM8XxNb"
set protocols rip traceoptions file rip-authentication-messages
set protocols rip traceoptions flag auth
set protocols rip traceoptions flag packets
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

Device R2

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set protocols rip authentication-type md5
set protocols rip authentication-key "$9$Lf1Xds2gJDHmoJCu1hKvoJGUjq"

```

```
set protocols rip traceoptions file rip-authentication-messages
set protocols rip traceoptions flag auth
set protocols rip traceoptions flag packets
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip authentication-type md5
set protocols rip authentication-key "$9$G.UkP5T39tOz3K87V4oz36/Cu"
set protocols rip traceoptions file rip-authentication-messages
set protocols rip traceoptions flag auth
set protocols rip traceoptions flag packets
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure RIP authentication:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

```
user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```


4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

5. Require MD5 authentication for RIP route queries received on an interface.

The passwords must match on neighboring RIP routers. If the password does not match, the packet is rejected. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.

Do not enter the password as shown here. The password shown here is the encrypted password that is displayed in the configuration after the actual password is already configured.

```
[edit protocols rip]
user@R1# set authentication-type md5
user@R1# set authentication-key "$9$ONLRBhreK87dsM8i.5FAAtM8XxNb"
```

6. Configure tracing operations to track authentication.

```
[edit protocols rip traceoptions]
user@R1# set file rip-authentication-messages
user@R1# set flag auth
user@R1# set flag packets
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 172.16.0.1/32;
      address 192.168.1.1/32;
    }
  }
}

user@R1# show protocols
rip {
  traceoptions {
    file rip-authentication-messages;
    flag auth;
    flag packets;
  }
}
```

```

authentication-type md5;
authentication-key "$9$ONLRBhreK87dsM8i.5FAtM8XxNb"; ## SECRET-DATA
group rip-group {
    export advertise-routes-through-rip;
    neighbor fe-1/2/0.1;
}
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
    term 1 {
        from protocol [ direct rip ];
        then accept;
    }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Checking for Authentication Failures on page 2932](#)
- [Verifying That MD5 Authentication Is Enabled in RIP Update Packets on page 2933](#)

Checking for Authentication Failures

Purpose Verify that there are no authentication failures.

Action From operational mode, enter the **show rip statistics** command.

```

user@R1> show rip statistics
RIPv2 info: port 520; holddown 120s.
      rts learned  rts held down  rqsts dropped  resps dropped
              5              0              0              0

fe-1/2/0.1: 5 routes learned; 2 routes advertised; timeout 180s; update interval
30s
Counter              Total    Last 5 min  Last minute
-----
Updates Sent          2669         10          2
Triggered Updates Sent      2          0          0
Responses Sent          0          0          0
Bad Messages           0          0          0
RIPv1 Updates Received      0          0          0
RIPv1 Bad Route Entries      0          0          0
RIPv1 Updates Ignored        0          0          0
RIPv2 Updates Received    2675         11          2
RIPv2 Bad Route Entries      0          0          0
RIPv2 Updates Ignored        0          0          0
Authentication Failures      0          0          0
RIP Requests Received        0          0          0
RIP Requests Ignored         0          0          0
none                      0          0          0

```

Meaning The output shows that there are no authentication failures.

Verifying That MD5 Authentication Is Enabled in RIP Update Packets

Purpose

Use tracing operations to verify that MD5 authentication is enabled in RIP updates.

Action

From operational mode, enter the **show log** command.

user@R1> **show log rip-authentication-messages | match md5**
Feb 15 15:45:13.969462 sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:45:43.229867 sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:46:13.174410 sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:46:42.716566 sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:47:11.425076 sending msg 0xb9a8c04, 3 rtes (needs MD5)
...

Meaning

The **(needs MD5)** output shows that all route updates require MD5 authentication.

Enabling Authentication with Plain-Text Passwords (CLI Procedure)

To configure authentication that requires a plain-text password to be included in the transmitted packet, enable simple authentication by performing these steps on all RIP devices in the network:

1. Navigate to the top of the configuration hierarchy.
2. Perform the configuration tasks described in [Table 239 on page 2933](#).
3. If you are finished configuring the router, commit the configuration.

Table 239: Configuring Simple RIP Authentication

Task	CLI Configuration Editor
Navigate to Rip level in the configuration hierarchy.	From the [edit] hierarchy level, enter edit protocols rip
Set the authentication type to simple .	Set the authentication type to simple : set authentication-type simple
Set the authentication key to a simple-text password. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.	Set the authentication key to a simple-text password: set authentication-key password

Enabling Authentication with MD5 Authentication (CLI Procedure)

To configure authentication that requires an MD5 password to be included in the transmitted packet, enable MD5 authentication by performing these steps on all RIP devices in the network:

1. Navigate to the top of the configuration hierarchy.
2. Perform the configuration tasks described in [Table 240 on page 2934](#).
3. If you are finished configuring the router, commit the configuration.

Table 240: Configuring MD5 RIP Authentication

Task	CLI Configuration Editor
Navigate to Rip level in the configuration hierarchy.	From the [edit] hierarchy level, enter edit protocols rip
Set the authentication type to MD5 .	Set the authentication type to md5 : set authentication-type md5
Set the MD5 authentication key (password). The key can be from 1 through 16 contiguous characters long and can include any ASCII strings.	Set the MD5 authentication key: set authentication-key password

Related Documentation

- [Example: Configuring RIP on page 2921](#)

Example: Configuring BFD for RIP

This example shows how to configure Bidirectional Forwarding Detection (BFD) for a RIP network.

- [Requirements on page 2934](#)
- [Overview on page 2934](#)
- [Configuration on page 2936](#)
- [Verification on page 2938](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

To enable failure detection, include the **bfd-liveness-detection** statement:

```

bfd-liveness-detection {
  detection-time {
    threshold milliseconds;
  }
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  no-adaptation;
  transmit-interval {
    threshold milliseconds;
    minimum-interval milliseconds;
  }
  version (1 | automatic);
}

```

Optionally, you can specify the threshold for the adaptation of the detection time by including the **threshold** statement. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a system log message are sent.

To specify the minimum transmit and receive interval for failure detection, include the **minimum-interval** statement. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval at which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. This examples sets a minimum interval of 600 milliseconds.



NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with nonstop active routing configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

You can optionally specify the minimum transmit and receive intervals separately.

To specify only the minimum receive interval for failure detection, include the **minimum-receive-interval** statement. This value represents the minimum interval at which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,00 milliseconds.

To specify only the minimum transmit interval for failure detection, include the **transmit-interval minimum-interval** statement. This value represents the minimum interval at which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds.

To specify the number of hello packets not received by a neighbor that causes the originating interface to be declared down, include the **multiplier** statement. The default is 3, and you can configure a value in the range from 1 through 255.

To specify the threshold for detecting the adaptation of the transmit interval, include the **transmit-interval threshold** statement. The threshold value must be greater than the transmit interval.


To specify the BFD version used for detection, include the **version** statement. The default is to have the version detected automatically.

You can trace BFD operations by including the **traceoptions** statement at the **[edit protocols bfd]** hierarchy level.

In Junos OS Release 9.0 and later, you can configure BFD sessions not to adapt to changing network conditions. To disable BFD adaptation, include the **no-adaptation** statement. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

Figure 35 on page 2936 shows the topology used in this example.

Figure 35: RIP BFD Network Topology



“CLI Quick Configuration” on page 2936 shows the configuration for all of the devices in Figure 35 on page 2936. The section “Step-by-Step Procedure” on page 2937 describes the steps on Device R1.

Configuration

CLI Quick Configuration

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

Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R2

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

Device R3

```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure a BFD for a RIP network:

1. Configure the network interfaces.

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```

[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1

```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```

[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept

```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```

[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip

```

5. Enable BFD.

```

[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection minimum-interval 600

```

6. Configure tracing operations to track BFD messages.

```
[edit protocols bfd traceoptions]
```

```
user@R1# set file bfd-trace
```

```
user@R1# set flag all
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
```

```
fe-1/2/0 {  
  unit 1 {  
    family inet {  
      address 10.0.0.1/30;  
    }  
  }  
}
```

```
user@R1# show protocols
```

```
bfd {  
  traceoptions {  
    file bfd-trace;  
    flag all;  
  }  
}  
rip {  
  group rip-group {  
    export advertise-routes-through-rip;  
    bfd-liveness-detection {  
      minimum-interval 600;  
    }  
    neighbor fe-1/2/0.1;  
  }  
}
```

```
user@R1# show policy-options
```

```
policy-statement advertise-routes-through-rip {  
  term 1 {  
    from protocol [ direct rip ];  
    then accept;  
  }  
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying That the BFD Sessions Are Up on page 2939](#)
- [Checking the BFD Trace File on page 2939](#)

Verifying That the BFD Sessions Are Up

Purpose Make sure that the BFD sessions are operating.

Action From operational mode, enter the **show bfd session** command.

```
user@R1> show bfd session
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.2	Up	fe-1/2/0.1	1.800	0.600	3

```
1 sessions, 1 clients
```

```
Cumulative transmit rate 1.7 pps, cumulative receive rate 1.7 pps
```

Meaning The output shows that there are no authentication failures.

Checking the BFD Trace File

Purpose Use tracing operations to verify that BFD packets are being exchanged.

Action From operational mode, enter the **show log** command.

```
user@R1> show log bfd-trace
```

```
Feb 16 10:26:32 PPM Trace: BFD periodic xmit to 10.0.0.2 (IFL 124, rtbl 53, single-hop port)
```

```
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 86:
```

```
Feb 16 10:26:32 IfIndex (3) len 4: 0
```

```
Feb 16 10:26:32 Protocol (1) len 1: BFD
```

```
Feb 16 10:26:32 Data (9) len 61: (hex) 42 46 44 20 70 61 63 6b 65 74 20 66 72 6f 6d 20 31 30 2e
```

```
Feb 16 10:26:32 PPM Trace: BFD packet from 10.0.0.1 (IFL 73, rtbl 56, ttl 255) absorbed
```

```
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 60:
```

```
Feb 16 10:26:32 IfIndex (3) len 4: 0
```

```
Feb 16 10:26:32 Protocol (1) len 1: BFD
```

```
Feb 16 10:26:32 Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20 78 6d 69 74 20 6f
```

```
...
```

Meaning The output shows the normal functioning of BFD.

Related Documentation

- [Understanding BFD for RIP](#)

Example: Configuring BFD Authentication for RIP

- [Understanding BFD Authentication for RIP on page 2939](#)
- [Example: Configuring BFD Authentication for RIP on page 2941](#)

Understanding BFD Authentication for RIP

BFD enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when running BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels. Beginning with Junos OS Release 9.6, Junos OS supports authentication

for BFD sessions running over RIP. BFD authentication is only supported in the domestic image and is not available in the export image.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and the level of authentication that can be configured:

- [BFD Authentication Algorithms on page 2940](#)
- [Security Authentication Keychains on page 2941](#)
- [Strict Versus Loose Authentication on page 2941](#)

BFD Authentication Algorithms

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



NOTE: Nonstop active routing is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

Security Authentication Keychains

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

Strict Versus Loose Authentication

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

Example: Configuring BFD Authentication for RIP

This example shows how to configure Bidirectional Forwarding Detection (BFD) authentication for a RIP network.

- [Requirements on page 2941](#)
- [Overview on page 2941](#)
- [Configuration on page 2942](#)
- [Verification on page 2946](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

The devices must be running Junos OS Release 9.6 or later.

Overview


Only three steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the RIP protocol.
2. Associate the authentication keychain with the RIP protocol.

3. Configure the related security authentication keychain.

Figure 36 on page 2942 shows the topology used in this example.

Figure 36: RIP BFD Authentication Network Topology



"CLI Quick Configuration" on page 2942 shows the configuration for all of the devices in Figure 36 on page 2942. The section "Step-by-Step Procedure" on page 2943 describes the steps on Device R1.

Configuration

CLI Quick Configuration

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

```
Device R1
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set protocols rip group rip-group bfd-liveness-detection authentication key-chain bfd-rip
set protocols rip group rip-group bfd-liveness-detection authentication algorithm
    keyed-md5
set protocols rip group rip-group bfd-liveness-detection authentication loose-check
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set security authentication-key-chains key-chain bfd-rip key 53 secret
    "$9$5d1V2aZGi.fzDiORSeXxDikqmT"
set security authentication-key-chains key-chain bfd-rip key 53 start-time
    "2012-2-16.12:00:00 -0800"

Device R2
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set protocols rip group rip-group bfd-liveness-detection authentication key-chain bfd-rip
set protocols rip group rip-group bfd-liveness-detection authentication algorithm
    keyed-md5
```

```

set protocols rip group rip-group bfd-liveness-detection authentication loose-check
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set security authentication-key-chains key-chain bfd-rip key 53 secret
  "$9$d1V2aZGi.fzDiORSeXxDikqmT"
set security authentication-key-chains key-chain bfd-rip key 53 start-time
  "2012-2-16.12:00:00 -0800"

```

Device R3

```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set protocols rip group rip-group bfd-liveness-detection authentication key-chain bfd-rip
set protocols rip group rip-group bfd-liveness-detection authentication algorithm
  keyed-md5
set protocols rip group rip-group bfd-liveness-detection authentication loose-check
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set security authentication-key-chains key-chain bfd-rip key 53 secret
  "$9$d1V2aZGi.fzDiORSeXxDikqmT"
set security authentication-key-chains key-chain bfd-rip key 53 start-time
  "2012-2-16.12:00:00 -0800"

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure a BFD authentication:

1. Configure the network interfaces.

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```

[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1

```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```

[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept

```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

5. Enable BFD.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection minimum-interval 600
```

6. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use.



NOTE: Nonstop active routing is not supported with **meticulous-keyed-md5** and **meticulous-keyed-sha-1** authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection authentication algorithm keyed-md5
```

7. Specify the keychain to be used to associate BFD sessions on RIP with the unique security authentication keychain attributes.

The keychain you specify must match a keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection authentication key-chain bfd-rip
```

8. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection authentication loose-check
```

9. Specify the unique security authentication information for BFD sessions:

- The matching keychain name as specified in Step 7.
- At least one key, a unique integer between 0 and 63. Creating multiple keys allows multiple clients to use the BFD session.
- The secret data used to allow access to the session.
- The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security authentication-key-chains key-chain bfd-rip]
user@R1# set key 53 secret "$9$d1V2aZGi.fzDiORSeXxDikqmT"
user@R1# set key 53 start-time "2012-2-16.12:00:00 -0800"
```

10. Configure tracing operations to track BFD authentication.

```
[edit protocols bfd traceoptions]
```

```

user@R1# set file bfd-trace
user@R1# set flag all

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show security** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}

user@R1# show protocols
bfd {
  traceoptions {
    file bfd-trace;
    flag all;
  }
}
rip {
  group rip-group {
    export advertise-routes-through-rip;
    bfd-liveness-detection {
      minimum-interval 600;
    }
    neighbor fe-1/2/0.1;
  }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}

user@R1# show security
authentication-key-chains {
  key-chain bfd-rip {
    key 53 {
      secret "$9$d1V2aZGi.fzDiORSeXxDikqmT"; ## SECRET-DATA
      start-time "2012-2-16.12:00:00 -0800";
    }
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying That the BFD Sessions Are Authenticated on page 2946](#)
- [Viewing Extensive Information About the BFD Authentication on page 2946](#)
- [Checking the BFD Trace File on page 2947](#)

Verifying That the BFD Sessions Are Authenticated

Purpose Make sure that the BFD sessions are authenticated.

Action From operational mode, enter the **show bfd session detail** command.

```
user@R1> show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.2	Up	fe-1/2/0.1	1.800	0.600	3

Client RIP, TX interval 0.600, RX interval 0.600, **Authenticate**
 Session up time 01:39:34
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Logical system 6, routing table index 53

1 sessions, 1 clients
 Cumulative transmit rate 1.7 pps, cumulative receive rate 1.7 pps

Meaning **Authenticate** is displayed to indicate that BFD authentication is configured.

Viewing Extensive Information About the BFD Authentication

Purpose View the keychain name, the authentication algorithm and mode for each client in the session, and the BFD authentication configuration status.

Action From operational mode, enter the **show bfd session extensive** command.

```
user@R1> show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.2	Up	fe-1/2/0.1	1.800	0.600	3

Client RIP, TX interval 0.600, RX interval 0.600, **Authenticate**
keychain bfd-rip, algo keyed-md5, mode loose
 Session up time 01:46:29
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Logical system 6, routing table index 53
 Min async interval 0.600, min slow interval 1.000
 Adaptive async TX interval 0.600, RX interval 0.600
 Local min TX interval 0.600, minimum RX interval 0.600, multiplier 3
 Remote min TX interval 0.600, min RX interval 0.600, multiplier 3
 Local discriminator 225, remote discriminator 226
 Echo mode disabled/inactive
Authentication enabled/active, keychain bfd-rip, algo keyed-md5, mode loose
 Session ID: 0x300501

1 sessions, 1 clients
 Cumulative transmit rate 1.7 pps, cumulative receive rate 1.7 pps

Meaning The output shows the keychain name, the authentication algorithm and mode for the client in the session, and the BFD authentication configuration status.

Checking the BFD Trace File

Purpose Use tracing operations to verify that BFD packets are being exchanged.

Action From operational mode, enter the **show log** command.

```
user@R1> show log bfd-trace
Feb 16 10:26:32 PPM Trace: BFD periodic xmit to 10.0.0.2 (IFL 124, rtbl 53,
single-hop port)
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 86:
Feb 16 10:26:32   IfIndex (3) len 4: 0
Feb 16 10:26:32   Protocol (1) len 1: BFD
Feb 16 10:26:32   Data (9) len 61: (hex) 42 46 44 20 70 61 63 6b 65 74 20 66 72
6f 6d 20 31 30 2e
Feb 16 10:26:32 PPM Trace: BFD packet from 10.0.0.1 (IFL 73, rtbl 56, ttl 255)
absorbed
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 60:
Feb 16 10:26:32   IfIndex (3) len 4: 0
Feb 16 10:26:32   Protocol (1) len 1: BFD
Feb 16 10:26:32   Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 6f
...
```

Meaning The output shows the normal functioning of BFD.

- Related Documentation**
- *Example: Configuring BFD for RIP*
 - [Example: Configuring Authentication for RIP Routes on page 2928](#)
 - [Example: Configuring RIP on page 2921](#)

Example: Configuring Point-to-Multipoint RIP Networks

- [Configuring Point-to-Multipoint RIP Networks Overview on page 2947](#)
- [Example: Configuring Point-to-Multipoint RIP Networks on page 2948](#)

Configuring Point-to-Multipoint RIP Networks Overview

A point-to-multipoint RIP network consists of a device having two or more peers on a single interface. All the devices forming a point-to-multipoint connection are placed in a single broadcast domain.

In a RIP network, a device can have a single peer or multiple peers for an interface. However, the demand circuit feature implementation in a RIP network requires the use of a single RIP peer. When you configure the following statements, a RIP network with demand circuits can also be configured to have multiple peers on an interface:

- Configuring the interface type to be a multipoint interface by using the **interface-type (Protocols RIP) p2mp** statement.
- Enabling dynamic peer discovery by using the **dynamic-peers** statement (SRX Series devices only).



NOTE: Before configuring the [dynamic-peers](#) statement, IPsec must be configured and IPsec tunnels must be set up by configuring IPsec parameters. Without IPsec configuration, the remote peers have to be explicitly configured at the RIP protocol level by using the [peer address](#) statement. See *Configuring Security Associations for IPsec on an ES PIC* for more details.

- Configuring peers by using the [peer address](#) statement.

```
[edit]
protocols {
  rip {
    group red {
      neighbor fe-0/1/3 {
        interface-type (Protocols RIP) p2mp;
        peer address; (or use dynamic-peers;)
      }
    }
  }
}
```

The [show rip statistics peer address](#) command can be used to display the RIP statistics at the peer level. The [clear rip statistics peer address](#) command can be used to clear the RIP statistics for a peer. Alternatively, you can use the [show rip statistics peer all](#) and [clear rip statistics peer all](#) command to display and clear RIP statistics for all peers.

Example: Configuring Point-to-Multipoint RIP Networks

This example shows how to configure a point-to-multipoint RIP network.

- [Requirements on page 2948](#)
- [Overview on page 2948](#)
- [Configuration on page 2950](#)
- [Verification on page 2952](#)

Requirements

This example uses the following hardware and software components:

- M Series routers, MX Series routers, T Series routers, or SRX Series devices
- Junos OS Release 12.1 or later

Overview

In a RIP network, a device can have a single peer or multiple peers for an interface. However, the demand circuit feature implementation in a RIP network requires the use of a single RIP peer.

When you include the following statements, the demand circuit implementation can have multiple peers for a given RIP neighbor.

- Configuring the interface type to be a multipoint interface by using the **interface-type (Protocols RIP) p2mp** statement.
- Enabling dynamic peer discovery by using the **dynamic-peers** statement (SRX Series devices only).



NOTE: To configure the **dynamic-peers** statement, IPsec tunnels must be set up by configuring IPsec parameters. See *Configuring Security Associations for IPsec on an ES PIC* for more details.

- Configuring peers by using the **peer address** statement.

```
[edit]
protocols {
  rip {
    group red {
      neighbor fe-0/1/3 {
        interface-type (Protocols RIP) p2mp;
        peer address; (or use dynamic-peers;)
      }
    }
  }
}
```


The **show rip statistics peer** command can be used to display the RIP statistics at the peer level.

Topology

In this example, Devices R1, R2, and R3 form a point-to-multipoint network. R1 is connected to R2 and to R3 as a point-to-multipoint connection through a switch that places all devices in the same broadcast domain. RIP demand circuits are configured on all three devices. The two peers to R1 are configured statically by using the **peer address** statement. The **dynamic-peers** statement is not used here.

Figure 37 on page 2949 shows the topology used in this example.

Figure 37: Configuring a Point-to-Multipoint RIP Network



Configuration

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

Device R1

```
set interfaces fe-0/1/3 unit 0 family inet address 23.1.1.1/24
set policy-options policy-statement accept-rip-routes term from-direct from protocol direct
set policy-options policy-statement accept-rip-routes term from-direct then accept
set policy-options policy-statement accept-rip-routes term from-rip from protocol rip
set policy-options policy-statement accept-rip-routes term from-rip then accept
set protocols rip traceoptions file R1.log size 4m world-readable
set protocols rip traceoptions flag all detail
set protocols rip group red export accept-rip-routes
set protocols rip group red neighbor fe-0/1/3.0 interface-type p2mp
set protocols rip group red neighbor fe-0/1/3.0 peer 23.1.1.2
set protocols rip group red neighbor fe-0/1/3.0 peer 23.1.1.3
set protocols rip group red neighbor fe-0/1/3.0 demand-circuit
set protocols rip group red neighbor fe-0/1/3.0 max-retrans-time 10
```

Similarly, configure Devices R2 and R3, omitting the **peer address** configuration statement.

Configuring a Point-to-Multipoint RIP Network (with Demand Circuits)

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure the point-to-multipoint feature across a RIP network:

1. Configure the device interface.

```
[edit interfaces fe-0/1/3 unit 0]
user@R1# set family inet address 23.1.1.1/24
```
2. Define a policy for exporting RIP routes from the routing table to the protocol for transmission through the network.

```
[edit policy-options policy-statement accept-rip-routes]
user@R1# set term from-direct from protocol direct
user@R1# set term from-direct then accept
user@R1# set term from-rip from protocol rip
user@R1# set term from-rip then accept
```
3. Configure RIP and a RIP group with the defined export policy and point-to-multipoint configuration statements.

```
[edit protocols rip]
user@R1# set traceoptions file R1.log size 4m world-readable
user@R1# set traceoptions flag all detail
user@R1# set group red export accept-rip-routes
user@R1# set group red neighbor fe-0/1/3.0 interface-type p2mp
user@R1# set group red neighbor fe-0/1/3.0 peer 23.1.1.2
user@R1# set group red neighbor fe-0/1/3.0 peer 23.1.1.3
user@R1# set group red neighbor fe-0/1/3.0 demand-circuit
```

```
user@R1# set group red neighbor fe-0/1/3.0 max-retrans-time 10
```

Similarly, configure Devices R2 and R3, omitting the **peer address** configuration statement.



NOTE: Configuring **max-retrans-time** is optional. In the absence of this configuration statement, the default retransmission time of 180 seconds is configured.

The configuration used in this example is for a RIP network with demand circuits. To configure RIP for networks without demand circuits, exclude the **demand-circuit** and **max-retrans-time** statements from the configuration and check the resulting output. For more information about configuring RIP demand circuits, see [“Example: Configuring RIP Demand Circuits” on page 2986](#).

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, and **show protocols rip** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@R1# show interfaces
fe-0/1/3 {
  unit 0 {
    family inet {
      address 23.1.1.1/24;
    }
  }
}

user@R1# show protocols rip
traceoptions {
  file R1.log size 4m world-readable;
  flag all detail;
}
group red {
  export accept-rip-routes;
  neighbor fe-0/1/3.0 {
    interface-type p2mp;
    peer 23.1.1.2;
    peer 23.1.1.3;
    demand-circuit;
    max-retrans-time 10;
  }
}

user@R1# show policy-options
policy-statement accept-rip-routes {
  term from-direct {
    from protocol direct;
    then accept;
  }
}
```

```
}  
term from-rip {  
    from protocol rip;  
    then accept;  
}  
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying the Point-to-Multipoint RIP Network

Purpose Verify that the RIP network is functional with the point-to-multipoint feature configured.

Action From operational mode, run the **show rip neighbor** command.

```
user@R1> show rip neighbor
```

Neighbor	Local State	Source Address	Destination Address	Send Mode	Receive Mode	In Met
fe-0/1/3.0(DC)	Up	23.1.1.1	23.1.1.2	unicast	unicast	1
fe-0/1/3.0(DC)	Up	23.1.1.1	23.1.1.3	unicast	unicast	1

From operational mode, run the **show rip statistics peer address** command.

```
user@R1> show rip statistics peer 23.1.1.2
```

RIPv2 info: port 520; holddown 120s.

rts learned	rts held down	rqsts dropped	resps dropped
3	0	0	0

fe-0/1/3.0 Peer-IP 23.1.1.2: 2 routes learned; 3 routes advertised; timeout 180s; update interval 0s

Counter	Total	Last 5 min	Last minute
Updates Sent	0	0	0
Triggered Updates Sent	3	0	0
Responses Sent	0	0	0
Bad Messages	0	0	0
RIPv1 Updates Received	0	0	0
RIPv1 Bad Route Entries	0	0	0
RIPv1 Updates Ignored	0	0	0
RIPv2 Updates Received	2	0	0
RIPv2 Bad Route Entries	0	0	0
RIPv2 Updates Ignored	0	0	0
Authentication Failures	0	0	0
RIP Requests Received	0	0	0
RIP Requests Ignored	0	0	0
none	3	0	0

```
user@R1> show rip statistics peer 23.1.1.3
```

RIPv2 info: port 520; holddown 120s.

rts learned	rts held down	rqsts dropped	resps dropped
3	0	0	0

fe-0/1/3.0 Peer-IP 23.1.1.3: 2 routes learned; 3 routes advertised; timeout 180s; update interval 0s

Counter	Total	Last 5 min	Last minute
Updates Sent	0	0	0
Triggered Updates Sent	3	0	0
Responses Sent	0	0	0
Bad Messages	0	0	0
RIPv1 Updates Received	0	0	0
RIPv1 Bad Route Entries	0	0	0
RIPv1 Updates Ignored	0	0	0
RIPv2 Updates Received	2	0	0
RIPv2 Bad Route Entries	0	0	0
RIPv2 Updates Ignored	0	0	0
Authentication Failures	0	0	0
RIP Requests Received	0	0	0
RIP Requests Ignored	0	0	0

none 3 0 0

Meaning The RIP network is up and running with the point-to-multipoint feature configured.

Related Documentation

- [Example: Configuring RIP on page 2921](#)

Example: Applying Policies to RIP Routes Imported from Neighbors

- [Understanding RIP Import Policy on page 2954](#)
- [Example: Applying Policies to RIP Routes Imported from Neighbors on page 2954](#)

Understanding RIP Import Policy

The default RIP import policy is to accept all received RIP routes that pass a sanity check. To filter routes being imported by the local routing device from its neighbors, include the **import** statement, and list the names of one or more policies to be evaluated. If you specify more than one policy, they are evaluated in order (first to last) and the first matching policy is applied to the route. If no match is found, the local routing device does not import any routes.

Example: Applying Policies to RIP Routes Imported from Neighbors

This example shows how to configure an import policy in a RIP network.

- [Requirements on page 2954](#)
- [Overview on page 2954](#)
- [Configuration on page 2955](#)
- [Verification on page 2958](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.


Overview

In this example, Device R1 has an import policy that accepts the 10/8 and 192.168/16 RIP routes and rejects all other RIP routes. This means that the 172.16/16 RIP routes are excluded from Device R1's routing table.

An export policy is also shown because an export policy is required as part of the minimum configuration for RIP.

[Figure 38 on page 2955](#) shows the topology used in this example.

Figure 38: RIP Import Policy Network Topology



"CLI Quick Configuration" on page 2955 shows the configuration for all of the devices in Figure 38 on page 2955. The section "Step-by-Step Procedure" on page 2956 describes the steps on Device R1.

Configuration

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

Device R1

```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip import rip-import
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set policy-options policy-statement rip-import term 1 from protocol rip
set policy-options policy-statement rip-import term 1 from route-filter 10.0.0.0/8 orlonger
set policy-options policy-statement rip-import term 1 from route-filter 192.168.0.0/16
  orlonger
set policy-options policy-statement rip-import term 1 then accept
set policy-options policy-statement rip-import term 2 then reject

```

Device R2

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

Device R3

```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32

```

```
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Step-by-Step
Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure a RIP import policy:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

```
user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled.

You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

5. Configure the import policy.

```
[edit policy-options policy-statement rip-import]
user@R1# set term 1 from protocol rip
user@R1# set term 1 from route-filter 10.0.0.0/8 orlonger
user@R1# set term 1 from route-filter 192.168.0.0/16 orlonger
user@R1# set term 1 then accept
user@R1# set term 2 then reject
```

6. Apply the import policy.

```
[edit protocols rip]
user@R1# set import rip-import
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 172.16.0.1/32;
      address 192.168.1.1/32;
    }
  }
}

user@R1# show protocols
rip {
  import rip-import;
  group rip-group {
    export advertise-routes-through-rip;
    neighbor fe-1/2/0.1;
  }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}
policy-statement rip-import {
  term 1 {
    from {
      protocol rip;
      route-filter 10.0.0.0/8 orlonger;
      route-filter 192.168.0.0/16 orlonger;
    }
    then accept;
  }
  term 2 {
    then reject;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Looking at the Routes That Device R2 Is Advertising to Device R1 on page 2958](#)
- [Looking at the Routes That Device R1 Is Receiving from Device R2 on page 2958](#)
- [Checking the Routing Table on page 2958](#)
- [Testing the Import Policy on page 2959](#)

Looking at the Routes That Device R2 Is Advertising to Device R1

Purpose Verify that Device R2 is sending the expected routes.

Action From operational mode, enter the **show route advertising-protocol rip** command.

```
user@R2> show route advertising-protocol rip 10.0.0.2
```

```
inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
10.0.0.4/30      *[Direct/0] 2d 01:17:44
                  > via fe-1/2/0.5
172.16.2.2/32    *[Direct/0] 2d 04:09:52
                  > via lo0.2
172.16.3.3/32    *[RIP/100] 23:40:02, metric 2, tag 0
                  > to 10.0.0.6 via fe-1/2/0.5
192.168.2.2/32   *[Direct/0] 2d 04:09:52
                  > via lo0.2
192.168.3.3/32   *[RIP/100] 23:40:02, metric 2, tag 0
                  > to 10.0.0.6 via fe-1/2/0.5
```

Meaning Device R2 is sending 172.16/16 routes to Device R1.

Looking at the Routes That Device R1 Is Receiving from Device R2

Purpose Verify that Device R1 is receiving the expected routes.

Action From operational mode, enter the **show route receive-protocol rip** command.

```
user@R1> show route receive-protocol rip 10.0.0.2
```

```
inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
10.0.0.4/30      *[RIP/100] 01:06:03, metric 2, tag 0
                  > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32    *[RIP/100] 01:06:03, metric 2, tag 0
                  > to 10.0.0.2 via fe-1/2/0.1
192.168.3.3/32    *[RIP/100] 01:06:03, metric 3, tag 0
                  > to 10.0.0.2 via fe-1/2/0.1
```

Meaning The output shows that the 172.16/16 routes are excluded.

Checking the Routing Table

Purpose Verify that the routing table is populated with the expected routes.

Action From operational mode, enter the **show route protocol rip** command.

```
user@R1> show route protocol rip

inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30      *[RIP/100] 00:54:34, metric 2, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32   *[RIP/100] 00:54:34, metric 2, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
192.168.3.3/32   *[RIP/100] 00:54:34, metric 3, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
224.0.0.9/32     *[RIP/100] 00:49:00, metric 1
                 MultiRecv
```

Meaning The output shows that the routes have been learned from Device R2 and Device R3.

If you delete or deactivate the import policy, the routing table contains the 172.16/16 routes.

Testing the Import Policy

Purpose By using the **test policy** command, monitor the number of rejected prefixes.

Action From operational mode, enter the **test policy rip-import 172.16/16** command.

```
user@R1> test policy rip-import 172.16/16
Policy rip-import: 0 prefix accepted, 1 prefix rejected
```

Meaning The output shows that the policy rejected one prefix.

Related Documentation

- [Example: Configuring RIP on page 2921](#)

Examples: Controlling Traffic with Metrics in a RIP Network

- [Understanding Traffic Control with Metrics in a RIP Network on page 2959](#)
- [Example: Controlling Traffic in a RIP Network with an Incoming Metric on page 2960](#)
- [Example: Controlling Traffic in a RIP Network with an Outgoing Metric on page 2962](#)
- [Example: Configuring the Metric Value Added to Imported RIP Routes on page 2963](#)

Understanding Traffic Control with Metrics in a RIP Network

To tune a RIP network and to control traffic flowing through the network, you increase or decrease the cost of the paths through the network. RIP provides two ways to modify the path cost: an incoming metric and an outgoing metric, which are each set to 1 by default. In other words, by default, the metric of routes that RIP imports from a neighbor or exports to a neighbor is incremented by 1. These routes include those learned from RIP as well as those learned from other protocols. The metrics are attributes that specify the cost of any route advertised through a host. By increasing or decreasing the metrics—and thus the cost—of links throughout the network, you can control packet transmission across the network.

The incoming metric modifies the cost of an individual segment when a route across the segment is imported into the routing table. For example, if you set the incoming metric on the segment to **3**, the individual segment cost along the link is changed from 1 to **3**. The increased cost affects all route calculations through that link. Other routes that were previously excluded because of a high hop count might now be included in the router's forwarding table.

The outgoing metric modifies the path cost for all the routes advertised out of a particular interface. Unlike the incoming metric, the outgoing metric modifies the routes that other routers are learning and thereby controls the way they send traffic.

If an exported route was learned from a member of the same RIP group, the metric associated with that route is the normal RIP metric. For example, a RIP route with a metric of 5 learned from a neighbor configured with an incoming metric of 2 is advertised with a combined metric of 7 when advertised to neighbors in the same group. However, if this route was learned from a RIP neighbor in a different group or from a different protocol, the route is advertised with the metric value configured in the outgoing metric for that group.

You might want to increase the metric of routes to decrease the likelihood that a particular route is selected and installed in the routing table. This process is sometimes referred to as *route poisoning*. Some reasons that you might want to poison a route are that the route is relatively expensive to use, or it has relatively low bandwidth.

A route with a higher metric than another route becomes the active route only when the lower-metric route becomes unavailable. In this way, the higher-metric route serves as a backup path.

One way to increase the metric of imported routes is to configure an import policy. Another way is to include the **metric-in** statement in the RIP neighbor configuration. One way to increase the metric of export routes is to configure an export policy. Another way is to include the **metric-out** statement in the RIP neighbor configuration.

Example: Controlling Traffic in a RIP Network with an Incoming Metric

This example shows how to control traffic with an incoming metric.

- [Requirements on page 2960](#)
- [Overview on page 2960](#)
- [Configuration on page 2961](#)
- [Verification on page 2961](#)

Requirements


Before you begin, define RIP groups, and add interfaces to the groups. Then configure a routing policy to export directly connected routes and routes learned through the RIP routing exchanges. See [“Example: Configuring a Basic RIP Network” on page 2922](#).

Overview

In this example, routes to Router D are received by Router A across both of its RIP-enabled interfaces as shown in [Figure 39 on page 2961](#). Because the route through Router B and the route through Router C have the same number of hops, both routes are imported into

the forwarding table. However, because the T3 link from Router B to Router D has a higher bandwidth than the T1 link from Router C to Router D, you want traffic to flow from Router A through Router B to Router D.

Figure 39: Controlling Traffic in a RIP Network with the Incoming Metric



To force this flow, you can modify the route metrics as they are imported into Router A's routing table. By setting the incoming metric on the interface from Router A to Router C, you modify the metric on all routes received through that interface. Setting the incoming route metric on Router A changes only the routes in Router A's routing table, and affects only how Router A sends traffic to Router D. Router D's route selection is based on its own routing table, which, by default, includes no adjusted metric values.

In the example, Router C receives a route advertisement from Router D and readvertises the route to Router A. When Router A receives the route, it applies the incoming metric on the interface. Instead of incrementing the metric by 1 (the default), Router A increments it by 3 (the configured incoming metric), giving the route from Router A to Router D through Router C a total path metric of 4. Because the route through Router B has a metric of 2, it becomes the preferred route for all traffic from Router A to Router D.

This example uses a RIP group called **alpha 1** on interface **g3-0/0/0**.

Configuration

Step-by-Step Procedure

To control traffic with an incoming metric:

1. Enable RIP on the interface.

```
[edit protocols rip]
user@host# set group alpha1 neighbor ge-0/0/0
```
2. Set the incoming metric.

```
[edit protocols rip]
user@host# set metric-in 3
```
3. If you are done configuring the device, commit the configuration.

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

Verification

To verify that the configuration is working properly, enter the **show route protocols rip** command.

Example: Controlling Traffic in a RIP Network with an Outgoing Metric

This example shows how to control traffic with an outgoing metric.

- [Requirements on page 2962](#)
- [Overview on page 2962](#)
- [Configuration on page 2963](#)
- [Verification on page 2963](#)

Requirements


Before you begin:

- Define RIP groups, and add interfaces to the groups. Then configure a routing policy to export directly connected routes and routes learned through RIP routing exchanges. See [“Example: Configuring a Basic RIP Network” on page 2922](#).
- Control traffic with an incoming metric. See [“Example: Controlling Traffic in a RIP Network with an Incoming Metric” on page 2960](#).

Overview

In this example, each route from Router A to Router D has two hops as shown in [Figure 40 on page 2962](#). However, because the link from Router A to Router B in the RIP group has a higher bandwidth than the link from Router A to Router C in RIP group Alpha 1, you want traffic from Router D to Router A to flow through Router B. To control the way Router D sends traffic to Router A, you can alter the routes that Router D receives by configuring the outgoing metric on Router A's interfaces in the Alpha 1 RIP group.

Figure 40: Controlling Traffic in a RIP Network with the Outgoing Metric



If the outgoing metric for the Alpha 1 RIP group—the A-to-C link—is changed to 3, Router D calculates the total path metric from Router A through Router C as 4. In contrast, the unchanged default total path metric to Router A through Router B in the RIP group is 2. The fact that Router A's interfaces belong to two different RIP groups allows you to configure two different outgoing metrics on its interfaces, because you configure path metrics at the group level.

By configuring the outgoing metric, you control the way Router A sends traffic to Router D. By configuring the outgoing metric on the same router, you control the way Router D sends traffic to Router A.

This example uses an outgoing metric of 3.

Configuration

Step-by-Step Procedure To control traffic with an outgoing metric:

1. Set the outgoing metric.

```
[edit protocols rip group alpha1]  
user@host# set metric-out 3
```
2. If you are done configuring the device, commit the configuration.

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

Verification

To verify that the configuration is working properly, enter the **show protocols rip** command.

Example: Configuring the Metric Value Added to Imported RIP Routes

This example shows how to change the default metric to be added to incoming routes to control the route selection process.

- [Requirements on page 2963](#)
- [Overview on page 2963](#)
- [Configuration on page 2964](#)
- [Verification on page 2967](#)

Requirements


No special configuration beyond device initialization is required before configuring this example.

Overview

Normally, when multiple routes are available, RIP selects the route with the lowest hop count. Changing the default metric enables you to control the route selection process such that a route with a higher hop count can be preferred over of a route with a lower hop count.

[Figure 41 on page 2964](#) shows the topology used in this example.

Figure 41: RIP Incoming Metrics Network Topology



Device R1 has two potential paths to reach 172.16.2.2/32. The default behavior is to send traffic out the 0.1/30 interface facing Device R2. Suppose, though, that the path through Device R3 is less expensive to use or has higher bandwidth links. This example shows how to use the **metric-in** statement to ensure that Device R1 uses the path through Device R3 to reach 172.16.2.2/32. “[CLI Quick Configuration](#)” on page 2964 shows the configuration for all of the devices in Figure 41 on page 2964. The section “[Step-by-Step Procedure](#)” on page 2965 describes the steps on Device R1.

Configuration

CLI Quick Configuration

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

Device R1

```
set interfaces fe-1/2/0 unit 1 description to-R2
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces ge-1/2/1 unit 10 description to-R3
set interfaces ge-1/2/1 unit 10 family inet address 10.0.0.10/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group primary export advertise-routes-through-rip
set protocols rip group primary neighbor ge-1/2/1.10
set protocols rip group secondary export advertise-routes-through-rip
set protocols rip group secondary neighbor fe-1/2/0.1 metric-in 4
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R2

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces ge-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor ge-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R3

```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces ge-1/2/1 unit 9 family inet address 10.0.0.9/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip group rip-group neighbor ge-1/2/1.9
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure a RIP metrics:

1. Configure the network interfaces.

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 1 description to-R2
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

user@R1# set ge-1/2/1 unit 10 description to-R3
user@R1# set ge-1/2/1 unit 10 family inet address 10.0.0.10/30

user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32

```

2. Create the RIP groups and add the interfaces.

To configure RIP in Junos OS, you must configure one or more groups that contain the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

For the interface that is facing Device R2, the **metric-in 4** setting causes this route to be less likely to be chosen as the active route.

```

[edit protocols rip]
user@R1# set group primary neighbor ge-1/2/1.10
user@R1# set group secondary neighbor fe-1/2/0.1 metric-in 4

```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```

[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept

```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```

[edit protocols rip]
user@R1# set group primary export advertise-routes-through-rip
user@R1# set group secondary export advertise-routes-through-rip

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    description to-R2;
    family inet {
      address 10.0.0.1/30;
    }
  }
}
ge-1/2/1 {
  unit 10 {
    description to-R3;
    family inet {
      address 10.0.0.10/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 172.16.0.1/32;
      address 192.168.1.1/32;
    }
  }
}

user@R1# show protocols
rip {
  group primary {
    export advertise-routes-through-rip;
    neighbor ge-1/2/1.10;
  }
  group secondary {
    export advertise-routes-through-rip;
    neighbor fe-1/2/0.1 {
      metric-in 4;
    }
  }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying That the Expected Route Is Active on page 2967](#)
- [Removing the metric-in Statement on page 2967](#)

Verifying That the Expected Route Is Active

Purpose Make sure that to reach 172.16.2.2/32, Device R1 uses the path through Device R3.

Action From operational mode, enter the **show route 172.16.2.2** command.

```
user@R1> show route 172.16.2.2
inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.2.2/32      *[RIP/100] 00:15:46, metric 3, tag 0
                  > to 10.0.0.9 via ge-1/2/1.10
```

Meaning The **to 10.0.0.9 via ge-1/2/1.10** output shows that Device R1 uses the path through Device R3 to reach 172.16.2.2/32. The metric for this route is 3.

Removing the metric-in Statement

Purpose Delete or deactivate the **metric-in** statement to see what happens to the 172.16.2.2/32 route.

Action 1. From configuration mode, deactivate the **metric-in** statement.

```
[edit protocols rip group secondary neighbor fe-1/2/0.1]
user@R1# deactivate metric-in
user@R1# commit
```

2. From operational mode, enter the **show route 172.16.2.2** command.

```
user@R1> show route 172.16.2.2
inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.2.2/32      *[RIP/100] 00:00:06, metric 2, tag 0
                  > to 10.0.0.2 via fe-1/2/0.1
```

Meaning The **to 10.0.0.2 via fe-1/2/0.1** output shows that Device R1 uses the path through Device R2 to reach 172.16.2.2/32. The metric for this route is 2.

Related Documentation

- [Example: Applying Policies to RIP Routes Imported from Neighbors on page 2954](#)

Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets

- [Understanding the Sending and Receiving of RIPv1 and RIPv2 Packets on page 2968](#)
- [Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 2968](#)

Understanding the Sending and Receiving of RIPv1 and RIPv2 Packets

RIP version 1 (RIPv1) and RIP version 2 (RIPv2) can run simultaneously. This might make sense when you are migrating a RIPv1 network to a RIPv2 network. This also allows interoperation with a device that supports RIPv1 but not RIPv2.

By default, when RIP is enabled on an interface, Junos OS receives both RIPv1 and RIPv2 packets and sends only RIPv2 packets. You can configure this behavior by including the [send](#) and [receive](#) statements in the RIP configuration.

Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets

This example shows how to configure whether the RIP update messages conform to RIP version 1 (RIPv1) only, to RIP version 2 (RIPv2) only, or to both versions. You can also disable the sending or receiving of update messages.

- [Requirements on page 2968](#)
- [Overview on page 2968](#)
- [Configuration on page 2969](#)
- [Verification on page 2971](#)

Requirements


No special configuration beyond device initialization is required before configuring this example.

Overview

By default, when RIP is enabled on an interface, Junos OS receives both RIPv1 and RIPv2 packets and sends only RIPv2 packets.

[Figure 42 on page 2968](#) shows the topology used in this example.

Figure 42: Sending and Receiving RIPv1 and RIPv2 Packets Network Topology



In this example, Device R1 is configured to receive only RIPv2 packets.

[“CLI Quick Configuration” on page 2969](#) shows the configuration for all of the devices in [Figure 42 on page 2968](#). The section [“Step-by-Step Procedure” on page 2969](#) describes the steps on Device R1.

Configuration

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

Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1 receive version-2
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R2

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure a RIP packet versions that can be received:

1. Configure the network interfaces.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
```

2. Create the RIP groups and add the interfaces.

To configure RIP in Junos OS, you must configure one or more groups that contain the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

For the interface that is facing Device R2, the **receive version-2** setting causes this interface to accept only RIPv2 packets.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1 receive version-2
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 172.16.0.1/32;
      address 192.168.1.1/32;
    }
  }
}

user@R1# show protocols
rip {
  group rip-group {
    export advertise-routes-through-rip;
    neighbor fe-1/2/0.1 {
      receive version-2;
    }
  }
}
```



```

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying That the Receive Mode Is Set to RIPv2 Only

Purpose Make sure that the interfacing Device R2 is configured to receive only RIPv2 packets, instead of both RIPv1 and RIPv2 packets.

Action From operational mode, enter the **show rip neighbor** command.

```
user@R1> show rip neighbor
```

Neighbor	Local State	Source Address	Destination Address	Send Mode	Receive Mode	In Met
fe-1/2/0.1	Up	10.0.0.1	224.0.0.9	mcast	v2 only	1

Meaning In the output, the **Receive Mode** field displays **v2 only**. The default **Receive Mode** is **both**.

Related Documentation

- [Example: Configuring RIP on page 2921](#)

Example: Redistributing Routes Among RIP Instances

- [Understanding Route Redistribution Among RIP instances on page 2971](#)
- [Example: Redistributing Routes Between Two RIP Instances on page 2972](#)

Understanding Route Redistribution Among RIP instances

You can redistribute routes among RIP processes. Another way to say this is to export RIP routes from one RIP instance to other RIP instances.

In Junos OS, route redistribution among routing instances is accomplished by using routing table groups, also called RIB groups. Routing table groups allow you to import and export routes from a protocol within one routing table into another routing table.



NOTE: In contrast, the policy-based import and export functions allow you import and export routes between different protocols within the same routing table.

Consider the following partial example:

```
protocols {
```

```
    rip {
      rib-group inet-to-voice;
    }
  }
  routing-instances {
    voice {
      protocols {
        rip {
          rib-group voice-to-inet;
        }
      }
    }
  }
  routing-options {
    rib-groups {
      inet-to-voice {
        import-rib [ inet.0 voice.inet.0 ];
      }
      voice-to-inet {
        import-rib [ voice.inet.0 inet.0 ];
      }
    }
  }
}
```

The way to read the **import-rib** statement is as follows. Take the routes from the protocol (RIP, in this case), and import them into the primary (or local) routing table and also into any other routing tables listed after this. The primary routing table is the routing table where the routing table group is being used. That would be either **inet.0** if used in the main routing instance or **voice.inet.0** if used within the routing instance. In the **inet-to-voice** routing table group, **inet.0** is listed first because this routing table group is used in the main routing instance. In the **voice-to-inet** routing table group, **voice.inet.0** is listed first because this routing table group is used in the voice routing instance.

Example: Redistributing Routes Between Two RIP Instances

This example shows how to configure a RIP routing instance and control the redistribution of RIP routes between the routing instance and the master instance.

- [Requirements on page 2972](#)
- [Overview on page 2972](#)
- [Configuration on page 2973](#)
- [Verification on page 2976](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.


Overview

When you create a routing instance called **voice**, Junos OS creates a routing table called **voice.inet.0**. The example shows how to install routes learned through the master RIP instance into the **voice.inet.0** routing table. The example also shows how to install routes learned through the voice routing instance into **inet.0**. This is done by configuring routing

table groups. RIP routes are installed into each routing table that belongs to a routing table group.

Figure 43 on page 2973 shows the topology used in this example.

Figure 43: Redistributing Routes Between RIP Instances Network Topology



"CLI Quick Configuration" on page 2973 shows the configuration for all of the devices in Figure 43 on page 2973. The section "Step-by-Step Procedure" on page 2974 describes the steps on Device R2.

Configuration

CLI Quick Configuration

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

```

Device R1
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group to-R2 export advertise-routes-through-rip
set protocols rip group to-R2 neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

Device R2
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip rib-group inet-to-voice
set protocols rip group to-R3 export advertise-routes-through-rip
set protocols rip group to-R3 neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set routing-instances voice protocols rip group to-R1 export advertise-routes-through-rip
set routing-instances voice interface fe-1/2/0.2
set routing-instances voice protocols rip rib-group voice-to-inet
set routing-instances voice protocols rip group to-R1 neighbor fe-1/2/0.2
set routing-options rib-groups inet-to-voice import-rib inet.0
  
```

```
set routing-options rib-groups inet-to-voice import-rib voice.inet.0
set routing-options rib-groups voice-to-inet import-rib voice.inet.0
set routing-options rib-groups voice-to-inet import-rib inet.0
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group to-R2 export advertise-routes-through-rip
set protocols rip group to-R2 neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To redistribute RIP routes between routing instances:

1. Configure the network interfaces.

```
[edit interfaces]
user@R2# set fe-1/2/0 unit 2 family inet address 10.0.0.2/30

user@R2# set fe-1/2/1 unit 5 family inet address 10.0.0.5/30

user@R2# set lo0 unit 2 family inet address 192.168.2.2/32
user@R2# set lo0 unit 2 family inet address 172.16.2.2/32
```
2. Create the routing instance, and add one or more interfaces to the routing instance.

```
[edit routing-instances voice]
user@R2# set interface fe-1/2/0.2
```
3. Create the RIP groups and add the interfaces.

```
[edit protocols rip group to-R3]
user@R2# set neighbor fe-1/2/1.5

[edit routing-instances voice protocols rip group to-R1]
user@R2# set neighbor fe-1/2/0.2
```
4. Create the routing table groups.

```
[edit routing-options rib-groups]
user@R2# set inet-to-voice import-rib inet.0
user@R2# set inet-to-voice import-rib voice.inet.0

user@R2# set voice-to-inet import-rib voice.inet.0
user@R2# set voice-to-inet import-rib inet.0
```
5. Apply the routing table groups.

```
[edit protocols rip]
user@R2# set rib-group inet-to-voice
```

```
[edit routing-instances voice protocols rip]
user@R2# set rib-group voice-to-inet
```

6. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R2# set from protocol direct
user@R2# set from protocol rip
user@R2# set then accept
```

7. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group to-R3]
user@R2# set export advertise-routes-through-rip
```

```
[edit routing-instances voice protocols rip group to-R1]
user@R2# set export advertise-routes-through-rip
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, **show routing-instances**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 2 {
    family inet {
      address 10.0.0.2/30;
    }
  }
}
fe-1/2/1 {
  unit 5 {
    family inet {
      address 10.0.0.5/30;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.168.2.2/32;
      address 172.16.2.2/32;
    }
  }
}

user@R2# show protocols
rip {
  rib-group inet-to-voice;
  group to-R3 {
    export advertise-routes-through-rip;
    neighbor fe-1/2/1.5;
  }
}
```

```
user@R2# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}

user@R2# show routing-instances
voice {
  interface fe-1/2/0.2;
  protocols {
    rip {
      rib-group voice-to-inet;
      group to-R1 {
        export advertise-routes-through-rip;
        neighbor fe-1/2/0.2;
      }
    }
  }
}

user@R2# show routing-options
rib-groups {
  inet-to-voice {
    import-rib [ inet.0 voice.inet.0 ];
  }
  voice-to-inet {
    import-rib [ voice.inet.0 inet.0 ];
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Checking the Routing Tables

Purpose Make sure that the routing tables contain the expected routes.

Action From operational mode, enter the **show route protocol rip** command.

```
user@R2> show route protocol rip
inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.1/32      *[RIP/100] 01:58:14, metric 2, tag 0
                  > to 10.0.0.1 via fe-1/2/0.2
172.16.3.3/32     *[RIP/100] 02:06:03, metric 2, tag 0
                  > to 10.0.0.6 via fe-1/2/0.5
192.168.1.1/32    *[RIP/100] 01:58:14, metric 2, tag 0
                  > to 10.0.0.1 via fe-1/2/0.2
192.168.3.3/32    *[RIP/100] 02:06:03, metric 2, tag 0
                  > to 10.0.0.6 via fe-1/2/0.5
224.0.0.9/32      *[RIP/100] 01:44:13, metric 1
                  MultiRecv
```

```
voice.inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
172.16.0.1/32      *[RIP/100] 02:06:03, metric 2, tag 0
                  > to 10.0.0.1 via fe-1/2/0.2
172.16.3.3/32      *[RIP/100] 01:58:14, metric 2, tag 0
                  > to 10.0.0.6 via fe-1/2/0.5
192.168.1.1/32     *[RIP/100] 02:06:03, metric 2, tag 0
                  > to 10.0.0.1 via fe-1/2/0.2
192.168.3.3/32     *[RIP/100] 01:58:14, metric 2, tag 0
                  > to 10.0.0.6 via fe-1/2/0.5
224.0.0.9/32       *[RIP/100] 01:44:13, metric 1
                  MultiRecv
```

Meaning The output shows that both routing tables contain all of the RIP routes.

Related Documentation

- [Example: Configuring RIP on page 2921](#)
- [Example: Applying Policies to RIP Routes Imported from Neighbors on page 2954](#)

Example: Configuring RIP Timers

- [Understanding RIP Timers on page 2977](#)
- [Example: Configuring RIP Timers on page 2978](#)

Understanding RIP Timers

RIP uses several timers to regulate its operation.

The update interval is the interval at which routes that are learned by RIP are advertised to neighbors. This timer controls the interval between routing updates. The update interval is set to 30 seconds, by default, with a small random amount of time added when the timer is reset. This added time prevents congestion that can occur if all routing devices update their neighbors simultaneously.

To configure the update time interval, include the **update-interval** statement:

```
update-interval seconds;
```

seconds can be a value from 10 through 60.

You can set a route timeout interval. If a route is not refreshed after being installed in the routing table by the specified time interval, the route is marked as invalid and is removed from the routing table after the hold-down period expires.

To configure the route timeout for RIP, include the **route-timeout** statement:

```
route-timeout seconds;
```

seconds can be a value from 30 through 360. The default value is 180 seconds.

RIP routes expire when either a route timeout limit is met or a route metric reaches infinity, and the route is no longer valid. However, the expired route is retained in the routing table for a specified period so that neighbors can be notified that the route has been dropped.

This time period is set by configuring the hold-down timer. Upon expiration of the hold-down timer, the route is removed from the routing table.

To configure the hold-down timer for RIP, include the **holddown** statement:

holddown *seconds*;

seconds can be a value from 10 through 180. The default value is 120 seconds.



NOTE: In Junos OS Release 11.1 and later, a retransmission timer is available for RIP demand circuits.

Generally, we recommend against changing the RIP timers, unless the effects of a change are well understood. The route timeout should be at least three times the update interval. The hold-down timer must be greater than the route timeout. Normally, the default values are best left in effect for standard operations.

Example: Configuring RIP Timers

This example shows how to configure the RIP update interval and how to monitor the impact of the change.

- [Requirements on page 2978](#)
- [Overview on page 2978](#)
- [Configuration on page 2979](#)
- [Verification on page 2981](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview


In this example, Device R2 has an update interval of 60 seconds for its neighbor, Device R1, and an update interval of 10 seconds for its neighbor, Device R3.

This example is not necessarily practical, but it is shown for demonstration purposes. Generally, we recommend against changing the RIP timers, unless the effects of a change are well understood. Normally, the default values are best left in effect for standard operations.

An export policy is also shown because an export policy is required as part of the minimum configuration for RIP.

[Figure 44 on page 2979](#) shows the topology used in this example.

Figure 44: RIP Timers Network Topology



"CLI Quick Configuration" on page 2979 shows the configuration for all of the devices in Figure 44 on page 2979. The section "Step-by-Step Procedure" on page 2980 describes the steps on Device R2.

Configuration

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

Device R1

```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

Device R2

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2 update-interval 60
set protocols rip group rip-group neighbor fe-1/2/1.5 update-interval 10
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

Device R3

```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip

```

```
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure the RIP update interval:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```
[edit interfaces]
user@R2# set fe-1/2/0 unit 2 family inet address 10.0.0.2/30

user@R2# set fe-1/2/1 unit 5 family inet address 10.0.0.5/30
```

```
user@R2# set lo0 unit 2 family inet address 192.168.2.2/32
user@R2# set lo0 unit 2 family inet address 172.16.2.2/32
```

2. Configure different update intervals for the two RIP neighbors.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R2# set neighbor fe-1/2/0.2 update-interval 60
user@R2# set neighbor fe-1/2/1.5 update-interval 10
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R2# set from protocol direct
user@R2# set from protocol rip
user@R2# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R2# set export advertise-routes-through-rip
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 2 {
    family inet {
      address 10.0.0.2/30;
    }
  }
}
```

```

fe-1/2/1 {
  unit 5 {
    family inet {
      address 10.0.0.5/30;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.168.2.2/32;
      address 172.16.2.2/32;
    }
  }
}

user@R2# show protocols
rip {
  group rip-group {
    export advertise-routes-through-rip;
    neighbor fe-1/2/0.2 {
      update-interval 60;
    }
    neighbor fe-1/2/1.5 {
      update-interval 10;
    }
  }
}

user@R2# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Checking the RIP Updates Sent by Device R2 on page 2981](#)
- [Checking the RIP Updates Received by Device R2 on page 2982](#)
- [Checking the RIP Updates Received by Device R3 on page 2983](#)

Checking the RIP Updates Sent by Device R2

Purpose Make sure that the RIP update packets are sent at the expected interval.

Action From operational mode, enter the **show rip statistics** command.

```

user@R2> show rip statistics
RIPv2 info: port 520; holddown 120s.
      rts learned  rts held down  rqsts dropped  resps dropped
              4              2              0              0

```

fe-1/2/0.2: 2 routes learned; 5 routes advertised; timeout 180s; update interval 60s

Counter	Total	Last 5 min	Last minute
-----	-----	-----	-----
Updates Sent	123	5	1
Triggered Updates Sent	0	0	0
Responses Sent	0	0	0
Bad Messages	0	0	0
RIPv1 Updates Received	0	0	0
RIPv1 Bad Route Entries	0	0	0
RIPv1 Updates Ignored	0	0	0
RIPv2 Updates Received	244	10	2
RIPv2 Bad Route Entries	0	0	0
RIPv2 Updates Ignored	0	0	0
Authentication Failures	0	0	0
RIP Requests Received	0	0	0
RIP Requests Ignored	0	0	0
none	0	0	0

fe-1/2/1.5: 2 routes learned; 5 routes advertised; timeout 180s; update interval 10s

Counter	Total	Last 5 min	Last minute
-----	-----	-----	-----
Updates Sent	734	32	6
Triggered Updates Sent	0	0	0
Responses Sent	0	0	0
Bad Messages	0	0	0
RIPv1 Updates Received	0	0	0
RIPv1 Bad Route Entries	0	0	0
RIPv1 Updates Ignored	0	0	0
RIPv2 Updates Received	245	11	2
RIPv2 Bad Route Entries	0	0	0
RIPv2 Updates Ignored	0	0	0
Authentication Failures	0	0	0
RIP Requests Received	0	0	0
RIP Requests Ignored	0	0	0
none	0	0	0

Meaning The **update interval** field shows that the interval is 60 seconds for Neighbor R1 and 10 seconds for Neighbor R3. The **Updates Sent** field shows that Device R2 is sending updates to Device R1 at roughly 1/6 of the rate that it is sending updates to Device R3.

Checking the RIP Updates Received by Device R2

Purpose Make sure that the RIP update packets are sent at the expected interval.

Action From operational mode, enter the **show rip statistics** command.

```
user@R1> show rip statistics
```

```
RIPv2 info: port 520; holddown 120s.
```

```
    rts learned  rts held down  rqsts dropped  resps dropped
          5             0             0             0
```

fe-1/2/0.1: 5 routes learned; 2 routes advertised; timeout 180s; update interval 30s

Counter	Total	Last 5 min	Last minute
-----	-----	-----	-----
Updates Sent	312	10	2

Triggered Updates Sent	2	0	0
Responses Sent	0	0	0
Bad Messages	0	0	0
RIPv1 Updates Received	0	0	0
RIPv1 Bad Route Entries	0	0	0
RIPv1 Updates Ignored	0	0	0
RIPv2 Updates Received	181	5	1
RIPv2 Bad Route Entries	0	0	0
RIPv2 Updates Ignored	0	0	0
Authentication Failures	0	0	0
RIP Requests Received	1	0	0
RIP Requests Ignored	0	0	0
none	0	0	0

Meaning The **RIPv2 Updates Received** field shows the number of updates received from Device R2.

Checking the RIP Updates Received by Device R3

Purpose Make sure that the RIP update packets are sent at the expected interval.

Action From operational mode, enter the **show rip statistics** command.

```
user@R3> show rip statistics
```

```
RIPv2 info: port 520; holddown 120s.
```

```
    rts learned  rts held down  rqsts dropped  resps dropped
           5             0             0             0
```

```
fe-1/2/0.6: 5 routes learned; 2 routes advertised; timeout 180s; update interval 30s
```

Counter	Total	Last 5 min	Last minute
-----	-----	-----	-----
Updates Sent	314	11	2
Triggered Updates Sent	1	0	0
Responses Sent	0	0	0
Bad Messages	0	0	0
RIPv1 Updates Received	0	0	0
RIPv1 Bad Route Entries	0	0	0
RIPv1 Updates Ignored	0	0	0
RIPv2 Updates Received	827	31	6
RIPv2 Bad Route Entries	0	0	0
RIPv2 Updates Ignored	0	0	0
Authentication Failures	0	0	0
RIP Requests Received	0	0	0
RIP Requests Ignored	0	0	0
none	0	0	0

Meaning The **RIPv2 Updates Received** field shows the number of updates received from Device R2.

Related Documentation

- [Example: Configuring RIP on page 2921](#)
- [Example: Configuring RIP Demand Circuits on page 2984](#)

Example: Configuring RIP Demand Circuits

- [RIP Demand Circuits Overview on page 2984](#)
- [Example: Configuring RIP Demand Circuits on page 2986](#)

RIP Demand Circuits Overview

RIP periodically sends routing information (RIP packets) to neighboring devices. These periodic broadcasts can consume bandwidth resources and interfere with network traffic by preventing WAN circuits from being closed. Demand circuits for RIP is defined in RFC 2091 and overcomes these issues by exchanging incremental updates on demand.

A demand circuit is a point-to-point connection between two neighboring interfaces configured for RIP. Demand circuits preserve bandwidth by establishing a link when data needs to be transferred, and terminating the link when the data transfer is complete. Demand circuits increase the efficiency of RIP on the configured interfaces by offering minimal network overhead in terms of messages passed between the demand circuit end points, thus conserving resources and reducing costs.

By configuring RIP demand circuits, a specific event triggers the device to send an update, thereby eliminating the periodic transmission of RIP packets over the neighboring interface. To save overhead, the device sends RIP information only when changes occur in the routing database, such as:

- The device is first powered on
- The device receives a request for route update information
- A change occurs in the network
- The demand circuit goes down or comes up

The device sends update requests, update responses, and acknowledgments. In addition, the device retransmits updates and requests until valid acknowledgments are received. The device dynamically learns RIP neighbors. If the neighboring interface goes down, RIP flushes routes learned from the neighbor's IP address.

Routes learned from demand circuits do not age like other RIP entries because demand circuits are in a permanent state. Routes in a permanent state are only removed under the following conditions:

- A formerly reachable route changes to unreachable in an incoming response
- The demand circuit is down due to an excessive number of unacknowledged retransmissions

You can also set the RIP hold-down timer and the RIP demand circuit retransmission timer to regulate performance. The demand circuit uses these timers to determine if there is a change that requires update messages to be sent. There is also a database timer that runs only when RIP flushes learned routes from the routing table.

This topic includes the following sections:

- [RIP Demand Circuit Packets on page 2985](#)

- [Timers Used by RIP Demand Circuits on page 2985](#)

RIP Demand Circuit Packets

When you configure an interface for RIP demand circuits, the supported command field packet types are different than those for RIP version 1 and RIP version 2. RIP packets for RIP demand circuits contain three additional packet types and an extended 4-byte update header. Both RIP version 1 and RIP version 2 support the three packet types and the extended 4-byte header. [Table 241 on page 2985](#) describes the three packet types.

Table 241: RIP Demand Circuit Packet Types

Packet Type	Description
Update Request	Update request messages seek information for the device's routing table. This message is sent when the device is first powered on or when a down demand circuit comes up. The device sends this message every 5 seconds (by default) until an update response message is received.
Update Response	Update response messages are sent in response to an update request message, which occurs when the device is first powered on or when a down demand circuit comes up. Each update response message contains a sequence number that the neighbor uses to acknowledge the update request.
Update Acknowledge	Update acknowledge messages are sent in response to every update response message received by the neighbor.



NOTE: These packets are only valid on interfaces configured for RIP demand circuits. If a demand circuit receives a RIP packet that does not contain these packet types, it silently discards the packet and logs an error message similar to the following:

Ignoring RIP packet with invalid version 0 from neighbor 10.0.0.0 and source 10.0.0.1

Timers Used by RIP Demand Circuits

RIP demand circuits use the RIP hold-down timer and the RIP demand circuit retransmission timer to regulate performance and to determine if there is a change in the network that requires the device to send update messages. The hold-down timer is a global RIP timer that affects the entire RIP configuration. Whatever range you configure for RIP applies to RIP demand circuits. The retransmission timer affects only RIP demand circuits. In addition, there is a database timer that runs only when RIP flushes learned routes from the routing table.

- **Hold-down timer (global RIP timer)**—Use the hold-down timer to configure the number of seconds that RIP waits before updating the routing table. The value of the hold-down timer affects the entire RIP configuration, not just the demand circuit interfaces. The hold-down timer starts when a route timeout limit is met, when a formerly reachable route is unreachable, or when a demand circuit interface is down. When the hold-down

timer is running, routes are advertised as unreachable on other interfaces. When the hold-down timer expires, the route is removed from the routing table if all destinations are aware that the route is unreachable or the remaining destinations are down. By default, RIP waits 120 seconds between routing table updates. The range is from 10 to 180 seconds.

- Retransmission timer (RIP demand circuit timer)—RIP demand circuits send update messages every 5 seconds to an unresponsive peer. Use the retransmission timer to limit the number of times a demand circuit resends update messages to an unresponsive peer. If the configured retransmission threshold is reached, routes from the next hop router are marked as unreachable and the hold-down timer starts. The value of the retransmission timer affects only the demand circuit interfaces. To determine the number of times to resend the update message, use the following calculation:

$$5 \text{ seconds} \times \text{number of retransmissions} = \text{retransmission seconds}$$

The retransmission range is from 5 through 180 seconds, which corresponds to sending an update message a minimum of 1 time (5 seconds) and a maximum of 36 times (180 seconds).

- Database timer (global timeout timer)—Routes learned from demand circuits do not age like other RIP entries because demand circuits are in a permanent state. On a RIP demand circuit, the database timer starts upon receipt of the update response message with the flush flag sent from a RIP demand circuit peer. When the neighbor receives this message, all routes from that peer are flushed, and the database timer starts and runs for the configured route timeout interval. When the database timer is running, routes are still advertised as reachable on other interfaces. When the database timer expires, the device advertises all routes from its peer as unreachable.

Example: Configuring RIP Demand Circuits

This example describes how to configure an interface as a RIP demand circuit.

- [Requirements on page 2986](#)
- [Overview on page 2986](#)
- [Configuration on page 2987](#)
- [Verification on page 2988](#)

Requirements

Before you begin, configure the device interfaces. See the *Router Interfaces* or the *Junos OS Interfaces Configuration Guide for Security Devices*.

Overview

A demand circuit is a point-to-point connection between two neighboring interfaces configured for RIP. Demand circuits increase the efficiency of RIP on the configured interfaces by eliminating the periodic transmission of RIP packets. Demand circuits preserve bandwidth by establishing a link when data needs to be transferred, and terminating the link when the data transfer is complete. In this example, two devices are connected using SONET/SDH interfaces.



NOTE: When you configure RIP demand circuits, any silent removal of the RIP configuration goes unnoticed by the RIP peer and leads to stale entries in the routing table. To clear the stale entries, deactivate and reactivate RIP on the neighboring devices.

In this example, you configure interface **so-0/1/0** with the following settings:

- **demand-circuit**—Configures the interface as a demand circuit. To complete the demand circuit, you must configure both ends of the pair as demand circuits.
- **max-retrans-time**—RIP demand circuits send update messages every 5 seconds to an unresponsive peer. Use the retransmission timer to limit the number of times a demand circuit resends update messages to an unresponsive peer. If the configured retransmission threshold is reached, routes from the next-hop router are marked as unreachable, and the hold-down timer starts. The value of the retransmission timer affects only the demand circuit interfaces. To determine the number of times to resend the update message, use the following calculation:

$$5 \text{ seconds} \times \text{retransmissions} = \text{retransmission seconds}$$

For example, if you want the demand circuit to send only two update messages to an unresponsive peer, the calculation is: $5 \times 2 = 10$. When you configure the retransmission timer, you enter 10 seconds.

The retransmission range is from 5 through 180 seconds, which corresponds to sending an update message a minimum of 1 time (5 seconds) and a maximum of 36 times (180 seconds).

Configuration

In the following example, you configure a neighboring interface to be a RIP demand circuit and save the configuration.

CLI Quick Configuration

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

```
set interfaces so-0/1/0 unit 0 family inet address 192.0.2.0/24
set protocols rip group group1 neighbor so-0/1/0 demand-circuit
set protocols rip group group1 neighbor so-0/1/0 max-retrans-time 10
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure a RIP demand circuit on one neighboring interface:

1. Configure the interface.

```
[edit interfaces]
```

```
user@host# set so-0/1/0 unit 0 family inet address 192.0.2.0/24
```

2. Configure the neighbor as a demand circuit.

```
[edit protocols rip]
user@host# set group group1 neighbor so-0/1/0 demand-circuit
```

3. Configure the demand circuit retransmission timer.

```
[edit protocols rip]
user@host# set group group1 neighbor so-0/1/0 max-retrans-time 10
```

4. If you are done configuring the device, commit the configuration.

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



NOTE: Repeat this entire configuration on the other neighboring interface.

Results

Confirm your configuration by entering the **show interfaces** and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
so-0/1/0 {
  unit 0 {
    family inet {
      address 192.0.2.0/24;
    }
  }
}

user@host# show protocols
rip {
  group group1 {
    neighbor so-0/1/0 {
      demand-circuit;
      max-retrans-time 10;
    }
  }
}
```

Verification

Confirm that the configuration is working properly.

Verifying a Demand Circuit Configuration

Purpose Verify that the demand circuit configuration is working.

Action To verify that the demand circuit configuration is in effect, use the **show rip neighbor** operational mode command.

```
user@host> show rip neighbor
```

Neighbor	State	Source Address	Destination Address	Send Mode	Receive Mode	In Met
so-0/1/0.0(DC)	Up	10.10.10.2	224.0.0.9	mcast	both	1

When you configure demand circuits, the **show rip neighbor** command displays a DC flag next to the neighboring interface configured for demand circuits.



NOTE: If you configure demand circuits at the [edit protocols rip group *group-name* neighbor *neighbor-name*] hierarchy level, the output shows only the neighboring interface that you specifically configured as a demand circuit. If you configure demand circuits at the [edit protocols rip group *group-name*] hierarchy level, all of the interfaces in the group are configured as demand circuits. Therefore, the output shows all of the interfaces in that group as demand circuits.

Related Documentation

- [Example: Configuring RIP Timers on page 2977](#)

Example: Tracing RIP Protocol Traffic

- [Understanding RIP Trace Operations on page 2989](#)
- [Example: Tracing RIP Protocol Traffic on page 2990](#)

Understanding RIP Trace Operations

You can trace various types of RIP protocol traffic to help debug RIP protocol issues.

To trace RIP protocol traffic, include the **traceoptions** statement at the [edit protocols rip] hierarchy level:

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

You can specify the following RIP protocol-specific trace options using the **flag** statement:

- **auth**—RIP authentication
- **error**—RIP error packets
- **expiration**—RIP route expiration processing
- **holddown**—RIP hold-down processing
- **nsr-synchronization**—Nonstop active routing synchronization events
- **packets**—All RIP packets
- **request**—RIP information packets
- **trigger**—RIP triggered updates
- **update**—RIP update packets

You can optionally specify one or more of the following flag modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted



NOTE: Use the **detail** flag modifier with caution as this may cause the CPU to become very busy.

Global tracing options are inherited from the configuration set by the **traceoptions** statement at the **[edit routing-options]** hierarchy level. You can override the following global trace options for the RIP protocol using the **traceoptions flag** statement included at the **[edit protocols rip]** hierarchy level:

- **all**—All tracing operations
- **general**—All normal operations and routing table changes (a combination of the normal and route trace operations)
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing



NOTE: Use the trace flag **all** with caution because this may cause the CPU to become very busy.

Example: Tracing RIP Protocol Traffic

This example shows how to trace RIP protocol operations.

- [Requirements on page 2990](#)
- [Overview on page 2990](#)
- [Configuration on page 2991](#)
- [Verification on page 2993](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.


Overview

In this example, Device R1 is set to trace routing information updates.

An export policy is also shown because an export policy is required as part of the minimum configuration for RIP.

[Figure 45 on page 2991](#) shows the topology used in this example.

Figure 45: RIP Trace Operations Network Topology



[“CLI Quick Configuration” on page 2991](#) shows the configuration for all of the devices in [Figure 45 on page 2991](#). The section [“Step-by-Step Procedure” on page 2992](#) describes the steps on Device R1.

Configuration

CLI Quick Configuration

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

- Device R1**
- ```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip traceoptions file rip-trace-file
set protocols rip traceoptions flag route
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```
- Device R2**
- ```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```
- Device R3**
- ```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32

```

```
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
 rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Step-by-Step  
Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure the RIP update interval:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

```
user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
```

2. Configure the RIP group, and add the interface to the group.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1
```

3. Configure RIP tracing operations.

```
[edit protocols rip traceoptions]
user@R1# set file rip-trace-file
user@R1# set flag route
```

4. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

5. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

user@R1# show interfaces
fe-1/2/0 {
 unit 1 {
 family inet {
 address 10.0.0.1/30;
 }
 }
}
lo0 {
 unit 1 {
 family inet {
 address 172.16.0.1/32;
 address 192.168.1.1/32;
 }
 }
}

user@R1# show protocols
rip {
 traceoptions {
 file rip-trace-file;
 flag route;
 }
 group rip-group {
 export advertise-routes-through-rip;
 neighbor fe-1/2/0.1;
 }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
 term 1 {
 from protocol [direct rip];
 then accept;
 }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

### Checking the Log File

**Purpose** Make sure that the RIP route updates are logged in the configured log file.

**Action** 1. Deactivate the extra loopback interface address on Device R3.

```

[edit interfaces lo0 unit 3 family inet]
user@R3# deactivate address 172.16.3.3/32
user@R3# commit

```

2. From operational mode on Device R1, enter the **show log rip-trace-file** command with the **| match 172.16.3.3** option.

```

user@R1> show log rip-trace-file | match 172.16.3.3

```

```

Mar 1 11:39:53.975192 Setting RIPv2 rtbit on route 172.16.3.3/32, tsi =
0xbb69228
Mar 1 11:39:59.847118 172.16.3.3/32: metric-in: 16, change: 3 -> 16; # gw:
1, pkt_upd_src 10.0.0.2, inx: 0, rte_upd_src 10.0.0.2
Mar 1 11:39:59.847568 CHANGE 172.16.3.3/32 nhid 591 gw 10.0.0.2
RIP pref 100/0 metric 3/0 fe-1/2/0.1 <Delete Int>
Mar 1 11:39:59.847629 Best route to 172.16.3.3/32 got deleted. Doing route calculation
on the stored rte-info

```

**Meaning** The output shows that the route to 172.16.3.3/32 was deleted.

**Related Documentation**

- [Example: Configuring RIP on page 2921](#)

### *Verifying a RIP Configuration*

To verify a RIP configuration, perform the following tasks:

- [Verifying the Exchange of RIP Messages on page 2994](#)
- [Verifying the RIP-Enabled Interfaces on page 2995](#)
- [Verifying Reachability of All Hosts in the RIP Network on page 2996](#)

### *Verifying the Exchange of RIP Messages*

**Purpose** Verify that RIP messages are being sent and received on all RIP-enabled interfaces.

**Action** From the CLI, enter the **show rip statistics** command.

## Sample Output

```

user@host> show rip statistics
RIPv2 info: port 520; holddown 120s.
 rts learned rts held down rqsts dropped resps dropped
 10 0 0 0

t1-0/0/2.0: 0 routes learned; 13 routes advertised; timeout 120s; update interval
45s
Counter Total Last 5 min Last minute

Updates Sent 2855 11 2
Triggered Updates Sent 5 0 0
Responses Sent 0 0 0
Bad Messages 0 0 0
RIPv1 Updates Received 0 0 0
RIPv1 Bad Route Entries 0 0 0
RIPv1 Updates Ignored 0 0 0
RIPv2 Updates Received 41 0 0
RIPv2 Bad Route Entries 0 0 0
RIPv2 Updates Ignored 0 0 0
Authentication Failures 0 0 0
RIP Requests Received 0 0 0
RIP Requests Ignored 0 0 0

ge-0/0/1.0: 10 routes learned; 3 routes advertised; timeout 180s; update interval
30s
Counter Total Last 5 min Last minute

```



|                         |      |    |   |
|-------------------------|------|----|---|
| Updates Sent            | 2855 | 11 | 2 |
| Triggered Updates Sent  | 3    | 0  | 0 |
| Responses Sent          | 0    | 0  | 0 |
| Bad Messages            | 1    | 0  | 0 |
| RIPv1 Updates Received  | 0    | 0  | 0 |
| RIPv1 Bad Route Entries | 0    | 0  | 0 |
| RIPv1 Updates Ignored   | 0    | 0  | 0 |
| RIPv2 Updates Received  | 2864 | 11 | 2 |
| RIPv2 Bad Route Entries | 14   | 0  | 0 |
| RIPv2 Updates Ignored   | 0    | 0  | 0 |
| Authentication Failures | 0    | 0  | 0 |
| RIP Requests Received   | 0    | 0  | 0 |
| RIP Requests Ignored    | 0    | 0  | 0 |

**Meaning** The output shows the number of RIP routes learned. It also shows the number of RIP updates sent and received on the RIP-enabled interfaces. Verify the following information:

- The number of RIP routes learned matches the number of expected routes learned. Subnets learned by direct connectivity through an outgoing interface are not listed as RIP routes.
- RIP updates are being sent on each RIP-enabled interface. If no updates are being sent, the routing policy might not be configured to export routes.
- RIP updates are being received on each RIP-enabled interface. If no updates are being received, the routing policy might not be configured to export routes on the host connected to that subnet. The lack of updates might also indicate an authentication error.

#### *Verifying the RIP-Enabled Interfaces*

**Purpose** Verify that all the RIP-enabled interfaces are available and active.

**Action** From the CLI, enter the **show rip neighbor** command.

### Sample Output

```
user@host> show rip neighbor
Source Destination Send Receive In
Neighbor State Address Address Mode Mode Met

ge-0/0/0.0 Dn (null) (null) mcast both 1
ge-0/0/1.0 Up 192.168.220.5 224.0.0.9 mcast both 1
```

**Meaning** The output shows a list of the RIP neighbors that are configured on the device. Verify the following information:

- Each configured interface is present. Interfaces are listed in alphabetical order.
- Each configured interface is up. The state of the interface is listed in the **Destination State** column. A state of **Up** indicates that the link is passing RIP traffic. A state of **Dn** indicates that the link is not passing RIP traffic. In a point-to-point link, this state generally means that either the end point is not configured for RIP or the link is unavailable.

### ***Verifying Reachability of All Hosts in the RIP Network***

**Purpose** By using the traceroute tool on each loopback address in the network, verify that all hosts in the RIP network are reachable from each Juniper Networks device.

**Action** For each device in the RIP network:

1. In the J-Web interface, select **Troubleshoot>Traceroute**.
2. In the Remote Host box, type the name of a host for which you want to verify reachability from the device.
3. Click **Start**. Output appears on a separate page.

### **Sample Output**

```
1 172.17.40.254 (172.17.40.254) 0.362 ms 0.284 ms 0.251 ms
2 routera-fxp0.eng1ab.mycompany.net (192.168.71.246) 0.251 ms 0.235 ms 0.200 ms
```

**Meaning** Each numbered row in the output indicates a routing hop in the path to the host. The three-time increments indicate the round-trip time (RTT) between the device and the hop for each traceroute packet.

To ensure that the RIP network is healthy, verify the following information:

- The final hop in the list is the host you want to reach.
- The number of expected hops to the host matches the number of hops in the traceroute output. The appearance of more hops than expected in the output indicates that a network segment is probably unreachable. It might also indicate that the incoming or outgoing metric on one or more hosts has been set unexpectedly.

#### **Related Documentation**

- *Junos OS Feature Support Reference for SRX Series and J Series Devices*
- [RIP Configuration Overview on page 2920](#)
- [show rip statistics on page 3060](#) in the *Junos OS Operational Mode Commands*
- [show rip neighbor on page 3058](#) in the *Junos OS Operational Mode Commands*
- [traceroute on page 4134](#) in the *Junos OS Operational Mode Commands*
- [RIP Overview on page 2915](#)

---

### **Configuration Statements**

- [\[edit protocols rip\] Hierarchy Level on page 2996](#)
- [\[edit protocols ripng\] Hierarchy Level on page 2998](#)

#### ***[edit protocols rip] Hierarchy Level***

The following statement hierarchy can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```
protocols {
```

```

rip {
 authentication-key password;
 authentication-type type;
 (check-zero | no-check-zero);
 graceful-restart {
 disable;
 restart-time seconds;
 }
 group group-name {
 ... the group subhierarchy appears after the main [edit protocols rip] hierarchy ...
 }
 holddown seconds;
 import [policy-names];
 message-size number;
 metric-in metric;
 receive (both | none | version-1 | version-2);
 rib-group group-name;
 route-timeout seconds;
 send (broadcast | multicast | none | version-1);
 traceoptions {
 file filename <files number> <size maximum-file-size> <world-readable |
 no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
 update-interval seconds;
}

rip {
 group group-name {
 bfd-liveness-detection {
 authentication {
 algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
 meticulous-keyed-sha-1 | simple-password);
 key-chain key-chain-name;
 loose-check;
 }
 detection-time {
 threshold milliseconds;
 }
 }
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 version (1 | automatic);
 }
 demand-circuit;
 export [policy-names];
 import [policy-names];
 max-retrans-time seconds;
 metric-out metric;
 neighbor interface-name {

```

```
... the neighbor subhierarchy appears after the main [edit protocols rip group
group-name] hierarchy level ...
}
preference preference;
route-timeout seconds;
update-interval seconds;
}

group group-name {
 neighbor neighbor-name {
 any-sender;
 authentication-key password;
 authentication-type type;
 bfd-liveness-detection {
 ... same statements as at the [edit protocols rip group group-name
 bfd-liveness-detection] hierarchy level ...
 }
 (check-zero | no-check-zero);
 demand-circuit;
 import [policy-names];
 max-retrans-time seconds;
 message-size number;
 metric-in metric;
 receive (both | none | version-1 | version-2);
 route-timeout seconds;
 send (broadcast | multicast | none | version-1);
 update-interval seconds;
 }
}
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
  - *[edit protocols] Hierarchy Level*

#### ***[edit protocols ripng] Hierarchy Level***

The following statement hierarchy can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```
protocols {
 ripng {
 graceful-restart {
 disable;
 restart-time seconds;
 }
 group group-name {
 export [policy-names];
 import [policy-names];
 metric-out metric;
 neighbor neighbor-name {
 import [policy-names];
 metric-in metric;
 receive <none>;
 route-timeout seconds;
 }
 }
 }
}
```

```

 send <none>;
 update-interval seconds;
 }
 preference number;
 route-timeout seconds;
 update-interval seconds;
}
holddown seconds;
import [policy-names];
metric-in metric;
receive <none>;
route-timeout seconds;
send <none>;
update-interval seconds;
traceoptions {
 file filename <files number> <size maximum-file-size> <world-readable |
 no-world-readable>;
 flag flag <flag-modifier> <disable>;
}
}
}


```

**Related  
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

## any-sender

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | any-sender;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ],<br>[edit protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Disable strict sender address checks.</p> <p>If the sender of a RIP message does not belong to the subnet of the interface, the message is discarded. This situation might cause problems with dropped packets when RIP is running on point-to-point interfaces, or when the addresses on the interfaces do not fall in the same subnet. You can resolve this by disabling strict address checks on the RIP traffic.</p>                                                                                                                 |
|                                 | <div><p><b>NOTE:</b> The <b>any-sender</b> statement is supported only for peer-to-peer interfaces.</p></div>                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring RIP on page 2921</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## authentication-key (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authentication-key password;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Require authentication for RIP route queries received on an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>password</b>—Authentication password. If the password does not match, the packet is rejected. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Route Authentication for RIP on page 2928</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## authentication-type (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authentication-type type;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the type of authentication for RIP route queries received on an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Default</b>                  | If you do not include this statement and the <b>authentication-key</b> statement, RIP authentication is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>type</b> —Authentication type: <ul style="list-style-type: none"><li>• <b>md5</b>—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing device uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme.</li><li>• <b>none</b>—Disable authentication. If <b>none</b> is configured, the configured authentication key is ignored.</li><li>• <b>simple</b>—Use a simple password. The password is included in the transmitted packet, which makes this method of authentication relatively insecure. The password can be from 1 through 16 contiguous letters or digits long.</li></ul>                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Route Authentication for RIP on page 2928</a></li><li>• <a href="#">authentication-key on page 3001</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |



## bfd-liveness-detection (Protocols RIP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> bfd-liveness-detection {     authentication {         algorithm <i>algorithm-name</i>;         key-chain <i>key-chain-name</i>;         loose-check;     }     detection-time {         threshold <i>milliseconds</i>;     }     minimum-interval <i>milliseconds</i>;     minimum-receive-interval <i>milliseconds</i>;     multiplier <i>number</i>;     no-adaptation;     transmit-interval {         minimum-interval <i>milliseconds</i>;         threshold <i>milliseconds</i>;     }     version (1   automatic); } </pre>                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols rip <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</p> <p>[edit protocols rip <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>]</p>                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.0.</p> <p>Options <b>detection-time threshold</b> and <b>transmit-interval threshold</b> introduced in Junos OS Release 8.2.</p> <p>Support for logical systems introduced in Junos OS Release 8.3.</p> <p>Option <b>no-adaptation</b> introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Options <b>authentication algorithm</b>, <b>authentication key-chain</b>, and <b>authentication loose-check</b> introduced in Junos OS Release 9.6.</p> <p>Options <b>authentication algorithm</b>, <b>authentication key-chain</b>, and <b>authentication loose-check</b> introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> |
| <b>Description</b>         | Configure bidirectional failure detection timers and authentication.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>             | <p><b>authentication algorithm <i>algorithm-name</i></b>—Configure the algorithm used to authenticate the specified BFD session: <b>simple-password</b>, <b>keyed-md5</b>, <b>keyed-sha-1</b>, <b>meticulous-keyed-md5</b>, or <b>meticulous-keyed-sha-1</b>.</p> <p><b>authentication key-chain <i>key-chain-name</i></b>—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the <b>authentication-key-chains key-chain</b> statement at the [edit security] hierarchy level.</p>                                                                                                                                                                                                                            |

**authentication loose-check**—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication is not configured at both ends of the BFD session.

**detection-time threshold *milliseconds***—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

**minimum-interval *milliseconds***—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.

**Range:** 1 through 255,000 milliseconds

**minimum-receive-interval *milliseconds***—Configure the minimum interval after which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

**Range:** 1 through 255,000 milliseconds

**multiplier *number***—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

**Range:** 1 through 255

**Default:** 3

**no-adaptation**—Configure BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

**transmit-interval threshold *milliseconds***—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ )

**transmit-interval minimum-interval *milliseconds***—Configure a minimum interval after which the local routing device transmits hello packets to a neighbor. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

**Range:** 1 through 255,000

**version**—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

**Default:** automatic

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

- Related Documentation**
- [Example: Configuring BFD for RIP on page 2934](#)
  - [Example: Configuring BFD Authentication for RIP on page 2941](#)

## check-zero

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (check-zero   no-check-zero);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>Some of the reserved fields in RIP version 1 packets must be zero, whereas in RIP version 2 packets, most of these reserved fields can contain nonzero values. By default, RIP discards version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications.</p> <p>If you find that you are receiving RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero, you can configure RIP to receive these packets even though they are being sent in violation of the specifications in RFC 1058 and RFC 2453.</p> <p>Check whether the reserved fields in a RIP packet are zero:</p> <ul style="list-style-type: none"> <li>• <b>check-zero</b>—Discard version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications.</li> <li>• <b>no-check-zero</b>—Receive RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero. This is in spite of the fact that they are being sent in violation of the specifications in RFC 1058 and RFC 2453.</li> </ul> |
| <b>Default</b>                  | check-zero                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIP on page 2921</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## demand-circuit (Protocols RIP)


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | demand-circuit;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ],<br>[edit protocols rip group <i>group-name</i> ],<br>[edit protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure a neighboring interface to act as a RIP demand circuit. To complete the demand circuit, you must configure both ends of the pair as demand circuits. When configured, the device sends RIP information only when changes occur in the routing database.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>                  | Disabled. You must explicitly configure two neighboring interfaces to act as a RIP demand circuit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIP Demand Circuits on page 2986</a></li> <li>• <a href="#">RIP Demand Circuits Overview on page 2984</a></li> <li>• <a href="#">max-retrans-time on page 3015</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## dynamic-peers

|                                 |                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | dynamic-peers;                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i> ],<br>[edit protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.4.                                                                                                                                                         |
| <b>Description</b>              | Configure an interface to have dynamic peers in a point-to-multipoint RIP network.                                                                                                                     |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Point-to-Multipoint RIP Networks on page 2948</a></li> </ul>                                                                 |

## export (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols rip <b>group</b> <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>rip <b>group</b> <i>group-name</i>],</code><br><code>[edit protocols rip <b>group</b> <i>group-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip <b>group</b> <i>group-name</i>]</code>                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Apply a policy to routes being exported to the neighbors.</p> <p>By default, RIP does not export routes it has learned to its neighbors. To enable RIP to export routes, apply one or more export policies.</p> <p>If no routes match the policies, the local routing device does not export any routes to its neighbors. Export policies override any metric values determined through calculations involving the values configured with the <a href="#">metric-in</a> and <a href="#">metric-out</a> statements.</p> <div><p><b>NOTE:</b> The export policy on RIP does not support manipulating routing information of the next hop.</p></div> |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring RIP on page 2921</a></li><li>• <a href="#">import on page 3013</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

---

## graceful-restart (Protocols RIP)

---

|                                 |                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>graceful-restart {<br/>  disable;<br/>  restart-time seconds;<br/>}</pre>                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> ],<br>[edit protocols <a href="#">rip</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>              | Configure graceful restart for RIP.                                                                                        |
| <b>Options</b>                  | <b>disable</b> —Disables graceful restart for RIP.<br><br>The remaining statement is explained separately.                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Junos OS High Availability Configuration Guide</i></li></ul>                    |

## group (Protocols RIP)

---

```
Syntax group group-name {
 bfd-liveness-detection {
 authentication {
 algorithm algorithm-name;
 key-chain key-chain-name;
 loose-check;
 }
 detection-time {
 threshold milliseconds;
 }
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 transmit-interval {
 threshold milliseconds;
 minimum-interval milliseconds;
 }
 multiplier number;
 version (0 | 1 | automatic);
 }
 demand-circuit;
 export policy;
 max-retrans-time seconds;
 metric-out metric;
 preference number;
 route-timeout seconds;
 update-interval seconds;
 neighbor neighbor-name {
 authentication-key password;
 authentication-type type;
 bfd-liveness-detection {
 authentication {
 algorithm algorithm-name;
 key-chain key-chain-name;
 loose-check;
 }
 detection-time {
 threshold milliseconds;
 }
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 transmit-interval {
 threshold milliseconds;
 minimum-interval milliseconds;
 }
 multiplier number;
 version (0 | 1 | automatic);
 }
 (check-zero | no-check-zero);
 demand-circuit;
 import policy-name;
 max-retrans-time seconds;
 message-size number;
```



```

metric-in metric;
metric-out metric;
receive receive-options;
route-timeout seconds;
send send-options;
update-interval seconds;
}
}

```

|                                 |                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                          |
| <b>Description</b>              | <p>Configure a set of RIP neighbors that share an export policy and metric. The export policy and metric govern what routes to advertise to neighbors in a given group. Each group must contain at least one neighbor. You should create a group for every export policy.</p>                                                                                          |
| <b>Options</b>                  | <p><b><i>group-name</i></b>—Name of a group, up to 16 characters long.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIP on page 2921</a></li> </ul>                                                                                                                                                                                                                                                              |

## holddown (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>holddown seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code><a href="#">rip</a>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                          |
| <b>Description</b>              | <p>Configure how long the expired route is retained in the routing table before being removed.</p> <p>When the hold-down timer runs on RIP demand circuits, routes are advertised as unreachable on other interfaces. When the hold-down timer expires, the route is removed from the routing table if all destinations detect that the route is unreachable or the remaining destinations are down.</p>                |
| <b>Options</b>                  | <b>seconds</b> —Estimated time to wait before making updates to the routing table.<br><b>Range:</b> 10 through 180 seconds<br><b>Default:</b> 180 seconds                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring RIP Timers on page 2978</a></li><li>• <a href="#">RIP Demand Circuits Overview on page 2984</a></li></ul>                                                                                                                                                                                                                                      |

## import (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Apply one or more policies to routes being imported by the local routing device from neighbors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Applying Policies to RIP Routes Imported from Neighbors on page 2954</a></li> <li>• <a href="#">Routing Policy Configuration Guide</a></li> <li>• <a href="#">export on page 3008</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## interface-type (Protocols RIP)


---

|                                 |                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | interface-type p2mp;                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i> ],<br>[edit protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.4.                                                                                                                                                         |
| <b>Description</b>              | Configure the type of interface in a RIP network.<br><br>This statement enables a RIP device to have single or multiple peers through an interface.                                                    |
| <b>Options</b>                  | <b>p2mp</b> —Configure an interface in a RIP network as a point-to-multipoint interface.                                                                                                               |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Point-to-Multipoint RIP Networks on page 2948</a></li></ul>                                                                   |

## max-retrans-time

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | max-retrans-time <i>seconds</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</p> <p>[edit protocols rip group <i>group-name</i>],</p> <p>[edit protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>RIP demand circuits send update messages every 5 seconds to an unresponsive peer. Configure the retransmission timer to limit the number of times the demand circuit resends update messages to an unresponsive peer. If the configured retransmission threshold is reached, routes from the next hop router are marked as unreachable and the hold-down timer starts. You must configure a pair of RIP demand circuits for this timer to take effect.</p> <p>To determine the number of times to resend the update message, use the following calculation:</p> $5 \text{ seconds} \times \text{number of retransmissions} = \text{retransmission seconds}$                                                                                                                                                      |
| <b>Options</b>                  | <p><b>seconds</b>—The total amount of time the demand circuit resends update messages to an unresponsive peer. The seconds range corresponds to sending an update message a minimum of 1 time (5 seconds) and a maximum of 36 times (180 seconds).</p> <p><b>Range:</b> 5 through 180 seconds</p> <p><b>Default:</b> 5 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIP Demand Circuits on page 2986</a></li> <li>• <a href="#">RIP Demand Circuits Overview on page 2984</a></li> <li>• <a href="#">demand-circuit on page 3007</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## message-size

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>message-size <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement for SRX Series devices introduced in Junos OS Release 9.5.</p> <p>Statement for J Series platform introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Specify the number of route entries to be included in every RIP update message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|                                 | <div>  <p><b>TIP:</b> To ensure interoperability with other vendors' equipment, use the standard of 25 route entries per message. Do not change the default number of route entries in a RIP update message.</p> </div>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><i>number</i>—Number of route entries per update message.</p> <p><b>Range:</b> 25 through 255 entries</p> <p><b>Default:</b> 25 entries</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIP on page 2921</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## metric-in (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>metric-in <i>metric</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Specify the metric to add to incoming routes when the routing device advertises into RIP routes that were learned from other protocols. Use this statement to configure the routing device to prefer RIP routes learned through a specific neighbor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><i>metric</i>—Metric value.</p> <p><b>Range:</b> 1 through 16</p> <p><b>Default:</b> 1</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring the Metric Value Added to Imported RIP Routes on page 2963</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## metric-out (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>metric-out <i>metric</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</code><br><code>[edit protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>]</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Specify the metric value to add to routes transmitted to the neighbor. Use this statement to control how other routing devices prefer RIP routes sent from this neighbor.</p> <p>If you have included the <b>export</b> statement, RIP exports routes it has learned to the neighbors configured by including the <b>neighbor</b> statement.</p> <p>The metric associated with a RIP route (unless modified by an export policy) is the normal RIP metric. For example, a RIP route with a metric of 5 learned from a neighbor configured with a <b>metric-in</b> value of 2 is advertised with a combined metric of 7 when advertised to RIP neighbors in the same group. However, if this route was learned from a RIP neighbor in a different group or from a different protocol, the route is advertised with the metric value configured for that group with the <b>metric-out</b> statement.</p> <p>The metric for a route can be modified with an export policy. That metric is seen when the route is exported to the next hop.</p> <p>To increase the metric for routes advertised outside a group, include the <b>metric-out</b> statement.</p> |
| <b>Options</b>                  | <b><i>metric</i></b> —Metric value.<br><b>Range:</b> 1 through 16<br><b>Default:</b> 1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Examples: Controlling Traffic with Metrics in a RIP Network on page 295950</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |



## neighbor (Protocols RIP)

**Syntax** `neighbor neighbor-name {`  
     `authentication-key password;`  
     `authentication-type type;`  
     `bfd-liveness-detection {`  
         `authentication {`  
             `algorithm algorithm-name;`  
             `key-chain key-chain-name;`  
             `loose-check;`  
         `}`  
         `detection-time {`  
             `threshold milliseconds;`  
         `}`  
     `minimum-interval milliseconds;`  
     `minimum-receive-interval milliseconds;`  
     `transmit-interval {`  
         `threshold milliseconds;`  
         `minimum-interval milliseconds;`  
     `}`  
     `multiplier number;`  
     `version (0 | 1 | automatic);`  
     `}`  
     `(check-zero | no-check-zero);`  
     `demand-circuit;`  
     `import policy-name;`  
     `max-retrans-time seconds;`  
     `message-size number;`  
     `metric-in metric;`  
     `metric-out metric;`  
     `receive receive-options;`  
     `route-timeout seconds;`  
     `send send-options;`  
     `update-interval seconds;`  
     `}`

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols rip **group** *group-name*],  
     [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
         rip **group** *group-name*],  
     [edit protocols rip **group** *group-name*],  
     [edit routing-instances *routing-instance-name* protocols rip **group** *group-name*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
     Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Configure neighbor-specific RIP parameters, thereby overriding the defaults set for the routing device.

**Options** *neighbor-name*—Name of an interface over which a routing device communicates to its neighbors.

The remaining statements are explained separately.

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

**Related Documentation**

- [Example: Configuring RIP on page 2921](#)

---

## peer (Protocols RIP)

---

**Syntax** `peer IP address;`

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],  
[edit protocols rip group *group-name* neighbor *neighbor-name*]

**Release Information** Statement introduced in Junos OS Release 11.4.

**Description** Configure a static peer for an interface in a point-to-multipoint RIP network.

**Options** *address*—IP address of the static peer to be configured.

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

**Related Documentation**

- [Example: Configuring Point-to-Multipoint RIP Networks on page 2948](#)

## preference (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>preference <i>preference</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols rip <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip <b>group</b> <i>group-name</i>],</p> <p>[edit protocols rip <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip <b>group</b> <i>group-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Specify the preference of external routes learned by RIP as compared to those learned from other routing protocols.</p> <p>By default, Junos OS assigns a preference of 100 to routes that originate from RIP. When Junos OS determines a route's preference to become the active route, the software selects the route with the lowest preference and installs this route into the forwarding table.</p>                       |
| <b>Options</b>                  | <p><b>preference</b>—Preference value. A lower value indicates a more preferred route.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> 100</p>                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Route Preferences Overview</i></li> </ul>                                                                                                                                                                                                                                                                                                                                              |

## receive (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>receive receive-options;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure RIP receive options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <i>receive-options</i> —One of the following: <ul style="list-style-type: none"><li>• <b>both</b>—Accept both RIP version 1 and version 2 packets.</li><li>• <b>none</b>—Do not receive RIP packets.</li><li>• <b>version-1</b>—Accept only RIP version 1 packets.</li><li>• <b>version-2</b>—Accept only RIP version 2 packets.</li></ul> <b>Default:</b> <b>both</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 2968</a></li><li>• <a href="#">send on page 3026</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## rib-group (Protocols RIP)

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rib-group group-name;</code>                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ],<br>[edit protocols <a href="#">rip</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                          |
| <b>Description</b>              | Install RIP routes into multiple routing tables by configuring a routing table group.                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <i>group-name</i> —Name of the routing table group.                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Redistributing Routes Between Two RIP Instances on page 2972</a></li> </ul>                                                                                                                                                                                                               |

## rip

|                                 |                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>rip {...}</code>                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],<br>[edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                      |
| <b>Description</b>              | Enable RIP routing on the routing device.                                                                                                                                                                                                                           |
| <b>Default</b>                  | RIP is disabled on the routing device.                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIP on page 2921</a></li> </ul>                                                                                                                                                           |

## route-timeout (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>route-timeout seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols rip <a href="#">group</a> <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols rip <a href="#">group</a> <i>group-name</i> neighbor</code><br><code>  <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  rip <a href="#">group</a> <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  rip <a href="#">group</a> <i>group-name</i> neighbor <i>neighbor-name</i>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit protocols rip <a href="#">group</a> <i>group-name</i>],</code><br><code>[edit protocols rip <a href="#">group</a> <i>group-name</i> neighbor <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip <a href="#">group</a> <i>group-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip <a href="#">group</a> <i>group-name</i> neighbor</code><br><code>  <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Configure the route timeout interval for RIP. If a route is not refreshed after being installed in the routing table by the specified timeout interval, the route is marked as invalid and is removed from the routing table after the hold-down period expires.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <b>seconds</b> —Estimated time to wait before making updates to the routing table.<br><b>Range:</b> 30 through 360 seconds<br><b>Default:</b> 180 seconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring RIP Timers on page 2978</a></li><li>• <a href="#">RIP Demand Circuits Overview on page 2984</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## routing-instances (Multiple Routing Entities)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>routing-instances <i>routing-instance-name</i> { ... }</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit],<br>[edit logical-systems <i>logical-system-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Configure an additional routing entity for a router. You can create multiple instances of BGP, IS-IS, OSPF, OSPFv3, and RIP for a router. You can also create multiple routing instances for separating routing tables, routing policies, and interfaces for individual wholesale subscribers (retailers) in a Layer 3 wholesale network.</p> <p>Each routing instance consist of the following:</p> <ul style="list-style-type: none"> <li>• A set of routing tables</li> <li>• A set of interfaces that belong to these routing tables</li> <li>• A set of routing option configurations</li> </ul> <p>Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name <b>my-instance</b>, its corresponding IP unicast table is my-instance.inet.0. All routes for <b>my-instance</b> are installed into my-instance.inet.0.</p> <p>Routes are installed into the default routing instance inet.0 by default, unless a routing instance is specified.</p> <p>In Junos OS Release 9.0 and later, you can no longer specify a routing-instance name of <i>master</i>, <i>default</i>, or <i>bgp</i> or include special characters within the name of a routing instance.</p> <p>In Junos OS Release 9.6 and later, you can include a slash (/) in a routing-instance name only if a logical system is not configured. That is, you cannot include the slash character in a routing-instance name if a logical system other than the default is explicitly configured. Routing-instance names, further, are restricted from having the form <code>__.*__</code> (beginning and ending with underscores). The colon : character cannot be used when multipotology routing (MTR) is enabled.</p> |
| <b>Default</b>                  | Routing instances are disabled for the router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b><i>routing-instance-name</i></b>—Name of the routing instance. This must be a non-reserved string of not more than 128 characters.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

- Related Documentation**
- [Example: Configuring Interprovider Layer 3 VPN Option A](#)
  - [Example: Configuring Interprovider Layer 3 VPN Option B](#)
  - [Example: Configuring Interprovider Layer 3 VPN Option C](#)
  - [Example: Configuring E-LINE and E-LAN Services for a PBB Network](#)

---

## send (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>send send-options;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure RIP send options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <b><i>send-options</i></b> —One of the following: <ul style="list-style-type: none"><li>• <b>broadcast</b>—Broadcast RIP version 2 packets (RIP version 1 compatible).</li><li>• <b>multicast</b>—Multicast RIP version 2 packets. This is the default.</li><li>• <b>none</b>—Do not send RIP updates.</li><li>• <b>version-1</b>—Broadcast RIP version 1 packets.</li></ul> <b>Default:</b> multicast                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 2968</a></li><li>• <a href="#">receive on page 3022</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |



## traceoptions (Protocols RIP)

|                            |                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;<i>flag-modifier</i>&gt; &lt;disable&gt;; }</pre>                                                                                                                                |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ],<br>[edit protocols <a href="#">rip</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ] |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                          |
| <b>Description</b>         | Set RIP protocol-level tracing options.                                                                                                                                                                                                                                                                                                                 |



**NOTE:** The `traceoptions` statement is not supported on QFabric systems.

**Default** The default RIP protocol-level trace options are inherited from the global `traceoptions` statement.

**Options** **disable**—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

**file *filename***—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place RIP tracing output in the file `/var/log/rip-log`.

**files *number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file.0***, then ***trace-file.1***, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

**Range:** 2 through 1000 files

**Default:** 10 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

### RIP Tracing Options

- **auth**—RIP authentication
- **error**—RIP error packets

- **expiration**—RIP route expiration processing
- **holddown**—RIP hold-down processing
- **nsr-synchronization**—Nonstop routing synchronization events
- **packets**—All RIP packets
- **request**—RIP information packets such as request, poll, and poll entry packets
- **trigger**—RIP triggered updates
- **update**—RIP update packets

#### Global Tracing Options

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **receive**—Trace the packets being received.
- **receive-detail**—Provide detailed trace information for packets being received.
- **send**—Trace the packets being transmitted.
- **send-detail**—Provide detailed trace information for packets being transmitted.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size** *size*—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.                                                                 |
|                                 | routing-control—To add this statement to the configuration.                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Tracing RIP Protocol Traffic on page 2990</a></li></ul> |

## update-interval (Protocols RIP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>update-interval seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> group <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> group <i>group-name</i> neighbor</code><br><code>    <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <a href="#">rip</a> group <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <a href="#">rip</a> group <i>group-name</i> neighbor <i>neighbor-name</i>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit protocols <a href="#">rip</a> group <i>group-name</i>],</code><br><code>[edit protocols <a href="#">rip</a> group <i>group-name</i> neighbor <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> group <i>group-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> group <i>group-name</i> neighbor</code><br><code>    <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the interval at which routes learned by RIP are sent to neighbors. This timer controls the interval between routing updates. This timer is set to 30 seconds, by default, with a small random amount of time added when the timer is reset. This added time prevents congestion that can happen if all routing devices update their neighbors simultaneously.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <b>seconds</b> —Estimated time to wait before making updates to the routing table.<br><b>Range:</b> 10 through 60 seconds<br><b>Default:</b> 30 seconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring RIP Timers on page 2978</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

### Administration

- [RIP Monitoring on page 3031](#)
- [Operational Commands on page 3039](#)

---

## RIP Monitoring

---

- [Routing Protocol Process Memory FAQs Overview on page 3031](#)
- [Routing Protocol Process Memory FAQs on page 3031](#)
- [Monitoring RIP Routing Information on page 3038](#)

### ***Routing Protocol Process Memory FAQs Overview***

Junos OS is based on the FreeBSD Unix operating system. The open source software is modified and hardened to operate in the device's specialized environment. For example, some executables have been deleted, while other utilities were de-emphasized. Additionally, certain software processes were added to enhance the routing functionality. The result of this transformation is the kernel, the heart of the Junos OS software.

The kernel is responsible for operating multiple processes that perform the actual functions of the device. Each process operates in its own protected memory space, while the communication among all the processes is still controlled by the kernel. This separation provides isolation between the processes, and resiliency in the event of a process failure. This is important in a core routing platform because a single process failure does not cause the entire device to cease functioning.

Some of the common software processes include the routing protocol process (rpd) that controls the device's protocols, the device control process (dcd) that controls the device's interfaces, the management process (mgd) that controls user access to the device, the chassis process (chassisd) that controls the device's properties itself, and the Packet Forwarding Engine process (pfed) that controls the communication between the device's Packet Forwarding Engine and the Routing Engine. The kernel also generates specialized processes as needed for additional functionality, such as SNMP, the Virtual Router Redundancy Protocol (VRRP), and Class of Service (CoS).

The routing protocol process is a software process within the Routing Engine software, which controls the routing protocols that run on the device. Its functionality includes all protocol messages, routing table updates, and implementation of routing policies.

The routing protocol process starts all configured routing protocols and handles all routing messages. It maintains one or more routing tables, which consolidate the routing information learned from all routing protocols. From this routing information, the routing protocol process determines the active routes to network destinations and installs these routes into the Routing Engine's forwarding table. Finally, it implements routing policy, which allows you to control the routing information that is transferred between the routing protocols and the routing table. Using routing policy, you can filter and limit the transfer of information as well as set properties associated with specific routes.

### **Related Documentation**

- [Routing Protocol Process Memory FAQs on page 3031](#)

### ***Routing Protocol Process Memory FAQs***

The following sections present the most frequently asked questions and answers related to the routing protocol process memory utilization, operation, interpretation of related command outputs, and troubleshooting the software process.

### ***Frequently Asked Questions: Routing Protocol Process Memory***

This section presents frequently asked questions and answers related to the memory usage of the routing protocol process.

#### **Why does the routing protocol process use excessive memory?**

The routing protocol process uses hundreds of megabytes of RAM in the Routing Engine to store information needed for the operation of routing and related protocols, such as BGP, OSPF, IS-IS, RSVP, LDP and MPLS. Such huge consumption of memory is common for the process, as the information it stores includes routes, next hops, interfaces, routing policies, labels, and label-switched paths (LSPs). Because access to the RAM memory is much faster than access to the hard disk, most of the routing protocol process information is stored in the RAM memory instead of using the hard disk space. This ensures that the performance of the routing protocol process is maximized.

#### **How can I check the amount of memory the routing protocol process is using?**

You can check routing protocol process memory usage by entering the **show system processes** and the **show task memory** Junos OS command-line interface (CLI) operational mode commands.

The **show system processes** command displays information about software processes that are running on the device and that have controlling terminals. The **show task memory** command displays memory utilization for routing protocol tasks on the Routing Engine.

You can check the routing protocol process memory usage by using the **show system processes** command with the **extensive** option. The **show task memory** command displays a report generated by the routing protocol process on its own memory usage. However, this report does not display all the memory used by the process. The value reported by the routing protocol process does not account for the memory used for the **TEXT** and **STACK** segments, or the memory used by the process's internal memory manager. Further, the Resident Set Size value includes shared library pages used by the routing protocol process.

For more information about checking the routing protocol process memory usage, see *Check Routing Protocol Process (rpd) Memory Usage*.

For more information, see the [show system processes](#) command and the **show task memory** command.

#### **I just deleted a large number of routes from the routing protocol process. Why is it still using so much memory?**

The **show system processes extensive** command displays a **RES** value measured in kilobytes. This value represents the amount of program memory resident in the physical memory. This is also known as RSS or Resident Set Size. The **RES** value includes shared library pages used by the process. Any amount of memory freed by the process might still be considered part of the **RES** value. Generally, the kernel delays the migrating of memory out of the **Inact** queue into the **Cache** or **Free** list unless there is a memory shortage. This can lead to large discrepancies between the values reported by the routing

protocol process and the kernel, even after the routing protocol process has freed a large amount of memory.

### ***Frequently Asked Questions: Interpreting Routing Protocol Process-Related Command Outputs***

This section presents frequently asked questions and answers about the routing protocol process-related Junos OS command-line interface (CLI) command outputs that are used to display the memory usage of the routing protocol process.

#### **How do I interpret memory numbers displayed in the show system processes extensive command output?**

The **show system processes extensive** command displays exhaustive system process information about software processes that are running on the device and have controlling terminals. This command is equivalent to the UNIX **top** command. However, the UNIX **top** command shows real-time memory usage, with the memory values constantly changing, while the **show system processes extensive** command provides a snapshot of memory usage in a given moment.

To check overall CPU and memory usage, enter the **show system processes extensive** command. Refer to [Table 242 on page 3034](#) for information about the **show system processes extensive** commands output fields.

```
user@host> show system processes extensive
last pid: 544; load averages: 0.00, 0.00, 0.00 18:30:33
37 processes: 1 running, 36 sleeping

Mem: 25M Active, 3968K Inact, 19M Wired, 184K Cache, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free
 PID USERNAME PRI NICE SIZE RES STATE TIME WCPU CPU COMMAND
 544 root 30 0 604K 768K RUN 0:00 0.00% 0.00% top
 3 root 28 0 0K 12K psleep 0:00 0.00% 0.00% vmdaemon
 4 root 28 0 0K 12K update 0:03 0.00% 0.00% update
 528 aviva 18 0 660K 948K pause 0:00 0.00% 0.00% tcsh
 204 root 18 0 300K 544K pause 0:00 0.00% 0.00% csh
 131 root 18 0 332K 532K pause 0:00 0.00% 0.00% cron
 186 root 18 0 196K 68K pause 0:00 0.00% 0.00% watchdog
 27 root 10 0 512M 16288K mfsidl 0:00 0.00% 0.00% mount_mfs
 1 root 10 0 620K 344K wait 0:00 0.00% 0.00% init
 304 root 3 0 884K 900K ttyin 0:00 0.00% 0.00% bash
 200 root 3 0 180K 540K ttyin 0:00 0.00% 0.00% getty
 203 root 3 0 180K 540K ttyin 0:00 0.00% 0.00% getty
 202 root 3 0 180K 540K ttyin 0:00 0.00% 0.00% getty
 201 root 3 0 180K 540K ttyin 0:00 0.00% 0.00% getty
 194 root 2 0 2248K 1640K select 0:11 0.00% 0.00% rpd
 205 root 2 0 964K 800K select 0:12 0.00% 0.00% tnp.chassisd
 189 root 2 -12 352K 740K select 0:03 0.00% 0.00% xntpd
 114 root 2 0 296K 612K select 0:00 0.00% 0.00% amd
 188 root 2 0 780K 600K select 0:00 0.00% 0.00% dcd
 527 root 2 0 176K 580K select 0:00 0.00% 0.00% rlogind
 195 root 2 0 212K 552K select 0:00 0.00% 0.00% inetd
 187 root 2 0 192K 532K select 0:00 0.00% 0.00% tnetd
 83 root 2 0 188K 520K select 0:00 0.00% 0.00% syslogd
 538 root 2 0 1324K 516K select 0:00 0.00% 0.00% mgd
 99 daemon 2 0 176K 492K select 0:00 0.00% 0.00% portmap
 163 root 2 0 572K 420K select 0:00 0.00% 0.00% nsrexecd
 192 root 2 0 560K 400K select 0:10 0.00% 0.00% snmpd
```

```

191 root 2 0 1284K 376K select 0:00 0.00% 0.00% mgd
537 aviva 2 0 636K 364K select 0:00 0.00% 0.00% cli
193 root 2 0 312K 204K select 0:07 0.00% 0.00% mib2d
 5 root 2 0 0K 12K pfesel 0:00 0.00% 0.00% if_pfe
 2 root -18 0 0K 12K psleep 0:00 0.00% 0.00% pagedaemon
 0 root -18 0 0K 0K sched 0:00 0.00% 0.00% swapper

```

Table 242 on page 3034 describes the output fields that represent the memory values for the **show system processes extensive** command. Output fields are listed in the approximate order in which they appear.

Table 242: show system processes extensive Output Fields

| Field Name    | Field Description                                                                                                                                                                                                                 |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Mem</b>    | Information about physical and virtual memory allocation.                                                                                                                                                                         |
| <b>Active</b> | Memory allocated and actively used by the program.                                                                                                                                                                                |
| <b>Inact</b>  | Memory allocated but not recently used or memory freed by the programs. Inactive memory remains mapped in the address space of one or more processes and, therefore, counts toward the RSS value of those processes.              |
| <b>Wired</b>  | Memory that is not eligible to be swapped, usually used for in-kernel memory structures and/or memory physically locked by a process.                                                                                             |
| <b>Cache</b>  | Memory that is not associated with any program and does not need to be swapped before being reused.                                                                                                                               |
| <b>Buf</b>    | Size of memory buffer used to hold data recently called from the disk.                                                                                                                                                            |
| <b>Free</b>   | Memory that is not associated with any programs. Memory freed by a process can become <b>Inactive</b> , <b>Cache</b> , or <b>Free</b> , depending on the method used by the process to free the memory.                           |
| <b>Swap</b>   | Information about swap memory. <ul style="list-style-type: none"> <li>• Total—Total memory available to be swapped to disk.</li> <li>• Used—Memory swapped to disk.</li> <li>• Free—Memory available for further swap.</li> </ul> |

The rest of the command output displays information about the memory usage of each process. The **SIZE** field indicates the size of the virtual address space, and the **RES** field indicates the amount of the program in physical memory, which is also known as RSS or Resident Set Size. For more information, see the [show system processes](#) command.

#### What is the difference between Active and Inact memory that is displayed by the show system processes extensive command?

When the system is under memory pressure, the pageout process reuses memory from the free, cache, inact and, if necessary, active pages. When the pageout process runs, it scans memory to see which pages are good candidates to be unmapped and freed up. Thus, the distinction between **Active** and **Inact** memory is only used by the pageout process to determine which pool of pages to free first at the time of a memory shortage.



The pageout process first scans the **Inact** list, and checks whether the pages on this list have been accessed since the time they have been listed here. The pages that have been accessed are moved from the **Inact** list to the **Active** list. On the other hand, pages that have not been accessed become prime candidates to be freed by the pageout process. If the pageout process cannot produce enough free pages from the **Inact** list, pages from the **Active** list get freed up.

Because the pageout process runs only when the system is under memory pressure, the pages on the **Inact** list remain untouched – even if they have not been accessed recently – when the amount of **Free** memory is adequate.

### How do I interpret memory numbers displayed in the show task memory command output?

The **show task memory** command provides a comprehensive picture of the memory utilization for routing protocol tasks on the Routing Engine. The routing protocol process is the main task that uses Routing Engine memory.

To check routing process memory usage, enter the **show task memory** command. Refer to [Table 243 on page 3035](#) for information about the **show task memory** command output fields.

```
user@host> show task memory
Memory Size (kB) %Available When
Currently In Use: 29417 3% now
Maximum Ever Used: 33882 4% 00/02/11 22:07:03
Available: 756281 100% now
```

[Table 243 on page 3035](#) describes the output fields for the **show task memory** command. Output fields are listed in the approximate order in which they appear.

**Table 243: show task memory Output Fields**

| Field Name                      | Field Description                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------|
| <b>Memory Currently In Use</b>  | Memory currently in use. Dynamically allocated memory plus the <b>DATA</b> segment memory in kilobytes. |
| <b>Memory Maximum Ever Used</b> | Maximum memory ever used.                                                                               |
| <b>Memory Available</b>         | Memory currently available.                                                                             |

The **show task memory** command does not display all the memory used by the routing protocol process. This value does not account for the memory used for the **TEXT** and **STACK** segments, or the memory used by the routing protocol process's internal memory manager.

### Why is the Currently In Use value less than the RES value?

The **show task memory** command displays a **Currently In Use** value measured in kilobytes. This value represents the memory currently in use. It is the dynamically allocated memory plus the **DATA** segment memory. The **show system processes extensive** command displays

a **RES** value measured in kilobytes. This value represents the amount of program memory resident in the physical memory. This is also known as RSS or Resident Set Size.

The **Currently In Use** value does not account for all of the memory that the routing protocol process uses. This value does not include the memory used for the **TEXT** and the **STACK** segments, and a small percentage of memory used by the routing protocol process's internal memory manager. Further, the **RES** value includes shared library pages used by the routing protocol process.

Any amount of memory freed by the routing protocol process might still be considered part of the **RES** value. Generally, the kernel delays the migrating of memory out of the **Inact** queue into the **Cache** or **Free** list unless there is a memory shortage. This can lead to large discrepancies between the **Currently In Use** value and the **RES** value.

### ***Frequently Asked Questions: Routing Protocol Process Memory Swapping***

This section presents frequently asked questions and answers related to the memory swapping of the routing protocol process from the Routing Engine memory to the hard disk memory.

#### **How do I monitor swap activity?**

When the system is under memory pressure, the pageout process reuses memory from the free, cache, inactive and, if necessary, active pages. You can monitor the swap activity by viewing the syslog message reported by the kernel during periods of high pageout activity.

The syslog message appears as follows:

```
Mar 3 20:08:02 olympic /kernel: High pageout rate!! 277 pages/sec.
```

You can use the **vmstat -s** command to print the statistics for the swapout activity. The displayed statistics appear as follows:

```
0 swap pager pageouts
0 swap pager pages paged out
```

The **swap pager pageouts** is the number of pageout operations to the swap device, and the **swap pager pages paged out** is the number of pages paged out to the swap device.

#### **Why does the system start swapping when I try to dump core using the request system core-dumps command?**

The **request system core-dumps** command displays a list of system core files created when the device has failed. This command can be useful for diagnostic purposes. Each list item includes the file permissions, number of links, owner, group, size, modification date, path, and filename. You can use the **core-filename** option and the **core-file-info**, **brief**, and **detail** options to display more information about the specified core-dump files.

You can use the **request system core-dumps** command to perform a non-fatal core-dump without aborting the routing protocol process. To do this, the routing protocol process is forked, generating a second copy, and then aborted. This process can double the memory consumed by the two copies of the routing protocol processes, pushing the system into swap.

**Why does the show system processes extensive command show that memory is swapped to disk although there is plenty of free memory?**

Memory can remain swapped out indefinitely if it is not accessed again. Therefore, the **show system processes extensive** command shows that memory is swapped to disk even though there is plenty of free memory, and such a situation is not unusual.

***Frequently Asked Questions: Troubleshooting the Routing Protocol Process***

This section presents frequently asked questions and answers related to a shortage of memory and memory leakage by the routing protocol process.

**What does the RPD\_OS\_MEMHIGH message mean?**

The **RPD\_OS\_MEMHIGH** message is written into the system message file if the routing protocol process is running out of memory. This message alerts you that the routing protocol process is using the indicated amount and percentage of Routing Engine memory, which is considered excessive. This message is generated either because the routing protocol process is leaking memory or the use of system resources is excessive, perhaps because routing filters are misconfigured or the configured network topology is very complex.

When the memory utilization for the routing protocol process is using all available Routing Engine DRAM memory (Routing Engines with maximum 2 GB DRAM) or reaches the limit of 2 GB of memory (Routing Engines with 4 GB DRAM), a message of the following form is written every minute in the syslog message file:

**RPD\_OS\_MEMHIGH:** Using 188830 KB of memory, 100 percent of available

This message includes the amount, in kilobytes and/or the percentage, of the available memory in use.

This message should not appear under normal conditions, as any further memory allocations usually require a portion of existing memory to be written to swap. As a recommended solution, increase the amount of RAM in the Routing Engine. For more information, go to <http://kb.juniper.net/InfoCenter/index?page=content&id=KB14186>.

**What can I do when there is a memory shortage even after a swap?**

It is not recommended for the system to operate in this state, notwithstanding the existence of swap. The protocols that run in the routing protocol process usually have a real-time requirement that cannot reliably withstand the latency of being swapped to hard disk. If the memory shortage has not resulted from a memory leak, then either a reduction in the memory usage or an upgrade to a higher memory-capacity Routing Engine is required.

**How do I determine whether there is a memory leak in the routing protocol process?**

Memory leaks are typically the result of a seemingly unbounded growth in the memory usage of a process as reported by the **show system processes extensive** command.

There are two classes of memory leaks that the routing protocol process can experience.

- The first class occurs when the allocated memory that is no longer in use is not freed. This class of leak can usually be fixed by taking several samples of the **show task memory detail** command over a period of time and comparing the deltas.
- The second class occurs when there is a late access to freed memory. If the access is not outside the mapped address space, the kernel backfills the accessed page with real memory. This backfill is done without the knowledge of the routing protocol process's internal memory allocator, which makes this class of leak much more difficult to resolve. If a memory leak of this class is suspected, writing the state of the system to a disk file (creating a core file) is suggested.

A large discrepancy between the **RES** value and the **Currently In Use** value might indicate a memory leak. However, large discrepancies can also occur for legitimate reasons. For example, the memory used for the **TEXT** and **STACK** segments or the memory used by the routing protocol process's internal memory manager might not be displayed. Further, the **RES** value includes shared library pages used by the process.

#### What is the task\_timer?

The source of a routing protocol process memory leak can usually be identified by dumping the timers for each task. You can use the **show task task-name** command to display routing protocol tasks on the Routing Engine. Tasks can be baseline tasks performed regardless of the device's configuration, and other tasks that depend on the device configuration.

For more information, see the **show task** command.

#### Related Documentation

- [Routing Protocol Process Memory FAQs Overview on page 3031](#)

#### Monitoring RIP Routing Information

**Purpose** Use the monitoring functionality to monitor RIP routing on routing devices.

**Action** To view RIP routing information in the J-Web interface, select **Monitor > Routing > RIP Information**.

To view RIP routing information in the CLI, enter the following CLI commands:

- **show rip statistics**
- **show rip neighbor**

**Meaning** [Table 244 on page 3038](#) summarizes key output fields in the RIP routing display in the J-Web interface.

Table 244: Summary of Key RIP Routing Output Fields

| Field          | Values | Additional Information |
|----------------|--------|------------------------|
| RIP Statistics |        |                        |

Table 244: Summary of Key RIP Routing Output Fields (*continued*)

| Field                    | Values                                                                                 | Additional Information                                                                                                |
|--------------------------|----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| Protocol Name            | The RIP protocol name.                                                                 |                                                                                                                       |
| Port number              | The port on which RIP is enabled.                                                      |                                                                                                                       |
| Hold down time           | The interval during which routes are neither advertised nor updated.                   |                                                                                                                       |
| Global routes learned    | Number of RIP routes learned on the logical interface.                                 |                                                                                                                       |
| Global routes held down  | Number of RIP routes that are not advertised or updated during the hold-down interval. |                                                                                                                       |
| Global request dropped   | Number of requests dropped.                                                            |                                                                                                                       |
| Global responses dropped | Number of responses dropped.                                                           |                                                                                                                       |
| <b>RIP Neighbors</b>     |                                                                                        |                                                                                                                       |
| Neighbor                 | Name of the RIP neighbor.                                                              | This value is the name of the interface on which RIP is enabled. Click the name to see the details for this neighbor. |
| State                    | State of the RIP connection: <b>Up</b> or <b>Dn</b> (Down).                            |                                                                                                                       |
| Source Address           | Local source address.                                                                  | This value is the configured address of the interface on which RIP is enabled.                                        |
| Destination Address      | Destination address.                                                                   | This value is the configured address of the immediate RIP adjacency.                                                  |
| Send Mode                | The mode of sending RIP messages.                                                      |                                                                                                                       |
| Receive Mode             | The mode in which messages are received.                                               |                                                                                                                       |
| In Metric                | Value of the incoming metric configured for the RIP neighbor.                          |                                                                                                                       |

- Related Documentation**
- *Configuring a RIP Network (J-Web Procedure)*
  - *Layer 3 Protocols Supported on EX Series Switches*

### Operational Commands

## clear rip general-statistics

---

|                                                   |                                                                                                                                                                                                         |
|---------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                     | clear rip general-statistics<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                    |
| <b>Syntax (EX Series Switches and QFX Series)</b> | clear rip general-statistics                                                                                                                                                                            |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                |
| <b>Description</b>                                | Clear RIP general statistics.                                                                                                                                                                           |
| <b>Options</b>                                    | <b>none</b> —Clear RIP general statistics.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                   | clear                                                                                                                                                                                                   |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">show rip general-statistics on page 3056</a></li></ul>                                                                                              |
| <b>List of Sample Output</b>                      | <a href="#">clear rip general-statistics on page 3040</a>                                                                                                                                               |
| <b>Output Fields</b>                              | When you enter this command, you are provided feedback on the status of your request.                                                                                                                   |

## Sample Output

### clear rip general-statistics

```
user@host> clear rip general-statistics
```

## clear rip statistics

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                     | clear rip statistics<br><instance (all   <i>instance-name</i> )><br><logical-system (all   <i>logical-system-name</i> )><br><neighbor><br><peer (all   <i>address</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (EX Series Switches and QFX Series)</b> | clear rip statistics<br><instance (all   <i>instance-name</i> )><br><neighbor>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>                                | Clear RIP statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                    | <p><b>none</b>—Reset RIP counters for all neighbors for all routing instances.</p> <p><b>instance (all   <i>instance-name</i>)</b>—(Optional) Clear RIP statistics for all instances or for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>neighbor</b>—(Optional) Clear RIP statistics for the specified neighbor only.</p> <p><b>peer (all   <i>address</i>)</b>—(Optional) Clear RIP statistics for a single peer or all peers.</p> |
| <b>Required Privilege Level</b>                   | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li>• <a href="#">show rip statistics on page 3060</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>                      | <a href="#">clear rip statistics on page 3041</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>                              | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## Sample Output

### clear rip statistics

```
user@host> clear rip statistics
```

## restart

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | <pre>restart &lt;adaptive-services   ancpd-service   application-identification   audit-process   auto-configuration   captive-portal-content-delivery   ce-l2tp-service   chassis-control   class-of-service   clksyncd-service   database-replication   datapath-trace-service   dhcp-service   diameter-service   disk-monitoring   dynamic-flow-capture   ecc-error-logging   ethernet-connectivity-fault-management   ethernet-link-fault-management   event-processing   firewall   general-authentication-service   gracefully   iccp-service   idp-policy   immediately   interface-control   ipsec-key-management   kernel-replication   l2-learning   l2cpd-service   l2tp-service   l2tp-universal-edge   lacp   license-service   link-management   local-policy-decision-function   mac-validation   mib-process   mobile-ip   mountd-service   mpls-traceroute   mspd   multicast-snooping   named-service   nfsd-service   packet-triggered-subscribers   peer-selection-service   pgcp-service   pgm   pic-services-logging   pki-service   ppp   ppp-service   pppoe   protected-system-domain-service   redundancy-interface-process   remote-operations   root-system-domain-service   routing &lt;logical-system <i>logical-system-name</i>&gt;   sampling   sbc-configuration-process   sdk-service   service-deployment   services   services pgcp gateway <i>gateway-name</i>   snmp   soft   static-subscribers   statistics-service   subscriber-management   subscriber-management-helper   tunnel-oamd   usb-control   vrrp   web-management&gt; &lt;gracefully   immediately   soft&gt;</pre> |
| <b>Syntax (ACX Series Routers)</b> | <pre>restart &lt;adaptive-services   audit-process   auto-configuration   autoinstallation   chassis-control   class-of-service   clksyncd-service   database-replication   dhcp-service   diameter-service   disk-monitoring   dynamic-flow-capture   ethernet-connectivity-fault-management   ethernet-link-fault-management   event-processing   firewall   general-authentication-service   gracefully   immediately   interface-control   ipsec-key-management   l2-learning   lacp   link-management   mib-process   mobile-ip   mountd-service   mpls-traceroute   mspd   named-service   nfsd-service   pgm   pki-service   ppp   pppoe   redundancy-interface-process   remote-operations   routing   sampling   sdk-service   secure-neighbor-discovery   service-deployment   services   snmp   soft   statistics-service   subscriber-management   subscriber-management-helper   tunnel-oamd   vrrp&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Syntax (EX Series Switches)</b> | <pre>restart &lt;autoinstallation   chassis-control   class-of-service   database-replication   dhcp   dhcp-service   diameter-service   dot1x-protocol   ethernet-link-fault-management   ethernet-switching   event-processing   firewall   general-authentication-service   interface-control   kernel-replication   l2-learning   lacp   license-service   link-management   lldpd-service   mib-process   mountd-service   multicast-snooping   pgm   redundancy-interface-process   remote-operations   routing   secure-neighbor-discovery   service-deployment   sflow-service   snmp   vrrp   web-management&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (Routing Matrix)</b>     | <pre>restart &lt;adaptive-services   audit-process   chassis-control   class-of-service   disk-monitoring   dynamic-flow-capture   ecc-error-logging   event-processing   firewall   interface-control   ipsec-key-management   kernel-replication   l2-learning   l2tp-service   lacp   link-management   mib-process   pgm   pic-services-logging   ppp   pppoe   redundancy-interface-process   remote-operations   routing &lt;logical-system <i>logical-system-name</i>&gt;   sampling   service-deployment   snmp&gt; &lt;all   all-lcc   lcc <i>number</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |



|                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                           | <gracefully   immediately   soft>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax (J Series Routing Platform)</b> | <p>restart</p> <p>&lt;adaptive-services   audit-process   chassis-control   class-of-service   dhcp   dialer-services   dlsw   event-processing   firewall   interface-control   ipsec-key-management   isdn-signaling   l2-learning   l2tp-service   mib-process   network-access-service   pgm   ppp   pppoe   remote-operations   routing &lt;logical-system <i>logical-system-name</i>&gt;   sampling   service-deployment   snmp   usb-control   web-management&gt;</p> <p>&lt;gracefully   immediately   soft&gt;</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (TX Matrix Routers)</b>         | <p>restart</p> <p>&lt;adaptive-services   audit-process   chassis-control   class-of-service   dhcp-service   diameter-service   disk-monitoring   dynamic-flow-capture   ecc-error-logging   event-processing   firewall   interface-control   ipsec-key-management   kernel-replication   l2-learning   l2tp-service   lacp   link-management   mib-process   pgm   pic-services-logging   ppp   pppoe   redundancy-interface-process   remote-operations   routing &lt;logical-system <i>logical-system-name</i>&gt;   sampling   service-deployment   snmp   statistics-service&gt;</p> <p>&lt;all-chassis   all-lcc   lcc <i>number</i>   scc&gt;</p> <p>&lt;gracefully   immediately   soft&gt;</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax (TX Matrix Plus Routers)</b>    | <p>restart</p> <p>&lt;adaptive-services   audit-process   chassis-control   class-of-service   dhcp-service   diameter-service   disk-monitoring   dynamic-flow-capture   ecc-error-logging   event-processing   firewall   interface-control   ipsec-key-management   kernel-replication   l2-learning   l2tp-service   lacp   link-management   mib-process   pgm   pic-services-logging   ppp   pppoe   redundancy-interface-process   remote-operations   routing &lt;logical-system <i>logical-system-name</i>&gt;   sampling   service-deployment   snmp   statistics-service&gt;</p> <p>&lt;all-chassis   all-lcc   all-sfc   lcc <i>number</i>   sfc <i>number</i>&gt;</p> <p>&lt;gracefully   immediately   soft&gt;</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax (MX Series Routers)</b>         | <p>restart</p> <p>&lt;adaptive-services   ancpd-service   application-identification   audit-process   auto-configuration   captive-portal-content-delivery   ce-l2tp-service   chassis-control   class-of-service   clksyncd-service   database-replication   datapath-trace-service   dhcp-service   diameter-service   disk-monitoring   dynamic-flow-capture   ecc-error-logging   ethernet-connectivity-fault-management   ethernet-link-fault-management   event-processing   firewall   general-authentication-service   gracefully   iccp-service   idp-policy   immediately   interface-control   ipsec-key-management   kernel-replication   l2-learning   l2cpd-service   l2tp-service   l2tp-universal-edge   lacp   license-service   link-management   local-policy-decision-function   mac-validation   mib-process   mobile-ip   mounstd-service   mpls-traceroute   mspd   multicast-snooping   named-service   nfsd-service   packet-triggered-subscribers   peer-selection-service   pgcp-service   pgm   pic-services-logging   pki-service   ppp   ppp-service   pppoe   protected-system-domain-service   redundancy-interface-process   remote-operations   root-system-domain-service   routing   routing &lt;logical-system <i>logical-system-name</i>&gt;   sampling   sbc-configuration-process   sdk-service   service-deployment   services   services pgcp gateway <i>gateway-name</i>   snmp   soft   static-subscribers   statistics-service   subscriber-management   subscriber-management-helper   tunnel-oamd   usb-control   vrrp   web-management&gt;</p> <p>&lt;all-members&gt;</p> <p>&lt;gracefully   immediately   soft&gt;</p> <p>&lt;local&gt;</p> <p>&lt;member <i>member-id</i>&gt;</p> |

|                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax (J Series Routers)</b> | <pre>restart &lt;adaptive-services   audit-process   chassis-control   class-of-service   dhcp   dhcp-service   dialer-services   diameter-service   dlsf   event-processing   firewall   interface-control   ipsec-key-management   isdn-signaling   l2ald   l2-learning   l2tp-service   mib-process   network-access-service   pgm   ppp   pppoe   remote-operations   routing &lt;logical-system logical-system-name&gt;   sampling   service-deployment   snmp   usb-control   web-management&gt; &lt;gracefully   immediately   soft&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Syntax (QFX Series)</b>       | <pre>restart &lt;adaptive-services   audit-process   chassis-control   class-of-service   dialer-services   diameter-service   dlsf   ethernet-connectivity   event-processing   fibre-channel   firewall   general-authentication-service   igmp-host-services   interface-control   ipsec-key-management   isdn-signaling   l2ald   l2-learning   l2tp-service   mib-process   named-service   network-access-service   nstrace-process   pgm   ppp   pppoe   redundancy-interface-process   remote-operations  logical-system-name&gt;   routing   sampling  secure-neighbor-discovery   service-deployment   snmp   usb-control   web-management&gt; &lt;gracefully   immediately   soft&gt;</pre>                                                                                                                                                                                                                                                        |
| <b>Release Information</b>       | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series routers.</p> <p>Options added:</p> <ul style="list-style-type: none"><li>• <b>dynamic-flow-capture</b> in Junos OS Release 7.4.</li><li>• <b>dlsf</b> in Junos OS Release 7.5.</li><li>• <b>event-processing</b> in Junos OS Release 7.5.</li><li>• <b>ppp</b> in Junos OS Release 7.5.</li><li>• <b>l2ald</b> in Junos OS Release 8.0.</li><li>• <b>link-management</b> in Release 8.0.</li><li>• <b>pgcp-service</b> in Junos OS Release 8.4.</li><li>• <b>sbc-configuration-process</b> in Junos OS Release 9.5.</li><li>• <b>services pgcp gateway</b> in Junos OS Release 9.6.</li><li>• <b>sfc</b> and <b>all-sfc</b> for the TX Matrix Router in Junos OS Release 9.6.</li></ul> |
| <b>Description</b>               | <p>Restart a Junos OS process.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|                                  | <div><p><b>CAUTION:</b> Never restart a software process unless instructed to do so by a customer support engineer. A restart might cause the router or switch to drop calls and interrupt transmission, resulting in possible loss of data.</p></div>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                   | <p><b>none</b>—Same as <b>gracefully</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

- adaptive-services**—(Optional) Restart the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.
- all-chassis**—(TX Matrix and TX Matrix Plus routers only) (Optional) Restart the software process on all chassis.
- all-lcc**—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process on all T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process on all T1600 routers connected to the TX Matrix Plus router.
- all-members**—(MX Series routers only) (Optional) Restart the software process for all members of the Virtual Chassis configuration.
- all-sfc**—(TX Matrix Plus routers only) (Optional) For a TX Matrix Plus router, restart the software processes for the TX Matrix Plus router (or switch-fabric chassis).
- ancpd-service**—(Optional) Restart the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.
- application-identification**—(Optional) Restart the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.
- audit-process**—(Optional) Restart the RADIUS accounting process that gathers statistical data that can be used for general network monitoring, analyzing, and tracking usage patterns, for billing a user based on the amount of time or type of services accessed.
- auto-configuration**—(Optional) Restart the Interface Auto-Configuration process.
- autoinstallation**—(EX Series switches only) (Optional) Restart the autoinstallation process.
- captive-portal-content-delivery**—(Optional) Restart the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.
- ce-l2tp-service**—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.
- chassis-control**—(Optional) Restart the chassis management process.
- class-of-service**—(Optional) Restart the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.
- clksyncd-service**—(Optional) Restart the external clock synchronization process, which uses synchronous Ethernet (SyncE).

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

**lcc number**—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process for a specific T640 router that is connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process for a specific T1600 router that is connected to the TX Matrix Plus router. Replace **number** with a value from 0 through 3.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

**secure-neighbor-discovery**—(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.

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

**service-deployment**—(Optional) Restart the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.

**services**—(Optional) Restart a service.

**services pgcp gateway gateway-name**—(Optional) Restart the pgcpd process for a specific border gateway function (BGF) running on an MS-PIC. This option does not restart the pgcpd process running on the Routing Engine. To restart the pgcpd process on the Routing Engine, use the **pgcp-service** option.

**sflow-service**—(EX Series switches only) (Optional) Restart the flow sampling (sFlow technology) process.

**snmp**—(Optional) Restart the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.

**soft**—(Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.

**static-subscribers**—(Optional) Restart the static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.

**statistics-service**—(Optional) Restart the process that manages the Packet Forwarding Engine statistics.

**subscriber-management**—(Optional) Restart the Subscriber Management process.

**subscriber-management-helper**—(Optional) Restart the Subscriber Management Helper process.

**tunnel-oamd**—(Optional) Restart the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol



tunneling (L2PT) allows service providers to send Layer 2 PDUs across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.

**usb-control**—(J Series routers and MX Series routers only) (Optional) Restart the USB control process.

**vrrp**—(ACX Series routers, EX Series switches, and MX Series routers only) (Optional) Restart the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

**web-management**—(J Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the Web management process.

**Required Privilege Level** reset

**Related Documentation**

- [Overview of Junos OS CLI Operational Mode Commands on page 3523](#)

**List of Sample Output** [restart interfaces on page 3051](#)

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

## Sample Output

### restart interfaces

```
user@host> restart interfaces
interfaces process terminated
interfaces process restarted
```

## show policy

|                                    |                                                                                                                                                                                                                                                                                                                          |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | show policy<br><logical-system (all   <i>logical-system-name</i> )><br>< <i>policy-name</i> >                                                                                                                                                                                                                            |
| <b>Syntax (EX Series Switches)</b> | show policy<br>< <i>policy-name</i> >                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                    |
| <b>Description</b>                 | Display information about configured routing policies.                                                                                                                                                                                                                                                                   |
| <b>Options</b>                     | <p><b>none</b>—List the names of all configured routing policies.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>policy-name</i></b>—(Optional) Show the contents of the specified policy.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"> <li><a href="#">show policy damping on page 2674</a></li> </ul>                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>       | <a href="#">show policy on page 3052</a><br><a href="#">show policy policy-name on page 3053</a><br><a href="#">show policy (Multicast Scoping) on page 3053</a>                                                                                                                                                         |
| <b>Output Fields</b>               | Table 245 on page 3052 lists the output fields for the <b>show policy</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                               |

**Table 245: show policy Output Fields**

| Field Name         | Field Description               |
|--------------------|---------------------------------|
| <i>policy-name</i> | Name of the policy listed.      |
| <i>term</i>        | Policy term listed.             |
| <i>from</i>        | Match condition for the policy. |
| <i>then</i>        | Action for the policy.          |

## Sample Output

### show policy

```
user@host> show policy
```

```
Configured policies:
__vrf-export-red-internal__
__vrf-import-red-internal__
red-export
all_routes
```

### **show policy policy-name**

```
user@host> show policy test-statics
Policy test-statics:
 from
 3.0.0.0/8 accept
 3.1.0.0/16 accept
 then reject
```

### **show policy (Multicast Scoping)**

```
user@host> show policy test-statics
Policy test-statics:
 from
 multicast-scoping == 8
```

## show policy conditions

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | <pre>show policy conditions &lt;condition-name&gt; &lt;detail&gt; &lt;dynamic&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax (EX Series Switches)</b> | <pre>show policy conditions &lt;condition-name&gt; &lt;detail&gt; &lt;dynamic&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>         | <p>Command introduced in Junos OS Release 9.0.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>                 | <p>Display all the configured conditions as well as the routing tables with which the configuration manager is interacting. If the <b>detail</b> keyword is included, the output also displays dependent routes for each condition.</p>                                                                                                                                                                                                                                                                            |
| <b>Options</b>                     | <p><b>none</b>—Display all configured conditions and associated routing tables.</p> <p><b>condition-name</b>—(Optional) Display information about the specified condition only.</p> <p><b>detail</b>—(Optional) Display the specified level of output.</p> <p><b>dynamic</b>—(Optional) Display information about the conditions in the dynamic database.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>       | <a href="#">show policy conditions detail on page 3055</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>               | <p><a href="#">Table 246 on page 3054</a> lists the output fields for the <b>show policy conditions</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                       |

**Table 246: show policy conditions Output Fields**

| Field Name              | Field Description                                                                                                                                  | Level of Output |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Condition</b>        | Name of configured condition.                                                                                                                      | All levels      |
| <b>event</b>            | Condition type. If the <b>if-route-exists</b> option is configured, the event type is:<br><b>Existence of a route in a specific routing table.</b> | All levels      |
| <b>Dependent routes</b> | List of routes dependent on the condition, along with the latest generation number.                                                                | <b>detail</b>   |
| <b>Condition tables</b> | List of routing tables associated with the condition, along with the latest generation number and number of dependencies.                          | All levels      |

Table 246: show policy conditions Output Fields (*continued*)

| Field Name                 | Field Description                                                         | Level of Output |
|----------------------------|---------------------------------------------------------------------------|-----------------|
| If-route-exists conditions | List of conditions configured to look for a route in the specified table. | All levels      |

## Sample Output

### show policy conditions detail

```
user@host> show policy conditions detail
Configured conditions:
Condition cond1, event: Existence of a route in a specific routing table
Dependent routes:
 4.4.4.4/32, generation 3
 6.6.6.6/32, generation 3
 10.10.10.10/32, generation 3

Condition cond2, event: Existence of a route in a specific routing table
Dependent routes:
None

Condition tables:
Table inet.0, generation 4, dependencies 3, If-route-exists conditions: cond1
(static) cond2 (static)
```

## show rip general-statistics

|                                                   |                                                                                                                                                                                                              |
|---------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                     | show rip general-statistics<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                          |
| <b>Syntax (EX Series Switches and QFX Series)</b> | show rip general-statistics                                                                                                                                                                                  |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                     |
| <b>Description</b>                                | Display brief RIP statistics.                                                                                                                                                                                |
| <b>Options</b>                                    | <p><b>none</b>—Display brief RIP statistics.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                                         |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li><a href="#">clear rip general-statistics on page 3040</a></li> </ul>                                                                                                  |
| <b>List of Sample Output</b>                      | <a href="#">show rip general-statistics on page 3056</a>                                                                                                                                                     |
| <b>Output Fields</b>                              | <a href="#">Table 247 on page 3056</a> lists the output fields for the <b>show rip general-statistics</b> command. Output fields are listed in the approximate order in which they appear.                   |

**Table 247: show rip general-statistics Output Fields**

| Field Name         | Field Description                                      |
|--------------------|--------------------------------------------------------|
| <b>bad msgs</b>    | Number of invalid messages received.                   |
| <b>no rcv intf</b> | Number of packets received with no matching interface. |
| <b>curr memory</b> | Amount of memory currently used by RIP.                |
| <b>max memory</b>  | Most memory used by RIP.                               |

## Sample Output

### show rip general-statistics

```

user@host> show rip general-statistics
RIPv2 I/O info:
 bad msgs : 0
 no rcv intf : 0

```

```
curr memory : 0
max memory : 0
```

## show rip neighbor

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                     | show rip neighbor<br><instance (all   <i>instance-name</i> )><br><logical-system (all   <i>logical-system-name</i> )><br>< <i>name</i> >                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax (EX Series Switches and QFX Series)</b> | show rip neighbor<br><instance (all   <i>instance-name</i> )><br>< <i>name</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>                                | Display information about RIP neighbors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                    | <p><b>none</b>—Display information about all RIP neighbors for all instances.</p> <p><b>instance (all   <i>instance-name</i>)</b>—(Optional) Display RIP neighbor information for all instances or for only the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>name</i></b>—(Optional) Display detailed information about only the specified RIP neighbor.</p> |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>                      | <a href="#">show rip neighbor on page 3059</a><br><a href="#">show rip neighbor (With Demand Circuits Configured) on page 3059</a>                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>                              | <a href="#">Table 248 on page 3058</a> lists the output fields for the <b>show rip neighbor</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                     |

**Table 248: show rip neighbor Output Fields**

| Field Name      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Neighbor</b> | <p>Name of the RIP neighbor.</p> <p><b>NOTE:</b> Beginning with Junos OS Release 11.1, when you configure demand circuits, the output displays a demand circuit (DC) flag next to neighbor interfaces configured for demand circuits.</p> <p>If you configure demand circuits at the <b>[edit protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>]</b> hierarchy level, the output shows only the neighboring interface that you specifically configured as a demand circuit. If you configure demand circuits at the <b>[edit protocols rip group <i>group-name</i>]</b> hierarchy level, all of the interfaces in the group are configured as demand circuits. Therefore, the output shows all of the interfaces in that group as demand circuits.</p> |



Table 248: show rip neighbor Output Fields (*continued*)

| Field Name                 | Field Description                                                                                        |
|----------------------------|----------------------------------------------------------------------------------------------------------|
| <b>State</b>               | State of the connection: <b>Up</b> or <b>Dn</b> (Down).                                                  |
| <b>Source Address</b>      | Address of the port on the local router.                                                                 |
| <b>Destination Address</b> | Address of the port on the remote router.                                                                |
| <b>Send Mode</b>           | Send options: <b>broadcast</b> , <b>multicast</b> , <b>none</b> , or <b>version 1</b> .                  |
| <b>Receive Mode</b>        | Type of packets to accept: <b>both</b> , <b>none</b> , <b>version 1</b> , or <b>version 2</b> .          |
| <b>In Met</b>              | Metric added to incoming routes when advertising into RIP routes that were learned from other protocols. |

## Sample Output

### show rip neighbor

```

user@host> show rip neighbor
Neighbor Local Source Destination Send Receive In
 State Address Address Mode Mode Met

ge-2/3/0.0 Up 192.168.9.105 192.168.9.107 bcast both 1
at-5/1/1.42 Dn (null) (null) mcast v2 only 3
at-5/1/0.42 Dn (null) (null) mcast both 3
at-5/1/0.0 Up 20.0.0.1 224.0.0.9 mcast both 3
so-0/0/0.0 Up 192.168.9.97 224.0.0.9 mcast both 3

```

### show rip neighbor (With Demand Circuits Configured)

```

user@host> show rip neighbor
Neighbor Local Source Destination Send Receive In
 State Address Address Mode Mode Met

so-0/1/0.0(DC) Up 10.10.10.2 224.0.0.9 mcast both 1
so-0/2/0.0(DC) Up 13.13.13.2 224.0.0.9 mcast both 1

```

## show rip statistics

---

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                     | <code>show rip statistics</code><br><code>&lt;instance (all   <i>instance-name</i>)&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;name&gt;</code><br><code>&lt;peer (all   <i>address</i>)&gt;</code>                                                                                                                                                                                                                                                                                                                                        |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <code>show rip statistics</code><br><code>&lt;instance (all   <i>instance-name</i>)&gt;</code><br><code>&lt;name&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                                | Display RIP statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                    | <b>none</b> —Display RIP statistics for all routing instances.<br><br><b>instance (all   <i>instance-name</i>)</b> —(Optional) Display RIP statistics for all instances or for only the specified routing instance.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b>name</b> —(Optional) Display detailed information about only the specified RIP neighbor.<br><br><b>peer (all   <i>address</i>)</b> —(Optional) Display RIP statistics for a single peer or all peers. |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">clear rip statistics on page 3041</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>                      | <a href="#">show rip statistics on page 3061</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>                              | <a href="#">Table 249 on page 3061</a> lists the output fields for the <b>show rip statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                               |

Table 249: show rip statistics Output Fields

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>RIP info</b>          | <p>Information about RIP on the specified interface:</p> <ul style="list-style-type: none"> <li>• <b>port</b>—UDP port number used for RIP.</li> <li>• <b>update interval</b>—Interval between routing table updates, in seconds.</li> <li>• <b>holddown</b>—Hold-down interval, in seconds.</li> <li>• <b>timeout</b>—Timeout interval, in seconds.</li> <li>• <b>restart in progress</b>—Graceful restart status. Displayed when RIP is or has been in the process of graceful restart.</li> <li>• <b>restart time</b>—Estimated time for the graceful restart to finish, in seconds.</li> <li>• <b>restart will complete in</b>—Remaining time for the graceful restart to finish, in seconds.</li> <li>• <b>rts learned</b>—Number of routes learned through RIP.</li> <li>• <b>rts held down</b>—Number of routes held down by RIP.</li> <li>• <b>rqsts dropped</b>—Number of received request packets that were dropped.</li> <li>• <b>resps dropped</b>—Number of received response packets that were dropped.</li> </ul>                                                                                                                                                   |
| <b>logical-interface</b> | <p>Name of the logical interface and its statistics:</p> <ul style="list-style-type: none"> <li>• <b>routes learned</b>—Number of routes learned on the logical interface.</li> <li>• <b>routes advertised</b>—Number of routes advertised by the logical interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Counter</b>           | <p>List of counter types:</p> <ul style="list-style-type: none"> <li>• <b>Updates Sent</b>—Number of update messages sent.</li> <li>• <b>Triggered Updates Sent</b>—Number of triggered update messages sent.</li> <li>• <b>Responses Sent</b>—Number of response messages sent.</li> <li>• <b>Bad Messages</b>—Number of invalid messages received.</li> <li>• <b>RIPv1 Updates Received</b>—Number of RIPv1 update messages received.</li> <li>• <b>RIPv1 Bad Route Entries</b>—Number of RIPv1 invalid route entry messages received.</li> <li>• <b>RIPv1 Updates Ignored</b>—Number of RIPv1 update messages ignored.</li> <li>• <b>RIPv2 Updates Received</b>—Number of RIPv2 update messages received.</li> <li>• <b>RIPv2 Bad Route Entries</b>—Number of RIPv2 invalid route entry messages received.</li> <li>• <b>RIPv2 Updates Ignored</b>—Number of RIPv2 update messages ignored.</li> <li>• <b>Authentication Failures</b>—Number of received update messages that failed authentication.</li> <li>• <b>RIP Requests Received</b>—Number of RIP request messages received.</li> <li>• <b>RIP Requests Ignored</b>—Number of RIP request messages ignored.</li> </ul> |
| <b>Total</b>             | Total number of packets for the selected counter.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Last 5 min</b>        | Number of packets for the selected counter in the most recent 5-minute period.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Last minute</b>       | Number of packets for the selected counter in the most recent 1-minute period.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## Sample Output

### show rip statistics

```
user@host> show rip statistics so-0/0/0.0
```

RIP info: port 520; update interval: 30s; holddown 180s; timeout 120s  
restart in progress: restart time 60s; restart will complete in 55s  
      rts learned  rts held down  rqsts dropped  resps dropped  
              0              0              0              0

so-0/0/0.0: 0 routes learned; 501 routes advertised

| Counter                 | Total | Last 5 min | Last minute |
|-------------------------|-------|------------|-------------|
| -----                   | ----- | -----      | -----       |
| Updates Sent            | 0     | 0          | 0           |
| Triggered Updates Sent  | 0     | 0          | 0           |
| Responses Sent          | 0     | 0          | 0           |
| Bad Messages            | 0     | 0          | 0           |
| RIPv1 Updates Received  | 0     | 0          | 0           |
| RIPv1 Bad Route Entries | 0     | 0          | 0           |
| RIPv1 Updates Ignored   | 0     | 0          | 0           |
| RIPv2 Updates Received  | 0     | 0          | 0           |
| RIPv2 Bad Route Entries | 0     | 0          | 0           |
| RIPv2 Updates Ignored   | 0     | 0          | 0           |
| Authentication Failures | 0     | 0          | 0           |
| RIP Requests Received   | 0     | 0          | 0           |
| RIP Requests Ignored    | 0     | 0          | 0           |

## show route

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | <pre>show route &lt;all&gt; &lt;destination-prefix&gt; &lt;logical-system (all   logical-system-name)&gt; &lt;private&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax (EX Series Switches)</b> | <pre>show route &lt;all&gt; &lt;destination-prefix&gt; &lt;private&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option <b>private</b> introduced in Junos OS Release 9.5.</p> <p>Option <b>private</b> introduced in Junos OS Release 9.5 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                 | Display the active entries in the routing tables.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                     | <p><b>none</b>—Display brief information about all active entries in the routing tables.</p> <p><b>all</b>—(Optional) Display information about all routing tables, including private, or internal, routing tables.</p> <p><b>destination-prefix</b>—(Optional) Display active entries for the specified address or range of addresses.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>private</b>—(Optional) Display information only about all private, or internal, routing tables.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIP on page 2921</a></li> <li>• <a href="#">Example: Configuring RIPv6</a></li> <li>• <a href="#">Example: Configuring IS-IS</a></li> <li>• <a href="#">Examples: Configuring Internal BGP Peering</a></li> <li>• <a href="#">Examples: Configuring External BGP Peering</a></li> <li>• <a href="#">Examples: Configuring OSPF Routing Policy</a></li> </ul>                                                                                                                                                                     |
| <b>List of Sample Output</b>       | <p><a href="#">show route on page 3066</a></p> <p><a href="#">show route destination-prefix on page 3066</a></p> <p><a href="#">show route extensive on page 3066</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**Output Fields** Table 250 on page 3064 describes the output fields for the **show route** command. Output fields are listed in the approximate order in which they appear.

**Table 250: show route Output Fields**

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>routing-table-name</i>       | Name of the routing table (for example, inet.0).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <i>number destinations</i>      | Number of destinations for which there are routes in the routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <i>number routes</i>            | <p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> <li>• <b>active</b> (routes that are active).</li> <li>• <b>holddown</b> (routes that are in the pending state before being declared inactive). A holddown route was once the active route and is no longer the active route. The route is in the holddown state because a protocol still has interest in the route, meaning that the interest bit is set. A protocol might have its interest bit set on the previously active route because the protocol is still advertising the route. The route will be deleted after all protocols withdraw their advertisement of the route and remove their interest bit. A persistent holddown state often means that the interested protocol is not releasing its interest bit properly.</li> </ul> <p>However, if you have configured advertisement of multiple routes (with the <b>add-path</b> or <b>advertise-inactive</b> statement), the holddown bit is most likely set because BGP is advertising the route as an active route. In this case, you can ignore the holddown state because nothing is wrong.</p> <ul style="list-style-type: none"> <li>• <b>hidden</b> (routes that are not used because of a routing policy).</li> </ul> |
| <i>destination-prefix</i>       | <p>Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as:</p> <ul style="list-style-type: none"> <li>• <b>MPLS-label</b> (for example, 80001).</li> <li>• <b>interface-name</b> (for example, ge-1/0/2).</li> <li>• <b>neighbor-address:control-word-status:encapsulation type:vc-id:source</b> (Layer 2 circuit only. For example, 10.1.1.195:NoCtrlWord:1:1:Local/96): <ul style="list-style-type: none"> <li>• <b>neighbor-address</b>—Address of the neighbor.</li> <li>• <b>control-word-status</b>—Whether the use of the control word has been negotiated for this virtual circuit: <b>NoCtrlWord</b> or <b>CtrlWord</b>.</li> <li>• <b>encapsulation type</b>—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.</li> <li>• <b>vc-id</b>—Virtual circuit identifier.</li> <li>• <b>source</b>—Source of the advertisement: <b>Local</b> or <b>Remote</b>.</li> </ul> </li> </ul>                                                                                                                                                                      |
| <b>[ protocol, preference ]</b> | <p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> <li>• <b>+</b>—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>-</b>—A hyphen indicates the last active route.</li> <li>• <b>*</b>—An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul> <p>In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.</p>                                                                                                                                                                                                                               |

Table 250: show route Output Fields (*continued*)

| Field Name                                        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>weeks:days</i><br><i>hours:minutes:seconds</i> | How long the route been known (for example, <b>2w4d 13:11:14</b> , or 2 weeks, 4 days, 13 hours, 11 minutes, and 14 seconds).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>metric</b>                                     | Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>localpref</b>                                  | Local preference value included in the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>from</b>                                       | Interface from which the route was received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>AS path</b>                                    | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• <b>I</b>—IGP.</li> <li>• <b>E</b>—EGP.</li> <li>• <b>?</b>—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• <b>[ ]</b>—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured.</li> <li>• <b>{ }</b>—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• <b>( )</b>—Parentheses enclose a confederation.</li> <li>• <b>( [ ] )</b>—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> |
| <b>validation-state</b>                           | <p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> <li>• <b>Invalid</b>—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.</li> <li>• <b>Unknown</b>—Indicates that the prefix is not among the prefixes or prefix ranges in the database.</li> <li>• <b>Valid</b>—Indicates that the prefix and autonomous system pair are found in the database.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>to</b>                                         | <p>Next hop to the destination. An angle bracket (&gt;) indicates that the route is the selected route.</p> <p>If the destination is <b>Discard</b>, traffic is dropped.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 250: show route Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>via</b> | <p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word <b>Selected</b>. This field can also contain the following information:</p> <ul style="list-style-type: none"> <li>• <b>Weight</b>—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.</li> <li>• <b>Balance</b>—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.</li> <li>• <b>lsp-path-name</b>—Name of the LSP used to reach the next hop.</li> <li>• <b>label-action</b>—MPLS label and operation occurring at the next hop. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).</li> </ul> |

## Sample Output

### show route

```

user@host> show route
inet.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:65500:1:10.0.0.20/240
 * [MVPN/70] 19:53:41, metric2 1
 Indirect
1:65500:1:10.0.0.40/240
 * [BGP/170] 19:53:29, localpref 100, from 10.0.0.30
 AS path: I
 > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
 [BGP/170] 19:53:26, localpref 100, from 10.0.0.33
 AS path: I
 > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
1:65500:1:10.0.0.60/240
 * [BGP/170] 19:53:29, localpref 100, from 10.0.0.30
 AS path: I
 > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
 [BGP/170] 19:53:25, localpref 100, from 10.0.0.33
 AS path: I
 > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF

```

### show route destination-prefix

```

user@host> show route 172.16.0.0/12

inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.0/12 * [Static/5] 2w4d 12:54:27
 > to 192.168.167.254 via fxp0.0

```

### show route extensive

```

user@host> show route extensive

```



```

v1.mvpn.0: 5 destinations, 8 routes (5 active, 1 holddown, 0 hidden)
1:65500:1:10.0.0.40/240 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 PMSI: Flags 0x0: Label[0:0:0]: PIM-SM: Sender 10.0.0.40 Group 225.1.1.1

 Next hop type: Indirect
 Address: 0x92455b8
 Next-hop reference count: 2
 Source: 10.0.0.30
 Protocol next hop: 10.0.0.40
 Indirect next hop: 2 no-forward
 State: <Active Int Ext>
 Local AS: 65500 Peer AS: 65500
 Age: 3 Metric2: 1
 Task: BGP_65500.10.0.0.30+179
 Announcement bits (2): 0-PIM.v1 1-mvpn global task
 AS path: I (Originator) Cluster list: 10.0.0.30
 AS path: Originator ID: 10.0.0.40
 Communities: target:65520:100
 Import Accepted
 Localpref: 100
 Router ID: 10.0.0.30
 Primary Routing Table bgp.mvpn.0
 Indirect next hops: 1
 Protocol next hop: 10.0.0.40 Metric: 1
 Indirect next hop: 2 no-forward
 Indirect path forwarding next hops: 1
 Next hop type: Router
 Next hop: 10.0.24.4 via lt-0/3/0.24 weight 0x1
 10.0.0.40/32 Originating RIB: inet.3
 Metric: 1 Node path count: 1
 Forwarding nexthops: 1
 Nexthop: 10.0.24.4 via lt-0/3/0.24

```

## show route active-path

|                                    |                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | show route active-path<br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                     |
| <b>Syntax (EX Series Switches)</b> | show route active-path<br><brief   detail   extensive   terse>                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>         | Command introduced in Junos OS Release 8.0.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                          |
| <b>Description</b>                 | Display all active routes for destinations. An active route is a route that is selected as the best path. Inactive routes are not displayed.                                                                                                                                                                                                                                               |
| <b>Options</b>                     | <p><b>none</b>—Display all active routes.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b>.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>       | <a href="#">show route active-path on page 3068</a><br><a href="#">show route active-path brief on page 3069</a><br><a href="#">show route active-path detail on page 3069</a><br><a href="#">show route active-path extensive on page 3070</a><br><a href="#">show route active-path terse on page 3072</a>                                                                               |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                                                                             |

## Sample Output

### show route active-path

```

user@host> show route active-path

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.70.19/32 *[Direct/0] 21:33:52
 > via lo0.0
10.255.71.50/32 *[IS-IS/15] 00:18:13, metric 10
 > to 100.1.2.1 via so-2/1/3.0
100.1.2.0/24 *[Direct/0] 00:18:36
 > via so-2/1/3.0
100.1.2.2/32 *[Local/0] 00:18:41
 Local via so-2/1/3.0
192.168.64.0/21 *[Direct/0] 21:33:52

```

```

> via fxp0.0
192.168.70.19/32 *Local/0] 21:33:52
 Local via fxp0.0

```

### show route active-path brief

The output for the **show route active-path brief** command is identical to that for the **show route active-path** command. For sample output, see [show route active-path on page 3068](#).

### show route active-path detail

```

user@host> show route active-path detail

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)

10.255.70.19/32 (1 entry, 1 announced)
 *Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 3
 Next hop: via lo0.0, selected
 State: <Active Int>
 Local AS: 200
 Age: 21:37:10
 Task: IF
 Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

 AS path: I

10.255.71.50/32 (1 entry, 1 announced)
 *IS-IS Preference: 15
 Level: 1
 Next hop type: Router, Next hop index: 397
 Next-hop reference count: 4
 Next hop: 100.1.2.1 via so-2/1/3.0, selected
 State: <Active Int>
 Local AS: 200
 Age: 21:31 Metric: 10
 Task: IS-IS
 Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
 AS path: I

100.1.2.0/24 (1 entry, 1 announced)
 *Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 3
 Next hop: via so-2/1/3.0, selected
 State: <Active Int>
 Local AS: 200
 Age: 21:54
 Task: IF
 Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

 AS path: I

100.1.2.2/32 (1 entry, 1 announced)
 *Local Preference: 0
 Next hop type: Local
 Next-hop reference count: 11
 Interface: so-2/1/3.0
 State: <Active NoReadvrt Int>

```

```
Local AS: 200
Age: 21:59
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.64.0/21 (1 entry, 1 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via fxp0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:37:10
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.70.19/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: fxp0.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:37:10
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I
```

#### show route active-path extensive

```
user@host> show route active-path extensive

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
10.255.70.19/32 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via lo0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
AS path: I

10.255.71.50/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.255.71.50/32 -> {100.1.2.1}
IS-IS level 2, LSP fragment 0
*IS-IS Preference: 15
Level: 1
Next hop type: Router, Next hop index: 397
Next-hop reference count: 4
Next hop: 100.1.2.1 via so-2/1/3.0, selected
State: <Active Int>
```

```

Local AS: 200
Age: 24:08 Metric: 10
Task: IS-IS
Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
AS path: I

100.1.2.0/24 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via so-2/1/3.0, selected
State: <Active Int>
Local AS: 200
Age: 24:31
Task: IF
Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
AS path: I

100.1.2.2/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: so-2/1/3.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 24:36
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.64.0/21 (1 entry, 1 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via fxp0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.70.19/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: fxp0.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

```

**show route active-path terse**

```
user@host> show route active-path terse
```

```
inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

| A | Destination      | P | Prf | Metric 1 | Metric 2 | Next hop    | AS path |
|---|------------------|---|-----|----------|----------|-------------|---------|
| * | 10.255.70.19/32  | D | 0   |          |          | >1o0.0      |         |
| * | 10.255.71.50/32  | I | 15  | 10       |          | >100.1.2.1  |         |
| * | 100.1.2.0/24     | D | 0   |          |          | >so-2/1/3.0 |         |
| * | 100.1.2.2/32     | L | 0   |          |          | Local       |         |
| * | 192.168.64.0/21  | D | 0   |          |          | >fxp0.0     |         |
| * | 192.168.70.19/32 | L | 0   |          |          | Local       |         |

## show route advertising-protocol

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show route advertising-protocol <i>protocol</i> <i>neighbor-address</i></code><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Display the routing information as it has been prepared for advertisement to a particular neighbor of a particular dynamic routing protocol.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>neighbor-address</i></b>—Address of the neighboring router to which the route entry is being transmitted.</p> <p><b><i>protocol</i></b>—Protocol transmitting the route:</p> <ul style="list-style-type: none"> <li>• <b>bgp</b>—Border Gateway Protocol</li> <li>• <b>dvmrp</b>—Distance Vector Multicast Routing Protocol</li> <li>• <b>msdp</b>—Multicast Source Discovery Protocol</li> <li>• <b>pim</b>—Protocol Independent Multicast</li> <li>• <b>rip</b>—Routing Information Protocol</li> <li>• <b>ripng</b>—Routing Information Protocol next generation</li> </ul> |
| <b>Additional Information</b>   | Routes displayed are routes that the routing table has exported into the routing protocol and that have been filtered by the associated protocol's <b>export</b> routing policy statements.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring the MED Attribute Directly</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>    | <a href="#">show route advertising-protocol bgp (Layer 3 VPN) on page 3075</a><br><a href="#">show route advertising-protocol bgp detail on page 3076</a><br><a href="#">show route advertising-protocol bgp detail (Layer 2 VPN) on page 3076</a><br><a href="#">show route advertising-protocol bgp detail (Layer 3 VPN) on page 3076</a><br><a href="#">show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address) on page 3076</a>                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>            | Table 251 on page 3074 lists the output fields for the <b>show route advertising-protocol</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

Table 251: show route advertising-protocol Output Fields

| Field Name                                   | Field Description                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output         |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <i>routing-table-name</i>                    | Name of the routing table—for example, inet.0.                                                                                                                                                                                                                                                                                                                                                                 | All levels              |
| <i>number destinations</i>                   | Number of destinations for which there are routes in the routing table.                                                                                                                                                                                                                                                                                                                                        | All levels              |
| <i>number routes</i>                         | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> <li>• <b>active</b> (routes that are active)</li> <li>• <b>holddown</b> (routes that are in the pending state before being declared inactive)</li> <li>• <b>hidden</b> (routes that are not used because of a routing policy)</li> </ul>                                          | All levels              |
| <b>Prefix</b>                                | Destination prefix.                                                                                                                                                                                                                                                                                                                                                                                            | <b>brief none</b>       |
| <i>destination-prefix (entry, announced)</i> | Destination prefix. The <b>entry</b> value is the number of routes for this destination, and the <b>announced</b> value is the number of routes being announced for this destination.                                                                                                                                                                                                                          | <b>detail extensive</b> |
| <b>BGP group and type</b>                    | BGP group name and type ( <b>Internal</b> or <b>External</b> ).                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b> |
| <b>Route Distinguisher</b>                   | Unique 64-bit prefix augmenting each IP subnet.                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b> |
| <b>Advertised Label</b>                      | Incoming label advertised by the LDP. When an IP packet enters a label-switched path (LSP), the ingress router examines the packet and assigns it a label based on its destination, placing the label in the packet's header. The label transforms the packet from one that is forwarded based on its IP routing information to one that is forwarded based on information associated with the label.          | <b>detail extensive</b> |
| <b>Label-Base, range</b>                     | First label in a block of labels and label block size. A remote PE router uses this first label when sending traffic toward the advertising PE router.                                                                                                                                                                                                                                                         | <b>detail extensive</b> |
| <b>VPN Label</b>                             | Virtual private network (VPN) label. Packets are sent between CE and PE routers by advertising VPN labels. VPN labels transit over either an RSVP or an LDP LSP tunnel.                                                                                                                                                                                                                                        | <b>detail extensive</b> |
| <b>Nexthop</b>                               | Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.<br><br>If the next-hop advertisement to the peer is <b>Self</b> , and the RIB-out next hop is a specific IP address, the RIB-out IP address is included in the extensive output. See <a href="#">show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address)</a> on page 3076. | All levels              |
| <b>MED</b>                                   | Multiple exit discriminator value included in the route.                                                                                                                                                                                                                                                                                                                                                       | <b>brief</b>            |
| <b>Lclpref or Localpref</b>                  | Local preference value included in the route.                                                                                                                                                                                                                                                                                                                                                                  | All levels              |



Table 251: show route advertising-protocol Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Level of Output         |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>AS path</b>             | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• <b>I</b>—IGP.</li> <li>• <b>E</b>—EGP.</li> <li>• <b>?</b>—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• <b>[ ]</b>—Brackets enclose the local AS number associated with the AS path if configured on the router, or if AS path prepending is configured.</li> <li>• <b>{ }</b>—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• <b>( )</b>—Parentheses enclose a confederation.</li> <li>• <b>( [ ] )</b>—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> | All levels              |
| <b>Communities</b>         | Community path attribute for the route. See the output field table for the <a href="#">show route detail</a> command for all possible values for this field.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b> |
| <b>AIGP</b>                | Accumulated interior gateway protocol (AIGP) BGP attribute.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b> |
| <b>Attrset AS</b>          | Number, local preference, and path of the autonomous system (AS) that originated the route. These values are stored in the <b>Attrset</b> attribute at the originating router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b> |
| <b>Layer2-info: encaps</b> | Layer 2 encapsulation (for example, VPLS).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b> |
| <b>control flags</b>       | Control flags: <b>none</b> or <b>Site Down</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail extensive</b> |
| <b>mtu</b>                 | Maximum transmission unit (MTU) of the Layer 2 circuit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b> |

## Sample Output

### show route advertising-protocol bgp (Layer 3 VPN)

```

user@host> show route advertising-protocol bgp 10.255.14.171
VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
10.255.14.172/32 Self 1 100 I
VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
10.255.14.181/32 Self 2 100 I

```

### show route advertising-protocol bgp detail

```
user@host> show route advertising-protocol bgp 111.222.1.3 detail
bgp20.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
111.222.1.11/32 (1 entry, 1 announced)
 BGP group pe-pe type Internal
 Route Distinguisher: 111.255.14.11:69
 Advertised Label: 100000
 next hop: Self
 Localpref: 100
 AS path: 2 I
 Communities: target:69:20
 AIGP 210
111.8.0.0/16 (1 entry, 1 announced)
 BGP group pe-pe type Internal
 Route Distinguisher: 111.255.14.11:69
 Advertised Label: 100000
 Next hop: Self
 Localpref: 100
 AS path: 2 I
 Communities: target:69:20
 AIGP 210
```

### show route advertising-protocol bgp detail (Layer 2 VPN)

```
user@host> show route advertising-protocol bgp 192.168.24.1 detail
vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
192.168.16.1:1:1:1/96 (1 entry, 1 announced)
 BGP group int type Internal
 Route Distinguisher: 192.168.16.1:1
 Label-base : 32768, range : 3
 Nexthop: Self
 Localpref: 100
 AS path: I
 Communities: target:65412:100
 AIGP 210
 Layer2-info: encaps:VLAN, control flags:, mtu:
```

### show route advertising-protocol bgp detail (Layer 3 VPN)

```
user@host> show route advertising-protocol bgp 10.255.14.176 detail
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
 BGP group ibgp type Internal
 Route Distinguisher: 10.255.14.174:2
 VPN Label: 101264
 Nexthop: Self
 Localpref: 100
 AS path: I
 Communities: target:200:100
 AIGP 210
 AttrSet AS: 100
 Localpref: 100
 AS path: I
 ...
```

### show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address)

```
user@host> show route advertising-protocol bgp 200.0.0.2 170.0.1.0/24 extensive all
inet.0: 13 destinations, 19 routes (13 active, 0 holddown, 6 hidden)
 170.0.1.0/24 (2 entries, 1 announced)
```

```
BGP group eBGP-INTEROP type External
 Nexthop: Self (rib-out 10.100.3.2)
 AS path: [4713] 200 I
...
```

## show route all

|                                    |                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | <code>show route all</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                                                         |
| <b>Syntax (EX Series Switches)</b> | <code>show route all</code>                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                         |
| <b>Description</b>                 | Display information about all routes in all routing tables, including private, or internal, tables.                                                                                                                                                                                                                                           |
| <b>Options</b>                     | <b>none</b> —Display information about all routes in all routing tables, including private, or internal, tables.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.                                                                 |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>       | <a href="#">show route all on page 3078</a>                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>               | In Junos OS Release 9.5 and later, only the output fields for the <b>show route all</b> command display all routing tables, including private, or hidden, routing tables. The output field table of the <a href="#">show route</a> command does not display entries for private, or hidden, routing tables in Junos OS Release 9.5 and later. |

## Sample Output

### show route all

The following example displays a snippet of output from the **show route** command and then displays the same snippet of output from the **show route all** command:

```

user@host> show route
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0 *[MPLS/0] 2d 02:24:39, metric 1
 Receive
1 *[MPLS/0] 2d 02:24:39, metric 1
 Receive
2 *[MPLS/0] 2d 02:24:39, metric 1
 Receive
800017 *[VPLS/7] 1d 14:00:16
 > via vt-3/2/0.32769, Pop
800018 *[VPLS/7] 1d 14:00:26
 > via vt-3/2/0.32772, Pop

user@host> show route all
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

```

+ = Active Route, - = Last Active, \* = Both

|                |                                                    |
|----------------|----------------------------------------------------|
| 0              | *[MPLS/0] 2d 02:19:12, metric 1<br>Receive         |
| 1              | *[MPLS/0] 2d 02:19:12, metric 1<br>Receive         |
| 2              | *[MPLS/0] 2d 02:19:12, metric 1<br>Receive         |
| 800017         | *[VPLS/7] 1d 13:54:49<br>> via vt-3/2/0.32769, Pop |
| 800018         | *[VPLS/7] 1d 13:54:59<br>> via vt-3/2/0.32772, Pop |
| vt-3/2/0.32769 | [VPLS/7] 1d 13:54:49<br>Unusable                   |
| vt-3/2/0.32772 | [VPLS/7] 1d 13:54:59<br>Unusable                   |

## show route best

---

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | <code>show route best <i>destination-prefix</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switches)</b> | <code>show route best <i>destination-prefix</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code>                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                            |
| <b>Description</b>                 | Display the route in the routing table that is the best route to the specified address or range of addresses. The best route is the longest matching route.                                                                                                                                                                                                                                                      |
| <b>Options</b>                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b> .<br><br><b><i>destination-prefix</i></b> —Address or range of addresses.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>       | <a href="#">show route best on page 3080</a><br><a href="#">show route best detail on page 3081</a><br><a href="#">show route best extensive on page 3082</a><br><a href="#">show route best terse on page 3082</a>                                                                                                                                                                                              |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                                                                                                   |

## Sample Output

### show route best

```
user@host> show route best 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32 *[OSPF/10] 1d 13:19:20, metric 2
 > to 10.31.1.6 via ge-3/1/0.0
 via so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32 *[RSVP/7] 1d 13:20:13, metric 2
 > via so-0/3/0.0, label-switched-path green-r1-r3

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
```

```

+ = Active Route, - = Last Active, * = Both
10.0.0.0/8 *[Direct/0] 2d 01:43:34
 > via fxp2.0
 [Direct/0] 2d 01:43:34
 > via fxp1.0

```

### show route best detail

```

user@host> show route best 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
 *OSPF Preference: 10
 Next-hop reference count: 9
 Next hop: 10.31.1.6 via ge-3/1/0.0, selected
 Next hop: via so-0/3/0.0
 State: <Active Int>
 Local AS: 69
 Age: 1d 13:20:06 Metric: 2
 Area: 0.0.0.0
 Task: OSPF
 Announcement bits (2): 0-KRT 3-Resolve tree 2
 AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
 State: <FlashAll>
 *RSVP Preference: 7
 Next-hop reference count: 5
 Next hop: via so-0/3/0.0 weight 0x1, selected
 Label-switched-path green-r1-r3
 Label operation: Push 100016
 State: <Active Int>
 Local AS: 69
 Age: 1d 13:20:59 Metric: 2
 Task: RSVP
 Announcement bits (1): 1-Resolve tree 2
 AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
10.0.0.0/8 (2 entries, 0 announced)
 *Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 1
 Next hop: via fxp2.0, selected
 State: <Active Int>
 Age: 2d 1:44:20
 Task: IF
 AS path: I
 Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 1
 Next hop: via fxp1.0, selected
 State: <NotBest Int>
 Inactive reason: No difference
 Age: 2d 1:44:20
 Task: IF
 AS path: I

```

## show route best extensive

The output for the **show route best extensive** command is identical to that for the **show route best detail** command. For sample output, see [show route best detail on page 3081](#).

## show route best terse

```
user@host> show route best 10.255.70.103 terse
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination P Prf Metric 1 Metric 2 Next hop AS path
* 10.255.70.103/32 0 10 2 >10.31.1.6
 so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination P Prf Metric 1 Metric 2 Next hop AS path
* 10.255.70.103/32 R 7 2 >so-0/3/0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination P Prf Metric 1 Metric 2 Next hop AS path
* 10.0.0.0/8 D 0 >fxp2.0
 D 0 >fxp1.0
```



## show route brief

|                                    |                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | show route brief<br><destination-prefix><br><logical-system (all   logical-system-name)>                                                                                                                                                                                                                                                     |
| <b>Syntax (EX Series Switches)</b> | show route brief<br><destination-prefix>                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                        |
| <b>Description</b>                 | Display brief information about the active entries in the routing tables.                                                                                                                                                                                                                                                                    |
| <b>Options</b>                     | <p><b>none</b>—Display all active entries in the routing table.</p> <p><b>destination-prefix</b>—(Optional) Display active entries for the specified address or range of addresses.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>       | <a href="#">show route brief on page 3083</a>                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>               | For information about output fields, see the Output Field table of the <a href="#">show route</a> command.                                                                                                                                                                                                                                   |

## Sample Output

### show route brief

```

user@host> show route brief
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0 *[Static/5] 1w5d 20:30:29
 Discard
10.255.245.51/32 *[Direct/0] 2w4d 13:11:14
 > via lo0.0
172.16.0.0/12 *[Static/5] 2w4d 13:11:14
 > to 192.168.167.254 via fxp0.0
192.168.0.0/18 *[Static/5] 1w5d 20:30:29
 > to 192.168.167.254 via fxp0.0
192.168.40.0/22 *[Static/5] 2w4d 13:11:14
 > to 192.168.167.254 via fxp0.0
192.168.64.0/18 *[Static/5] 2w4d 13:11:14
 > to 192.168.167.254 via fxp0.0
192.168.164.0/22 *[Direct/0] 2w4d 13:11:14
 > via fxp0.0
192.168.164.51/32 *[Local/0] 2w4d 13:11:14
 Local via fxp0.0
207.17.136.192/32 *[Static/5] 2w4d 13:11:14

```

```

> to 192.168.167.254 via fxp0.0
green.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100.101.0.0/16 *[Direct/0] 1w5d 20:30:28
 > via fe-0/0/3.0
100.101.2.3/32 *[Local/0] 1w5d 20:30:28
 Local via fe-0/0/3.0
224.0.0.5/32 *[OSPF/10] 1w5d 20:30:29, metric 1
 MultiRecv
```

## show route detail

|                                    |                                                                                                                                                                                                                                                                                                                                                             |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | show route detail<br><destination-prefix><br><logical-system (all   logical-system-name)>                                                                                                                                                                                                                                                                   |
| <b>Syntax (EX Series Switches)</b> | show route detail<br><destination-prefix>                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                       |
| <b>Description</b>                 | Display detailed information about the active entries in the routing tables.                                                                                                                                                                                                                                                                                |
| <b>Options</b>                     | <p><b>none</b>—Display all active entries in the routing table on all systems.</p> <p><b>destination-prefix</b>—(Optional) Display active entries for the specified address or range of addresses.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>       | <a href="#">show route detail on page 3093</a><br><a href="#">show route detail (with BGP Multipath) on page 3099</a>                                                                                                                                                                                                                                       |
| <b>Output Fields</b>               | <p><a href="#">Table 252 on page 3085</a> describes the output fields for the <b>show route detail</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                 |

**Table 252: show route detail Output Fields**

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                               |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>routing-table-name</i>  | Name of the routing table (for example, inet.0).                                                                                                                                                                                                                                                                                                                |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table.                                                                                                                                                                                                                                                                                         |
| <i>number routes</i>       | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> <li><b>active</b> (routes that are active)</li> <li><b>holddown</b> (routes that are in the pending state before being declared inactive)</li> <li><b>hidden</b> (routes that are not used because of a routing policy)</li> </ul> |

Table 252: show route detail Output Fields (*continued*)

| Field Name                                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>route-destination</i><br>(entry, announced) | <p>Route destination (for example:10.0.0.1/24). The <b>entry</b> value is the number of routes for this destination, and the <b>announced</b> value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> <li>• <b>MPLS-label</b> (for example, 80001).</li> <li>• <b>interface-name</b> (for example, ge-1/0/2).</li> <li>• <b>neighbor-address:control-word-status:encapsulation type:vc-id:source</b> (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> <li>• <b>neighbor-address</b>—Address of the neighbor.</li> <li>• <b>control-word-status</b>—Whether the use of the control word has been negotiated for this virtual circuit: <b>NoCtrlWord</b> or <b>CtrlWord</b>.</li> <li>• <b>encapsulation type</b>—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.</li> <li>• <b>vc-id</b>—Virtual circuit identifier.</li> <li>• <b>source</b>—Source of the advertisement: <b>Local</b> or <b>Remote</b>.</li> </ul> </li> </ul> |
| label stacking                                 | <p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> <li>• <b>S=0 route</b> indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed).</li> <li>• If there is no <b>S=</b> information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| [ <i>protocol, preference</i> ]                | <p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> <li>• <b>+</b>—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>-</b>—A hyphen indicates the last active route.</li> <li>• <b>*</b>—An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul> <p>In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.</p>                                                                                                                                                                                                                            |
| Level                                          | <p>(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Route Distinguisher                            | IP subnet augmented with a 64-bit prefix.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Next-hop type                                  | Type of next hop. For a description of possible values for this field, see <a href="#">Table 253 on page 3089</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 252: show route detail Output Fields (*continued*)

| Field Name                                           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Next-hop reference count</b>                      | Number of references made to the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Flood nexthop branches exceed maximum message</b> | Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Source</b>                                        | IP address of the route source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Next hop</b>                                      | Network layer address of the directly reachable neighboring system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>via</b>                                           | <p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word <b>Selected</b>. This field can also contain the following information:</p> <ul style="list-style-type: none"> <li>• <b>Weight</b>—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.</li> <li>• <b>Balance</b>—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.</li> </ul> |
| <b>Label-switched-path<br/>lsp-path-name</b>         | Name of the LSP used to reach the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Label operation</b>                               | MPLS label and operation occurring at this routing device. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Interface</b>                                     | (Local only) Local interface name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Protocol next hop</b>                             | Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Indirect next hop</b>                             | Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>State</b>                                         | State of the route (a route can be in more than one state). See <a href="#">Table 254 on page 3091</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Local AS</b>                                      | AS number of the local routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Age</b>                                           | How long the route has been known.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>AIGP</b>                                          | Accumulated interior gateway protocol (AIGP) BGP attribute.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Metricn</b>                                       | Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

Table 252: show route detail Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>MED-plus-IGP</b>            | Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>TTL-Action</b>              | For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.<br><br>For sample output, see <a href="#">show route table</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Task</b>                    | Name of the protocol that has added the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Announcement bits</b>       | List of protocols that announce this route. <b>n-Resolve inet</b> indicates that the route is used for route resolution for next hops found in the routing table. <b>n</b> is an index used by Juniper Networks customer support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>AS path</b>                 | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• <b>I</b>—IGP.</li> <li>• <b>E</b>—EGP.</li> <li>• <b>Recorded</b>—The AS path is recorded by the sample process (sampled).</li> <li>• <b>?</b>—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• <b>[ ]</b>—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893.</li> <li>• <b>[ ]</b>—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path.</li> <li>• <b>{ }</b>—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• <b>( )</b>—Parentheses enclose a confederation.</li> <li>• <b>( [ ] )</b>—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> |
| <b>VC Label</b>                | MPLS label assigned to the Layer 2 circuit virtual connection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>MTU</b>                     | Maximum transmission unit (MTU) of the Layer 2 circuit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>VLAN ID</b>                 | VLAN identifier of the Layer 2 circuit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Prefixes bound to route</b> | Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Communities</b>             | Community path attribute for the route. See <a href="#">Table 255 on page 3093</a> for all possible values for this field.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Layer2-info: encaps</b>     | Layer 2 encapsulation (for example, VPLS).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

Table 252: show route detail Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                                                      |
|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>control flags</b>             | Control flags: <b>none</b> or <b>Site Down</b> .                                                                                                                       |
| <b>mtu</b>                       | Maximum transmission unit (MTU) information.                                                                                                                           |
| <b>Label-Base, range</b>         | First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device. |
| <b>status vector</b>             | Layer 2 VPN and VPLS network layer reachability information (NLRI).                                                                                                    |
| <b>Accepted Multipath</b>        | Current active path when BGP multipath is configured.                                                                                                                  |
| <b>Accepted MultipathContrib</b> | Path currently contributing to BGP multipath.                                                                                                                          |
| <b>Localpref</b>                 | Local preference value included in the route.                                                                                                                          |
| <b>Router ID</b>                 | BGP router ID as advertised by the neighbor in the open message.                                                                                                       |
| <b>Primary Routing Table</b>     | In a routing table group, the name of the primary routing table in which the route resides.                                                                            |
| <b>Secondary Tables</b>          | In a routing table group, the name of one or more secondary tables in which the route resides.                                                                         |

[Table 253 on page 3089](#) describes all possible values for the **Next-hop Types** output field.

Table 253: Next-hop Types Output Field Values

| Next-Hop Type            | Description                                                                                                                                                                                                                                                                    |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Broadcast (bcast)</b> | Broadcast next hop.                                                                                                                                                                                                                                                            |
| <b>Deny</b>              | Deny next hop.                                                                                                                                                                                                                                                                 |
| <b>Discard</b>           | Discard next hop.                                                                                                                                                                                                                                                              |
| <b>Flood</b>             | Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast. |
| <b>Hold</b>              | Next hop is waiting to be resolved into a unicast or multicast type.                                                                                                                                                                                                           |
| <b>Indexed (idxd)</b>    | Indexed next hop.                                                                                                                                                                                                                                                              |

Table 253: Next-hop Types Output Field Values (*continued*)

| Next-Hop Type                   | Description                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Indirect (indr)</b>          | Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.                                                                                                                                                                               |
| <b>Interface</b>                | Used for a network address assigned to an interface. Unlike the <b>router</b> next hop, the <b>interface</b> next hop does not reference any specific node on the network.                                                                                                                                                                                                                                 |
| <b>Local (locl)</b>             | Local address on an interface. This next-hop type causes packets with this destination address to be received locally.                                                                                                                                                                                                                                                                                     |
| <b>Multicast (mcst)</b>         | Wire multicast next hop (limited to the LAN).                                                                                                                                                                                                                                                                                                                                                              |
| <b>Multicast discard (mdsc)</b> | Multicast discard.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Multicast group (mgrp)</b>   | Multicast group member.                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Receive (recv)</b>           | Receive.                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Reject (rjct)</b>            | Discard. An ICMP unreachable message was sent.                                                                                                                                                                                                                                                                                                                                                             |
| <b>Resolve (rslv)</b>           | Resolving next hop.                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Routed multicast (mcrt)</b>  | Regular multicast next hop.                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Router</b>                   | <p>A specific node or set of nodes to which the routing device forwards packets that match the route prefix.</p> <p>To qualify as next-hop type router, the route must meet the following criteria:</p> <ul style="list-style-type: none"> <li>• Must not be a direct or local subnet for the routing device.</li> <li>• Must have a next hop that is directly connected to the routing device.</li> </ul> |
| <b>Table</b>                    | Routing table next hop.                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Unicast (ucst)</b>           | Unicast.                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Unilist (ulst)</b>           | List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.                                                                                                                                                                                                                                                                                                                |

Table 254 on page 3091 describes all possible values for the **State** output field. A route can be in more than one state (for example, **<Active NoReadvrt Int Ext>**).



Table 254: State Output Field Values

| Value                                       | Description                                                                                                                                                                          |
|---------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Accounting                                  | Route needs accounting.                                                                                                                                                              |
| Active                                      | Route is active.                                                                                                                                                                     |
| Always Compare MED                          | Path with a lower multiple exit discriminator (MED) is available.                                                                                                                    |
| AS path                                     | Shorter AS path is available.                                                                                                                                                        |
| Cisco Non-deterministic MED selection       | Cisco nondeterministic MED is enabled, and a path with a lower MED is available.                                                                                                     |
| Clone                                       | Route is a clone.                                                                                                                                                                    |
| Cluster list length                         | Length of cluster list sent by the route reflector.                                                                                                                                  |
| Delete                                      | Route has been deleted.                                                                                                                                                              |
| Ex                                          | Exterior route.                                                                                                                                                                      |
| Ext                                         | BGP route received from an external BGP neighbor.                                                                                                                                    |
| FlashAll                                    | Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes. |
| Hidden                                      | Route not used because of routing policy.                                                                                                                                            |
| IfCheck                                     | Route needs forwarding RPF check.                                                                                                                                                    |
| IGP metric                                  | Path through next hop with lower IGP metric is available.                                                                                                                            |
| Inactive reason                             | Flags for this route, which was not selected as best for a particular destination.                                                                                                   |
| Initial                                     | Route being added.                                                                                                                                                                   |
| Int                                         | Interior route.                                                                                                                                                                      |
| Int Ext                                     | BGP route received from an internal BGP peer or a BGP confederation peer.                                                                                                            |
| Interior > Exterior > Exterior via Interior | Direct, static, IGP, or EBGp path is available.                                                                                                                                      |
| Local Preference                            | Path with a higher local preference value is available.                                                                                                                              |
| Martian                                     | Route is a martian (ignored because it is obviously invalid).                                                                                                                        |

Table 254: State Output Field Values (*continued*)

| Value                                 | Description                                                                                                                                                                                                                       |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>MartianOK</b>                      | Route exempt from martian filtering.                                                                                                                                                                                              |
| <b>Next hop address</b>               | Path with lower metric next hop is available.                                                                                                                                                                                     |
| <b>No difference</b>                  | Path from neighbor with lower IP address is available.                                                                                                                                                                            |
| <b>NoReadvrt</b>                      | Route not to be advertised.                                                                                                                                                                                                       |
| <b>NotBest</b>                        | Route not chosen because it does not have the lowest MED.                                                                                                                                                                         |
| <b>Not Best in its group</b>          | Incoming BGP AS is not the best of a group (only one AS can be the best).                                                                                                                                                         |
| <b>NotInstall</b>                     | Route not to be installed in the forwarding table.                                                                                                                                                                                |
| <b>Number of gateways</b>             | Path with a greater number of next hops is available.                                                                                                                                                                             |
| <b>Origin</b>                         | Path with a lower origin code is available.                                                                                                                                                                                       |
| <b>Pending</b>                        | Route pending because of a hold-down configured on another route.                                                                                                                                                                 |
| <b>Release</b>                        | Route scheduled for release.                                                                                                                                                                                                      |
| <b>RIB preference</b>                 | Route from a higher-numbered routing table is available.                                                                                                                                                                          |
| <b>Route Distinguisher</b>            | 64-bit prefix added to IP subnets to make them unique.                                                                                                                                                                            |
| <b>Route Metric or MED comparison</b> | Route with a lower metric or MED is available.                                                                                                                                                                                    |
| <b>Route Preference</b>               | Route with lower preference value is available                                                                                                                                                                                    |
| <b>Router ID</b>                      | Path through a neighbor with lower ID is available.                                                                                                                                                                               |
| <b>Secondary</b>                      | Route not a primary route.                                                                                                                                                                                                        |
| <b>Unusable path</b>                  | Path is not usable because of one of the following conditions: <ul style="list-style-type: none"> <li>• The route is damped.</li> <li>• The route is rejected by an import policy.</li> <li>• The route is unresolved.</li> </ul> |
| <b>Update source</b>                  | Last tiebreaker is the lowest IP address value.                                                                                                                                                                                   |

Table 255 on page 3093 describes the possible values for the **Communities** output field.

Table 255: Communities Output Field Values

| Value                                                   | Description                                                                                                                                                                                                                                                                                             |
|---------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>area-number</i>                                      | 4 bytes, encoding a 32-bit area number. For AS-external routes, the value is <b>0</b> . A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.                                                  |
| <i>bandwidth: local AS number:link-bandwidth-number</i> | Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.                   |
| <i>domain-id</i>                                        | Unique configurable number that identifies the OSPF domain.                                                                                                                                                                                                                                             |
| <i>domain-id-vendor</i>                                 | Unique configurable number that further identifies the OSPF domain.                                                                                                                                                                                                                                     |
| <i>link-bandwidth-number</i>                            | Link-bandwidth number: from <b>0</b> through <b>4,294,967,295</b> (bytes per second).                                                                                                                                                                                                                   |
| <i>local AS number</i>                                  | Local AS number: from <b>1</b> through <b>65,535</b> .                                                                                                                                                                                                                                                  |
| <i>options</i>                                          | 1 byte. Currently this is only used if the route type is <b>5</b> or <b>7</b> . Setting the least significant bit in the field indicates that the route carries a type 2 metric.                                                                                                                        |
| <i>origin</i>                                           | (Used with VPNs) Identifies where the route came from.                                                                                                                                                                                                                                                  |
| <i>ospf-route-type</i>                                  | 1 byte, encoded as <b>1</b> or <b>2</b> for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); <b>3</b> for summary routes; <b>5</b> for external routes (area number must be <b>0</b> ); <b>7</b> for NSSA routes; or <b>129</b> for sham link endpoint addresses. |
| <i>route-type-vendor</i>                                | Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute <b>0x8000</b> . The format is <i>area-number:ospf-route-type:options</i> .                                                                                            |
| <i>rte-type</i>                                         | Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute <b>0x0306</b> . The format is <i>area-number:ospf-route-type:options</i> .                                                                                            |
| <i>target</i>                                           | Defines which VPN the route participates in; <b>target</b> has the format <i>32-bit IP address:16-bit number</i> . For example, 10.19.0.0:100.                                                                                                                                                          |
| <i>unknown IANA</i>                                     | Incoming IANA codes with a value between <b>0x1</b> and <b>0x7fff</b> . This code of the BGP extended community attribute is accepted, but it is not recognized.                                                                                                                                        |
| <i>unknown OSPF vendor community</i>                    | Incoming IANA codes with a value above <b>0x8000</b> . This code of the BGP extended community attribute is accepted, but it is not recognized.                                                                                                                                                         |

## Sample Output

### show route detail

```
user@host> show route detail
```

```
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
```

```
10.10.0.0/16 (1 entry, 1 announced)
 *Static Preference: 5
 Next-hop reference count: 29
 Next hop: 192.168.71.254 via fxp0.0, selected
 State: <Active NoReadvrt Int Ext>
 Local AS: 69
 Age: 1:31:43
 Task: RT
 Announcement bits (2): 0-KRT 3-Resolve tree 2
 AS path: I

10.31.1.0/30 (2 entries, 1 announced)
 *Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 2
 Next hop: via so-0/3/0.0, selected
 State: <Active Int>
 Local AS: 69
 Age: 1:30:17
 Task: IF
 Announcement bits (1): 3-Resolve tree 2
 AS path: I
 OSPF Preference: 10
 Next-hop reference count: 1
 Next hop: via so-0/3/0.0, selected
 State: <Int>
 Inactive reason: Route Preference
 Local AS: 69
 Age: 1:30:17 Metric: 1
 Area: 0.0.0.0
 Task: OSPF
 AS path: I

10.31.1.1/32 (1 entry, 1 announced)
 *Local Preference: 0
 Next hop type: Local
 Next-hop reference count: 7
 Interface: so-0/3/0.0
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:30:20
 Task: IF
 Announcement bits (1): 3-Resolve tree 2
 AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
 *OSPF Preference: 10
 Next-hop reference count: 9
 Next hop: via so-0/3/0.0
 Next hop: 10.31.1.6 via ge-3/1/0.0, selected
 State: <Active Int>
 Local AS: 69
 Age: 1:29:56 Metric: 2
 Area: 0.0.0.0
 Task: OSPF
 Announcement bits (2): 0-KRT 3-Resolve tree 2
 AS path: I

...
```

```

224.0.0.2/32 (1 entry, 1 announced)
 *PIM Preference: 0
 Next-hop reference count: 18
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:31:45
 Task: PIM Recv
 Announcement bits (2): 0-KRT 3-Resolve tree 2
 AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
 *IGMP Preference: 0
 Next-hop reference count: 18
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:31:43
 Task: IGMP
 Announcement bits (2): 0-KRT 3-Resolve tree 2
 AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
 State: <FlashAll>
 *RSVP Preference: 7
 Next-hop reference count: 6
 Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
 Label-switched-path green-r1-r3
 Label operation: Push 100096
 State: <Active Int>
 Local AS: 69
 Age: 1:25:49 Metric: 2
 Task: RSVP
 Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
 AS path: I

10.255.71.238/32 (1 entry, 1 announced)
 State: <FlashAll>
 *RSVP Preference: 7
 Next-hop reference count: 6
 Next hop: via so-0/3/0.0 weight 0x1, selected
 Label-switched-path green-r1-r2
 State: <Active Int>
 Local AS: 69
 Age: 1:25:49 Metric: 1
 Task: RSVP
 Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
 AS path: I

private__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
 *Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 1
 Next hop: via lo0.0, selected

```

```
State: <Active Int>
Local AS: 69
Age: 1:31:44
Task: IF
AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
0 (1 entry, 1 announced)
 *MPLS Preference: 0
 Next hop type: Receive
 Next-hop reference count: 6
 State: <Active Int>
 Local AS: 69
 Age: 1:31:45 Metric: 1
 Task: MPLS
 Announcement bits (1): 0-KRT
 AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299776 (1 entry, 1 announced)
TSI:
KRT in-kernel 299776 /52 -> {Flood}
 *RSVP Preference: 7
 Next hop type: Flood
 Next-hop reference count: 130
 Flood nexthop branches exceed maximum
 Address: 0x8ea65d0

...

800010 (1 entry, 1 announced)
 *VPLS Preference: 7
 Next-hop reference count: 2
 Next hop: via vt-3/2/0.32769, selected
 Label operation: Pop
 State: <Active Int>
 Age: 1:29:30
 Task: Common L2 VC
 Announcement bits (1): 0-KRT
 AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
 *VPLS Preference: 7
 Next-hop reference count: 2
 Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
 Label-switched-path green-r1-r3
 Label operation: Push 800012, Push 100096(top)
 Protocol next hop: 10.255.70.103
 Push 800012
 Indirect next hop: 87272e4 1048574
 State: <Active Int>
 Age: 1:29:30 Metric2: 2
 Task: Common L2 VC
 Announcement bits (2): 0-KRT 1-Common L2 VC
 AS path: I
 Communities: target:11111:1 Layer2-info: encaps:VPLS,
 control flags:, mtu: 0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
```

```

abcd::10:255:71:52/128 (1 entry, 0 announced)
 *Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 1
 Next hop: via lo0.0, selected
 State: <Active Int>
 Local AS: 69
 Age: 1:31:44
 Task: IF
 AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
 *Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 1
 Next hop: via lo0.0, selected
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:31:44
 Task: IF
 AS path: I

ff02::2/128 (1 entry, 1 announced)
 *PIM Preference: 0
 Next-hop reference count: 18
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:31:45
 Task: PIM Recv6
 Announcement bits (1): 0-KRT
 AS path: I

ff02::d/128 (1 entry, 1 announced)
 *PIM Preference: 0
 Next-hop reference count: 18
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:31:45
 Task: PIM Recv6
 Announcement bits (1): 0-KRT
 AS path: I

ff02::16/128 (1 entry, 1 announced)
 *MLD Preference: 0
 Next-hop reference count: 18
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:31:43
 Task: MLD
 Announcement bits (1): 0-KRT
 AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
 *Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 1
 Next hop: via lo0.16385, selected
 State: <Active NoReadvrt Int>
 Age: 1:31:44

```

```
Task: IF
AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 10.255.70.103:1
 Next-hop reference count: 7
 Source: 10.255.70.103
 Protocol next hop: 10.255.70.103
 Indirect next hop: 2 no-forward
 State: <Secondary Active Int Ext>
 Local AS: 69 Peer AS: 69
 Age: 1:25:49 Metric2: 1
 AIGP 210
 Task: BGP_69.10.255.70.103+179
 Announcement bits (1): 0-green-l2vpn
 AS path: I
 Communities: target:11111:1 Layer2-info: encaps:VPLS,
 control flags:, mtu: 0
 Label-base: 800008, range: 8
 Localpref: 100
 Router ID: 10.255.70.103
 Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
 *L2VPN Preference: 170/-1
 Next-hop reference count: 5
 Protocol next hop: 10.255.71.52
 Indirect next hop: 0 -
 State: <Active Int Ext>
 Age: 1:31:40 Metric2: 1
 Task: green-l2vpn
 Announcement bits (1): 1-BGP.0.0.0.0+179
 AS path: I
 Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
 mtu: 0
 Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
 *L2VPN Preference: 170/-101
 Next-hop reference count: 5
 Protocol next hop: 10.255.71.52
 Indirect next hop: 0 -
 State: <Active Int Ext>
 Age: 1:31:40 Metric2: 1
 Task: green-l2vpn
 Announcement bits (1): 1-BGP.0.0.0.0+179
 AS path: I
 Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
 Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
 *L2CKT Preference: 7
 Next hop: via so-1/1/2.0 weight 1, selected
 Label-switched-path my-lsp
 Label operation: Push 100000[0]
```



```

Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
State: <Active Int>
Local AS: 99
Age: 10:21
Task: 12 circuit
Announcement bits (1): 0-LDP
AS path: I
VC Label 100000, MTU 1500, VLAN ID 512

```

### show route detail (with BGP Multipath)

```
user@host> show route detail
```

```

10.1.1.8/30 (2 entries, 1 announced)
 *BGP Preference: 170/-101
 Next hop type: Router, Next hop index: 262142
 Address: 0x901a010
 Next-hop reference count: 2
 Source: 10.1.1.2
 Next hop: 10.1.1.2 via ge-0/3/0.1, selected
 Next hop: 10.1.1.6 via ge-0/3/0.5
 State: <Active Ext>
 Local AS: 1 Peer AS: 2
 Age: 5:04:43
 Task: BGP_2.10.1.1.2+59955
 Announcement bits (1): 0-KRT
 AS path: 2 I
 Accepted Multipath
 Localpref: 100
 Router ID: 1.1.1.2
 BGP Preference: 170/-101
 Next hop type: Router, Next hop index: 678
 Address: 0x8f97520
 Next-hop reference count: 9
 Source: 10.1.1.6
 Next hop: 10.1.1.6 via ge-0/3/0.5, selected
 State: <NotBest Ext>
 Inactive reason: Not Best in its group - Active preferred
 Local AS: 1 Peer AS: 2
 Age: 5:04:43
 Task: BGP_2.10.1.1.6+58198
 AS path: 2 I
 Accepted MultipathContrib
 Localpref: 100
 Router ID: 1.1.1.3

```

## show route exact

---

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | <code>show route exact <i>destination-prefix</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                                     |
| <b>Syntax (EX Series Switches)</b> | <code>show route exact <i>destination-prefix</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code>                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                            |
| <b>Description</b>                 | Display only the routes that exactly match the specified address or range of addresses.                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b> .<br><br><b><i>destination-prefix</i></b> —Address or range of addresses.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>       | <a href="#">show route exact on page 3100</a><br><a href="#">show route exact detail on page 3100</a><br><a href="#">show route exact extensive on page 3101</a><br><a href="#">show route exact terse on page 3101</a>                                                                                                                                                                                          |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                                                                                                   |

## Sample Output

### show route exact

```
user@host> show route exact 207.17.136.0/24

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
207.17.136.0/24 *[Static/5] 2d 03:30:22
 > to 192.168.71.254 via fxp0.0
```

### show route exact detail

```
user@host> show route exact 207.17.136.0/24 detail

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
207.17.136.0/24 (1 entry, 1 announced)
 *Static Preference: 5
```

```

Next-hop reference count: 29
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2d 3:30:26
Task: RT
Announcement bits (2): 0-KRT 3-Resolve tree 2
AS path: I

```

#### show route exact extensive

```

user@host> show route exact 207.17.136.0/24 extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
 *Static Preference: 5
 Next-hop reference count: 29
 Next hop: 192.168.71.254 via fxp0.0, selected
 State: <Active NoReadvrt Int Ext>
 Local AS: 69
 Age: 1:25:18
 Task: RT
 Announcement bits (2): 0-KRT 3-Resolve tree 2
 AS path: I

```

#### show route exact terse

```

user@host> show route exact 207.17.136.0/24 terse

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination P Prf Metric 1 Metric 2 Next hop AS path
* 207.17.136.0/24 S 5 >192.168.71.254

```

## show route export

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | show route export<br><brief   detail><br><instance <instance-name>   routing-table-name><br><logical-system (all   logical-system-name)>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switches)</b> | show route export<br><brief   detail><br><instance <instance-name>   routing-table-name>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                 | Display policy-based route export information. Policy-based export simplifies the process of exchanging route information between routing instances.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                     | <p><b>none</b>—(Same as <b>brief</b>.) Display standard information about policy-based export for all instances and routing tables on all systems.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>instance &lt;instance-name&gt;</b>—(Optional) Display a particular routing instance for which policy-based export is currently enabled.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>routing-table-name</b>—(Optional) Display information about policy-based export for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the <b>show route export inet</b> command).</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>       | <a href="#">show route export on page 3103</a><br><a href="#">show route export detail on page 3103</a><br><a href="#">show route export instance detail on page 3103</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>               | Table 256 on page 3102 lists the output fields for the <b>show route export</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

Table 256: show route export Output Fields

| Field Name                        | Field Description                                                                                                                                           | Level of Output   |
|-----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| <b>Table</b> or <b>table-name</b> | Name of the routing tables that either import or export routes.                                                                                             | All levels        |
| <b>Routes</b>                     | Number of routes exported from this table into other tables. If a particular route is exported to different tables, the counter will only increment by one. | <b>brief</b> none |
| <b>Export</b>                     | Whether the table is currently exporting routes to other tables: <b>Y</b> or <b>N</b> (Yes or No).                                                          | <b>brief</b> none |

Table 256: show route export Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                           | Level of Output |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Import</b>        | Tables currently importing routes from the originator table. (Not displayed for tables that are not exporting any routes.)                                                                                                                                                                                                                                                  | <b>detail</b>   |
| <b>Flags</b>         | ( <b>instance</b> keyword only) Flags for this feature on this instance: <ul style="list-style-type: none"> <li><b>config auto-policy</b>—The policy was deduced from the configured IGP export policies.</li> <li><b>cleanup</b>—Configuration information for this instance is no longer valid.</li> <li><b>config</b>—The instance was explicitly configured.</li> </ul> | <b>detail</b>   |
| <b>Options</b>       | ( <b>instance</b> keyword only) Configured option displays the type of routing tables the feature handles: <ul style="list-style-type: none"> <li><b>unicast</b>—Indicates <i>instance.inet.0</i>.</li> <li><b>multicast</b>—Indicates <i>instance.inet.2</i>.</li> <li><b>unicast multicast</b>—Indicates <i>instance.inet.0</i> and <i>instance.inet.2</i>.</li> </ul>    | <b>detail</b>   |
| <b>Import policy</b> | ( <b>instance</b> keyword only) Policy that <b>route export</b> uses to construct the import-export matrix. Not displayed if the instance type is <b>vrf</b> .                                                                                                                                                                                                              | <b>detail</b>   |
| <b>Instance</b>      | ( <b>instance</b> keyword only) Name of the routing instance.                                                                                                                                                                                                                                                                                                               | <b>detail</b>   |
| <b>Type</b>          | ( <b>instance</b> keyword only) Type of routing instance: <b>forwarding</b> , <b>non-forwarding</b> , or <b>vrf</b> .                                                                                                                                                                                                                                                       | <b>detail</b>   |

## Sample Output

### show route export

```

user@host> show route export
Table Export Routes
inet.0 N 0
black.inet.0 Y 3
red.inet.0 Y 4

```

### show route export detail

```

user@host> show route export detail
inet.0 Routes: 0
black.inet.0 Routes: 3
 Import: [inet.0]
red.inet.0 Routes: 4
 Import: [inet.0]

```

### show route export instance detail

```

user@host> show route export instance detail
Instance: master Type: forwarding
Flags: <config auto-policy> Options: <unicast multicast>
Import policy: [(ospf-master-from-red || isis-master-from-black)]

```

Instance: black  
Instance: red

Type: non-forwarding  
Type: non-forwarding

## show route extensive

|                                    |                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | show route extensive<br><destination-prefix><br><logical-system (all   logical-system-name)>                                                                                                                                                                                                                                                 |
| <b>Syntax (EX Series Switches)</b> | show route extensive<br><destination-prefix>                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                        |
| <b>Description</b>                 | Display extensive information about the active entries in the routing tables.                                                                                                                                                                                                                                                                |
| <b>Options</b>                     | <p><b>none</b>—Display all active entries in the routing table.</p> <p><b>destination-prefix</b>—(Optional) Display active entries for the specified address or range of addresses.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>       | <a href="#">show route extensive on page 3110</a><br><a href="#">show route extensive (Access Route) on page 3116</a><br><a href="#">show route extensive (Route Reflector) on page 3117</a><br><a href="#">show route extensive (FRR and LFA) on page 3117</a><br><a href="#">show route extensive (FRR and LFA) on page 3118</a>           |
| <b>Output Fields</b>               | Table 89 on page 1219 describes the output fields for the <b>show route extensive</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                       |

Table 257: show route extensive Output Fields

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>routing-table-name</i>  | Name of the routing table (for example, inet.0).                                                                                                                                                                                                                                                                                                                   |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table.                                                                                                                                                                                                                                                                                            |
| <i>number routes</i>       | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> <li><b>active</b> (routes that are active).</li> <li><b>holddown</b> (routes that are in the pending state before being declared inactive).</li> <li><b>hidden</b> (routes that are not used because of a routing policy).</li> </ul> |

Table 257: show route extensive Output Fields (*continued*)

| Field Name                                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>route-destination</b><br>(entry, announced) | <p>Route destination (for example: 10.0.0.1/24). The <b>entry</b> value is the number of route for this destination, and the <b>announced</b> value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> <li>• <b>MPLS-label</b> (for example, 80001).</li> <li>• <b>interface-name</b> (for example, ge-1/0/2).</li> <li>• <b>neighbor-address:control-word-status:encapsulation type:vc-id:source</b> (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96).</li> <li>• <b>neighbor-address</b>—Address of the neighbor.</li> <li>• <b>control-word-status</b>—Whether the use of the control word has been negotiated for this virtual circuit: <b>NoCtrlWord</b> or <b>CtrlWord</b>.</li> <li>• <b>encapsulation type</b>—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.</li> <li>• <b>vc-id</b>—Virtual circuit identifier.</li> <li>• <b>source</b>—Source of the advertisement: <b>Local</b> or <b>Remote</b>.</li> </ul> |
| <b>TSI</b>                                     | Protocol header information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>label stacking</b>                          | <p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> <li>• <b>S=0 route</b> indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed).</li> <li>• If there is no <b>S=</b> information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>[protocol, preference]</b>                  | <p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> <li>• <b>+</b>—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>-</b>—A hyphen indicates the last active route.</li> <li>• <b>*</b>—An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul> <p>In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.</p>                                                                                                                                                                                  |
| <b>Level</b>                                   | <p>(IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Route Distinguisher</b>                     | IP subnet augmented with a 64-bit prefix.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |



Table 257: show route extensive Output Fields (*continued*)

| Field Name                                           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Next-hop type</b>                                 | Type of next hop. For a description of possible values for this field, see the Output Field table in the <a href="#">show route detail</a> command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Next-hop reference count</b>                      | Number of references made to the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Flood nexthop branches exceed maximum message</b> | Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Source</b>                                        | IP address of the route source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Next hop</b>                                      | Network layer address of the directly reachable neighboring system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>via</b>                                           | <p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word <b>Selected</b>. This field can also contain the following information:</p> <ul style="list-style-type: none"> <li>• <b>Weight</b>—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.</li> <li>• <b>Balance</b>—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.</li> </ul> |
| <b>Label-switched-path <i>lsp-path-name</i></b>      | Name of the LSP used to reach the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Label operation</b>                               | MPLS label and operation occurring at this routing device. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Offset</b>                                        | Whether the metric has been increased or decreased by an offset value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Interface</b>                                     | (Local only) Local interface name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Protocol next hop</b>                             | Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b><i>label-operation</i></b>                        | MPLS label and operation occurring at this routing device. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Indirect next hops</b>                            | When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>State</b>                                         | State of the route (a route can be in more than one state). See the Output Field table in the <a href="#">show route detail</a> command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

Table 257: show route extensive Output Fields (*continued*)

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Session ID</b>      | The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Inactive reason</b> | <p>If the route is inactive, the reason for its current state is indicated. Typical reasons include:</p> <ul style="list-style-type: none"> <li>• <b>Active preferred</b>—Currently active route was selected over this route.</li> <li>• <b>Always compare MED</b>—Path with a lower multiple exit discriminator (MED) is available.</li> <li>• <b>AS path</b>—Shorter AS path is available.</li> <li>• <b>Cisco Non-deterministic MED selection</b>—Cisco nondeterministic MED is enabled and a path with a lower MED is available.</li> <li>• <b>Cluster list length</b>—Path with a shorter cluster list length is available.</li> <li>• <b>Forwarding use only</b>—Path is only available for forwarding purposes.</li> <li>• <b>IGP metric</b>—Path through the next hop with a lower IGP metric is available.</li> <li>• <b>IGP metric type</b>—Path with a lower OSPF link-state advertisement type is available.</li> <li>• <b>Interior &gt; Exterior &gt; Exterior via Interior</b>—Direct, static, IGP, or EBGP path is available.</li> <li>• <b>Local preference</b>—Path with a higher local preference value is available.</li> <li>• <b>Next hop address</b>—Path with a lower metric next hop is available.</li> <li>• <b>No difference</b>—Path from a neighbor with a lower IP address is available.</li> <li>• <b>Not Best in its group</b>—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed).</li> <li>• <b>Number of gateways</b>—Path with a higher number of next hops is available.</li> <li>• <b>Origin</b>—Path with a lower origin code is available.</li> <li>• <b>OSPF version</b>—Path does not support the indicated OSPF version.</li> <li>• <b>RIB preference</b>—Route from a higher-numbered routing table is available.</li> <li>• <b>Route distinguisher</b>—64-bit prefix added to IP subnets to make them unique.</li> <li>• <b>Route metric or MED comparison</b>—Route with a lower metric or MED is available.</li> <li>• <b>Route preference</b>—Route with a lower preference value is available.</li> <li>• <b>Router ID</b>—Path through a neighbor with a lower ID is available.</li> <li>• <b>Unusable path</b>—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved.</li> <li>• <b>Update source</b>—Last tiebreaker is the lowest IP address value.</li> </ul> |
| <b>Local AS</b>        | Autonomous system (AS) number of the local routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Age</b>             | How long the route has been known.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>AIGP</b>            | Accumulated interior gateway protocol (AIGP) BGP attribute.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Metric</b>          | Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>MED-plus-IGP</b>    | Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

Table 257: show route extensive Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TTL-Action              | <p>For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.</p> <p>For sample output, see <a href="#">show route table</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Task                    | Name of the protocol that has added the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Announcement bits       | List of protocols that announce this route. <b>n-Resolve inet</b> indicates that the route is used for route resolution for next hops found in the routing table. <b>n</b> is an index used by Juniper Networks customer support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| AS path                 | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• <b>I</b>—IGP.</li> <li>• <b>E</b>—EGP.</li> <li>• <b>Recorded</b>—The AS path is recorded by the sample process (sampled).</li> <li>• <b>?</b>—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• <b>[ ]</b>—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured.</li> <li>• <b>{ }</b>—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• <b>( )</b>—Parentheses enclose a confederation.</li> <li>• <b>( [ ] )</b>—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> |
| AS path: I <Originator> | (For route reflected output only) Originator ID attribute set by the route reflector.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| VC Label                | MPLS label assigned to the Layer 2 circuit virtual connection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| MTU                     | Maximum transmission unit (MTU) of the Layer 2 circuit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| VLAN ID                 | VLAN identifier of the Layer 2 circuit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Cluster list            | (For route reflected output only) Cluster ID sent by the route reflector.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Originator ID           | (For route reflected output only) Address of router that originally sent the route to the route reflector.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Prefixes bound to route | Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Communities             | Community path attribute for the route. See the Output Field table in the <a href="#">show route detail</a> command for all possible values for this field.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 257: show route extensive Output Fields (*continued*)

| Field Name                   | Field Description                                                                                                                                                                                                                                                                                                |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Layer2-info: encaps</b>   | Layer 2 encapsulation (for example, VPLS).                                                                                                                                                                                                                                                                       |
| <b>control flags</b>         | Control flags: <b>none</b> or Site Down.                                                                                                                                                                                                                                                                         |
| <b>mtu</b>                   | Maximum transmission unit (MTU) information.                                                                                                                                                                                                                                                                     |
| <b>Label-Base, range</b>     | First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.                                                                                                                                           |
| <b>status vector</b>         | Layer 2 VPN and VPLS network layer reachability information (NLRI).                                                                                                                                                                                                                                              |
| <b>Localpref</b>             | Local preference value included in the route.                                                                                                                                                                                                                                                                    |
| <b>Router ID</b>             | BGP router ID as advertised by the neighbor in the open message.                                                                                                                                                                                                                                                 |
| <b>Primary Routing Table</b> | In a routing table group, the name of the primary routing table in which the route resides.                                                                                                                                                                                                                      |
| <b>Secondary Tables</b>      | In a routing table group, the name of one or more secondary tables in which the route resides.                                                                                                                                                                                                                   |
| <b>Originating RIB</b>       | Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix. |
| <b>Node path count</b>       | Number of nodes in the path.                                                                                                                                                                                                                                                                                     |
| <b>Forwarding nexthops</b>   | Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.                                                                                                                            |

## Sample Output

### show route extensive

```

user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
 *Static Preference: 5
 Next-hop reference count: 29
 Next hop: 192.168.71.254 via fxp0.0, selected
 State: <Active NoReadvrt Int Ext>
 Local AS: 69
 Age: 1:34:06
 Task: RT
 Announcement bits (2): 0-KRT 3-Resolve tree 2
 AS path: I

10.31.1.0/30 (2 entries, 1 announced)
 *Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 2

```

```

 Next hop: via so-0/3/0.0, selected
 State: <Active Int>
 Local AS: 69
 Age: 1:32:40
 Task: IF
 Announcement bits (1): 3-Resolve tree 2
 AS path: I
OSPF Preference: 10
 Next-hop reference count: 1
 Next hop: via so-0/3/0.0, selected
 State: <Int>
 Inactive reason: Route Preference
 Local AS: 69
 Age: 1:32:40 Metric: 1
 Area: 0.0.0.0
 Task: OSPF
 AS path: I

10.31.1.1/32 (1 entry, 1 announced)
*Local Preference: 0
 Next hop type: Local
 Next-hop reference count: 7
 Interface: so-0/3/0.0
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:32:43
 Task: IF
 Announcement bits (1): 3-Resolve tree 2
 AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.31.2.0/30 -> {10.31.1.6}
*OSPF Preference: 10
 Next-hop reference count: 9
 Next hop: via so-0/3/0.0
 Next hop: 10.31.1.6 via ge-3/1/0.0, selected
 State: <Active Int>
 Local AS: 69
 Age: 1:32:19 Metric: 2
 Area: 0.0.0.0
 Task: OSPF
 Announcement bits (2): 0-KRT 3-Resolve tree 2
 AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 224.0.0.2/32 -> {}
*PIM Preference: 0
 Next-hop reference count: 18
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:34:08
 Task: PIM Recv
 Announcement bits (2): 0-KRT 3-Resolve tree 2
 AS path: I

```

```
...

224.0.0.22/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 224.0.0.22/32 -> {}
 *IGMP Preference: 0
 Next-hop reference count: 18
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:34:06
 Task: IGMP
 Announcement bits (2): 0-KRT 3-Resolve tree 2
 AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
State: <FlashAll>
 *RSVP Preference: 7
 Next-hop reference count: 6
 Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
 Label-switched-path green-r1-r3
 Label operation: Push 100096
 State: <Active Int>
 Local AS: 69
 Age: 1:28:12 Metric: 2
 Task: RSVP
 Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
 AS path: I

10.255.71.238/32 (1 entry, 1 announced)
State: <FlashAll>
 *RSVP Preference: 7
 Next-hop reference count: 6
 Next hop: via so-0/3/0.0 weight 0x1, selected
 Label-switched-path green-r1-r2
 State: <Active Int>
 Local AS: 69
 Age: 1:28:12 Metric: 1
 Task: RSVP
 Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
 AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
 *Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 1
 Next hop: via lo0.0, selected
 State: <Active Int>
 Local AS: 69
 Age: 1:34:07
 Task: IF
 AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
```

```

0 (1 entry, 1 announced)
TSI:
KRT in-kernel 0 /36 -> {}
 *MPLS Preference: 0
 Next hop type: Receive
 Next-hop reference count: 6
 State: <Active Int>
 Local AS: 69
 Age: 1:34:08 Metric: 1
 Task: MPLS
 Announcement bits (1): 0-KRT
 AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299776 (1 entry, 1 announced)
TSI:
KRT in-kernel 299776 /52 -> {Flood}
 *RSVP Preference: 7
 Next hop type: Flood
 Next-hop reference count: 130
 Flood nexthop branches exceed maximum
 Address: 0x8ea65d0

...

800010 (1 entry, 1 announced)
TSI:
KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
 *VPLS Preference: 7
 Next-hop reference count: 2
 Next hop: via vt-3/2/0.32769, selected
 Label operation: Pop
 State: <Active Int>
 Age: 1:31:53
 Task: Common L2 VC
 Announcement bits (1): 0-KRT
 AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
TSI:
KRT in-kernel vt-3/2/0.32769.0 /16 -> {indirect(1048574)}
 *VPLS Preference: 7
 Next-hop reference count: 2
 Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
 Label-switched-path green-r1-r3
 Label operation: Push 800012, Push 100096(top)
 Protocol next hop: 10.255.70.103
 Push 800012
 Indirect next hop: 87272e4 1048574
 State: <Active Int>
 Age: 1:31:53 Metric2: 2
 Task: Common L2 VC
 Announcement bits (2): 0-KRT 1-Common L2 VC
 AS path: I
 Communities: target:11111:1 Layer2-info: encaps:VPLS,
 control flags:, mtu: 0
 Indirect next hops: 1
 Protocol next hop: 10.255.70.103 Metric: 2

```

```

 Push 800012
 Indirect next hop: 87272e4 1048574
 Indirect path forwarding next hops: 1
 Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1
 10.255.70.103/32 Originating RIB: inet.3
 Metric: 2 Node path count: 1
 Forwarding nexthops: 1
 Nexthop: 10.31.1.6 via ge-3/1/0.0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
 *Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 1
 Next hop: via lo0.0, selected
 State: <Active Int>
 Local AS: 69
 Age: 1:34:07
 Task: IF
 AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
 *Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 1
 Next hop: via lo0.0, selected
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:34:07
 Task: IF
 AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
 *PIM Preference: 0
 Next-hop reference count: 18
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:34:08
 Task: PIM Recv6
 Announcement bits (1): 0-KRT
 AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
 *PIM Preference: 0
 Next-hop reference count: 18
 State: <Active NoReadvrt Int>
 Local AS: 69
 Age: 1:34:08
 Task: PIM Recv6
 Announcement bits (1): 0-KRT
 AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
 *MLD Preference: 0
```



```

Next-hop reference count: 18
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:34:06
Task: MLD
Announcement bits (1): 0-KRT
AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
*Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 1
 Next hop: via lo0.16385, selected
 State: <Active NoReadvrt Int>
 Age: 1:34:07
 Task: IF
 AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
*BGP Preference: 170/-101
 Route Distinguisher: 10.255.70.103:1
 Next-hop reference count: 7
 Source: 10.255.70.103
 Protocol next hop: 10.255.70.103
 Indirect next hop: 2 no-forward
 State: <Secondary Active Int Ext>
 Local AS: 69 Peer AS: 69
 Age: 1:28:12 Metric2: 1
 Task: BGP_69.10.255.70.103+179
 Announcement bits (1): 0-green-l2vpn
 AS path: I
 Communities: target:11111:1 Layer2-info: encaps:VPLS,
control flags:, mtu: 0
 Label-base: 800008, range: 8
 Localpref: 100
 Router ID: 10.255.70.103
 Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699540
*L2VPN Preference: 170/-1
 Next-hop reference count: 5
 Protocol next hop: 10.255.71.52
 Indirect next hop: 0 -
 State: <Active Int Ext>
 Age: 1:34:03 Metric2: 1
 Task: green-l2vpn
 Announcement bits (1): 1-BGP.0.0.0.0+179
 AS path: I
 Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
mtu: 0
 Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699528

```

```
*L2VPN Preference: 170/-101
Next-hop reference count: 5
Protocol next hop: 10.255.71.52
Indirect next hop: 0 -
State: <Active Int Ext>
Age: 1:34:03 Metric2: 1
Task: green-l2vpn
Announcement bits (1): 1-BGP.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

TSI:

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
*L2CKT Preference: 7
Next hop: via so-1/1/2.0 weight 1, selected
Label-switched-path my-lsp
Label operation: Push 100000[0]
Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
State: <Active Int>
Local AS: 99
Age: 10:21
Task: l2 circuit
Announcement bits (1): 0-LDP
AS path: I
VC Label 100000, MTU 1500, VLAN ID 512

55.0.0.0/24 (1 entry, 1 announced)
TSI:
KRT queued (pending) add
55.0.0.0/24 -> {Push 300112}
*BGP Preference: 170/-101
Next hop type: Router
Address: 0x925c208
Next-hop reference count: 2
Source: 10.0.0.9
Next hop: 10.0.0.9 via ge-1/2/0.15, selected
Label operation: Push 300112
Label TTL action: prop-ttl
State: <Active Ext>
Local AS: 7019 Peer AS: 13979
Age: 1w0d 23:06:56
AIGP: 25
Task: BGP_13979.10.0.0.9+56732
Announcement bits (1): 0-KRT
AS path: 13979 7018 I
Accepted
Route Label: 300112
Localpref: 100
Router ID: 10.9.9.1
```

#### show route extensive (Access Route)

```
user@host> show route 13.160.0.102 extensive
```

```

inet.0: 39256 destinations, 39258 routes (39255 active, 0 holddown, 1 hidden)
13.160.0.102/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.102/32 -> {13.160.0.2}
OSPF area : 0.0.0.0, LSA ID : 13.160.0.102, LSA type : Extern
 *Access Preference: 13
 Next-hop reference count: 78472
 Next hop: 13.160.0.2 via fe-0/0/0.0, selected
 State: <Active Int>
 Age: 12
 Task: RPD Unix Domain Server./var/run/rpd_serv.local
 Announcement bits (2): 0-KRT 1-OSPFv2
 AS path: I

```

### show route extensive (Route Reflector)

```

user@host> show route extensive
1.0.0.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 1.0.0.0/8 -> {indirect(40)}
 *BGP Preference: 170/-101
 Source: 192.168.4.214
 Protocol next hop: 207.17.136.192 Indirect next hop: 84ac908 40
 State: <Active Int Ext>
 Local AS: 10458 Peer AS: 10458
 Age: 3:09 Metric: 0 Metric2: 0
 Task: BGP_10458.192.168.4.214+1033
 Announcement bits (2): 0-KRT 4-Resolve inet.0
 AS path: 3944 7777 I <Originator>
 Cluster list: 1.1.1.1
 Originator ID: 10.255.245.88
 Communities: 7777:7777
 Localpref: 100
 Router ID: 4.4.4.4
 Indirect next hops: 1
 Protocol next hop: 207.17.136.192 Metric: 0
 Indirect next hop: 84ac908 40
 Indirect path forwarding next hops: 0
 Next hop type: Discard

```

### show route extensive (FRR and LFA)

```

user@host> show route 20:31:2:0 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
20.31.2.0/24 (2 entries, 1 announced)
 State: FlashAll

TSI:
KRT in-kernel 20.31.2.0/24 -> {Push 299776, Push 299792}
 *RSVP Preference: 7/1
 Next hop type: Router, Next hop index: 1048574
 Address: 0xbbbc010
 Next-hop reference count: 5
 Next hop: 10.31.1.2 via ge-2/1/8.0 weight 0x1, selected
 Label-switched-path europa-d-to-europa-e
 Label operation: Push 299776
 Label TTL action: prop-ttl
 Session Id: 0x201
 Next hop: 10.31.2.2 via ge-2/1/4.0 weight 0x4001
 Label-switched-path europa-d-to-europa-e
 Label operation: Push 299792

```

```
Label TTL action: prop-ttl
Session Id: 0x202
State: Active Int
Local AS: 100
Age: 5:31 Metric: 2
Task: RSVP
Announcement bits (1): 0-KRT
AS path: I
OSPF Preference: 10
Next hop type: Router, Next hop index: 615
Address: 0xb9d78c4
Next-hop reference count: 7
Next hop: 10.31.1.2 via ge-2/1/8.0, selected
Session Id: 0x201
State: Int
Inactive reason: Route Preference
Local AS: 100
Age: 5:35 Metric: 3
Area: 0.0.0.0
Task: OSPF
AS path: I
```

#### show route extensive (FRR and LFA)

```
user@host> show route 20:31:2:0 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
20.31.2.0/24 (2 entries, 1 announced)
State: FlashAll
TSI:
KRT in-kernel 20.31.2.0/24 -> {Push 299776, Push 299792}
*RSVP Preference: 7/1
Next hop type: Router, Next hop index: 1048574
Address: 0xbbbc010
Next-hop reference count: 5
Next hop: 10.31.1.2 via ge-2/1/8.0 weight 0x1, selected
Label-switched-path europa-d-to-europa-e
Label operation: Push 299776
Label TTL action: prop-ttl
Session Id: 0x201
Next hop: 10.31.2.2 via ge-2/1/4.0 weight 0x4001
Label-switched-path europa-d-to-europa-e
Label operation: Push 299792
Label TTL action: prop-ttl
Session Id: 0x202
State: Active Int
Local AS: 100
Age: 5:31 Metric: 2
Task: RSVP
Announcement bits (1): 0-KRT
AS path: I
OSPF Preference: 10
Next hop type: Router, Next hop index: 615
Address: 0xb9d78c4
Next-hop reference count: 7
Next hop: 10.31.1.2 via ge-2/1/8.0, selected
Session Id: 0x201
State: Int
Inactive reason: Route Preference
Local AS: 100
Age: 5:35 Metric: 3
Area: 0.0.0.0
```

Task: OSPF  
AS path: I

## show route forwarding-table

---

|                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                     | <pre>show route forwarding-table &lt;detail   extensive   summary&gt; &lt;all&gt; &lt;ccc interface-name&gt; &lt;destination destination-prefix&gt; &lt;family family   matching matching&gt; &lt;interface-name interface-name&gt; &lt;label name&gt; &lt;matching matching&gt; &lt;multicast&gt; &lt;table (default   logical-system-name/routing-instance-name   routing-instance-name)&gt; &lt;vlan (all   vlan-name)&gt; &lt;vpn vpn&gt;</pre>                                                                                     |
| <b>Syntax (MX Series Routers)</b> | <pre>show route forwarding-table &lt;detail   extensive   summary&gt; &lt;all&gt; &lt;bridge-domain (all   domain-name)&gt; &lt;ccc interface-name&gt; &lt;destination destination-prefix&gt; &lt;family family   matching matching&gt; &lt;interface-name interface-name&gt; &lt;label name&gt; &lt;learning-vlan-id learning-vlan-id&gt; &lt;matching matching&gt; &lt;multicast&gt; &lt;table (default   logical-system-name/routing-instance-name   routing-instance-name)&gt; &lt;vlan (all   vlan-name)&gt; &lt;vpn vpn&gt;</pre> |
| <b>Syntax (Routing Matrix)</b>    | <pre>show route forwarding-table &lt;detail   extensive   summary&gt; &lt;all&gt; &lt;ccc interface-name&gt; &lt;destination destination-prefix&gt; &lt;family family   matching matching&gt; &lt;interface-name interface-name&gt; &lt;matching matching&gt; &lt;label name&gt; &lt;lcc number&gt; &lt;multicast&gt; &lt;table routing-instance-name&gt; &lt;vpn vpn&gt;</pre>                                                                                                                                                         |
| <b>Release Information</b>        | <p>Command introduced before Junos OS Release 7.4.</p> <p>Option <b>bridge-domain</b> introduced in Junos OS Release 7.5</p> <p>Option <b>learning-vlan-id</b> introduced in Junos OS Release 8.4</p> <p>Options <b>all</b> and <b>vlan</b> introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                              |

**Description** Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.



**NOTE:** The Routing Engine copies the forwarding table to the Packet Forwarding Engine, the part of the router that is responsible for forwarding packets. To display the entries in the Packet Forwarding Engine's forwarding table, use the **show pfe route** command.

**Options** **none**—Display the routes in the forwarding tables. By default, the **show route forwarding-table** command does not display information about private, or internal, forwarding tables.

**detail | extensive | summary**—(Optional) Display the specified level of output.

**all**—(Optional) Display routing table entries for all forwarding tables, including private, or internal, tables.

**bridge-domain (all | *bridge-domain-name*)**—(MX Series routers only) (Optional) Display route entries for all bridge domains or the specified bridge domain.

**ccc *interface-name***—(Optional) Display route entries for the specified circuit cross-connect interface.

**destination *destination-prefix***—(Optional) Destination prefix.

**family *family***—(Optional) Display routing table entries for the specified family: **fibre-channel**, **fmembers**, **inet**, **inet6**, **iso**, **mpls**, **tnp**, **unix**, **vpls**, or **vlan-classification**.

**interface-name *interface-name***—(Optional) Display routing table entries for the specified interface.

**label *name***—(Optional) Display route entries for the specified label.

**lcc *number***—(Routing Matrix only) (Optional) On a routing matrix composed of a TX Matrix Plus router and T640 routers configured in the routing matrix, display information for the specified T640 router (or line-card chassis) connected to the TX Matrix router. On a routing matrix composed of the TX Matrix Plus router and T1600 routers configured in the routing matrix, display information for the specified T1600 router (or line-card chassis) connected to the TX Matrix Plus router. Replace ***number*** with a value from 0 through 3.

**learning-vlan-id *learning-vlan-id***—(MX Series routers only) (Optional) Display learned information for all VLANs or for the specified VLAN.

**matching *matching***—(Optional) Display routing table entries matching the specified prefix or prefix length.

**multicast**—(Optional) Display routing table entries for multicast routes.

**table** (**default** | *logical-system-name/routing-instance-name* |

*routing-instance-name*)—(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing instance. If your device supports logical systems, you can also display route entries for the specified logical system and routing instance. To view the routing instances on your device, use the [show route instance](#) command.

**vlan** (**all** | *vlan-name*)—(Optional) Display information for all VLANs or for the specified VLAN.

**vpn** *vpn*—(Optional) Display routing table entries for a specified VPN.

**Required Privilege Level** view

**List of Sample Output** [show route forwarding-table on page 3125](#)  
[show route forwarding-table detail on page 3126](#)  
[show route forwarding-table destination extensive \(Weights and Balances\) on page 3126](#)  
[show route forwarding-table extensive on page 3127](#)  
[show route forwarding-table extensive \(RPF\) on page 3128](#)  
[show route forwarding-table family mpls on page 3129](#)  
[show route forwarding-table family vpls on page 3129](#)  
[show route forwarding-table family vpls extensive on page 3129](#)  
[show route forwarding-table table default on page 3131](#)  
[show route forwarding-table table](#)  
[logical-system-name/routing-instance-name on page 3132](#)  
[show route forwarding-table vpn on page 3132](#)

**Output Fields** [Table 258 on page 3122](#) lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified, or when the **detail** keyword is used instead of the **extensive** keyword.

**Table 258: show route forwarding-table Output Fields**

| Field Name     | Field Description                                                                                                                                                                                            | Level of Output                |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------|
| Logical system | Name of the logical system. This field is displayed if you specify the <b>table</b> <i>logical-system-name/routing-instance-name</i> option on a device that is configured for and supports logical systems. | All levels                     |
| Routing table  | Name of the routing table (for example, inet, inet6, mpls).                                                                                                                                                  | All levels                     |
| Address family | Address family (for example, IP, IPv6, ISO, MPLS, and VPLS).                                                                                                                                                 | All levels                     |
| Destination    | Destination of the route.                                                                                                                                                                                    | <b>detail</b> <b>extensive</b> |



Table 258: show route forwarding-table Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output         |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Route Type (Type)</b>       | How the route was placed into the forwarding table. When the <b>detail</b> keyword is used, the route type might be abbreviated (as shown in parentheses): <ul style="list-style-type: none"> <li>• <b>cloned (clon)</b>—(TCP or multicast only) Cloned route.</li> <li>• <b>destination (dest)</b>—Remote addresses directly reachable through an interface.</li> <li>• <b>destination down (iddn)</b>—Destination route for which the interface is unreachable.</li> <li>• <b>interface cloned (ifcl)</b>—Cloned route for which the interface is unreachable.</li> <li>• <b>route down (ifdn)</b>—Interface route for which the interface is unreachable.</li> <li>• <b>ignore (ignr)</b>—Ignore this route.</li> <li>• <b>interface (intf)</b>—Installed as a result of configuring an interface.</li> <li>• <b>permanent (perm)</b>—Routes installed by the kernel when the routing table is initialized.</li> <li>• <b>user</b>—Routes installed by the routing protocol process or as a result of the configuration.</li> </ul> | All levels              |
| <b>Route Reference (RtRef)</b> | Number of routes to reference.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b> |
| <b>Flags</b>                   | Route type flags: <ul style="list-style-type: none"> <li>• <b>none</b>—No flags are enabled.</li> <li>• <b>accounting</b>—Route has accounting enabled.</li> <li>• <b>cached</b>—Cache route.</li> <li>• <b>incoming-iface interface-number</b>—Check against incoming interface.</li> <li>• <b>prefix load balance</b>—Load balancing is enabled for this prefix.</li> <li>• <b>rt nh decoupled</b>—Route has been decoupled from the next hop to the destination.</li> <li>• <b>sent to PFE</b>—Route has been sent to the Packet Forwarding Engine.</li> <li>• <b>static</b>—Static route.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>extensive</b>        |
| <b>Next hop</b>                | IP address of the next hop to the destination.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b> |

Table 258: show route forwarding-table Output Fields (*continued*)

| Field Name                        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output              |
|-----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Next hop Type (Type)</b>       | <p>Next-hop type. When the <b>detail</b> keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):</p> <ul style="list-style-type: none"> <li>• <b>broadcast (bcst)</b>—Broadcast.</li> <li>• <b>deny</b>—Deny.</li> <li>• <b>discard (dscd)</b>—Discard.</li> <li>• <b>hold</b>—Next hop is waiting to be resolved into a unicast or multicast type.</li> <li>• <b>indexed (idxd)</b>—Indexed next hop.</li> <li>• <b>indirect (indr)</b>—Indirect next hop.</li> <li>• <b>local (locl)</b>—Local address on an interface.</li> <li>• <b>routed multicast (mcrst)</b>—Regular multicast next hop.</li> <li>• <b>multicast (mcst)</b>—Wire multicast next hop (limited to the LAN).</li> <li>• <b>multicast discard (mdsc)</b>—Multicast discard.</li> <li>• <b>multicast group (mgrp)</b>—Multicast group member.</li> <li>• <b>receive (rcv)</b>—Receive.</li> <li>• <b>reject (rjct)</b>—Discard. An ICMP unreachable message was sent.</li> <li>• <b>resolve (rslv)</b>—Resolving the next hop.</li> <li>• <b>unicast (ucst)</b>—Unicast.</li> <li>• <b>unilist (ulst)</b>—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.</li> </ul> | <b>detail extensive</b>      |
| <b>Index</b>                      | Software index of the next hop that is used to route the traffic for a given prefix.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b> |
| <b>Route interface-index</b>      | Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>extensive</b>             |
| <b>Reference (NhRef)</b>          | Number of routes that refer to this next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive none</b> |
| <b>Next-hop interface (Netif)</b> | Interface used to reach the next hop.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive none</b> |
| <b>Weight</b>                     | Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible (see the <b>Balance</b> field description).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>extensive</b>             |
| <b>Balance</b>                    | Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a router is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>extensive</b>             |
| <b>RPF interface</b>              | List of interfaces from which the prefix can be accepted. Reverse path forwarding (RPF) information is displayed only when <b>rpf-check</b> is configured on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>extensive</b>             |

## Sample Output

### show route forwarding-table

```

user@host> show route forwarding-table
Routing table: default.inet
Internet:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 46 4
0.0.0.0/32 perm 0 dscd 44 1
1.1.1.0/24 ifdn 0 rslv 608 1 ge-2/0/1.0
1.1.1.0/32 iddn 0 1.1.1.0 recv 606 1 ge-2/0/1.0
1.1.1.1/32 user 0 rjct 46 4
1.1.1.1/32 intf 0 1.1.1.1 locl 607 2
1.1.1.1/32 iddn 0 1.1.1.1 locl 607 2
1.1.1.255/32 iddn 0 ff:ff:ff:ff:ff:ff bcst 605 1 ge-2/0/1.0
10.0.0.0/24 intf 0 rslv 616 1 ge-2/0/0.0
10.0.0.0/32 dest 0 10.0.0.0 recv 614 1 ge-2/0/0.0
10.0.0.1/32 intf 0 10.0.0.1 locl 615 2
10.0.0.1/32 dest 0 10.0.0.1 locl 615 2
10.0.0.255/32 dest 0 10.0.0.255 bcst 613 1 ge-2/0/0.0
10.1.1.0/24 ifdn 0 rslv 612 1 ge-2/0/1.0
10.1.1.0/32 iddn 0 10.1.1.0 recv 610 1 ge-2/0/1.0
10.1.1.1/32 user 0 rjct 46 4
10.1.1.1/32 intf 0 10.1.1.1 locl 611 2
10.1.1.1/32 iddn 0 10.1.1.1 locl 611 2
10.1.1.255/32 iddn 0 ff:ff:ff:ff:ff:ff bcst 609 1 ge-2/0/1.0
10.209.0.0/16 user 0 10.209.63.254 ucst 419 20 fxp0.0
10.209.0.0/16 user 1 0:12:1e:ca:98:0 ucst 419 20 fxp0.0
10.209.0.0/18 intf 0 rslv 418 1 fxp0.0
10.209.0.0/32 dest 0 10.209.0.0 recv 416 1 fxp0.0
10.209.2.131/32 intf 0 10.209.2.131 locl 417 2
10.209.2.131/32 dest 0 10.209.2.131 locl 417 2
10.209.17.55/32 dest 0 0:30:48:5b:78:d2 ucst 435 1 fxp0.0
10.209.63.42/32 dest 0 0:23:7d:58:92:ca ucst 434 1 fxp0.0
10.209.63.254/32 dest 0 0:12:1e:ca:98:0 ucst 419 20 fxp0.0
10.209.63.255/32 dest 0 10.209.63.255 bcst 415 1 fxp0.0
10.227.0.0/16 user 0 10.209.63.254 ucst 419 20 fxp0.0

...

Routing table: iso
ISO:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 27 1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf 0 locl 28 1

Routing table: inet6
Internet6:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 6 1
ff00::/8 perm 0 mdsc 4 1
ff02::1/128 perm 0 ff02::1 mcst 3 1

Routing table: ccc
MPLS:
Interface.Label Type RtRef Next hop Type Index NhRef Netif

```

```

default perm 0 rjct 16 1
100004(top) fe-0/0/1.0

```

### show route forwarding-table detail

```

user@host> show route forwarding-table detail
Routing table: inet
Internet:
Destination Type RtRef Next hop Type Index NhRef Netif
default user 2 0:90:69:8e:b1:1b ucst 132 4 fxp0.0
default perm 0 rjct 14 1
10.1.1.0/24 intf 0 ff.3.0.21 ucst 322 1 so-5/3/0.0
10.1.1.0/32 dest 0 10.1.1.0 recv 324 1 so-5/3/0.0
10.1.1.1/32 intf 0 10.1.1.1 locl 321 1
10.1.1.255/32 dest 0 10.1.1.255 bcst 323 1 so-5/3/0.0
10.21.21.0/24 intf 0 ff.3.0.21 ucst 326 1 so-5/3/0.0
10.21.21.0/32 dest 0 10.21.21.0 recv 328 1 so-5/3/0.0
10.21.21.1/32 intf 0 10.21.21.1 locl 325 1
10.21.21.255/32 dest 0 10.21.21.255 bcst 327 1 so-5/3/0.0
127.0.0.1/32 intf 0 127.0.0.1 locl 320 1
172.17.28.19/32 clon 1 192.168.4.254 ucst 132 4 fxp0.0
172.17.28.44/32 clon 1 192.168.4.254 ucst 132 4 fxp0.0
...

Routing table: private1__inet
Internet:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 46 1
10.0.0.0/8 intf 0 rslv 136 1 fxp1.0
10.0.0.0/32 dest 0 10.0.0.0 recv 134 1 fxp1.0
10.0.0.4/32 intf 0 10.0.0.4 locl 135 2
10.0.0.4/32 dest 0 10.0.0.4 locl 135 2
...

Routing table: iso
ISO:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 38 1

Routing table: inet6
Internet6:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 22 1
ff00::/8 perm 0 mdsc 21 1
ff02::1/128 perm 0 ff02::1 mcst 17 1
...

Routing table: mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 28 1

```

### show route forwarding-table destination extensive (Weights and Balances)

```

user@host> show route forwarding-table destination 3.4.2.1 extensive
Routing table: inet [Index 0]
Internet:

```

```

Destination: 3.4.2.1/32
Route type: user
Route reference: 0
Flags: sent to PFE
Next-hop type: unicast
Nexthop: 4.4.4.4
Index: 262143 Reference: 1
Next-hop type: unicast
Next-hop interface: so-1/1/0.0
Index: 335 Reference: 2
Weight: 22 Balance: 3
Nexthop: 145.12.1.2
Next-hop type: unicast
Index: 337 Reference: 2
Next-hop interface: so-0/1/2.0
Weight: 33 Balance: 33

```

### show route forwarding-table extensive

```
user@host> show route forwarding-table extensive
```

```
Routing table: inet [Index 0]
```

```
Internet:
```

```

Destination: default
Route type: user
Route reference: 2
Flags: sent to PFE
Nexthop: 0:90:69:8e:b1:1b
Index: 132 Reference: 4
Next-hop type: unicast
Next-hop interface: fxp0.0

```

```

Destination: default
Route type: permanent
Route reference: 0
Flags: none
Next-hop type: reject
Index: 14 Reference: 1

```

```

Destination: 127.0.0.1/32
Route type: interface
Route reference: 0
Flags: sent to PFE
Nexthop: 127.0.0.1
Index: 320 Reference: 1
Next-hop type: local

```

```
...
```

```
Routing table: private1__inet [Index 1]
```

```
Internet:
```

```

Destination: default
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Index: 46 Reference: 1

```

```

Destination: 10.0.0.0/8
Route type: interface
Route reference: 0
Flags: sent to PFE
Next-hop type: resolve
Index: 136 Reference: 1
Next-hop interface: fxp1.0

```

```
...
```

```
Routing table: iso [Index 0]
```

```
ISO:
```

```
Destination: default
 Route type: permanent
 Route reference: 0
 Flags: sent to PFE
 Next-hop type: reject
 Route interface-index: 0
 Index: 38 Reference: 1

Routing table: inet6 [Index 0]
Internet6:

Destination: default
 Route type: permanent
 Route reference: 0
 Flags: sent to PFE
 Next-hop type: reject
 Route interface-index: 0
 Index: 22 Reference: 1

Destination: ff00::/8
 Route type: permanent
 Route reference: 0
 Flags: sent to PFE
 Next-hop type: multicast discard
 Route interface-index: 0
 Index: 21 Reference: 1

...

Routing table: private1__inet6 [Index 1]
Internet6:

Destination: default
 Route type: permanent
 Route reference: 0
 Flags: sent to PFE
 Next-hop type: reject
 Route interface-index: 0
 Index: 54 Reference: 1

Destination: fe80::2a0:a5ff:fe3d:375/128
 Route type: interface
 Route reference: 0
 Flags: sent to PFE
 Nexthop: fe80::2a0:a5ff:fe3d:375
 Next-hop type: local
 Route interface-index: 0
 Index: 75 Reference: 1

...
```

### show route forwarding-table extensive (RPF)

The next example is based on the following configuration, which enables an RPF check on all routes that are learned from this interface, including the interface route:

```
so-1/1/0 {
 unit 0 {
 family inet {
 rpf-check;
 address 15.95.1.2/30;
 }
 }
}

user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:
...
...
```

```

Destination: 15.95.1.3/32
Route type: destination
Route reference: 0
Flags: sent to PFE
Next-hop type: broadcast
Next-hop interface: so-1/1/0.0
RPF interface: so-1/1/0.0
Route interface-index: 67
Index: 328
Reference: 1

```

### show route forwarding-table family mpls

```

user@host> show route forwarding-table family mpls
Routing table: mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 rjct 19 1
0 user 0 recv 18 3
1 user 0 recv 18 3
2 user 0 recv 18 3
100000 user 0 10.31.1.6 swap 100001 fe-1/1/0.0
800002 user 0 Pop vt-0/3/0.32770

vt-0/3/0.32770 (VPLS)
 user 0 indr 351 4
 Push 800000, Push 100002(top)

so-0/0/0.0

```

### show route forwarding-table family vpls

```

user@host> show route forwarding-table family vpls
Routing table: green.vpls
VPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default dynm 0 flood 353 1
default perm 0 rjct 298 1
fe-0/1/0.0 dynm 0 flood 355 1
00:90:69:0c:20:1f/48 <<<<<Remote CE
 dynm 0 indr 351 4
 Push 800000, Push 100002(top)

so-0/0/0.0
00:90:69:85:b0:1f/48 <<<<<Local CE
 dynm 0 ucst 354 2 fe-0/1/0.0

```

### show route forwarding-table family vpls extensive

```

user@host> show route forwarding-table family vpls extensive
Routing table: green.vpls [Index 2]
VPLS:

Destination: default
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Next-hop type: unicast
Next-hop interface: fe-0/1/3.0
Next-hop type: unicast
Next-hop interface: fe-0/1/2.0
Route interface-index: 72
Index: 289
Index: 291
Index: 290
Reference: 1
Reference: 3
Reference: 3

Destination: default

```

```

Route type: permanent
Route reference: 0
Flags: none
Next-hop type: discard
Route interface-index: 0
Index: 341 Reference: 1

Destination: fe-0/1/2.0
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Next-hop type: indirect
Next-hop type: Push 800016
Next-hop interface: at-1/0/1.0
Next-hop type: indirect
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast
Next-hop interface: fe-0/1/3.0
Index: 293 Reference: 1
Index: 363 Reference: 4
Index: 301 Reference: 5
Index: 291 Reference: 3

Destination: fe-0/1/3.0
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Next-hop type: indirect
Next-hop type: Push 800016
Next-hop interface: at-1/0/1.0
Next-hop type: indirect
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast
Next-hop interface: fe-0/1/2.0
Index: 292 Reference: 1
Index: 363 Reference: 4
Index: 301 Reference: 5
Index: 290 Reference: 3

Destination: 10:00:00:01:01:01/48
Route type: dynamic
Route reference: 0
Flags: sent to PFE, prefix load balance
Next-hop type: unicast
Next-hop interface: fe-0/1/3.0
Route used as destination:
 Packet count: 6640 Byte count: 675786
Route used as source:
 Packet count: 6894 Byte count: 696424
Route interface-index: 70
Index: 291 Reference: 3

Destination: 10:00:00:01:01:04/48
Route type: dynamic
Route reference: 0
Flags: sent to PFE, prefix load balance
Next-hop type: unicast
Next-hop interface: fe-0/1/2.0
Route used as destination:
 Packet count: 96 Byte count: 8079
Route used as source:
 Packet count: 296 Byte count: 24955
Route interface-index: 69
Index: 290 Reference: 3

Destination: 10:00:00:01:03:05/48
Route type: dynamic
Route reference: 0
Flags: sent to PFE, prefix load balance
Route interface-index: 74

```



```

Next-hop type: indirect Index: 301 Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0

```

### show route forwarding-table table default

```
user@host> show route forwarding-table table default
```

```
Routing table: default.inet
```

```
Internet:
```

| Destination    | Type | RtRef | Next hop        | Type | Index | NhRef | Netif      |
|----------------|------|-------|-----------------|------|-------|-------|------------|
| default        | perm | 0     |                 | rjct | 36    | 2     |            |
| 0.0.0.0/32     | perm | 0     |                 | dscd | 34    | 1     |            |
| 10.0.60.0/30   | user | 0     | 10.0.60.13      | ucst | 713   | 5     | fe-0/1/3.0 |
| 10.0.60.12/30  | intf | 0     |                 | rslv | 688   | 1     | fe-0/1/3.0 |
| 10.0.60.12/32  | dest | 0     | 10.0.60.12      | recv | 686   | 1     | fe-0/1/3.0 |
| 10.0.60.13/32  | dest | 0     | 0:5:85:8b:bc:22 | ucst | 713   | 5     | fe-0/1/3.0 |
| 10.0.60.14/32  | intf | 0     | 10.0.60.14      | loc1 | 687   | 2     |            |
| 10.0.60.14/32  | dest | 0     | 10.0.60.14      | loc1 | 687   | 2     |            |
| 10.0.60.15/32  | dest | 0     | 10.0.60.15      | bcst | 685   | 1     | fe-0/1/3.0 |
| 10.0.67.12/30  | user | 0     | 10.0.60.13      | ucst | 713   | 5     | fe-0/1/3.0 |
| 10.0.80.0/30   | ifdn | 0     | ff.3.0.21       | ucst | 676   | 1     | so-0/0/1.0 |
| 10.0.80.0/32   | dest | 0     | 10.0.80.0       | recv | 678   | 1     | so-0/0/1.0 |
| 10.0.80.2/32   | user | 0     |                 | rjct | 36    | 2     |            |
| 10.0.80.2/32   | intf | 0     | 10.0.80.2       | loc1 | 675   | 1     |            |
| 10.0.80.3/32   | dest | 0     | 10.0.80.3       | bcst | 677   | 1     | so-0/0/1.0 |
| 10.0.90.12/30  | intf | 0     |                 | rslv | 684   | 1     | fe-0/1/0.0 |
| 10.0.90.12/32  | dest | 0     | 10.0.90.12      | recv | 682   | 1     | fe-0/1/0.0 |
| 10.0.90.14/32  | intf | 0     | 10.0.90.14      | loc1 | 683   | 2     |            |
| 10.0.90.14/32  | dest | 0     | 10.0.90.14      | loc1 | 683   | 2     |            |
| 10.0.90.15/32  | dest | 0     | 10.0.90.15      | bcst | 681   | 1     | fe-0/1/0.0 |
| 10.5.0.0/16    | user | 0     | 192.168.187.126 | ucst | 324   | 15    | fxp0.0     |
| 10.10.0.0/16   | user | 0     | 192.168.187.126 | ucst | 324   | 15    | fxp0.0     |
| 10.13.10.0/23  | user | 0     | 192.168.187.126 | ucst | 324   | 15    | fxp0.0     |
| 10.84.0.0/16   | user | 0     | 192.168.187.126 | ucst | 324   | 15    | fxp0.0     |
| 10.150.0.0/16  | user | 0     | 192.168.187.126 | ucst | 324   | 15    | fxp0.0     |
| 10.157.64.0/19 | user | 0     | 192.168.187.126 | ucst | 324   | 15    | fxp0.0     |
| 10.209.0.0/16  | user | 0     | 192.168.187.126 | ucst | 324   | 15    | fxp0.0     |

```
...
```

```
Routing table: default.iso
```

```
ISO:
```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default     | perm | 0     |          | rjct | 60    | 1     |       |

```
Routing table: default.inet6
```

```
Internet6:
```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default     | perm | 0     |          | rjct | 44    | 1     |       |
| ::/128      | perm | 0     |          | dscd | 42    | 1     |       |
| ff00::/8    | perm | 0     |          | mdsc | 43    | 1     |       |
| ff02::1/128 | perm | 0     | ff02::1  | mcst | 39    | 1     |       |

```
Routing table: default.mpls
```

```
MPLS:
```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default     | perm | 0     |          | dscd | 50    | 1     |       |

**show route forwarding-table table logical-system-name/routing-instance-name**

```
user@host> show route forwarding-table table R4/vpn-red
```

```
Logical system: R4
```

```
Routing table: vpn-red.inet
```

```
Internet:
```

| Destination        | Type | RtRef | Next hop                                       | Type | Index | NhRef | Netif      |
|--------------------|------|-------|------------------------------------------------|------|-------|-------|------------|
| default            | perm | 0     |                                                | rjct | 563   | 1     |            |
| 0.0.0.0/32         | perm | 0     |                                                | dscd | 561   | 2     |            |
| 1.0.0.1/32         | user | 0     |                                                | dscd | 561   | 2     |            |
| 2.0.2.0/24         | intf | 0     |                                                | rslv | 771   | 1     | ge-1/2/0.3 |
| 2.0.2.0/32         | dest | 0     | 2.0.2.0                                        | recv | 769   | 1     | ge-1/2/0.3 |
| 2.0.2.1/32         | intf | 0     | 2.0.2.1                                        | loc1 | 770   | 2     |            |
| 2.0.2.1/32         | dest | 0     | 2.0.2.1                                        | loc1 | 770   | 2     |            |
| 2.0.2.2/32         | dest | 0     | 0.4.80.3.0.1b.c0.d5.e4.bd.0.1b.c0.d5.e4.bc.8.0 | ucst | 789   | 1     | ge-1/2/0.3 |
| 2.0.2.255/32       | dest | 0     | 2.0.2.255                                      | bcst | 768   | 1     | ge-1/2/0.3 |
| 224.0.0.0/4        | perm | 1     |                                                | mdsc | 562   | 1     |            |
| 224.0.0.1/32       | perm | 0     | 224.0.0.1                                      | mcst | 558   | 1     |            |
| 255.255.255.255/32 | perm | 0     |                                                | bcst | 559   | 1     |            |

```
Logical system: R4
```

```
Routing table: vpn-red.iso
```

```
ISO:
```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default     | perm | 0     |          | rjct | 608   | 1     |       |

```
Logical system: R4
```

```
Routing table: vpn-red.inet6
```

```
Internet6:
```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default     | perm | 0     |          | rjct | 708   | 1     |       |
| ::/128      | perm | 0     |          | dscd | 706   | 1     |       |
| ff00::/8    | perm | 0     |          | mdsc | 707   | 1     |       |
| ff02::1/128 | perm | 0     | ff02::1  | mcst | 704   | 1     |       |

```
Logical system: R4
```

```
Routing table: vpn-red.mpls
```

```
MPLS:
```

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default     | perm | 0     |          | dscd | 638   |       |       |

**show route forwarding-table vpn**

```
user@host> show route forwarding-table vpn VPN-A
```

```
Routing table:: VPN-A.inet
```

```
Internet:
```

| Destination            | Type | RtRef | Next hop    | Type | Index  | NhRef | Netif |
|------------------------|------|-------|-------------|------|--------|-------|-------|
| default                | perm | 0     |             | rjct | 4      | 4     |       |
| 10.39.10.20/30         | intf | 0     | ff.3.0.21   | ucst | 40     | 1     |       |
| so-0/0/0.0             |      |       |             |      |        |       |       |
| 10.39.10.21/32         | intf | 0     | 10.39.10.21 | loc1 | 36     | 1     |       |
| 10.255.14.172/32       | user | 0     |             | ucst | 69     | 2     |       |
| so-0/0/0.0             |      |       |             |      |        |       |       |
| 10.255.14.175/32       | user | 0     |             | indr | 81     | 3     |       |
|                        |      |       |             | Push | 100004 |       | Push  |
| 100004(top) so-1/0/0.0 |      |       |             |      |        |       |       |
| 224.0.0.0/4            | perm | 2     |             | mdsc | 5      | 3     |       |
| 224.0.0.1/32           | perm | 0     | 224.0.0.1   | mcst | 1      | 8     |       |

|                    |      |   |           |      |   |   |
|--------------------|------|---|-----------|------|---|---|
| 224.0.0.5/32       | user | 1 | 224.0.0.5 | mcst | 1 | 8 |
| 255.255.255.255/32 | perm | 0 |           | bcst | 2 | 3 |

## show route hidden

---

|                                 |                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show route hidden</code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Display only hidden route information. A hidden route is unusable, even if it is the best path.                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b> .<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>    | <a href="#">show route hidden on page 3134</a><br><a href="#">show route hidden detail on page 3135</a><br><a href="#">show route hidden extensive on page 3135</a><br><a href="#">show route hidden terse on page 3135</a>                                                                                                              |
| <b>Output Fields</b>            | For information about output fields, see the output field table for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                            |

## Sample Output

### show route hidden

```
user@host> show route hidden
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
127.0.0.1/32 [Direct/0] 04:26:38
 > via lo0.0

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.5.5.5/32 [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
 AS path: 100 I
 Unusable
10.12.1.0/24 [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
 AS path: 100 I
 Unusable
10.12.80.4/30 [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
 AS path: I
```

Unusable

...

### show route hidden detail

```
user@host> show route hidden detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
127.0.0.1/32 (1 entry, 0 announced)
 Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 1
 Next hop: via lo0.0, selected
 State: <Hidden Martian Int>
 Local AS: 1
 Age: 4:27:37
 Task: IF
 AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.5.5.5/32 (1 entry, 0 announced)
 BGP Preference: 170/-101
 Route Distinguisher: 10.4.4.4:4
 Next hop type: Unusable
 Next-hop reference count: 6
 State: <Secondary Hidden Int Ext>
 Local AS: 1 Peer AS: 1
 Age: 3:45:09
 Task: BGP_1.10.4.4.4+2493
 AS path: 100 I
 Communities: target:1:999
 VPN Label: 100064
 Localpref: 100
 Router ID: 10.4.4.4
 Primary Routing Table bgp.13vpn.0

...
```

### show route hidden extensive

The output for the **show route hidden extensive** command is identical to that of the **show route hidden detail** command. For sample output, see [show route hidden detail on page 3135](#).

### show route hidden terse

```
user@host> show route hidden terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination P Prf Metric 1 Metric 2 Next hop AS path
127.0.0.1/32 D 0 >100.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
```

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)

Restart Complete

+ = Active Route, - = Last Active, \* = Both

| A Destination | P Prf | Metric 1 | Metric 2 | Next hop | AS path |
|---------------|-------|----------|----------|----------|---------|
| 10.5.5.5/32   | B 170 | 100      |          | Unusable | 100 I   |
| 10.12.1.0/24  | B 170 | 100      |          | Unusable | 100 I   |
| 10.12.80.4/30 | B 170 | 100      |          | Unusable | I       |

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)

Restart Complete

+ = Active Route, - = Last Active, \* = Both

| A Destination            | P Prf | Metric 1 | Metric 2 | Next hop | AS path |
|--------------------------|-------|----------|----------|----------|---------|
| 10.4.4.4:4:10.5.5.5/32   | B 170 | 100      |          | Unusable | 100 I   |
| 10.4.4.4:4:10.12.1.0/24  | B 170 | 100      |          | Unusable | 100 I   |
| 10.4.4.4:4:10.12.80.4/30 | B 170 | 100      |          | Unusable | I       |

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

Restart Complete

private1\_\_\_.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

## show route inactive-path

|                                    |                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | show route inactive-path<br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                     |
| <b>Syntax (EX Series Switches)</b> | show route inactive-path<br><brief   detail   extensive   terse>                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                        |
| <b>Description</b>                 | Display routes for destinations that have no active route. An inactive route is a route that was not selected as the best path.                                                                                                                                                                                                                                                              |
| <b>Options</b>                     | <p><b>none</b>—Display all inactive routes.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b>.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>       | <a href="#">show route inactive-path on page 3137</a><br><a href="#">show route inactive-path detail on page 3138</a><br><a href="#">show route inactive-path extensive on page 3139</a><br><a href="#">show route inactive-path terse on page 3139</a>                                                                                                                                      |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                                                                               |

## Sample Output

### show route inactive-path

```

user@host> show route inactive-path

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.100.12/30 [OSPF/10] 03:57:28, metric 1
> via so-0/3/0.0

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/8 [Direct/0] 04:39:56
> via fxp1.0

```

```
red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.80.0/30 [BGP/170] 04:38:17, localpref 100
 AS path: 100 I
 > to 10.12.80.1 via ge-6/3/2.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

#### show route inactive-path detail

```
user@host> show route inactive-path detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete

10.12.100.12/30 (2 entries, 1 announced)
 OSPF Preference: 10
 Next-hop reference count: 1
 Next hop: via so-0/3/0.0, selected
 State: <Int>
 Inactive reason: Route Preference
 Local AS: 1
 Age: 3:58:24 Metric: 1
 Area: 0.0.0.0
 Task: OSPF
 AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

10.0.0.0/8 (2 entries, 0 announced)
 Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 1
 Next hop: via fxp1.0, selected
 State: <NotBest Int>
 Inactive reason: No difference
 Age: 4:40:52
 Task: IF
 AS path: I

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.12.80.0/30 (2 entries, 1 announced)
 BGP Preference: 170/-101
 Next-hop reference count: 6
 Source: 10.12.80.1
```



```

Next hop: 10.12.80.1 via ge-6/3/2.0, selected
State: <Ext>
Inactive reason: Route Preference
Peer AS: 100
Age: 4:39:13
Task: BGP_100.10.12.80.1+179
AS path: 100 I
Localpref: 100
Router ID: 10.0.0.0

```

### show route inactive-path extensive

The output for the **show route inactive-path extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see [show route inactive-path detail on page 3138](#).

### show route inactive-path terse

```
user@host> show route inactive-path terse
```

```
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
```

```
Restart Complete
```

```
+ = Active Route, - = Last Active, * = Both
```

| A Destination   | P Prf | Metric 1 | Metric 2 | Next hop    | AS path |
|-----------------|-------|----------|----------|-------------|---------|
| 10.12.100.12/30 | 0 10  | 1        |          | >so-0/3/0.0 |         |

```
private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

| A Destination | P Prf | Metric 1 | Metric 2 | Next hop | AS path |
|---------------|-------|----------|----------|----------|---------|
| 10.0.0.0/8    | D 0   |          |          | >fxp1.0  |         |

```
red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
```

```
Restart Complete
```

```
+ = Active Route, - = Last Active, * = Both
```

| A Destination | P Prf | Metric 1 | Metric 2 | Next hop    | AS path |
|---------------|-------|----------|----------|-------------|---------|
| 10.12.80.0/30 | B 170 | 100      |          | >10.12.80.1 | 100 I   |

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

```
Restart Complete
```

```
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
```

```
Restart Complete
```

```
bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
```

```
Restart Complete
```

```
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

```
Restart Complete
```

```
private1__inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

## show route instance

|                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                     | show route instance<br><brief   detail   summary><br><instance-name><br><logical-system (all   <i>logical-system-name</i> )><br><operational>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax (EX Series Switches and QFX Series)</b> | show route instance<br><brief   detail   summary><br><instance-name><br><operational>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                                | Display routing instance information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                    | <p><b>none</b>—(Same as <b>brief</b>) Display standard information about all routing instances.</p> <p><b>brief   detail   summary</b>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b>. (These options are not available with the <b>operational</b> keyword.)</p> <p><b>instance-name</b>—(Optional) Display information for all routing instances whose name begins with this string (for example, <b>cust1</b>, <b>cust11</b>, and <b>cust111</b> are all displayed when you run the <b>show route instance cust1</b> command).</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>operational</b>—(Optional) Display operational routing instances.</p> |
| <b>Required Privilege Level</b>                   | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>                      | <a href="#">show route instance on page 3141</a><br><a href="#">show route instance detail (Graceful Restart Complete) on page 3142</a><br><a href="#">show route instance detail (Graceful Restart Incomplete) on page 3143</a><br><a href="#">show route instance detail (VPLS Routing Instance) on page 3145</a><br><a href="#">show route instance operational on page 3146</a><br><a href="#">show route instance summary on page 3146</a>                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>                              | Table 259 on page 3140 lists the output fields for the <b>show route instance</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

Table 259: show route instance Output Fields

| Field Name                       | Field Description             | Level of Output |
|----------------------------------|-------------------------------|-----------------|
| Instance or <i>instance-name</i> | Name of the routing instance. | All levels      |

Table 259: show route instance Output Fields (*continued*)

| Field Name                           | Field Description                                                                                                                                                                                                                                                  | Level of Output           |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------|
| <b>Operational Routing Instances</b> | ( <b>operational</b> keyword only) Names of all operational routing instances.                                                                                                                                                                                     | —                         |
| <b>Type</b>                          | Type of routing instance: <b>forwarding</b> , <b>l2vpn</b> , <b>no-forwarding</b> , <b>vpls</b> , <b>virtual-router</b> , or <b>vrf</b> .                                                                                                                          | All levels                |
| <b>State</b>                         | State of the routing instance: <b>active</b> or <b>inactive</b> .                                                                                                                                                                                                  | <b>brief detail none</b>  |
| <b>Interfaces</b>                    | Name of interfaces belonging to this routing instance.                                                                                                                                                                                                             | <b>brief detail none</b>  |
| <b>Restart State</b>                 | Status of graceful restart for this instance: <b>Pending</b> or <b>Complete</b> .                                                                                                                                                                                  | <b>detail</b>             |
| <b>Path selection timeout</b>        | Maximum amount of time, in seconds, remaining until graceful restart is declared complete. The default is <b>300</b> .                                                                                                                                             | <b>detail</b>             |
| <b>Tables</b>                        | Tables (and number of routes) associated with this routing instance.                                                                                                                                                                                               | <b>brief detail none</b>  |
| <b>Route-distinguisher</b>           | Unique route distinguisher associated with this routing instance.                                                                                                                                                                                                  | <b>detail</b>             |
| <b>Vrf-import</b>                    | VPN routing and forwarding instance import policy name.                                                                                                                                                                                                            | <b>detail</b>             |
| <b>Vrf-export</b>                    | VPN routing and forwarding instance export policy name.                                                                                                                                                                                                            | <b>detail</b>             |
| <b>Vrf-import-target</b>             | VPN routing and forwarding instance import target community name.                                                                                                                                                                                                  | <b>detail</b>             |
| <b>Vrf-export-target</b>             | VPN routing and forwarding instance export target community name.                                                                                                                                                                                                  | <b>detail</b>             |
| <b>Fast-reroute-priority</b>         | Fast reroute priority setting for a VPLS routing instance: <b>high</b> , <b>medium</b> , or <b>low</b> . The default is <b>low</b> .                                                                                                                               | <b>detail</b>             |
| <b>Restart State</b>                 | Restart state: <ul style="list-style-type: none"> <li><b>Pending:protocol-name</b>—List of protocols that have not yet completed graceful restart for this routing table.</li> <li><b>Complete</b>—All protocols have restarted for this routing table.</li> </ul> | <b>detail</b>             |
| <b>Primary rib</b>                   | Primary table for this routing instance.                                                                                                                                                                                                                           | <b>brief none summary</b> |
| <b>Active/holddown/hidden</b>        | Number of active, hold-down, and hidden routes.                                                                                                                                                                                                                    | All levels                |

## Sample Output

### show route instance

```

user@host> show route instance
Instance Type
Primary RIB
master forwarding
inet.0 16/0/1
iso.0 1/0/0

```

```

mpls.0 0/0/0
inet6.0 2/0/0
l2circuit.0 0/0/0
__juniper_private1__ forwarding
__juniper_private1__.inet.0 12/0/0
__juniper_private1__.inet6.0 1/0/0

```

### show route instance detail (Graceful Restart Complete)

```

user@host> show route instance detail
master:
 Router ID: 10.255.14.176
 Type: forwarding State: Active
 Restart State: Complete Path selection timeout: 300
 Tables:
 inet.0 : 17 routes (15 active, 0 holddown, 1 hidden)
 Restart Complete
 inet.3 : 2 routes (2 active, 0 holddown, 0 hidden)
 Restart Complete
 iso.0 : 1 routes (1 active, 0 holddown, 0 hidden)
 Restart Complete
 mpls.0 : 19 routes (19 active, 0 holddown, 0 hidden)
 Restart Complete
 bgp.l3vpn.0 : 10 routes (10 active, 0 holddown, 0 hidden)
 Restart Complete
 inet6.0 : 2 routes (2 active, 0 holddown, 0 hidden)
 Restart Complete
 bgp.l2vpn.0 : 1 routes (1 active, 0 holddown, 0 hidden)
 Restart Complete
 BGP-INET:
 Router ID: 10.69.103.1
 Type: vrf State: Active
 Restart State: Complete Path selection timeout: 300
 Interfaces:
 t3-0/0/0.103
 Route-distinguisher: 10.255.14.176:103
 Vrf-import: [BGP-INET-import]
 Vrf-export: [BGP-INET-export]
 Tables:
 BGP-INET.inet.0 : 4 routes (4 active, 0 holddown, 0 hidden)
 Restart Complete
 BGP-L:
 Router ID: 10.69.104.1
 Type: vrf State: Active
 Restart State: Complete Path selection timeout: 300
 Interfaces:
 t3-0/0/0.104
 Route-distinguisher: 10.255.14.176:104
 Vrf-import: [BGP-L-import]
 Vrf-export: [BGP-L-export]
 Tables:
 BGP-L.inet.0 : 4 routes (4 active, 0 holddown, 0 hidden)
 Restart Complete
 BGP-L.mpls.0 : 3 routes (3 active, 0 holddown, 0 hidden)
 Restart Complete
 L2VPN:
 Router ID: 0.0.0.0
 Type: l2vpn State: Active
 Restart State: Complete Path selection timeout: 300
 Interfaces:
 t3-0/0/0.512

```

```

Route-distinguisher: 10.255.14.176:512
Vrf-import: [L2VPN-import]
Vrf-export: [L2VPN-export]
Tables:
 L2VPN.l2vpn.0 : 2 routes (2 active, 0 holddown, 0 hidden)
 Restart Complete
LDP:
 Router ID: 10.69.105.1
 Type: vrf State: Active
 Restart State: Complete Path selection timeout: 300
 Interfaces:
 t3-0/0/0.105
 Route-distinguisher: 10.255.14.176:105
 Vrf-import: [LDP-import]
 Vrf-export: [LDP-export]
 Tables:
 LDP.inet.0 : 5 routes (4 active, 0 holddown, 0 hidden)
 Restart Complete
OSPF:
 Router ID: 10.69.101.1
 Type: vrf State: Active
 Restart State: Complete Path selection timeout: 300
 Interfaces:
 t3-0/0/0.101
 Route-distinguisher: 10.255.14.176:101
 Vrf-import: [OSPF-import]
 Vrf-export: [OSPF-export]
 Vrf-import-target: [target:11111
 Tables:
 OSPF.inet.0 : 8 routes (7 active, 0 holddown, 0 hidden)
 Restart Complete
RIP:
 Router ID: 10.69.102.1
 Type: vrf State: Active
 Restart State: Complete Path selection timeout: 300
 Interfaces:
 t3-0/0/0.102
 Route-distinguisher: 10.255.14.176:102
 Vrf-import: [RIP-import]
 Vrf-export: [RIP-export]
 Tables:
 RIP.inet.0 : 6 routes (6 active, 0 holddown, 0 hidden)
 Restart Complete
STATIC:
 Router ID: 10.69.100.1
 Type: vrf State: Active
 Restart State: Complete Path selection timeout: 300
 Interfaces:
 t3-0/0/0.100
 Route-distinguisher: 10.255.14.176:100
 Vrf-import: [STATIC-import]
 Vrf-export: [STATIC-export]
 Tables:
 STATIC.inet.0 : 4 routes (4 active, 0 holddown, 0 hidden)
 Restart Complete

```

#### show route instance detail (Graceful Restart Incomplete)

```

user@host> show route instance detail
master:
 Router ID: 10.255.14.176

```

```

Type: forwarding State: Active
Restart State: Pending Path selection timeout: 300
Tables:
 inet.0 : 17 routes (15 active, 1 holddown, 1 hidden)
 Restart Pending: OSPF LDP
 inet.3 : 2 routes (2 active, 0 holddown, 0 hidden)
 Restart Pending: OSPF LDP
 iso.0 : 1 routes (1 active, 0 holddown, 0 hidden)
 Restart Complete
 mpls.0 : 23 routes (23 active, 0 holddown, 0 hidden)
 Restart Pending: LDP VPN
 bgp.l3vpn.0 : 10 routes (10 active, 0 holddown, 0 hidden)
 Restart Pending: BGP VPN
 inet6.0 : 2 routes (2 active, 0 holddown, 0 hidden)
 Restart Complete
 bgp.l2vpn.0 : 1 routes (1 active, 0 holddown, 0 hidden)
 Restart Pending: BGP VPN
BGP-INET:
 Router ID: 10.69.103.1
 Type: vrf State: Active
 Restart State: Pending Path selection timeout: 300
 Interfaces:
 t3-0/0/0.103
 Route-distinguisher: 10.255.14.176:103
 Vrf-import: [BGP-INET-import]
 Vrf-export: [BGP-INET-export]
 Tables:
 BGP-INET.inet.0 : 6 routes (5 active, 0 holddown, 0 hidden)
 Restart Pending: VPN
BGP-L:
 Router ID: 10.69.104.1
 Type: vrf State: Active
 Restart State: Pending Path selection timeout: 300
 Interfaces:
 t3-0/0/0.104
 Route-distinguisher: 10.255.14.176:104
 Vrf-import: [BGP-L-import]
 Vrf-export: [BGP-L-export]
 Tables:
 BGP-L.inet.0 : 6 routes (5 active, 0 holddown, 0 hidden)
 Restart Pending: VPN
 BGP-L.mpls.0 : 2 routes (2 active, 0 holddown, 0 hidden)
 Restart Pending: VPN
L2VPN:
 Router ID: 0.0.0.0
 Type: l2vpn State: Active
 Restart State: Pending Path selection timeout: 300
 Interfaces:
 t3-0/0/0.512
 Route-distinguisher: 10.255.14.176:512
 Vrf-import: [L2VPN-import]
 Vrf-export: [L2VPN-export]
 Tables:
 L2VPN.l2vpn.0 : 2 routes (2 active, 0 holddown, 0 hidden)
 Restart Pending: VPN L2VPN
LDP:
 Router ID: 10.69.105.1
 Type: vrf State: Active
 Restart State: Pending Path selection timeout: 300
 Interfaces:
 t3-0/0/0.105

```

```

Route-distinguisher: 10.255.14.176:105
Vrf-import: [LDP-import]
Vrf-export: [LDP-export]
Tables:
 LDP.inet.0 : 5 routes (4 active, 1 holddown, 0 hidden)
Restart Pending: OSPF LDP VPN
OSPF:
Router ID: 10.69.101.1
Type: vrf State: Active
Restart State: Pending Path selection timeout: 300
Interfaces:
 t3-0/0/0.101
Route-distinguisher: 10.255.14.176:101
Vrf-import: [OSPF-import]
Vrf-export: [OSPF-export]
Tables:
 OSPF.inet.0 : 8 routes (7 active, 1 holddown, 0 hidden)
Restart Pending: OSPF VPN
RIP:
Router ID: 10.69.102.1
Type: vrf State: Active
Restart State: Pending Path selection timeout: 300
Interfaces:
 t3-0/0/0.102
Route-distinguisher: 10.255.14.176:102
Vrf-import: [RIP-import]
Vrf-export: [RIP-export]
Tables:
 RIP.inet.0 : 8 routes (6 active, 2 holddown, 0 hidden)
Restart Pending: RIP VPN
STATIC:
Router ID: 10.69.100.1
Type: vrf State: Active
Restart State: Pending Path selection timeout: 300
Interfaces:
 t3-0/0/0.100
Route-distinguisher: 10.255.14.176:100
Vrf-import: [STATIC-import]
Vrf-export: [STATIC-export]
Tables:
 STATIC.inet.0 : 4 routes (4 active, 0 holddown, 0 hidden)
Restart Pending: VPN

```

### show route instance detail (VPLS Routing Instance)

```

user@host> show route instance detail test-vpls
test-vpls:
Router ID: 0.0.0.0
Type: vpls State: Active
Interfaces:
 lsi.1048833
 lsi.1048832
 fe-0/1/0.513
Route-distinguisher: 10.255.37.65:1
Vrf-import: [__vrf-import-test-vpls-internal__]
Vrf-export: [__vrf-export-test-vpls-internal__]
Vrf-import-target: [target:300:1]
Vrf-export-target: [target:300:1]
Fast-reroute-priority: high

```

## Tables:

test-vpls.l2vpn.0 : 3 routes (3 active, 0 holddown, 0 hidden)

**show route instance operational**

```
user@host> show route instance operational
Operational Routing Instances:
```

```
master
default
```

**show route instance summary**

```
user@host> show route instance summary
```

| Instance | Type       | Primary rib      | Active/holddown/hidden |
|----------|------------|------------------|------------------------|
| master   | forwarding |                  |                        |
|          |            | inet.0           | 15/0/1                 |
|          |            | iso.0            | 1/0/0                  |
|          |            | mpls.0           | 35/0/0                 |
|          |            | l3vpn.0          | 0/0/0                  |
|          |            | inet6.0          | 2/0/0                  |
|          |            | l2vpn.0          | 0/0/0                  |
|          |            | l2circuit.0      | 0/0/0                  |
| BGP-INET | vrf        |                  |                        |
|          |            | BGP-INET.inet.0  | 5/0/0                  |
|          |            | BGP-INET.iso.0   | 0/0/0                  |
|          |            | BGP-INET.inet6.0 | 0/0/0                  |
| BGP-L    | vrf        |                  |                        |
|          |            | BGP-L.inet.0     | 5/0/0                  |
|          |            | BGP-L.iso.0      | 0/0/0                  |
|          |            | BGP-L.mpls.0     | 4/0/0                  |
|          |            | BGP-L.inet6.0    | 0/0/0                  |
| L2VPN    | l2vpn      |                  |                        |
|          |            | L2VPN.inet.0     | 0/0/0                  |
|          |            | L2VPN.iso.0      | 0/0/0                  |
|          |            | L2VPN.inet6.0    | 0/0/0                  |
|          |            | L2VPN.l2vpn.0    | 2/0/0                  |
| LDP      | vrf        |                  |                        |
|          |            | LDP.inet.0       | 4/0/0                  |
|          |            | LDP.iso.0        | 0/0/0                  |
|          |            | LDP.mpls.0       | 0/0/0                  |
|          |            | LDP.inet6.0      | 0/0/0                  |
|          |            | LDP.l2circuit.0  | 0/0/0                  |
| OSPF     | vrf        |                  |                        |
|          |            | OSPF.inet.0      | 7/0/0                  |
|          |            | OSPF.iso.0       | 0/0/0                  |
|          |            | OSPF.inet6.0     | 0/0/0                  |
| RIP      | vrf        |                  |                        |
|          |            | RIP.inet.0       | 6/0/0                  |
|          |            | RIP.iso.0        | 0/0/0                  |
|          |            | RIP.inet6.0      | 0/0/0                  |
| STATIC   | vrf        |                  |                        |
|          |            | STATIC.inet.0    | 4/0/0                  |
|          |            | STATIC.iso.0     | 0/0/0                  |
|          |            | STATIC.inet6.0   | 0/0/0                  |



## show route next-hop

|                                    |                                                                                                                                                                                                                                                                                                            |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | <code>show route next-hop <i>next-hop</i></code><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                           |
| <b>Syntax (EX Series Switches)</b> | <code>show route next-hop <i>next-hop</i></code><br><brief   detail   extensive   terse>                                                                                                                                                                                                                   |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                      |
| <b>Description</b>                 | Display the entries in the routing table that are being sent to the specified next-hop address.                                                                                                                                                                                                            |
| <b>Options</b>                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b><i>next-hop</i></b> —Next-hop address. |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>       | <a href="#">show route next-hop on page 3147</a><br><a href="#">show route next-hop detail on page 3148</a><br><a href="#">show route next-hop extensive on page 3150</a><br><a href="#">show route next-hop terse on page 3151</a>                                                                        |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                             |

## Sample Output

### show route next-hop

```

user@host> show route next-hop 192.168.71.254

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16 *[Static/5] 06:26:25
 > to 192.168.71.254 via fxp0.0
10.209.0.0/16 *[Static/5] 06:26:25
 > to 192.168.71.254 via fxp0.0
172.16.0.0/12 *[Static/5] 06:26:25
 > to 192.168.71.254 via fxp0.0
192.168.0.0/16 *[Static/5] 06:26:25
 > to 192.168.71.254 via fxp0.0
192.168.102.0/23 *[Static/5] 06:26:25
 > to 192.168.71.254 via fxp0.0

```

```
207.17.136.0/24 *[Static/5] 06:26:25
 > to 192.168.71.254 via fxp0.0
207.17.136.192/32 *[Static/5] 06:26:25
 > to 192.168.71.254 via fxp0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

#### show route next-hop detail

```
user@host> show route next-hop 192.168.71.254 detail

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
10.10.0.0/16 (1 entry, 1 announced)
 *Static Preference: 5
 Next-hop reference count: 36
 Next hop: 192.168.71.254 via fxp0.0, selected
 State: <Active NoReadvrt Int Ext>
 Local AS: 1
 Age: 6:27:41
 Task: RT
 Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
 AS path: I

10.209.0.0/16 (1 entry, 1 announced)
 *Static Preference: 5
 Next-hop reference count: 36
 Next hop: 192.168.71.254 via fxp0.0, selected
 State: <Active NoReadvrt Int Ext>
 Local AS: 1
 Age: 6:27:41
 Task: RT
 Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
 AS path: I

172.16.0.0/12 (1 entry, 1 announced)
 *Static Preference: 5
 Next-hop reference count: 36
 Next hop: 192.168.71.254 via fxp0.0, selected
 State: <Active NoReadvrt Int Ext>
 Local AS: 1
 Age: 6:27:41
 Task: RT
 Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
 AS path: I

192.168.0.0/16 (1 entry, 1 announced)
```

```

*Static Preference: 5
 Next-hop reference count: 36
 Next hop: 192.168.71.254 via fxp0.0, selected
 State: <Active NoReadvrt Int Ext>
 Local AS: 1
 Age: 6:27:41
 Task: RT
 Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
 AS path: I

192.168.102.0/23 (1 entry, 1 announced)
 *Static Preference: 5
 Next-hop reference count: 36
 Next hop: 192.168.71.254 via fxp0.0, selected
 State: <Active NoReadvrt Int Ext>
 Local AS: 1
 Age: 6:27:41
 Task: RT
 Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
 AS path: I

207.17.136.0/24 (1 entry, 1 announced)
 *Static Preference: 5
 Next-hop reference count: 36
 Next hop: 192.168.71.254 via fxp0.0, selected
 State: <Active NoReadvrt Int Ext>
 Local AS: 1
 Age: 6:27:41
 Task: RT
 Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
 AS path: I

207.17.136.192/32 (1 entry, 1 announced)
 *Static Preference: 5
 Next-hop reference count: 36
 Next hop: 192.168.71.254 via fxp0.0, selected
 State: <Active NoReadvrt Int Ext>
 Local AS: 1
 Age: 6:27:41
 Task: RT
 Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
 AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

### show route next-hop extensive

```
user@host> show route next-hop 192.168.71.254 extensive
```

```
inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
```

```
10.10.0.0/16 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
10.209.0.0/16 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
172.16.0.0/12 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
192.168.0.0/16 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 192.168.0.0/16 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
192.168.102.0/23 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 192.168.102.0/23 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```

Next-hop reference count: 22
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2:02:28
Task: RT
Announcement bits (1): 0-KRT
AS path: I

207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
 *Static Preference: 5
 Next-hop reference count: 22
 Next hop: 192.168.71.254 via fxp0.0, selected
 State: <Active NoReadvrt Int Ext>
 Local AS: 69
 Age: 2:02:28
 Task: RT
 Announcement bits (1): 0-KRT
 AS path: I

207.17.136.192/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.192/32 -> {192.168.71.254}
 *Static Preference: 5
 Next-hop reference count: 22
 Next hop: 192.168.71.254 via fxp0.0, selected
 State: <Active NoReadvrt Int Ext>
 Local AS: 69
 Age: 2:02:28
 Task: RT
 Announcement bits (1): 0-KRT
 AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

### show route next-hop terse

```

user@host> show route next-hop 192.168.71.254 terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination P Prf Metric 1 Metric 2 Next hop AS path
* 10.10.0.0/16 S 5 0 0 >192.168.71.254
* 10.209.0.0/16 S 5 0 0 >192.168.71.254
* 172.16.0.0/12 S 5 0 0 >192.168.71.254

```

```
* 192.168.0.0/16 S 5 >192.168.71.254
* 192.168.102.0/23 S 5 >192.168.71.254
* 207.17.136.0/24 S 5 >192.168.71.254
* 207.17.136.192/32 S 5 >192.168.71.254

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

## show route output

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | show route output (address <i>ip-address</i>   interface <i>interface-name</i> )<br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax (EX Series Switches)</b> | show route output (address <i>ip-address</i>   interface <i>interface-name</i> )<br><brief   detail   extensive   terse>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>                 | <p>Display the entries in the routing table learned through static routes and interior gateway protocols that are to be sent out the interface with either the specified IP address or specified name.</p> <p>To view routes advertised to a neighbor or received from a neighbor for the BGP protocol, use the <b>show route advertising-protocol bgp</b> and <b>show route receive-protocol bgp</b> commands instead.</p>                                                                                                                                                                                                                      |
| <b>Options</b>                     | <p><b>address <i>ip-address</i></b>—Display entries in the routing table that are to be sent out the interface with the specified IP address.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b>.</p> <p><b>interface <i>interface-name</i></b>—Display entries in the routing table that are to be sent out the interface with the specified name.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>       | <a href="#">show route output address on page 3154</a><br><a href="#">show route output address detail on page 3154</a><br><a href="#">show route output address extensive on page 3155</a><br><a href="#">show route output address terse on page 3155</a><br><a href="#">show route output interface on page 3155</a><br><a href="#">show route output interface detail on page 3156</a><br><a href="#">show route output interface extensive on page 3156</a><br><a href="#">show route output interface terse on page 3156</a>                                                                                                               |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                                                                                                                                                                                                                                                                                                                                   |

## Sample Output

### show route output address

```
user@host> show route output address 36.1.1.1/24

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

36.1.1.0/24 *[Direct/0] 00:19:56
 > via so-0/1/2.0
 [OSPF/10] 00:19:55, metric 1
 > via so-0/1/2.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

### show route output address detail

```
user@host> show route output address 36.1.1.1 detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
36.1.1.0/24 (2 entries, 0 announced)
 *Direct Preference: 0
 Next hop type: Interface
 Next-hop reference count: 1
 Next hop: via so-0/1/2.0, selected
 State: <Active Int>
 Age: 23:00
 Task: IF
 AS path: I
 OSPF Preference: 10
 Next-hop reference count: 1
 Next hop: via so-0/1/2.0, selected
 State: <Int>
 Inactive reason: Route Preference
 Age: 22:59 Metric: 1
 Area: 0.0.0.0
 Task: OSPF
 AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```



### show route output address extensive

The output for the **show route output address extensive** command is identical to that of the **show route output address detail** command. For sample output, see [show route output address detail on page 3154](#).

### show route output address terse

```
user@host> show route output address 36.1.1.1 terse

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination P Prf Metric 1 Metric 2 Next hop AS path
* 36.1.1.0/24 D 0 >so-0/1/2.0
 O 10 1 >so-0/1/2.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

### show route output interface

```
user@host> show route output interface so-0/1/2.0

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.71.240/32 * [OSPF/10] 00:13:00, metric 2
 via so-0/1/2.0
 > via so-0/3/2.0
10.255.71.241/32 * [OSPF/10] 00:13:10, metric 1
 > via so-0/1/2.0
14.1.1.0/24 * [OSPF/10] 00:05:11, metric 3
 to 35.1.1.2 via ge-3/1/0.0
 > via so-0/1/2.0
 via so-0/3/2.0
16.1.1.0/24 * [OSPF/10] 00:13:10, metric 2
 > via so-0/1/2.0
36.1.1.0/24 * [Direct/0] 00:13:21
 > via so-0/1/2.0
 [OSPF/10] 00:13:20, metric 1
 > via so-0/1/2.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

```
private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

### show route output interface detail

```
user@host> show route output interface so-0/1/2.0 detail
```

```
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.255.71.240/32 (1 entry, 1 announced)
```

```
*OSPF Preference: 10
 Next-hop reference count: 2
 Next hop: via so-0/1/2.0
 Next hop: via so-0/3/2.0, selected
 State: <Active Int>
 Age: 14:52 Metric: 2
 Area: 0.0.0.0
 Task: OSPF
 Announcement bits (1): 0-KRT
 AS path: I
```

```
10.255.71.241/32 (1 entry, 1 announced)
```

```
*OSPF Preference: 10
 Next-hop reference count: 4
 Next hop: via so-0/1/2.0, selected
 State: <Active Int>
 Age: 15:02 Metric: 1
 Area: 0.0.0.0
 Task: OSPF
 Announcement bits (1): 0-KRT
 AS path: I
```

```
...
```

### show route output interface extensive

The output for the **show route output interface extensive** command is identical to that of the **show route output interface detail** command. For sample output, see [show route output interface detail on page 3156](#).

### show route output interface terse

```
user@host> show route output interface so-0/1/2.0 terse
```

```
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

| A | Destination      | P | Prf | Metric 1 | Metric 2 | Next hop    | AS path |
|---|------------------|---|-----|----------|----------|-------------|---------|
| * | 10.255.71.240/32 | 0 | 10  | 2        |          | so-0/1/2.0  |         |
|   |                  |   |     |          |          | >so-0/3/2.0 |         |
| * | 10.255.71.241/32 | 0 | 10  | 1        |          | >so-0/1/2.0 |         |
| * | 14.1.1.0/24      | 0 | 10  | 3        |          | 35.1.1.2    |         |
|   |                  |   |     |          |          | >so-0/1/2.0 |         |
|   |                  |   |     |          |          | so-0/3/2.0  |         |
| * | 16.1.1.0/24      | 0 | 10  | 2        |          | >so-0/1/2.0 |         |
| * | 36.1.1.0/24      | D | 0   |          |          | >so-0/1/2.0 |         |
|   |                  | 0 | 10  | 1        |          | >so-0/1/2.0 |         |

```
private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

```
mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1\_\_inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

## show route protocol

---

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | <code>show route protocol <i>protocol</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Syntax (EX Series Switches)</b> | <code>show route protocol <i>protocol</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>ospf2</b> and <b>ospf3</b> options introduced in Junos OS Release 9.2.<br><b>ospf2</b> and <b>ospf3</b> options introduced in Junos OS Release 9.2 for EX Series switches.<br><b>flow</b> option introduced in Junos OS Release 10.0.<br><b>flow</b> option introduced in Junos OS Release 10.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                 | Display the route entries in the routing table that were learned from a particular protocol.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b> .<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b><i>protocol</i></b> —Protocol from which the route was learned: <ul style="list-style-type: none"><li>• <b>access</b>—Access route for use by DHCP application</li><li>• <b>access-internal</b>—Access-internal route for use by DHCP application</li><li>• <b>aggregate</b>—Locally generated aggregate route</li><li>• <b>arp</b>—Route learned through the Address Resolution Protocol</li><li>• <b>atmvpn</b>—Asynchronous Transfer Mode virtual private network</li><li>• <b>bgp</b>—Border Gateway Protocol</li><li>• <b>ccc</b>—Circuit cross-connect</li><li>• <b>direct</b>—Directly connected route</li><li>• <b>dvmrp</b>—Distance Vector Multicast Routing Protocol</li><li>• <b>esis</b>—End System-to-Intermediate System</li><li>• <b>flow</b>—Locally defined flow-specification route</li><li>• <b>frr</b>—Precomputed protection route or backup route used when a link goes down</li><li>• <b>isis</b>—Intermediate System-to-Intermediate System</li><li>• <b>ldp</b>—Label Distribution Protocol</li><li>• <b>l2circuit</b>—Layer 2 circuit</li><li>• <b>l2vpn</b>—Layer 2 virtual private network</li></ul> |

- **local**—Local address
- **mpls**—Multiprotocol Label Switching
- **msdp**—Multicast Source Discovery Protocol
- **ospf**—Open Shortest Path First versions 2 and 3
- **ospf2**—Open Shortest Path First versions 2 only
- **ospf3**—Open Shortest Path First version 3 only
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation
- **rsvp**—Resource Reservation Protocol
- **rtarget**—Local route target virtual private network
- **static**—Statically defined route
- **tunnel**—Dynamic tunnel
- **vpn**—Virtual private network



**NOTE:** EX Series switches run a subset of these protocols. See the switch CLI for details.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>    | <a href="#">show route protocol access on page 3160</a><br><a href="#">show route protocol access-internal extensive on page 3160</a><br><a href="#">show route protocol arp on page 3160</a><br><a href="#">show route protocol bgp on page 3161</a><br><a href="#">show route protocol bgp detail on page 3161</a><br><a href="#">show route protocol bgp extensive on page 3161</a><br><a href="#">show route protocol bgp terse on page 3162</a><br><a href="#">show route protocol direct on page 3162</a><br><a href="#">show route protocol frr on page 3163</a><br><a href="#">show route protocol l2circuit detail on page 3163</a><br><a href="#">show route protocol l2vpn extensive on page 3164</a><br><a href="#">show route protocol ldp on page 3165</a><br><a href="#">show route protocol ldp extensive on page 3165</a><br><a href="#">show route protocol ospf (Layer 3 VPN) on page 3166</a><br><a href="#">show route protocol ospf detail on page 3167</a><br><a href="#">show route protocol rip on page 3167</a><br><a href="#">show route protocol rip detail on page 3167</a><br><a href="#">show route protocol ripng table inet6 on page 3168</a><br><a href="#">show route protocol static detail on page 3168</a> |

**Output Fields** For information about output fields, see the output field tables for the [show route](#) command, the [show route detail](#) command, the [show route extensive](#) command, or the [show route terse](#) command.

## Sample Output

### show route protocol access

```
user@host> show route protocol access
inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32 *[Access/13] 00:00:09
 > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32 *[Access/13] 00:00:09
 > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32 *[Access/13] 00:00:09
 > to 13.160.0.2 via fe-0/0/0.0
```

### show route protocol access-internal extensive

```
user@host> show route protocol access-internal 13.160.0.19 extensive
inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.19/32 -> {13.160.0.2}
 *Access-internal Preference: 12
 Next-hop reference count: 200000
 Next hop: 13.160.0.2 via fe-0/0/0.0, selected
 State: <Active Int>
 Age: 36
 Task: RPD Unix Domain Server./var/run/rpd_serv.local
 Announcement bits (1): 0-KRT
 AS path: I
```

### show route protocol arp

```
user@host> show route protocol arp
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32 [ARP/4294967293] 00:04:35, from 20.20.1.1
 Unusable
20.20.1.4/32 [ARP/4294967293] 00:04:35, from 20.20.1.1
 Unusable
20.20.1.5/32 [ARP/4294967293] 00:04:32, from 20.20.1.1
 Unusable
20.20.1.6/32 [ARP/4294967293] 00:04:34, from 20.20.1.1
 Unusable
20.20.1.7/32 [ARP/4294967293] 00:04:35, from 20.20.1.1
 Unusable
20.20.1.8/32 [ARP/4294967293] 00:04:35, from 20.20.1.1
 Unusable
20.20.1.9/32 [ARP/4294967293] 00:04:35, from 20.20.1.1
 Unusable
20.20.1.10/32 [ARP/4294967293] 00:04:35, from 20.20.1.1
```

```

Unusable
20.20.1.11/32 [ARP/4294967293] 00:04:33, from 20.20.1.1
Unusable
20.20.1.12/32 [ARP/4294967293] 00:04:33, from 20.20.1.1
Unusable
20.20.1.13/32 [ARP/4294967293] 00:04:33, from 20.20.1.1
Unusable
...

```

### show route protocol bgp

```

user@host> show route protocol bgp 192.168.64.0/21
inet.0: 335832 destinations, 335833 routes (335383 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21 *[BGP/170] 6d 10:41:16, localpref 100, from 192.168.69.71
AS path: 10458 14203 2914 4788 4788 I
> to 192.168.167.254 via fxp0.0

```

### show route protocol bgp detail

```

user@host> show route protocol bgp 66.117.63.0/24 detail
inet.0: 335805 destinations, 335806 routes (335356 active, 0 holddown, 450 hidden)
66.117.63.0/24 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Next hop type: Indirect
 Next-hop reference count: 1006436
 Source: 192.168.69.71
 Next hop type: Router, Next hop index: 324
 Next hop: 192.168.167.254 via fxp0.0, selected
 Protocol next hop: 192.168.69.71
 Indirect next hop: 8e166c0 342
 State: <Active Ext>
 Local AS: 69 Peer AS: 10458
 Age: 6d 10:42:42 Metric2: 0
 Task: BGP_10458.192.168.69.71+179
 Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree

1
 AS path: 10458 14203 2914 4788 4788 I
 Communities: 2914:410 2914:2403 2914:3400
 Accepted
 Localpref: 100
 Router ID: 207.17.136.192

```

### show route protocol bgp extensive

```

user@host> show route protocol bgp 192.168.64.0/21 extensive

inet.0: 335827 destinations, 335828 routes (335378 active, 0 holddown, 450 hidden)
192.168.64.0/21 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.9.0.0/16 -> {indirect(342)}
Page 0 idx 1 Type 1 val db31a80
 Nexthop: Self
 AS path: [69] 10458 14203 2914 4788 4788 I
 Communities: 2914:410 2914:2403 2914:3400
Path 1.9.0.0 from 192.168.69.71 Vector len 4. Val: 1
 *BGP Preference: 170/-101
 Next hop type: Indirect
 Next-hop reference count: 1006502
 Source: 192.168.69.71
 Next hop type: Router, Next hop index: 324

```

```

Next hop: 192.168.167.254 via fxp0.0, selected
Protocol next hop: 192.168.69.71
Indirect next hop: 8e166c0 342
State: <Active Ext>
Local AS: 69 Peer AS: 10458
Age: 6d 10:44:45 Metric2: 0
Task: BGP_10458.192.168.69.71+179
Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree

1
AS path: 10458 14203 2914 4788 4788 I
Communities: 2914:410 2914:2403 2914:3400
Accepted
Localpref: 100
Router ID: 207.17.136.192
Indirect next hops: 1
 Protocol next hop: 192.168.69.71
 Indirect next hop: 8e166c0 342
 Indirect path forwarding next hops: 1
 Next hop type: Router
 Next hop: 192.168.167.254 via fxp0.0
 192.168.0.0/16 Originating RIB: inet.0
 Node path count: 1
 Forwarding nexthops: 1
 Nexthop: 192.168.167.254 via fxp0.0

```

### show route protocol bgp terse

```

user@host> show route protocol bgp 192.168.64.0/21 terse

inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination P Prf Metric 1 Metric 2 Next hop AS path
192.168.64.0/21 B 170 100 >100.1.3.2 10023 21 I

```

### show route protocol direct

```

user@host> show route protocol direct

inet.0: 335843 destinations, 335844 routes (335394 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

8.8.8.0/24 *[Direct/0] 17w0d 10:31:49
> via fe-1/3/1.0
10.255.165.1/32 *[Direct/0] 25w4d 04:13:18
> via lo0.0
30.30.30.0/24 *[Direct/0] 17w0d 23:06:26
> via fe-1/3/2.0
192.168.164.0/22 *[Direct/0] 25w4d 04:13:20
> via fxp0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.0102.5516.5001/152
*[Direct/0] 25w4d 04:13:21
> via lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```



```

abcd::10:255:165:1/128
 *[Direct/0] 25w4d 04:13:21
 > via lo0.0
fe80::2a0:a5ff:fe12:ad7/128
 *[Direct/0] 25w4d 04:13:21
 > via lo0.0

```

### show route protocol frr

```

user@host> show route protocol frr
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32 *[FRR/200] 00:05:38, from 20.20.1.1
 > to 20.20.1.3 via ge-4/1/0.0
 to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.4/32 *[FRR/200] 00:05:38, from 20.20.1.1
 > to 20.20.1.4 via ge-4/1/0.0
 to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.5/32 *[FRR/200] 00:05:35, from 20.20.1.1
 > to 20.20.1.5 via ge-4/1/0.0
 to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.6/32 *[FRR/200] 00:05:37, from 20.20.1.1
 > to 20.20.1.6 via ge-4/1/0.0
 to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.7/32 *[FRR/200] 00:05:38, from 20.20.1.1
 > to 20.20.1.7 via ge-4/1/0.0
 to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.8/32 *[FRR/200] 00:05:38, from 20.20.1.1
 > to 20.20.1.8 via ge-4/1/0.0
 to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.9/32 *[FRR/200] 00:05:38, from 20.20.1.1
 > to 20.20.1.9 via ge-4/1/0.0
 to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.10/32 *[FRR/200] 00:05:38, from 20.20.1.1
...

```

### show route protocol l2circuit detail

```

user@host> show route protocol l2circuit detail

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)
 *L2CKT Preference: 7
 Next hop: via ge-2/0/0.0, selected
 Label operation: Pop Offset: 4
 State: <Active Int>
 Local AS: 99
 Age: 9:52
 Task: Common L2 VC
 Announcement bits (1): 0-KRT
 AS path: I

ge-2/0/0.0 (1 entry, 1 announced)
 *L2CKT Preference: 7
 Next hop: via so-1/1/2.0 weight 1, selected
 Label-switched-path my-lsp

```

```

Label operation: Push 100000, Push 100000(top)[0] Offset: -4
Protocol next hop: 10.245.255.63
Push 100000 Offset: -4
 Indirect next hop: 86af0c0 298
State: <Active Int>
Local AS: 99
Age: 9:52
Task: Common L2 VC
Announcement bits (2): 0-KRT 1-Common L2 VC
AS path: I

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
 *L2CKT Preference: 7
 Next hop: via so-1/1/2.0 weight 1, selected
 Label-switched-path my-lsp
 Label operation: Push 100000[0]
 Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
 State: <Active Int>
 Local AS: 99
 Age: 10:21
 Task: l2 circuit
 Announcement bits (1): 0-LDP
 AS path: I
 VC Label 100000, MTU 1500, VLAN ID 512

```

#### show route protocol l2vpn extensive

```

user@host> show route protocol l2vpn extensive

inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)
TSI:
KRT in-kernel 800001 /36 -> {so-0/0/0.0}
 *L2VPN Preference: 7
 Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
 Label operation: Pop Offset: 4
 State: <Active Int>
 Local AS: 69
 Age: 7:48
 Task: Common L2 VC
 Announcement bits (1): 0-KRT
 AS path: I

so-0/0/0.0 (1 entry, 1 announced)
TSI:
KRT in-kernel so-0/0/0.0 /16 -> {indirect(288)}
 *L2VPN Preference: 7
 Next hop: via so-0/0/1.0, selected
 Label operation: Push 800000 Offset: -4
 Protocol next hop: 10.255.14.220
 Push 800000 Offset: -4
 Indirect next hop: 85142a0 288
 State: <Active Int>

```

```

Local AS: 69
Age: 7:48
Task: Common L2 VC
Announcement bits (2): 0-KRT 1-Common L2 VC
AS path: I
Communities: target:69:1 Layer2-info: encaps:PPP,
control flags:2, mtu: 0

```

### show route protocol ldp

```

user@host> show route protocol ldp
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.16.1/32 *[LDP/9] 1d 23:03:35, metric 1
 > via t1-4/0/0.0, Push 100000
192.168.17.1/32 *[LDP/9] 1d 23:03:35, metric 1
 > via t1-4/0/0.0

private1___.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100064 *[LDP/9] 1d 23:03:35, metric 1
 > via t1-4/0/0.0, Pop
100064(S=0) *[LDP/9] 1d 23:03:35, metric 1
 > via t1-4/0/0.0, Pop
100080 *[LDP/9] 1d 23:03:35, metric 1
 > via t1-4/0/0.0, Swap 100000

```

### show route protocol ldp extensive

```

user@host> show route protocol ldp extensive
192.168.16.1/32 (1 entry, 1 announced)
 State: <FlashAll>
 *LDP Preference: 9
 Next-hop reference count: 3
 Next hop: via t1-4/0/0.0, selected
 Label operation: Push 100000
 State: <Active Int>
 Local AS: 65500
 Age: 1d 23:03:58 Metric: 1
 Task: LDP
 Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
 AS path: I

192.168.17.1/32 (1 entry, 1 announced)
 State: <FlashAll>
 *LDP Preference: 9
 Next-hop reference count: 3
 Next hop: via t1-4/0/0.0, selected
 State: <Active Int>
 Local AS: 65500
 Age: 1d 23:03:58 Metric: 1
 Task: LDP
 Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
 AS path: I

```

```
private1__inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

```
mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
```

```
100064 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 100064 /36 -> {t1-4/0/0.0}
```

```
*LDP Preference: 9
 Next-hop reference count: 2
 Next hop: via t1-4/0/0.0, selected
 State: <Active Int>
 Local AS: 65500
 Age: 1d 23:03:58 Metric: 1
 Task: LDP
 Announcement bits (1): 0-KRT
 AS path: I
 Prefixes bound to route: 192.168.17.1/32
```

```
100064(S=0) (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 100064 /40 -> {t1-4/0/0.0}
```

```
*LDP Preference: 9
 Next-hop reference count: 2
 Next hop: via t1-4/0/0.0, selected
 Label operation: Pop
 State: <Active Int>
 Local AS: 65500
 Age: 1d 23:03:58 Metric: 1
 Task: LDP
 Announcement bits (1): 0-KRT
 AS path: I
```

```
100080 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 100080 /36 -> {t1-4/0/0.0}
```

```
*LDP Preference: 9
 Next-hop reference count: 2
 Next hop: via t1-4/0/0.0, selected
 Label operation: Swap 100000
 State: <Active Int>
 Local AS: 65500
 Age: 1d 23:03:58 Metric: 1
 Task: LDP
 Announcement bits (1): 0-KRT
 AS path: I
 Prefixes bound to route: 192.168.16.1/32
```

### show route protocol ospf (Layer 3 VPN)

```
user@host> show route protocol ospf
```

```
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

```
10.39.1.4/30 *[OSPF/10] 00:05:18, metric 4
 > via t3-3/2/0.0
10.39.1.8/30 [OSPF/10] 00:05:18, metric 2
 > via t3-3/2/0.0
10.255.14.171/32 *[OSPF/10] 00:05:18, metric 4
 > via t3-3/2/0.0
10.255.14.179/32 *[OSPF/10] 00:05:18, metric 2
 > via t3-3/2/0.0
```

```

224.0.0.5/32 *[OSPF/10] 20:25:55, metric 1

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 [OSPF/10] 00:05:43, metric 1
 > via so-0/2/2.0
10.255.14.173/32 *[OSPF/10] 00:05:43, metric 1
 > via so-0/2/2.0
224.0.0.5/32 *[OSPF/10] 20:26:20, metric 1

```

### show route protocol ospf detail

```

user@host> show route protocol ospf detail
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
 OSPF Preference: 10
 Nexthop: via so-0/2/2.0, selected
 State: <Int>
 Inactive reason: Route Preference
 Age: 6:25 Metric: 1
 Area: 0.0.0.0
 Task: VPN-AB-OSPF
 AS path: I
 Communities: Route-Type:0.0.0.0:1:0

...

```

### show route protocol rip

```

user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 *[RIP/100] 20:24:34, metric 2
 > to 10.39.1.22 via t3-0/2/2.0
224.0.0.9/32 *[RIP/100] 00:03:59, metric 1

```

### show route protocol rip detail

```

user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
 *RIP Preference: 100
 Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
 State: <Active Int>
 Age: 20:25:02 Metric: 2
 Task: VPN-AB-RIPv2
 Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179
 AS path: I
 Route learned from 10.39.1.22 expires in 96 seconds

```

**show route protocol ripng table inet6**

```

user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128 *[RIPng/100] 02:13:33, metric 2
 > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128 *[RIPng/100] 02:13:33, metric 2
 > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128 *[RIPng/100] 02:13:33, metric 2
 > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::4/128 *[RIPng/100] 02:13:33, metric 2
 > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128 *[RIPng/100] 02:13:33, metric 2
 > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::6/128 *[RIPng/100] 02:13:33, metric 2
 > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0

```

**show route protocol static detail**

```

user@host> show route protocol static detail
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
10.5.0.0/16 (1 entry, 1 announced)
 *Static Preference: 5
 Next hop type: Router, Next hop index: 324
 Address: 0x9274010
 Next-hop reference count: 27
 Next hop: 192.168.187.126 via fxp0.0, selected
 Session Id: 0x0
 State: <Active NoReadvrt Int Ext>
 Age: 7w3d 21:24:25
 Validation State: unverified
 Task: RT
 Announcement bits (1): 0-KRT
 AS path: I

10.10.0.0/16 (1 entry, 1 announced)
 *Static Preference: 5
 Next hop type: Router, Next hop index: 324
 Address: 0x9274010
 Next-hop reference count: 27
 Next hop: 192.168.187.126 via fxp0.0, selected
 Session Id: 0x0
 State: <Active NoReadvrt Int Ext>
 Age: 7w3d 21:24:25
 Validation State: unverified
 Task: RT
 Announcement bits (1): 0-KRT
 AS path: I

10.13.10.0/23 (1 entry, 1 announced)
 *Static Preference: 5
 Next hop type: Router, Next hop index: 324
 Address: 0x9274010
 Next-hop reference count: 27
 Next hop: 192.168.187.126 via fxp0.0, selected
 Session Id: 0x0
 State: <Active NoReadvrt Int Ext>
 Age: 7w3d 21:24:25
 Validation State: unverified

```

Task: RT  
Announcement bits (1): 0-KRT  
AS path: I

## show route receive-protocol

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |  |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--|
| <b>Syntax</b>                      | show route receive-protocol <i>protocol neighbor-address</i><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |  |
| <b>Syntax (EX Series Switches)</b> | show route receive-protocol <i>protocol neighbor-address</i><br><brief   detail   extensive   terse>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |  |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |  |
| <b>Description</b>                 | Display the routing information as it was received through a particular neighbor using a particular dynamic routing protocol.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |  |
| <b>Options</b>                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b><i>protocol neighbor-address</i></b> —Protocol transmitting the route ( <b>bgp</b> , <b>dvmrp</b> , <b>msdp</b> , <b>pim</b> , <b>rip</b> , or <b>ripng</b> ) and address of the neighboring router from which the route entry was received.                                                                                                                                                                                                                                        |  |
| <b>Additional Information</b>      | The output displays the selected routes and the attributes with which they were received, but does not show the effects of import policy on the routing attributes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |  |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |  |
| <b>List of Sample Output</b>       | <a href="#">show route receive-protocol bgp on page 3173</a><br><a href="#">show route receive-protocol bgp extensive on page 3173</a><br><a href="#">show route receive-protocol bgp table extensive on page 3173</a><br><a href="#">show route receive-protocol bgp logical-system extensive on page 3174</a><br><a href="#">show route receive-protocol bgp detail (Layer 2 VPN) on page 3175</a><br><a href="#">show route receive-protocol bgp extensive (Layer 2 VPN) on page 3175</a><br><a href="#">show route receive-protocol bgp (Layer 3 VPN) on page 3176</a><br><a href="#">show route receive-protocol bgp detail (Layer 3 VPN) on page 3176</a><br><a href="#">show route receive-protocol bgp extensive (Layer 3 VPN) on page 3177</a> |  |
| <b>Output Fields</b>               | Table 260 on page 3170 describes the output fields for the <b>show route receive-protocol</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |  |

Table 260: show route receive-protocol Output Fields

| Field Name                 | Field Description                                                       | Level of Output |
|----------------------------|-------------------------------------------------------------------------|-----------------|
| <i>routing-table-name</i>  | Name of the routing table—for example, inet.0.                          | All levels      |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table. | All levels      |



Table 260: show route receive-protocol Output Fields (*continued*)

| Field Name                                   | Field Description                                                                                                                                                                                                                                                                                                                        | Level of Output         |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>number routes</b>                         | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> <li>• <b>active</b></li> <li>• <b>holddown</b> (routes that are in pending state before being declared inactive)</li> <li>• <b>hidden</b> (routes that are not used because of a routing policy)</li> </ul> | All levels              |
| <b>Prefix</b>                                | Destination prefix.                                                                                                                                                                                                                                                                                                                      | none <b>brief</b>       |
| <b>MED</b>                                   | Multiple exit discriminator value included in the route.                                                                                                                                                                                                                                                                                 | none <b>brief</b>       |
| <b>destination-prefix (entry, announced)</b> | Destination prefix. The <b>entry</b> value is the number of routes for this destination, and the <b>announced</b> value is the number of routes being announced for this destination.                                                                                                                                                    | <b>detail extensive</b> |
| <b>Route Distinguisher</b>                   | 64-bit prefix added to IP subnets to make them unique.                                                                                                                                                                                                                                                                                   | <b>detail extensive</b> |
| <b>Label-Base, range</b>                     | First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.                                                                                                                                                                   | <b>detail extensive</b> |
| <b>VPN Label</b>                             | Virtual private network (VPN) label. Packets are sent between CE and PE routing devices by advertising VPN labels. VPN labels transit over either an RSVP or an LDP label-switched path (LSP) tunnel.                                                                                                                                    | <b>detail extensive</b> |
| <b>Next hop</b>                              | Next hop to the destination. An angle bracket ( > ) indicates that the route is the selected route.                                                                                                                                                                                                                                      | All levels              |
| <b>Localpref or Lclpref</b>                  | Local preference value included in the route.                                                                                                                                                                                                                                                                                            | All levels              |

Table 260: show route receive-protocol Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output         |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>AS path</b>             | <p>Autonomous system (AS) path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• <b>I</b>—IGP.</li> <li>• <b>E</b>—EGP.</li> <li>• <b>?</b>—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• <b>[ ]</b>—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used the AS-path merge process, as defined in RFC 4893.</li> <li>• <b>[ ]</b>—If more than one AS number is configured on the router, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path.</li> <li>• <b>{ }</b>—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• <b>( )</b>—Parentheses enclose a confederation.</li> <li>• <b>( [ ] )</b>—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> | All levels              |
| <b>Cluster list</b>        | (For route reflected output only) Cluster ID sent by the route reflector.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b> |
| <b>Originator ID</b>       | (For route reflected output only) Address of routing device that originally sent the route to the route reflector.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b> |
| <b>Communities</b>         | Community path attribute for the route. See the Output Field table in the <a href="#">show route detail</a> command for all possible values for this field.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>AIGP</b>                | Accumulated interior gateway protocol (AIGP) BGP attribute.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Attrset AS</b>          | Number, local preference, and path of the AS that originated the route. These values are stored in the <b>Attrset</b> attribute at the originating routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b> |
| <b>Layer2-info: encaps</b> | Layer 2 encapsulation (for example, VPLS).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b> |
| <b>control flags</b>       | Control flags: <b>none</b> or <b>Site Down</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b> |
| <b>mtu</b>                 | Maximum transmission unit (MTU) of the Layer 2 circuit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b> |

## Sample Output

### show route receive-protocol bgp

```
user@host> show route receive-protocol bgp 10.255.245.215

inet.0: 28 destinations, 33 routes (27 active, 0 holddown, 1 hidden)
Prefix Next hop MED Lclpref AS path
10.22.1.0/24 10.255.245.215 0 100 I
10.22.2.0/24 10.255.245.215 0 100 I
```

### show route receive-protocol bgp extensive

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
Prefix Next hop MED Lclpref AS path
1.1.1.0/24 (1 entry, 1 announced)
 Next hop: 10.0.50.3
 Localpref: 100
 AS path: I <Originator>
 Cluster list: 10.2.3.1
 Originator ID: 10.255.245.45
165.3.0.0/16 (1 entry, 1 announced)
 Next hop: 111.222.5.254
 Localpref: 100
 AS path: I <Originator>
 Cluster list: 10.2.3.1
 Originator ID: 10.255.245.68
165.4.0.0/16 (1 entry, 1 announced)
 Next hop: 111.222.5.254
 Localpref: 100
 AS path: I <Originator>
 Cluster list: 10.2.3.1
 Originator ID: 10.255.245.45
195.1.2.0/24 (1 entry, 1 announced)
 Next hop: 111.222.5.254
 Localpref: 100
 AS path: I <Originator>
 Cluster list: 10.2.3.1
 Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
Prefix Next hop MED Lclpref AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix Next hop MED Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix Next hop MED Lclpref AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

### show route receive-protocol bgp table extensive

```
user@host> show route receive-protocol bgp 207.17.136.192 table inet.0 66.117.68.0/24 extensive
inet.0: 227315 destinations, 227316 routes (227302 active, 0 holddown, 13 hidden)
* 66.117.63.0/24 (1 entry, 1 announced)
 Nexthop: 207.17.136.29
 Localpref: 100
 AS path: AS2 PA[6]: 14203 2914 3356 29748 33437 AS_TRANS
 AS path: AS4 PA[2]: 33437 393219
 AS path: Merged[6]: 14203 2914 3356 29748 33437 393219 I
 Communities: 2914:420
```

**show route receive-protocol bgp logical-system extensive**

```
user@host> show route receive-protocol bgp 10.0.0.9 logical-system PE4 extensive
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)
* 10.0.0.0/30 (1 entry, 1 announced)
 Accepted
 Route Label: 3
 Nexthop: 10.0.0.9
 AS path: 13979 I

* 10.0.0.4/30 (1 entry, 1 announced)
 Accepted
 Route Label: 3
 Nexthop: 10.0.0.9
 AS path: 13979 I

10.0.0.8/30 (2 entries, 1 announced)
 Accepted
 Route Label: 3
 Nexthop: 10.0.0.9
 AS path: 13979 I

* 10.9.9.1/32 (1 entry, 1 announced)
 Accepted
 Route Label: 3
 Nexthop: 10.0.0.9
 AS path: 13979 I

* 10.100.1.1/32 (1 entry, 1 announced)
 Accepted
 Route Label: 3
 Nexthop: 10.0.0.9
 AS path: 13979 I

* 44.0.0.0/24 (1 entry, 1 announced)
 Accepted
 Route Label: 300096
 Nexthop: 10.0.0.9
 AS path: 13979 I
 AIGP: 203

* 55.0.0.0/24 (1 entry, 1 announced)
 Accepted
 Route Label: 300112
 Nexthop: 10.0.0.9
 AS path: 13979 7018 I
 AIGP: 25

* 66.0.0.0/24 (1 entry, 1 announced)
 Accepted
 Route Label: 300144
 Nexthop: 10.0.0.9
 AS path: 13979 7018 I

* 99.0.0.0/24 (1 entry, 1 announced)
 Accepted
 Route Label: 300160
 Nexthop: 10.0.0.9
 AS path: 13979 7018 I
```

**show route receive-protocol bgp detail (Layer 2 VPN)**

```

user@host> show route receive-protocol bgp 10.255.14.171 detail
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix Nexthop MED Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
Prefix Nexthop MED Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
 Route Distinguisher: 10.255.245.35:1
 Label-base : 800000, range : 4, status-vector : 0x0
 Nexthop: 10.255.245.35
 Localpref: 100
 AS path: I
 Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags: 0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
 Route Distinguisher: 10.255.245.35:1
 Label-base : 800000, range : 4, status-vector : 0x0
 Nexthop: 10.255.245.35
 Localpref: 100
 AS path: I
 Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0

```

**show route receive-protocol bgp extensive (Layer 2 VPN)**

```

user@host> show route receive-protocol bgp 10.255.14.171 extensive
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix Nexthop MED Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
 Route Distinguisher: 10.255.245.35:1
 Label-base : 800000, range : 4, status-vector : 0x0
 Nexthop: 10.255.245.35
 Localpref: 100
 AS path: I
 Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
 Route Distinguisher: 10.255.245.35:1
 Label-base : 800000, range : 4, status-vector : 0x0
 Nexthop: 10.255.245.35
 Localpref: 100

```

```

AS path: I
Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0

```

### show route receive-protocol bgp (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.171
inet.0: 33 destinations, 33 routes (32 active, 0 holddown, 1 hidden)
Prefix Nexthop MED Lclpref AS path
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
10.255.14.175/32 10.255.14.171 100 2 I
10.255.14.179/32 10.255.14.171 2 100 I
VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
10.255.14.175/32 10.255.14.171 100 2 I
10.255.14.177/32 10.255.14.171 100 I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
10.255.14.171:300:10.255.14.177/32
 10.255.14.171 100 I
10.255.14.171:100:10.255.14.179/32
 10.255.14.171 2 100 I
10.255.14.171:200:10.255.14.175/32
 10.255.14.171 100 2 I

```

### show route receive-protocol bgp detail (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.174 detail
inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
 Route Distinguisher: 10.255.14.176:2
 VPN Label: 101264
 Nexthop: 10.255.14.174
 Localpref: 100
 AS path: I
 Communities: target:200:100
 AttrSet AS: 100
 Localpref: 100
 AS path: I
* 10.255.14.172/32 (1 entry, 1 announced)
 Route Distinguisher: 10.255.14.176:2
 VPN Label: 101280
 Nexthop: 10.255.14.174
 Localpref: 100
 AS path: I
 Communities: target:200:100
 AttrSet AS: 100
 Localpref: 100
 AS path: I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

```

```

* 10.255.14.174:2:10.49.0.0/30 (1 entry, 0 announced)
 Route Distinguisher: 10.255.14.174:2
 VPN Label: 101264
 Nexthop: 10.255.14.174
 Localpref: 100
 AS path: I
 Communities: target:200:100
 AttrSet AS: 100
 Localpref: 100
 AS path: I
* 10.255.14.174:2:10.255.14.172/32 (1 entry, 0 announced)
 Route Distinguisher: 10.255.14.174:2
 VPN Label: 101280
 Nexthop: 10.255.14.174
 Localpref: 100
 AS path: I
 Communities: target:200:100
 AttrSet AS: 100
 Localpref: 100
 AS path: I
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

```

#### show route receive-protocol bgp extensive (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
 Prefix Nexthop MED Lclpref AS path
 1.1.1.0/24 (1 entry, 1 announced)
 Nexthop: 10.0.50.3
 Localpref: 100
 AS path: I <Originator>
 Cluster list: 10.2.3.1
 Originator ID: 10.255.245.45
 165.3.0.0/16 (1 entry, 1 announced)
 Nexthop: 111.222.5.254
 Localpref: 100
 AS path: I <Originator>
 Cluster list: 10.2.3.1
 Originator ID: 10.255.245.68
 165.4.0.0/16 (1 entry, 1 announced)
 Nexthop: 111.222.5.254
 Localpref: 100
 AS path: I <Originator>
 Cluster list: 10.2.3.1
 Originator ID: 10.255.245.45
 195.1.2.0/24 (1 entry, 1 announced)
 Nexthop: 111.222.5.254
 Localpref: 100
 AS path: I <Originator>
 Cluster list: 10.2.3.1
 Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
 Prefix Nexthop MED Lclpref AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
 Prefix Nexthop MED Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
 Prefix Nexthop MED Lclpref AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)

```

## show route table

---

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | <code>show route table <i>routing-table-name</i></code><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Syntax (EX Series Switches)</b> | <code>show route table <i>routing-table-name</i></code><br><brief   detail   extensive   terse>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                 | Display the route entries in a particular routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b><i>routing-table-name</i></b> —Display route entries for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the <b>show route table inet</b> command).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"><li>• <a href="#">show route summary</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>       | <a href="#">show route table bgp.l2.vpn on page 3179</a><br><a href="#">show route table bgp.l3vpn.0 on page 3179</a><br><a href="#">show route table bgp.l3vpn.0 detail on page 3179</a><br><a href="#">show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured) on page 3181</a><br><a href="#">show route table inet.0 on page 3181</a><br><a href="#">show route table inet6.0 on page 3181</a><br><a href="#">show route table inet6.3 on page 3182</a><br><a href="#">show route table inetflow detail on page 3182</a><br><a href="#">show route table l2circuit.0 on page 3182</a><br><a href="#">show route table mpls on page 3183</a><br><a href="#">show route table mpls extensive on page 3183</a><br><a href="#">show route table mpls.0 on page 3183</a><br><a href="#">show route table mpls.0 (RSVP Route—Transit LSP) on page 3184</a><br><a href="#">show route table vpls_1 detail on page 3184</a><br><a href="#">show route table vpn-a on page 3184</a><br><a href="#">show route table vpn-a.mdt.0 on page 3185</a><br><a href="#">show route table VPN-A detail on page 3185</a><br><a href="#">show route table VPN-AB.inet.0 on page 3185</a><br><a href="#">show route table VPN_blue.mvpn-inet6.0 on page 3186</a><br><a href="#">show route table VPN-A detail on page 3186</a> |



[show route table inetflow detail on page 3187](#)

**Output Fields** For information about output fields, see the output field tables for the [show route](#) command, the [show route detail](#) command, the [show route extensive](#) command, or the [show route terse](#) command.

## Sample Output

### show route table bgp.l2vpn

```
user@host> show route table bgp.l2vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
 *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
 AS path: I
 > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

### show route table bgp.l3vpn.0

```
user@host> show route table bgp.l3vpn.0
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.71.15:100:10.255.71.17/32
 *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
 AS path: I
 > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.15:200:10.255.71.18/32
 *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
 AS path: I
 > via so-2/1/0.0, Push 100021, Push 100011(top)
```

### show route table bgp.l3vpn.0 detail

```
user@host> show route table bgp.l3vpn.0 detail
bgp.l3vpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

10.255.245.12:1:4.0.0.0/8 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 10.255.245.12:1
 Source: 10.255.245.12
 Next hop: 192.168.208.66 via fe-0/0/0.0, selected
 Label operation: Push 182449
 Protocol next hop: 10.255.245.12
 Push 182449
 Indirect next hop: 863a630 297
 State: <Active Int Ext>
 Local AS: 35 Peer AS: 35
 Age: 12:19 Metric2: 1
 Task: BGP_35.10.255.245.12+179
 Announcement bits (1): 0-BGP.0.0.0.0+179
 AS path: 30 10458 14203 2914 3356 I (Atomic) Aggregator: 3356 4.68.0.11

 Communities: 2914:420 target:11111:1 origin:56:78
 VPN Label: 182449
 Localpref: 100
```

```
Router ID: 10.255.245.12

10.255.245.12:1:4.17.225.0/24 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 10.255.245.12:1
 Source: 10.255.245.12
 Next hop: 192.168.208.66 via fe-0/0/0.0, selected
 Label operation: Push 182465
 Protocol next hop: 10.255.245.12
 Push 182465
 Indirect next hop: 863a8f0 305
 State: <Active Int Ext>
 Local AS: 35 Peer AS: 35
 Age: 12:19 Metric2: 1
 Task: BGP_35.10.255.245.12+179
 Announcement bits (1): 0-BGP.0.0.0.0+179
 AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496 6496 I
 Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
 VPN Label: 182465
 Localpref: 100
 Router ID: 10.255.245.12

10.255.245.12:1:4.17.226.0/23 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 10.255.245.12:1
 Source: 10.255.245.12
 Next hop: 192.168.208.66 via fe-0/0/0.0, selected
 Label operation: Push 182465
 Protocol next hop: 10.255.245.12
 Push 182465
 Indirect next hop: 86bd210 330
 State: <Active Int Ext>
 Local AS: 35 Peer AS: 35
 Age: 12:19 Metric2: 1
 Task: BGP_35.10.255.245.12+179
 Announcement bits (1): 0-BGP.0.0.0.0+179
 AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496
6496 I
 Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
 VPN Label: 182465
 Localpref: 100
 Router ID: 10.255.245.12

10.255.245.12:1:4.17.251.0/24 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 10.255.245.12:1
 Source: 10.255.245.12
 Next hop: 192.168.208.66 via fe-0/0/0.0, selected
 Label operation: Push 182465
 Protocol next hop: 10.255.245.12
 Push 182465
 Indirect next hop: 86bd210 330
 State: <Active Int Ext>
 Local AS: 35 Peer AS: 35
 Age: 12:19 Metric2: 1
 Task: BGP_35.10.255.245.12+179
 Announcement bits (1): 0-BGP.0.0.0.0+179
 AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496
6496 I
```

```

Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100

```

### show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured)

```

user@host> show route table bgp.rtarget.0
bgp.rtarget.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100:100:100/96
 *[RTarget/5] 00:03:14
 Type Proxy
 for 10.255.165.103
 for 10.255.166.124
 Local

```

### show route table inet.0

```

user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0 *[Static/5] 00:51:57
 > to 111.222.5.254 via fxp0.0
1.0.0.1/32 *[Direct/0] 00:51:58
 > via at-5/3/0.0
1.0.0.2/32 *[Local/0] 00:51:58
 Local
12.12.12.21/32 *[Local/0] 00:51:57
 Reject
13.13.13.13/32 *[Direct/0] 00:51:58
 > via t3-5/2/1.0
13.13.13.14/32 *[Local/0] 00:51:58
 Local
13.13.13.21/32 *[Local/0] 00:51:58
 Local
13.13.13.22/32 *[Direct/0] 00:33:59
 > via t3-5/2/0.0
127.0.0.1/32 [Direct/0] 00:51:58
 > via lo0.0
111.222.5.0/24 *[Direct/0] 00:51:58
 > via fxp0.0
111.222.5.81/32 *[Local/0] 00:51:58
 Local

```

### show route table inet6.0

```

user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64 *[Direct/0] 00:01:34
>via fe-0/1/0.0

fec0:0:0:3::/128 *[Local/0] 00:01:34
>Local

fec0:0:0:4::/64 *[Static/5] 00:01:34
>to fec0:0:0:3::ffff via fe-0/1/0.0

```

### show route table inet6.3

```
user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
 *[LDP/9] 00:00:22, metric 1
 > via so-1/0/0.0
::10.255.245.196/128
 *[LDP/9] 00:00:08, metric 1
 > via so-1/0/0.0, Push 100008
```

### show route table inetflow detail

```
user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Next-hop reference count: 2
 State: **Active Ext>
 Local AS: 65002 Peer AS: 65000
 Age: 4
 Task: BGP_65000.10.12.99.5+3792
 Announcement bits (1): 0-Flow
 AS path: 65000 I
 Communities: traffic-rate:0:0
 Validation state: Accept, Originator: 10.12.99.5
 Via: 10.12.44.0/24, Active
 Localpref: 100
 Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
 *Flow Preference: 5
 Next-hop reference count: 2
 State: **Active>
 Local AS: 65002
 Age: 6:30
 Task: RT Flow
 Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
 AS path: I
 Communities: 1:1
```

### show route table l2circuit.0

```
user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
 *[L2CKT/7] 00:50:47
 > via so-0/1/2.0, Push 100049
 > via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
 *[LDP/9] 00:50:14
 Discard
10.1.1.195:CtrlWord:1:2:Local/96
 *[L2CKT/7] 00:50:47
 > via so-0/1/2.0, Push 100049
 > via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
```

```
*[LDP/9] 00:50:14
Discard
```

### show route table mpls

```
user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0 *[MPLS/0] 00:13:55, metric 1
 Receive
1 *[MPLS/0] 00:13:55, metric 1
 Receive
2 *[MPLS/0] 00:13:55, metric 1
 Receive
1024 *[VPN/0] 00:04:18
 to table red.inet.0, Pop
```

### show route table mpls extensive

```
user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}
 *LDP Preference: 9
 Next hop: via so-1/0/0.0, selected
 Pop
 State: <Active Int>
 Age: 29:50 Metric: 1
 Task: LDP
 Announcement bits (1): 0-KRT
 AS path: I
 Prefixes bound to route: 10.0.0.194/32
```

### show route table mpls.0

```
user@host> show route table mpls.0
mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0 *[MPLS/0] 00:45:09, metric 1
 Receive
1 *[MPLS/0] 00:45:09, metric 1
 Receive
2 *[MPLS/0] 00:45:09, metric 1
 Receive
100000 *[L2VPN/7] 00:43:04
 > via so-0/1/0.1, Pop
100001 *[L2VPN/7] 00:43:03
 > via so-0/1/0.2, Pop Offset: 4
100002 *[LDP/9] 00:43:22, metric 1
 via so-0/1/2.0, Pop
 > via so-0/1/3.0, Pop
100002(S=0) *[LDP/9] 00:43:22, metric 1
 via so-0/1/2.0, Pop
 > via so-0/1/3.0, Pop
100003 *[LDP/9] 00:43:22, metric 1
 > via so-0/1/2.0, Swap 100002
 via so-0/1/3.0, Swap 100002
100004 *[LDP/9] 00:43:16, metric 1
 via so-0/1/2.0, Swap 100049
 > via so-0/1/3.0, Swap 100049
```

```

so-0/1/0.1 *[L2VPN/7] 00:43:04
 > via so-0/1/2.0, Push 100001, Push 100049(top)
 via so-0/1/3.0, Push 100001, Push 100049(top)
so-0/1/0.2 *[L2VPN/7] 00:43:03
 > via so-0/1/2.0, Push 100000, Push 100049(top) Offset: -4
 > via so-0/1/3.0, Push 100000, Push 100049(top) Offset: -4

```

### show route table mpls.0 (RSVP Route—Transit LSP)

```
user@host> show route table mpls.0
```

```

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

0 *[MPLS/0] 00:37:31, metric 1
 Receive
1 *[MPLS/0] 00:37:31, metric 1
 Receive
2 *[MPLS/0] 00:37:31, metric 1
 Receive
13 *[MPLS/0] 00:37:31, metric 1
 Receive
300352 *[RSVP/7/1] 00:08:00, metric 1
 > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300352(S=0) *[RSVP/7/1] 00:08:00, metric 1
 > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384 *[RSVP/7/2] 00:05:20, metric 1
 > to 8.64.1.106 via ge-1/0/0.0, Pop
300384(S=0) *[RSVP/7/2] 00:05:20, metric 1
 > to 8.64.1.106 via ge-1/0/0.0, Pop

```

### show route table vpls\_1 detail

```
user@host> show route table vpls_1 detail
```

```

vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

```

```

1.1.1.11:1000:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Receive table: vpls_1.l2vpn.0
Next-hop reference count: 2
State: <Active Int Ext>
Age: 4:29:47 Metric2: 1
Task: vpls_1-l2vpn
Announcement bits (1): 1-BGP.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF

```

### show route table vpn-a

```
user@host> show route table vpn-a
```

```
vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

```

192.168.16.1:1:1/96
 *[VPN/7] 05:48:27
 Discard
192.168.24.1:1:2/96
 *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
 AS path: I
 > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am

```

```

192.168.24.1:1:3:1/96
 *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
 AS path: I
 > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am

```

#### show route table vpn-a.mdt.0

```

user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
 *[MVPN/70] 01:23:05, metric2 1
 Indirect
1:1:1:10.255.14.218:232.1.1.1/144
 *[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
 AS path: I
 > via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
 *[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
 AS path: I
 > via so-0/0/1.0, label-switched-path r0-to-r2

```

#### show route table VPN-A detail

```

user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 10.255.179.13:200
 Next hop type: Indirect
 Next-hop reference count: 5
 Source: 10.255.179.13
 Next hop type: Router, Next hop index: 732
 Next hop: 10.39.1.14 via fe-0/3/0.0, selected
 Label operation: Push 299824, Push 299824(top)
 Protocol next hop: 10.255.179.13
 Push 299824
 Indirect next hop: 8f275a0 1048574
 State: (Secondary Active Int Ext)
 Local AS: 1 Peer AS: 1
 Age: 3:41:06 Metric: 1 Metric2: 1
 Task: BGP_1.10.255.179.13+64309
 Announcement bits (2): 0-KRT 1-BGP RT Background
 AS path: I
 Communities: target:1:200 rte-type:0.0.0.0:1:0
 Import Accepted
 VPN Label: 299824 TTL Action: vrf-ttl-propagate
 Localpref: 100
 Router ID: 10.255.179.13
 Primary Routing Table bgp.13vpn.0

```

#### show route table VPN-AB.inet.0

```

user@host> show route table VPN-AB.inet.0
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.0/30 *[OSPF/10] 00:07:24, metric 1
 > via so-7/3/1.0
10.39.1.4/30 *[Direct/0] 00:08:42
 > via so-5/1/0.0

```

```

10.39.1.6/32 *[Local/0] 00:08:46
 Local
10.255.71.16/32 *[Static/5] 00:07:24
 > via so-2/0/0.0
10.255.71.17/32 *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
 AS path: I
 > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.18/32 *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
 AS path: I
 > via so-2/1/0.0, Push 100021, Push 100011(top)
10.255.245.245/32 *[BGP/170] 00:08:35, localpref 100
 AS path: 2 I
 > to 10.39.1.5 via so-5/1/0.0
10.255.245.246/32 *[OSPF/10] 00:07:24, metric 1
 > via so-7/3/1.0

```

#### show route table VPN\_blue.mvpn-inet6.0

```

user@host> show route table VPN_blue.mvpn-inet6.0
vpn_blue.mvpn-inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.255.2.202:65535:10.255.2.202/432
 *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
 AS path: I
 > via so-0/1/3.0
1:10.255.2.203:65535:10.255.2.203/432
 *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
 AS path: I
 > via so-0/1/0.0
1:10.255.2.204:65535:10.255.2.204/432
 *[MVPN/70] 00:57:23, metric2 1
 Indirect
5:10.255.2.202:65535:128::192.168.90.2:128:ffff::1/432
 *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
 AS path: I
 > via so-0/1/3.0
6:10.255.2.203:65535:65000:128::10.12.53.12:128:ffff::1/432
 *[PIM/105] 00:02:37
 Multicast (IPv6)
7:10.255.2.202:65535:65000:128::192.168.90.2:128:ffff::1/432
 *[MVPN/70] 00:02:37, metric2 1
 Indirect

```

#### show route table VPN-A detail

```

user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 10.255.179.13:200
 Next hop type: Indirect
 Next-hop reference count: 5
 Source: 10.255.179.13
 Next hop type: Router, Next hop index: 732
 Next hop: 10.39.1.14 via fe-0/3/0.0, selected
 Label operation: Push 299824, Push 299824(top)
 Protocol next hop: 10.255.179.13
 Push 299824

```



```

Indirect next hop: 8f275a0 1048574
State: (Secondary Active Int Ext)
Local AS: 1 Peer AS: 1
Age: 3:41:06 Metric: 1 Metric2: 1
Task: BGP_1.10.255.179.13+64309
Announcement bits (2): 0-KRT 1-BGP RT Background
AS path: I
Communities: target:1:200 rte-type:0.0.0.0:1:0
Import Accepted
VPN Label: 299824 TTL Action: vrf-ttl-propagate
Localpref: 100
Router ID: 10.255.179.13
Primary Routing Table bgp.l3vpn.0

```

### show route table inetflow detail

```

user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Next-hop reference count: 2
 State: **Active Ext>
 Local AS: 65002 Peer AS: 65000
 Age: 4
 Task: BGP_65000.10.12.99.5+3792
 Announcement bits (1): 0-Flow
 AS path: 65000 I
 Communities: traffic-rate:0:0
 Validation state: Accept, Originator: 10.12.99.5
 Via: 10.12.44.0/24, Active
 Localpref: 100
 Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
 *Flow Preference: 5
 Next-hop reference count: 2
 State: **Active>
 Local AS: 65002
 Age: 6:30
 Task: RT Flow
 Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
 AS path: I
 Communities: 1:1

user@PE1> show route table green.l2vpn.0 (VPLS Multihoming with FEC 129)
green.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.2:100:1.1.1.2/96 AD
 *[VPLS/170] 1d 03:11:03, metric2 1
 Indirect
1.1.1.4:100:1.1.1.4/96 AD
 *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
 AS path: I, validation-state: unverified
 > via ge-1/2/1.5
1.1.1.2:100:1.0/96 MH
 *[VPLS/170] 1d 03:11:03, metric2 1
 Indirect
1.1.1.4:100:1.0/96 MH
 *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
 AS path: I, validation-state: unverified

```

```

> via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.2:1.1.1.4/176
*[VPLS/7] 1d 03:11:02, metric2 1
> via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.4:1.1.1.2/176
*[LDP/9] 1d 03:11:02
Discard
```

## show route terse


|                                                                                                                                                                                                                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                 |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                            | show route terse<br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                        |
| <b>Syntax (EX Series Switches)</b>                                                                                                                                                                                                                                                                                                                                                                                       | show route terse                                                                                                                                                                                                                                |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                               | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                           |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                       | Display a high-level summary of the routes in the routing table.                                                                                                                                                                                |
| <div>  <p><b>NOTE:</b> For BGP routes, the <b>show route terse</b> command displays the local preference attribute and MED instead of the metric1 and metric2 values. This is mostly due to historical reasons.</p> <p>To display the metric1 and metric2 value of a BGP route, use the <b>show route extensive</b> command.</p> </div> |                                                                                                                                                                                                                                                 |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                           | <p><b>none</b>—Display a high-level summary of the routes in the routing table.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                          | view                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>                                                                                                                                                                                                                                                                                                                                                                                             | <a href="#">show route terse on page 3191</a>                                                                                                                                                                                                   |
| <b>Output Fields</b>                                                                                                                                                                                                                                                                                                                                                                                                     | Table 261 on page 3189 describes the output fields for the <b>show route terse</b> command. Output fields are listed in the approximate order in which they appear.                                                                             |

Table 261: show route terse Output Fields

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                               |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>routing-table-name</i>  | Name of the routing table (for example, inet.0).                                                                                                                                                                                                                                                                                                                |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table.                                                                                                                                                                                                                                                                                         |
| <i>number routes</i>       | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> <li><b>active</b> (routes that are active)</li> <li><b>holddown</b> (routes that are in the pending state before being declared inactive)</li> <li><b>hidden</b> (routes that are not used because of a routing policy)</li> </ul> |

Table 261: show route terse Output Fields (*continued*)

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>route key</i>   | Key for the state of the route: <ul style="list-style-type: none"> <li>• <b>+</b>—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>-</b>—A hyphen indicates the last active route.</li> <li>• <b>*</b>—An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul>                                                                                                                                                                                                                |
| <b>A</b>           | Active route. An asterisk (*) indicates this is the active route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>V</b>           | Validation status of the route: <ul style="list-style-type: none"> <li>• <b>?</b>—Not evaluated. Indicates that the route was not learned through BGP.</li> <li>• <b>I</b>—Invalid. Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.</li> <li>• <b>N</b>—Unknown. Indicates that the prefix is not among the prefixes or prefix ranges in the database.</li> <li>• <b>V</b>—Valid. Indicates that the prefix and autonomous system pair are found in the database.</li> </ul> |
| <b>Destination</b> | Destination of the route.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>P</b>           | Protocol through which the route was learned: <ul style="list-style-type: none"> <li>• <b>A</b>—Aggregate</li> <li>• <b>B</b>—BGP</li> <li>• <b>C</b>—CCC</li> <li>• <b>D</b>—Direct</li> <li>• <b>G</b>—GMPLS</li> <li>• <b>I</b>—IS-IS</li> <li>• <b>L</b>—L2CKT, L2VPN, LDP, Local</li> <li>• <b>K</b>—Kernel</li> <li>• <b>M</b>—MPLS, MSDP</li> <li>• <b>O</b>—OSPF</li> <li>• <b>P</b>—PIM</li> <li>• <b>R</b>—RIP, RIPng</li> <li>• <b>S</b>—Static</li> <li>• <b>T</b>—Tunnel</li> </ul>                                                                                                                                                                                     |
| <b>Prf</b>         | Preference value of the route. In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.                                                                                              |
| <b>Metric 1</b>    | First metric value in the route. For routes learned from BGP, this is the MED metric.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Metric 2</b>    | Second metric value in the route. For routes learned from BGP, this is the IGP metric.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

Table 261: show route terse Output Fields (*continued*)

| Field Name      | Field Description                                                                                                                                                                                                                                                                                                                                          |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Next hop</b> | Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.                                                                                                                                                                                                                                                          |
| <b>AS path</b>  | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>I—IGP.</li> <li>E—EGP.</li> <li>?—Incomplete; typically, the AS path was aggregated.</li> </ul> |

## Sample Output

### show route terse

```

user@host> show route terse
inet.0: 10 destinations, 12 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A V Destination P Prf Metric 1 Metric 2 Next hop AS path
* ? 1.0.1.1/32 0 10 1 >10.0.0.2
? B 170 100 I
 unverified
* ? 1.1.1.1/32 D 0 >10.0.0.2
* ? 1.1.1.1/32 D 0 >10.0.0.2
* V 2.2.0.2/32 B 170 110 200 I
 valid
* ? 10.0.0.0/30 D 0 >10.0.0.2
* ? 10.0.0.0/30 B 170 100 >1t-1/2/0.1
? I
 unverified
* ? 10.0.0.1/32 L 0 Local
* ? 10.0.0.4/30 B 170 100 I
 unverified
* ? 10.0.0.8/30 B 170 100 >10.0.0.2
 unverified
* I 172.16.1.1/32 B 170 90 >10.0.0.2
 invalid
* N 192.168.2.3/32 B 170 100 >10.0.0.2
 unknown
* ? 224.0.0.5/32 O 10 1 >10.0.0.2
 MultiRecv

```

## test policy

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>test policy <i>policy-name</i> <i>prefix</i></code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Test a policy configuration to determine which prefixes match routes in the routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <i>policy-name</i> —Name of a policy.<br><i>prefix</i> —Destination prefix to match.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Additional Information</b>   | All prefixes in the default unicast routing table (inet.0) that match prefixes that are the same as or longer than the specific prefix are processed by the <b>from</b> clause in the specified policy. All prefixes accepted by the policy are displayed. The <b>test policy</b> command evaluates a policy differently from the BGP import process. When testing a policy that contains an <b>interface</b> match condition in the <b>from</b> clause, the <b>test policy</b> command uses the match condition. In contrast, BGP does not use the <b>interface</b> match condition when evaluating the policy against routes learned from internal BGP (IBGP) or external BGP (EGBP) multihop peers. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">show policy damping on page 2674</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>    | <a href="#">test policy on page 3192</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>            | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## Sample Output

### test policy

```

user@host> test policy test-statics 3.0.0.1/8
inet.0: 44 destinations, 44 routes (44 active, 0 holddown, 0 hidden)
Prefixes passing policy:

3.0.0.0/8 *[BGP/170] 16:22:46, localpref 100, from 10.255.255.41
 AS Path: 50888 I
 > to 10.11.4.32 via en0.2, label-switched-path 12
3.3.3.1/32 *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
 > to 10.0.4.7 via fxp0.0
3.3.3.2/32 *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
 > to 10.0.4.7 via fxp0.0
3.3.3.3/32 *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
 > to 10.0.4.7 via fxp0.0
3.3.3.4/32 *[IS-IS/18] 2d 00:21:46, metric 0, tag 2

```

```
> to 10.0.4.7 via fxp0.0
Policy test-statics: 5 prefixes accepted, 0 prefixes rejected
```

## Routing Options

---

- [Overview on page 3193](#)
- [Configuration on page 3193](#)
- [Administration on page 3232](#)

### Overview

- [Routing Properties Overview on page 3193](#)

#### Routing Properties Overview

---

- [Protocol-Independent Routing Properties Overview on page 3193](#)

##### ***Protocol-Independent Routing Properties Overview***

In Junos OS, routing capabilities and features that are not specific to any particular routing protocol are collectively called protocol-independent routing properties. These features often interact with routing protocols. In many cases, you combine protocol-independent properties and routing policy to achieve a goal. For example, you define a static route using protocol-independent properties, and then, using a routing policy, you can redistribute the static route into a routing protocol, such as BGP, OSPF, or IS-IS.

Protocol-independent routing properties include:

- Static, aggregate, and generated routes
- Bidirectional Forwarding Detection on static routes
- Global preference
- Martian routes
- Routing tables and routing information base (RIB) groups

#### **Related Documentation**

- [Examples: Configuring Static Routes](#)
- [Examples: Creating a Routing Table and Populating It with Routes](#)

## Configuration

- [Configuration Tasks on page 3193](#)
- [Configuration Statements on page 3216](#)

#### Configuration Tasks

---

- [Examples: Configuring BFD for Static Routes on page 3194](#)
- [Example: Configuring BFD Authentication for Static Routes on page 3208](#)

### **Examples: Configuring BFD for Static Routes**

- [Understanding BFD for Static Routes on page 3194](#)
- [Example: Configuring BFD for Static Routes on page 3197](#)
- [Example: Enabling BFD on Qualified Next Hops in Static Routes on page 3203](#)

### **Understanding BFD for Static Routes**

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchanges BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the static route failure detection mechanisms, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

By default, BFD is supported on single-hop static routes. In Junos OS Release 8.2 and later, BFD also supports multihop static routes.

To enable failure detection, include the **bfd-liveness-detection** statement in the static route configuration.

In Junos OS Release 9.1 and later, the BFD protocol is supported for IPv6 static routes. Global unicast and link-local IPv6 addresses are supported for static routes. The BFD protocol is not supported on multicast or anycast IPv6 addresses. For IPv6, the BFD protocol supports only static routes and only in Junos OS Release 9.3 and later. IPv6 for BFD is not supported for any other protocol.

To configure the BFD protocol for IPv6 static routes, include the **bfd-liveness-detection** statement at the **[edit routing-options rib inet6.0 static route destination-prefix]** hierarchy level.

In Junos OS Release 8.5 and later, you can configure a hold-down interval to specify how long the BFD session must remain up before a state change notification is sent.

To specify the hold-down interval, include the **holddown-interval** statement in the BFD configuration.



You can configure a number in the range from 0 through 255,000 milliseconds. The default is 0. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.



**NOTE:** If a single BFD session includes multiple static routes, the hold-down interval with the highest value is used.

To specify the minimum transmit and receive intervals for failure detection, include the **minimum-interval** statement in the BFD configuration.

This value represents both the minimum interval after which the local routing device transmits hello packets and the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.



**NOTE:** BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.



**NOTE:** SRX Series devices do not support distributed BFD.

To specify the minimum receive interval for failure detection, include the **minimum-receive-interval** statement in the BFD configuration. This value represents the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you

can configure the minimum receive interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.

To specify the number of hello packets not received by the neighbor that causes the originating interface to be declared down, include the **multiplier** statement in the BFD configuration.

The default value is 3. You can configure a number in the range from 1 through 255.

To specify a threshold for detecting the adaptation of the detection time, include the **threshold** statement in the BFD configuration.

When the BFD session detection time adapts to a value equal to or higher than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the **minimum-interval** or the **minimum-receive-interval** value. The threshold must be a higher value than the multiplier for either of these configured values. For example if the **minimum-receive-interval** is 300 ms and the **multiplier** is 3, the total detection time is 900 ms. Therefore, the detection time threshold must have a value higher than 900.

To specify the minimum transmit interval for failure detection, include the **transmit-interval** **minimum-interval** statement in the BFD configuration.

This value represents the minimum interval after which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.

To specify the threshold for the adaptation of the transmit interval, include the **transmit-interval threshold** statement in the BFD configuration.

The threshold value must be greater than the transmit interval. When the BFD session transmit time adapts to a value greater than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the value for the **minimum-interval** or the **minimum-receive-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level. The threshold must be a higher value than the multiplier for either of these configured values.

To specify the BFD version, include the **version** statement in the BFD configuration. The default is to have the version detected automatically.

To include an IP address for the next hop of the BFD session, include the **neighbor** statement in the BFD configuration.



**NOTE:** You must configure the **neighbor** statement if the next hop specified is an interface name. If you specify an IP address as the next hop, that address is used as the neighbor address for the BFD session.

---

In Junos OS Release 9.0 and later, you can configure BFD sessions not to adapt to changing network conditions.

To disable BFD adaptation, include the **no-adaptation** statement in the BFD configuration.



**NOTE:** We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation in your network.



**NOTE:** If BFD is configured only on one end of a static route, the route is removed from the routing table. BFD establishes a session when BFD is configured on both ends of the static route.

BFD is not supported on ISO address families in static routes. BFD does support IS-IS.

If you configure graceful Routing Engine switchover (GRES) at the same time as BFD, GRES does not preserve the BFD state information during a failover.

Junos OS also supports BFD over multihop static routes. For example, you can configure BFD over a Layer 3 path to provide path integrity over that path. You can limit the number of hops by specifying the time to live (TTL).

To configure BFD over multihop static routes, include the following statements:

```
static route destination-prefix {
 bfd-liveness-detection {
 local-address ip-address;
 minimum-receive-ttl number;
 }
}
```

To specify the source address for the multihop static route and to enable multihop BFD support, include the **local-address** statement.

To specify the number of hops, include the **minimum-receive-ttl** statement. You must configure this statement for a multihop BFD session. You can configure a value in the range from 1 through 255. It is optional for a single-hop BFD session. If you configure the **minimum-receive-ttl** statement for a single-hop session, the value must be 255.

#### **Example: Configuring BFD for Static Routes**

This example shows how to configure Bidirectional Forwarding Detection (BFD) for static routes.

- [Requirements on page 3198](#)
- [Overview on page 3198](#)
- [Configuration on page 3198](#)
- [Verification on page 3201](#)

### Requirements

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

### Overview


There are many practical applications for static routes. Static routing is often used at the network edge to support attachment to stub networks, which, given their single point of entry and egress, are well suited to the simplicity of a static route. In Junos OS, static routes have a global preference of 5. Static routes are activated if the specified next hop is reachable.

In this example, you configure the static route 192.168.47.0/24 from the provider network to the customer network, using the next-hop address of 172.16.1.2. You also configure a static default route of 0.0.0.0/0 from the customer network to the provider network, using a next-hop address of 172.16.1.1.

For demonstration purposes, some loopback interfaces are configured on Device B and Device D. These loopback interfaces provide addresses to ping and thus verify that the static routes are working.

Figure 46 on page 3198 shows the sample network.

**Figure 46: Customer Routes Connected to a Service Provider**



### Configuration

#### CLI Quick Configuration

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

#### Device B

```
set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set interfaces lo0 unit 57 family inet address 10.0.0.1/32
```

```

set interfaces lo0 unit 57 family inet address 10.0.0.2/32
set routing-options static route 192.168.47.0/24 next-hop 172.16.1.2
set routing-options static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
 1000
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all

```

**Device D**

```

set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set interfaces lo0 unit 2 family inet address 192.168.47.5/32
set interfaces lo0 unit 2 family inet address 192.168.47.6/32
set routing-options static route 0.0.0.0/0 next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all

```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure BFD for static routes:

1. On Device B, configure the interfaces.
 

```

[edit interfaces]
user@B# set ge-1/2/0 unit 0 description B->D
user@B# set ge-1/2/0 unit 0 family inet address 172.16.1.1/24
user@B# set lo0 unit 57 family inet address 10.0.0.1/32
user@B# set lo0 unit 57 family inet address 10.0.0.2/32

```
2. On Device B, create a static route and set the next-hop address.
 

```

[edit routing-options]
user@B# set static route 192.168.47.0/24 next-hop 172.16.1.2

```
3. On Device B, configure BFD for the static route.
 

```

[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
 1000

```
4. On Device B, configure tracing operations for BFD.
 

```

[edit protocols]
user@B# set bfd traceoptions file bfd-trace
user@B# set bfd traceoptions flag all

```
5. If you are done configuring Device B, commit the configuration.
 

```

[edit]
user@B# commit

```
6. On Device D, configure the interfaces.
 

```

[edit interfaces]
user@D# set ge-1/2/0 unit 1 description D->B
user@D# set ge-1/2/0 unit 1 family inet address 172.16.1.2/24
user@D# set lo0 unit 2 family inet address 192.168.47.5/32
user@D# set lo0 unit 2 family inet address 192.168.47.6/32

```

7. On Device D, create a static route and set the next-hop address.  
[edit routing-options]  
user@D# set static route 0.0.0.0/0 next-hop 172.16.1.1
8. On Device D, configure BFD for the static route.  
[edit routing-options]  
user@D# set static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
9. On Device D, configure tracing operations for BFD.  
[edit protocols]  
user@D# set bfd traceoptions file bfd-trace  
user@D# set bfd traceoptions flag all
10. If you are done configuring Device D, commit the configuration.  
[edit]  
user@D# commit

### Results

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

```
Device B user@B# show interfaces
ge-1/2/0 {
 unit 0 {
 description B->D;
 family inet {
 address 172.16.1.1/24;
 }
 }
}
lo0 {
 unit 57 {
 family inet {
 address 10.0.0.1/32;
 address 10.0.0.2/32;
 }
 }
}

user@D# show protocols
bfd {
 traceoptions {
 file bfd-trace;
 flag all;
 }
}

user@B# show routing-options
static {
 route 192.168.47.0/24 {
 next-hop 172.16.1.2;
 bfd-liveness-detection {
```

```

 minimum-interval 1000;
 }
}
}

Device D user@D# show interfaces
ge-1/2/0 {
 unit 1 {
 description D->B;
 family inet {
 address 172.16.1.2/24;
 }
 }
}
lo0 {
 unit 2 {
 family inet {
 address 192.168.47.5/32;
 address 192.168.47.6/32;
 }
 }
}

user@D# show routing-options
static {
 route 0.0.0.0/0 {
 next-hop 172.16.1.1;
 bfd-liveness-detection {
 minimum-interval 1000;
 }
 }
}
}

```

### Verification

Confirm that the configuration is working properly.

- [Verifying That BFD Sessions Are Up on page 3201](#)
- [Viewing Detailed BFD Events on page 3202](#)

### Verifying That BFD Sessions Are Up

**Purpose** Verify that the BFD sessions are up, and view details about the BFD sessions.

**Action** From operational mode, enter the `show bfd session extensive` command.

```

user@B> show bfd session extensive

```

| Address    | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|------------|-------|------------|-------------|-------------------|------------|
| 172.16.1.2 | Up    | lt-1/2/0.0 | 3.000       | 1.000             | 3          |

```

Client Static, TX interval 1.000, RX interval 1.000
Session up time 00:14:30
Local diagnostic None, remote diagnostic None
Remote state Up, version 1
Replicated, routing table index 172
Min async interval 1.000, min slow interval 1.000
Adaptive async TX interval 1.000, RX interval 1.000

```

```

Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
Local discriminator 2, remote discriminator 1
Echo mode disabled/inactive

```

```
1 sessions, 1 clients
```

```
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

```
user@D> show bfd session extensive
```

| Address    | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|------------|-------|------------|-------------|-------------------|------------|
| 172.16.1.1 | Up    | lt-1/2/0.1 | 3.000       | 1.000             | 3          |

```

Client Static, TX interval 1.000, RX interval 1.000
Session up time 00:14:35
Local diagnostic None, remote diagnostic None
Remote state Up, version 1
Replicated, routing table index 170
Min async interval 1.000, min slow interval 1.000
Adaptive async TX interval 1.000, RX interval 1.000
Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
Local discriminator 1, remote discriminator 2
Echo mode disabled/inactive

```

```
1 sessions, 1 clients
```

```
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

**Meaning** The TX interval 1.000, RX interval 1.000 output represents the setting configured with the **minimum-interval** statement. All of the other output represents the default settings for BFD. To modify the default settings, include the optional statements under the **bfd-liveness-detection** statement.

### Viewing Detailed BFD Events

**Purpose** View the contents of the BFD trace file to assist in troubleshooting, if needed.

**Action** From operational mode, enter the **file show /var/log/bfd-trace** command.

```

user@B> file show /var/log/bfd-trace
Nov 23 14:26:55 Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 72
Nov 23 14:26:55 PPM Trace: BFD periodic xmit rt tbl index 172
Nov 23 14:26:55 Received Downstream TraceMsg (22) len 108:
Nov 23 14:26:55 IfIndex (3) len 4: 0
Nov 23 14:26:55 Protocol (1) len 1: BFD
Nov 23 14:26:55 Data (9) len 83: (hex) 70 70 6d 64 5f 62 66 64 5f 73 65 6e 64
6d 73 67 20 3a 20
Nov 23 14:26:55 PPM Trace: pcmd_bfd_sendmsg : socket 12 len 24, ifl 78 src
172.16.1.1 dst 172.16.1.2 errno 65
Nov 23 14:26:55 Received Downstream TraceMsg (22) len 93:
Nov 23 14:26:55 IfIndex (3) len 4: 0
Nov 23 14:26:55 Protocol (1) len 1: BFD
Nov 23 14:26:55 Data (9) len 68: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 74

```

**Meaning** BFD messages are being written to the trace file.



**Example: Enabling BFD on Qualified Next Hops in Static Routes**

This example shows how to configure a static route with multiple possible next hops. Each next hop has Bidirectional Forwarding Detection (BFD) enabled.

- [Requirements on page 3203](#)
- [Overview on page 3203](#)
- [Configuration on page 3203](#)
- [Verification on page 3206](#)

**Requirements**

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

**Overview**


In this example, Device B has the static route **192.168.47.0/24** with two possible next hops. The two next hops are defined using two **qualified-next-hop** statements. Each next hop has BFD enabled.

BFD is also enabled on Device D because BFD must be enabled on both ends of the connection.

A next hop is included in the routing table if the BFD session is up. The next hop is removed from the routing table if the BFD session is down.

See [Figure 47 on page 3203](#).

**Figure 47: BFD Enabled on Qualified Next Hops**

**Configuration**

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device B**

```
set interfaces fe-0/1/0 unit 2 description secondary-B->D
set interfaces fe-0/1/0 unit 2 family inet address 192.168.2.1/24
set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set routing-options static route 192.168.47.0/24 qualified-next-hop 192.168.2.2
 bfd-liveness-detection minimum-interval 60
set routing-options static route 192.168.47.0/24 qualified-next-hop 172.16.1.2
 bfd-liveness-detection minimum-interval 60
```

**Device D**

```
set interfaces fe-0/1/0 unit 3 description secondary-D->B
set interfaces fe-0/1/0 unit 3 family inet address 192.168.2.2/24
set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set routing-options static route 0.0.0.0/0 qualified-next-hop 192.168.2.1
set routing-options static route 0.0.0.0/0 qualified-next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 60
```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure a static route with two possible next hops, both with BFD enabled:

1. On Device B, configure the interfaces.

```
[edit interfaces fe-0/1/0]
user@B# set unit 2 description secondary-B->D
user@B# set unit 2 family inet address 192.168.2.1/24

[edit interfaces ge-1/2/0]
user@B# set unit 0 description B->D
user@B# set unit 0 family inet address 172.16.1.1/24
```
2. On Device B, configure the static route with two next hops, both with BFD enabled.

```
[edit routing-options static route 192.168.47.0/24]
user@B# set qualified-next-hop 192.168.2.2 bfd-liveness-detection minimum-interval
60
user@B# set qualified-next-hop 172.16.1.2 bfd-liveness-detection minimum-interval
60
```
3. On Device D, configure the interfaces.

```
[edit interfaces fe-0/1/0]
user@D# set unit 3 description secondary-D->B
user@D# set unit 3 family inet address 192.168.2.2/24

[edit interfaces ge-1/2/0]
user@D# set unit 1 description D->B
user@D# set unit 1 family inet address 172.16.1.2/24
```
4. On Device D, configure a BFD-enabled default static route with two next hops to the provider network.

In this case, BFD is enabled on the route, not on the next hops.

```
[edit routing-options static route 0.0.0.0/0]
user@D# set qualified-next-hop 192.168.2.1
user@D# set qualified-next-hop 172.16.1.1
user@D# set bfd-liveness-detection minimum-interval 60
```

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

```
user@B# show interfaces
fe-0/1/0 {
 unit 2 {
 description secondary-B->D;
 family inet {
 address 192.168.2.1/24;
 }
 }
}
ge-1/2/0 {
 unit 0 {
 description B->D;
 family inet {
 address 172.16.1.1/24;
 }
 }
}

user@B# show routing-options
static {
 route 192.168.47.0/24 {
 qualified-next-hop 192.168.2.2 {
 bfd-liveness-detection {
 minimum-interval 60;
 }
 }
 qualified-next-hop 172.16.1.2 {
 bfd-liveness-detection {
 minimum-interval 60;
 }
 }
 }
}

user@D# show interfaces
fe-0/1/0 {
 unit 3 {
 description secondary-D->B;
 family inet {
 address 192.168.2.2/24;
 }
 }
}
ge-1/2/0 {
 unit 1 {
```

```
description D->B;
family inet {
 address 172.16.1.2/24;
}
}
}

user@D# show routing-options
static {
 route 0.0.0.0/0 {
 qualified-next-hop 192.168.2.1;
 qualified-next-hop 172.16.1.1;
 bfd-liveness-detection {
 minimum-interval 60;
 }
 }
}
```

If you are done configuring the devices, enter **commit** from configuration mode.

### **Verification**

Confirm that the configuration is working properly.

- [Checking the Routing Tables on page 3206](#)
- [Verifying the BFD Sessions on page 3207](#)
- [Removing BFD from Device D on page 3207](#)
- [Removing BFD from One Next Hop on page 3207](#)

### **Checking the Routing Tables**

**Purpose** Make sure that the static route appears in the routing table on Device B with two possible next hops.

**Action**

```
user@B> show route 192.168.47.0 extensive
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
192.168.47.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.47.0/24 -> {192.168.2.2}
 *Static Preference: 5
 Next hop type: Router
 Address: 0x9334010
 Next-hop reference count: 1
 Next hop: 172.16.1.2 via ge-1/2/0.0
 Next hop: 192.168.2.2 via fe-0/1/0.2, selected
 State: <Active Int Ext>
 Age: 9
 Task: RT
 Announcement bits (1): 3-KRT
 AS path: I
```

**Meaning** Both next hops are listed. The next hop 192.168.2.2 is the selected route.

**Verifying the BFD Sessions**

**Purpose** Make sure that the BFD sessions are up.

**Action** user@B> show bfd session

| Address     | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|-------------|-------|------------|-------------|-------------------|------------|
| 172.16.1.2  | Up    | ge-1/2/0.0 | 0.720       | 0.240             | 3          |
| 192.168.2.2 | Up    | fe-0/1/0.2 | 0.720       | 0.240             | 3          |

2 sessions, 2 clients

Cumulative transmit rate 8.3 pps, cumulative receive rate 8.3 pps

**Meaning** The output shows that the BFD sessions are up.

**Removing BFD from Device D**

**Purpose** Demonstrate what happens when the BFD session is down for both next hops.

**Action** 1. Deactivate BFD on Device D.

```
[edit routing-options static route 0.0.0.0/0]
user@D# deactivate bfd-liveness-detection
user@D# commit
```

2. Rerun the **show bfd session** command on Device B.

user@B> show bfd session

| Address     | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|-------------|-------|------------|-------------|-------------------|------------|
| 172.16.1.2  | Down  | ge-1/2/0.0 | 3.000       | 1.000             | 3          |
| 192.168.2.2 | Down  | fe-0/1/0.2 | 3.000       | 1.000             | 3          |

2 sessions, 2 clients

Cumulative transmit rate 2.0 pps, cumulative receive rate 2.0 pps

3. Rerun the **show route 192.168.47.0** command on Device B.

user@B> show route 192.168.47.0

**Meaning** As expected, when the BFD sessions are down, the static route is removed from the routing table.

**Removing BFD from One Next Hop**

**Purpose** Demonstrate what happens when only one next hop has BFD enabled.

**Action** 1. If it is not already deactivated, deactivate BFD on Device D.

```
[edit routing-options static route 0.0.0.0/0]
user@D# deactivate bfd-liveness-detection
user@D# commit
```

2. Deactivate BFD on one of the next hops on Device B.

```
[edit routing-options static route 192.168.47.0/24 qualified-next-hop 172.16.1.2]
user@B# deactivate bfd-liveness-detection
user@B# commit
```

3. Rerun the **show bfd session** command on Device B.

```
user@B> show bfd session
```

| Address     | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|-------------|-------|------------|-------------|-------------------|------------|
| 192.168.2.2 | Down  | fe-0/1/0.2 | 3.000       | 1.000             | 3          |

4. Rerun the **show route 192.168.47.0 extensive** command on Device B.

```
user@B> show route 192.168.47.0 extensive
```

```
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
192.168.47.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.47.0/24 -> {172.16.1.2}
 *Static Preference: 5
 Next hop type: Router, Next hop index: 624
 Address: 0x92f0178
 Next-hop reference count: 3
 Next hop: 172.16.1.2 via ge-1/2/0.0, selected
 State: <Active Int Ext>
 Age: 2:36
 Task: RT
 Announcement bits (1): 3-KRT
 AS path: I
```

**Meaning** As expected, the BFD session is down for the 192.168.2.2 next hop. The 172.16.1.2 next hop remains in the routing table, and the route remains active, because BFD is not a condition for this next hop to remain valid.

#### Related Documentation

- [Example: Configuring BFD Authentication for Static Routes on page 3208](#)
- [Example: Configuring BFD for OSPF](#)
- [Example: Configuring BFD for BGP](#)
- [Example: Configuring BFD for IS-IS](#)
- [Configuring PIM and the Bidirectional Forwarding Detection \(BFD\) Protocol on page 3725](#)

#### **Example: Configuring BFD Authentication for Static Routes**

- [Understanding BFD Authentication for Static Routes on page 3208](#)
- [Example: Configuring BFD Authentication for Static Routes on page 3210](#)

#### **Understanding BFD Authentication for Static Routes**

Bidirectional Forwarding Detection (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when you run BFD over Network Layer protocols, the risk of service attacks can be significant.



**NOTE:** We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels.

Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over IPv4 and IPv6 static routes. BFD authentication is not supported on MPLS OAM sessions. BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

- [BFD Authentication Algorithms on page 3209](#)
- [Security Authentication Keychains on page 3210](#)
- [Strict Versus Loose Authentication on page 3210](#)

#### ***BFD Authentication Algorithms***

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with

every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



**NOTE:** Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

---

### **Security Authentication Keychains**

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

### **Strict Versus Loose Authentication**

By default, strict authentication is enabled, and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

### **Example: Configuring BFD Authentication for Static Routes**

This example shows how to configure Bidirectional Forwarding Detection (BFD) authentication for static routes.

- [Requirements on page 3210](#)
- [Overview on page 3211](#)
- [Configuration on page 3211](#)
- [Verification on page 3215](#)

### **Requirements**

Junos OS Release 9.6 or later (Canda and United States version).

BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.



### Overview

You can configure authentication for BFD sessions running over IPv4 and IPv6 static routes. Routing instances and logical systems are also supported.

The following steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the static route.
2. Associate the authentication keychain with the static route.
3. Configure the related security authentication keychain. This must be configured on the main router.




**TIP:** We recommend that you specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

[edit]

```
user@host> set routing-options static route ipv4 bfd-liveness-detection
authentication loose-check
```

Figure 48 on page 3211 shows the sample network.

**Figure 48: Customer Routes Connected to a Service Provider**



### Configuration

#### CLI Quick Configuration

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

```
Device B set interfaces ge-1/2/0 unit 0 description B->D
Device B set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
```

```
set interfaces lo0 unit 57 family inet address 10.0.0.1/32
set interfaces lo0 unit 57 family inet address 10.0.0.2/32
set routing-options static route 192.168.47.0/24 next-hop 172.16.1.2
set routing-options static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
 1000
set routing-options static route 192.168.47.0/24 bfd-liveness-detection authentication
 key-chain bfd-kc4
set routing-options static route 192.168.47.0/24 bfd-liveness-detection authentication
 algorithm keyed-sha-1
set security authentication-key-chains key-chain bfd-kc4 key 5 secret
 "9JhZHmn6Ap0In/9ApOcSs24oaZikPft3wY24ZG.mz36AtOIeyMWxSrlKvM-dbs2a
 DkP5FtOIQFclev7N"
set security authentication-key-chains key-chain bfd-kc4 key 5 start-time
 "2011-1-1.12:00:00 -0800"
```

Device D

```
set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set interfaces lo0 unit 2 family inet address 192.168.47.5/32
set interfaces lo0 unit 2 family inet address 192.168.47.6/32
set routing-options static route 0.0.0.0/0 next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
set routing-options static route 0.0.0.0/0 bfd-liveness-detection authentication key-chain
 bfd-kc4
set routing-options static route 0.0.0.0/0 bfd-liveness-detection authentication algorithm
 keyed-sha-1
set security authentication-key-chains key-chain bfd-kc4 key 5 secret
 "9JhZHmn6Ap0In/9ApOcSs24oaZikPft3wY24ZG.mz36AtOIeyMWxSrlKvM-dbs2a
 DkP5FtOIQFclev7N"
set security authentication-key-chains key-chain bfd-kc4 key 5 start-time
 "2011-1-1.12:00:00 -0800"
```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure BFD for static routes:

1. On Device B, configure the interfaces.

```
[edit interfaces]
user@B# set ge-1/2/0 unit 0 description B->D
user@B# set ge-1/2/0 unit 0 family inet address 172.16.1.1/24

user@B# set lo0 unit 57 family inet address 10.0.0.1/32
user@B# set lo0 unit 57 family inet address 10.0.0.2/32
```

2. On Device B, create a static route and set the next-hop address.

```
[edit routing-options]
user@B# set static route 192.168.47.0/24 next-hop 172.16.1.2
```

3. On Device B, configure BFD for the static route.

```
[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
 1000
```

4. On Device B, specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on the static route.

```
[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection authentication
algorithm keyed-sha-1
```



**NOTE:** Nonstop active routing (NSR) is not supported with the **meticulous-keyed-md5** and **meticulous-keyed-sha-1** authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

5. On Device B, specify the keychain to be used to associate BFD sessions on the specified route with the unique security authentication keychain attributes.

This should match the keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

```
[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection authentication
key-chain bfd-kc4
```

6. On Device B, specify the unique security authentication information for BFD sessions:

- The matching keychain name as specified in Step 5.
- At least one key, a unique integer between **0** and **63**. Creating multiple keys allows multiple clients to use the BFD session.
- The secret data used to allow access to the session.
- The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security authentication-key-chains key-chain bfd-kc4]
user@B# set key 5 secret
"9JhZHmn6Ap0In/9ApOcSs24oaZikPft3wY24ZG.mz36AtOIEyMWxSrlKvM-dbs2a
DkP5Ft0IQFclev7N"
user@B# set key 5 start-time "2011-1-1.12:00:00 -0800"
```

7. If you are done configuring Device B, commit the configuration.

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

8. Repeat the configuration on Device D.

The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

### Results

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

```
Device B user@B# show interfaces
ge-1/2/0 {
 unit 0 {
 description B->D;
 family inet {
 address 172.16.1.1/24;
 }
 }
}
lo0 {
 unit 57 {
 family inet {
 address 10.0.0.1/32;
 address 10.0.0.2/32;
 }
 }
}

user@B# show routing-options
static {
 route 192.168.47.0/24 {
 next-hop 172.16.1.2;
 bfd-liveness-detection {
 minimum-interval 1000;
 authentication {
 key-chain bfd-kc4;
 algorithm keyed-sha-1;
 }
 }
 }
}

user@B# show security
authentication-key-chains {
 key-chain bfd-kc4 {
 key 5 {
 secret
 "9JhZHmn6Ap0ln/9ApOcSs24oaZikPft3wY24ZG.mz36AtOIeyMWxSrlKvM-dbs2a
 DkP5FtOIQFclev7N"; ## SECRET-DATA
 start-time "2011-1-1.12:00:00 -0800";
 }
 }
}
```

**Verification**

Confirm that the configuration is working properly.

- [Verifying That BFD Sessions Are Up on page 3215](#)
- [Viewing Details About the BFD Session on page 3215](#)
- [Viewing Extensive BFD Session Information on page 3215](#)

**Verifying That BFD Sessions Are Up**

**Purpose** Verify that the BFD sessions are up.

**Action** From operational mode, enter the `show bfd session` command.

```
user@B> show bfd session
```

| Address    | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|------------|-------|------------|-------------|-------------------|------------|
| 172.16.1.2 | Up    | ge-1/2/0.0 | 3.000       | 1.000             | 3          |

```
1 sessions, 1 clients
```

```
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

**Meaning** The command output shows that the BFD session is up.

**Viewing Details About the BFD Session**

**Purpose** View details about the BFD sessions and make sure that authentication is configured.

**Action** From operational mode, enter the `show bfd session detail` command.

```
user@B> show bfd session detail
```

| Address    | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|------------|-------|------------|-------------|-------------------|------------|
| 172.16.1.2 | Up    | ge-1/2/0.0 | 3.000       | 1.000             | 3          |

```
Client Static, TX interval 1.000, RX interval 1.000, Authenticate
Session up time 00:53:58
Local diagnostic NbrSignal, remote diagnostic None
Remote state Up, version 1
Logical system 9, routing table index 22
```

```
1 sessions, 1 clients
```

```
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

**Meaning** In the command output, **Authenticate** is displayed to indicate that BFD authentication is configured.

**Viewing Extensive BFD Session Information**

**Purpose** View more detailed information about the BFD sessions.

**Action** From operational mode, enter the `show bfd session extensive` command.

```
user@B> show bfd session extensive
```

| Address    | State | Interface  | Time  | Interval | Multiplier |
|------------|-------|------------|-------|----------|------------|
| 172.16.1.2 | Up    | ge-1/2/0.0 | 3.000 | 1.000    | 3          |

```
Client Static, TX interval 1.000, RX interval 1.000, Authenticate
```

```
keychain bfd-kc4, algo keyed-sha-1, mode strict
Session up time 01:39:45
Local diagnostic NbrSignal, remote diagnostic None
Remote state Up, version 1
Logical system 9, routing table index 22
Min async interval 1.000, min slow interval 1.000
Adaptive async TX interval 1.000, RX interval 1.000
Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
Local discriminator 3, remote discriminator 4
Echo mode disabled/inactive
Authentication enabled/active, keychain bfd-kc4, algo keyed-sha-1, mode strict

1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

**Meaning** In the command output, **Authenticate** is displayed to indicate that BFD authentication is configured. The output for the **extensive** command provides the keychain name, the authentication algorithm, and the mode for each client in the session.

**Related Documentation**

- [Examples: Configuring BFD for Static Routes on page 3194](#)

---

### Configuration Statements

- [\[edit routing-options\] Hierarchy Level on page 3216](#)

#### *[edit routing-options] Hierarchy Level*

Several statements in the **[edit routing-options]** hierarchy are valid at numerous locations within the hierarchy. To make the complete hierarchy easier to read, the repeated statements are listed in “[Common Routing Options](#)” on page 412 and that section is referenced at the appropriate locations in “[Complete \[edit routing-options\] Hierarchy](#)” on page 413.

- [Common Routing Options on page 3216](#)
- [Complete \[edit routing-options\] Hierarchy on page 3218](#)

#### *Common Routing Options*

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit routing-options\] Hierarchy](#)” on page 413 instead of the statements being repeated.

- **[edit routing-options aggregate defaults]**
- **[edit routing-options aggregate route *ip-prefix* </prefix-length>]**
- **[edit routing-options generate defaults]**
- **[edit routing-options generate route *ip-prefix* </prefix-length>]**
- **[edit routing-options static defaults]**
- **[edit routing-options static route *ip-prefix* </prefix-length>]**

The common routing options are as follows:

```
(active | passive);
as-path {
 aggregator as-number address;
 atomic-aggregate;
 origin (egp | igp | incomplete);
 path path-identifier;
}
color metric <type metric-type>;
color2 metric <type metric-type>;
community [community-id no-advertise no-export no-export-subconfed];
metric metric <type metric-type>;
metric2 metric <type metric-type>;
metric3 metric <type metric-type>;
metric4 metric <type metric-type>;
passive;
preference preference-value <type metric-type>;
preference2 preference-value <type metric-type>;
tag metric <type metric-type>;
tag2 metric <type metric-type>;
```

### ***Complete [edit routing-options] Hierarchy***

The statement hierarchy in this section can also be included at the [edit logical-systems *logical-system-name*] hierarchy level.

```
routing-options {
 access {
 route ip-prefix</prefix-length> {
 metric metric;
 next-hop [addresses];
 preference preference-value;
 qualified-next-hop address;
 tag route-tag;
 }
 }
 access-internal {
 route ip-prefix</prefix-length> {
 next-hop [addresses];
 qualified-next-hop address;
 tag route-tag;
 }
 }
 admin-groups-extended group-name {
 group-value group-identifier;
 }
 admin-groups-extended-range {
 maximum maximum-number;
 minimum minimum-number;
 }
 aggregate {
 defaults {
 ... statements in Common Routing Options on page 412 PLUS ...
 (brief | full);
 discard;
 }
 route ip-prefix</prefix-length> {
 ... statements in Common Routing Options on page 412 PLUS ...
 (brief | full);
 discard;
 policy [policy-names];
 }
 }
 auto-export {
 disable;
 family inet {
 disable;
 flow {
 disable;
 rib-group rib-group;
 }
 multicast {
 disable;
 rib-group rib-group;
 }
 }
 unicast {
 disable;
 }
 }
}
```



```

 rib-group rib-group;
 }
}
family inet6 {
 disable;
 multicast {
 disable;
 rib-group rib-group;
 }
 unicast {
 disable;
 rib-group rib-group;
 }
}
family iso {
 disable;
 unicast {
 disable;
 rib-group rib-group;
 }
}
}
traceoptions {
 file filename <files number> <size maximum-file-size> <world-readable |
 no-world-readable>;
 flag flag <flag-modifier> <disable>;
}
}
autonomous-system autonomous-system <asdot-notation> <loops number>;
no-bfd-triggered-local-repair;
bgp-orf-cisco-mode;
bmp {
 memory-limit bytes;
 station-address (ip-address | name);
 station-port-number port-number;
 statistics-timeout seconds;
}
confederation as-number members [as-numbers];
dynamic-tunnels tunnel-name {
 destination-networks prefix;
 gre;
 rsvp-te entry-name {
 destination-networks network-prefix;
 label-switched-path-template {
 default-template;
 template-name;
 }
 }
}
source-address address;
}
fate-sharing {
 group group-name {
 cost value;
 from {
 address <to address>;
 }
 }
}

```

```

}
flow {
 firewall-install-disable;
 route name {
 match {
 destination address;
 destination-port [afs bgp biff bootpc bootps cmd cvspserver dhcp domain eklogin
 ekshell exec finger ftp ftp-data http https ident imap kerberos-sec klogin kpasswd
 krb-prop krbupdate kshell ldap ldap login mobileip-agent mobilip-mn msdp
 netbios-dgm netbios-ns netbios-ssn nfsd nntp ntalk ntp pop3 pptp printer radacct
 radius rip rkinit smtp snmp snmptrap snpp socks ssh sunrpc syslog tacacs
 tacacs-ds talk telnet tftp timed who xdmcp];
 dscp [code-points];
 fragment [don't-fragment first-fragment is-fragment last-fragment
 not-a-fragment];
 icmp-code [communication-prohibited-by-filtering destination-host-prohibited
 destination-host-unknown fragmentation-needed host-precedence-violation
 host-unreachable host-unreachable-for-tos ip-header-bad network-unreachable
 network-unreachable-for-tos port-unreachable precedence-cutoff-in-effect
 protocol-unreachable redirect-for-host redirect-for-network
 redirect-for-tos-and-host redirect-for-tos-and-net required-option-missing
 source-host-isolated source-route-failed ttl-eq-zero-during-reassembly
 ttl-eq-zero-during-transit];
 icmp-type [echo-reply echo-request info-reply info-request mask-reply
 mask-request parameter-problem redirect router-advertisement router-solicit
 source-quench time-exceeded timestamp timestamp-reply unreachable];
 packet-length [values];
 port [... same values as for the preceding destination-port statement ...];
 protocol [ah esp gre icmp igmp ipip ospf pim rsvp sctp tcp udp];
 source address;
 source-port [... same values as for the preceding destination-port statement ...];
 tcp-flags [ack fin push rst syn urgent];
 }
 then {
 (accept | discard);
 community community-name;
 next-term;
 rate-limit value;
 routing-instance routing-instance-name;
 sample;
 }
 }
}
term order (legacy | standard);
validation {
 traceoptions {
 file filename <files number> <size maximum-file-size> <world-readable |
 no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
}
}
forwarding-table {
 chained-composite-next-hop {
 ingress {
 l3vpn {
 extended-space;

```

```

 }
 }
}
export [policy-name];
indexed-next-hop;
(indirect-next-hop | no-indirect-next-hop);
(indirect-next-hop-change-acknowledgements |
 no-indirect-next-hop-change-acknowledgements);
krt-nexthop-ack-timeout interval;
unicast-reverse-path (active-paths | feasible-paths);
}
generate {
 defaults {
 ... statements in Common Routing Options on page 412 PLUS ...
 (brief | full);
 discard;
 }
 route ip-prefix</prefix-length> {
 ... statements in Common Routing Options on page 412 PLUS ...
 (brief | full);
 discard;
 policy [policy-names];
 }
}
graceful-restart {
 disable;
 restart-duration seconds;
}
host-fast-reroute {
 global-arp-prefix-limit number;
 global-supplementary-blackout-timer minutes;
}
instance-export [policy-names];
instance-import [policy-names];
interface interface-name { # In the routing-instance only
 arp-prefix-limit number;
 link-protection;
 supplementary-blackout-timer minutes;
}
interface-routes {
 family (inet | inet6) {
 export {
 lan;
 point-to-point;
 }
 import [policy-names];
 }
 rib-group {
 inet group-name;
 inet6 group-name;
 }
}
martians {
 ip-prefix</prefix-length> (exact | longer | orlonger |
 prefix-length-range /minimum-prefix-length–/maximum-prefix-length |
 through ip-prefix</prefix-length> | upto /prefix-length> <allow>;

```

```

}
maximum-paths path-limit <log-only | threshold value> <log-interval seconds>;
maximum-prefixes prefix-limit <log-only | threshold value> <log-interval seconds>;
med-igp-update-interval minutes;
multicast {
 ... the multicast subhierarchy appears after the main [edit routing-options] hierarchy ...
}
nonstop-routing;
options {
 mark seconds;
 syslog {
 level level;
 upto level;
 }
}
ppm {
 no-delegate-processing;
}
resolution {
 rib routing-table-name {
 import [policy-names];
 resolution-ribs [routing-table-names];
 }
 tracefilter [filter-policy-names];
 traceoptions {
 file filename <files number> <size maximum-file-size> <world-readable |
 no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
}
rib routing-table-name {
 access {
 ... same statements as at the [edit routing-options access] hierarchy level ...
 }
 access-internal {
 ... same statements as at the [edit routing-options access-internal] hierarchy level ...
 }
 aggregate {
 ... same statements as at the [edit routing-options aggregate] hierarchy level ...
 }
 generate {
 ... same statements as at the [edit routing-options generate] hierarchy level ...
 }
 martians {
 ip-prefix</prefix-length> (exact | longer | orlonger |
 prefix-length-range /minimum-prefix-length–/maximum-prefix-length |
 through ip-prefix</prefix-length> | upto /prefix-length) <allow>;
 }
 maximum-paths path-limit <log-only | threshold value> <log-interval seconds>;
 maximum-prefixes prefix-limit <log-only | threshold value> <log-interval seconds>;
 static {
 ... same statements as at the [edit routing-options static] hierarchy level ...
 }
}
rib-groups {
 group-name {

```

```

 export-rib table-name;
 import-policy [policy-names];
 import-rib [table-names];
 }
}
route-distinguisher-id address;
route-record;
router-id address;
source-routing {
 ip;
 ipv6;
}
srlg {
 srlg-name {
 srlg-cost srlg-cost;
 srlg-value srlg-value;
 }
}
static {
 ... the static subhierarchy appears after the main [edit routing-options] hierarchy ...
}
topologies {
 family (inet | inet6) {
 topology topology-name;
 }
}
traceoptions {
 file filename <files number> <size maximum-file-size> <world-readable |
 no-world-readable>;
 flag flag <disable>;
}
validation {
 group group-name {
 max-sessions number;
 session address {
 hold-time seconds;
 local-address local-ip-address;
 port port-number;
 preference number;
 record-lifetime seconds;
 refresh-time seconds;
 }
 }
}
notification-rib value;
static {
 record destination {
 maximum-length prefix-length {
 origin-autonomous-system as-number {
 validation-state (invalid | valid);
 }
 }
 }
}
}
traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag;
}

```

```
 }
 }
}

routing-options {
 multicast {
 asm-override-ssm;
 backup-pe-group group-name {
 backups [addresses];
 local-address address;
 }
 flow-map flow-map-name {
 bandwidth <bps> <adaptive>;
 forwarding-cache {
 timeout (never <non-discard-entry-only> | minutes);
 }
 policy [policy-names];
 redundant-sources [addresses];
 }
 forwarding-cache {
 family (inet | inet6) {
 threshold {
 log-warning value;
 suppress value <reuse value>;
 }
 threshold {
 log-warning value;
 suppress value <reuse value>;
 }
 timeout minutes;
 }
 }
 interface interface-name {
 maximum-bandwidth bps;
 no-qos-adjust;
 reverse-oif-mapping {
 no-qos-adjust;
 }
 subscriber-leave-timer seconds;
 }
 }
 pim-to-igmp-proxy {
 upstream-interface [interface-names];
 }
 pim-to-mld-proxy {
 upstream-interface [interface-names];
 }
 rpf-check-policy [policy-names];
 scope scope-name {
 interface [interface-names];
 prefix ip-prefix </prefix-length>;
 }
 scope-policy [policy-names];
 ssm-groups [ip-prefix </prefix-length>];
 ssm-map ssm-map-name {
 policy [policy-names];
 source [addresses];
 }
}
```

```

 traceoptions {
 file filename <files number> <size maximum-file-size> <world-readable |
 no-world-readable>;
 flag flag <disable>;
 }
}

routing-options {
 static {
 defaults {
 ... statements in Common Routing Options on page 412 PLUS ...
 (install | no-install);
 (readvertise | no-readvertise);
 (resolve | no-resolve);
 (retain | no-retain);
 }
 rib-group group-name;
 route destination-prefix {
 ... statements in Common Routing Options on page 412 PLUS ...
 backup-pe-group group-name;
 bfd-liveness-detection {
 detection-time {
 threshold milliseconds;
 }
 holddown-interval milliseconds;
 local-address ip-address;
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 minimum-receive-ttl milliseconds;
 multiplier number;
 neighbor address;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 version (1 | automatic);
 }
 (discard | next-hop [addresses] | next-table address | receive | reject);
 (install | no-install);
 lsp-next-hop {
 metric metric;
 preference preference;
 }
 p2mp-lsp-next-hop lsp-name {
 metric metric;
 preference preference;
 }
 (readvertise | no-readvertise);
 (resolve | no-resolve);
 (retain | no-retain);
 static-lsp-next-hop lsp-name {
 metric metric;
 preference preference-value;
 }
 }
 }
}

```

```
}
}
}
```

**Related  
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*



## bfd

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> bfd {   traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;match <i>regular-expression</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable         no-world-readable&gt;;     flag <i>flag</i> &lt;<i>flag-modifier</i>&gt; &lt;disable&gt;;   } } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],</p> <p>[edit protocols],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>         | Configure trace options for Bidirectional Forwarding Protocol (BFD) traffic.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>             | If you do not include this statement, no BFD tracing operations are performed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the BFD tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the <b>/var/log</b> directory. We recommend that you place global routing protocol tracing output in the <b>routing-log</b> file.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 2 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. These are the BFD protocol tracing options:</p> <ul style="list-style-type: none"> <li>• <b>adjacency</b>—Trace adjacency messages.</li> <li>• <b>all</b>—Trace all options for BFD.</li> <li>• <b>error</b>—Trace all errors.</li> <li>• <b>event</b>—Trace all events.</li> <li>• <b>issu</b>—Trace in-service software upgrade (ISSU) packet activity.</li> </ul> |

- **nsr-packet**—Trace non-stop-routing (NSR) packet activity.
- **nsr-synchronization**—Trace NSR synchronization events.
- **packet**—Trace all packets.
- **pipe**—Trace pipe messages.
- **pipe-detail**—Trace pipe messages in detail.
- **ppm-packet**—Trace packet activity by periodic packet management (PPM).
- **state**—Trace state transitions.
- **timer**—Trace timer processing.

**match *regular-expression***—(Optional) Regular expression for lines to be logged.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size *size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the trace file again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing and trace—To view this statement in the configuration.                                                            |
|                                 | routing-control and trace-control—To add this statement to the configuration.                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring BFD for Static Routes on page 3197</a></li></ul> |

## bfd-liveness-detection (Routing Options Static Route)

**Syntax**

```

bfd-liveness-detection {
 authentication {
 algorithm algorithm-name;
 key-chain key-chain-name;
 loose-check;
 }
 detection-time {
 threshold milliseconds;
 }
 holddown-interval milliseconds;
 local-address ip-address;
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 minimum-receive-ttl number;
 multiplier number;
 neighbor address;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 version (1 | automatic);
}

```

**Hierarchy Level**

```

[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options rib routing-table-name static route destination-prefix],
[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options rib routing-table-name static route destination-prefix qualified-next-hop
 (interface-name | address)],
[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options static route destination-prefix],
[edit logical-systems logical-system-name routing-instances routing-instance-name
 routing-options static route destination-prefix qualified-next-hop (interface-name |
 address)],
[edit logical-systems logical-system-name routing-options rib routing-table-name static
 route destination-prefix],
[edit logical-systems logical-system-name routing-options rib routing-table-name static
 route destination-prefix qualified-next-hop (interface-name | address)],
[edit logical-systems logical-system-name routing-options static route destination-prefix],
[edit logical-systems logical-system-name routing-options static route destination-prefix
 qualified-next-hop (interface-name | address)],
[edit routing-instances routing-instance-name routing-options rib routing-table-name static
 route destination-prefix],
[edit routing-instances routing-instance-name routing-options rib routing-table-name static
 route destination-prefix qualified-next-hop (interface-name | address)],
[edit routing-instances routing-instance-name routing-options static route destination-prefix],
[edit routing-instances routing-instance-name routing-options static route destination-prefix
 qualified-next-hop (interface-name | address)],
[edit routing-options rib routing-table-name static route destination-prefix],
[edit routing-options rib routing-table-name static route destination-prefix qualified-next-hop
 (interface-name | address)],
[edit routing-options static route destination-prefix],

```

[edit routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)]

- Release Information** Statement introduced before Junos OS Release 7.4.  
**detection-time threshold** and **transmit-interval threshold** options introduced in Junos OS Release 8.2.  
**local-address** statement introduced in Junos OS Release 8.2.  
**minimum-receive-ttl** statement introduced in Junos OS Release 8.2.  
Support for logical routers introduced in Junos OS Release 8.3.  
**holddown-interval** statement introduced in Junos OS Release 8.5.  
**no-adaptation** statement introduced in Junos OS Release 9.0.  
Support for IPv6 static routes introduced in Junos OS Release 9.1.  
**authentication algorithm**, **authentication key-chain**, and **authentication loose-check** statements introduced in Junos OS Release 9.6.  
Statement introduced in Junos OS Release 12.1 for the QFX Series.  
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
- Description** Configure bidirectional failure detection timers and authentication criteria for static routes.

- Options** **authentication algorithm** *algorithm-name*—Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.
- authentication key-chain** *key-chain-name*—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.
- authentication loose-check**—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.
- detection-time threshold** *milliseconds*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.
- holddown-interval** *milliseconds*—Configure an interval specifying how long a BFD session must remain up before a state change notification is sent. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.  
**Range:** 0 through 255,000  
**Default:** 0
- local-address** *ip-address*—Enable a multihop BFD session and configure the source address for the BFD session.
- minimum-interval** *milliseconds*—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval** **minimum-interval** and **minimum-receive-interval** statements.  
**Range:** 1 through 255,000
- minimum-receive-interval** *milliseconds*—Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.  
**Range:** 1 through 255,000
- minimum-receive-ttl** *number*—Configure the time to live (TTL) for the multihop BFD session.  
**Range:** 1 through 255  
**Default:** 255
- multiplier** *number*—Configure number of hello packets not received by the neighbor that causes the originating interface to be declared down.  
**Range:** 1 through 255  
**Default:** 3

**neighbor address**—Configure a next-hop address for the BFD session for a next hop specified as an interface name.

**no-adaptation**—Specify for BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

**transmit-interval threshold milliseconds**—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

**Range:** 0 through 4,294,967,295

**transmit-interval minimum-interval milliseconds**—Configure the minimum interval at which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.

**Range:** 1 through 255,000

**version**—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

**Default:** automatic

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

**Related Documentation**

- [Example: Configuring BFD for Static Routes on page 3197](#)
- [Example: Configuring BFD Authentication for Static Routes on page 3210](#)

## Administration

- [Operational Commands on page 3232](#)

### Operational Commands

- [BFD Operational Mode Commands on page 3232](#)



#### BFD Operational Mode Commands

[Table 262 on page 3232](#) summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot Bidirectional Forwarding Detection (BFD) sessions. Commands are listed in alphabetical order.

**Table 262: BFD Operational Mode Commands**

| Task                  | Command                           |
|-----------------------|-----------------------------------|
| Clear BFD parameters. | <code>clear bfd adaptation</code> |
| Clear BFD sessions.   | <code>clear bfd session</code>    |

Table 262: BFD Operational Mode Commands (*continued*)

| Task                                                                                                                                                                                                                                          | Command                       |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------|
| Display BFD session statistics.                                                                                                                                                                                                               | <code>show bfd session</code> |
| <div><div></div><div><p><b>NOTE:</b> The protocol client for which the BFD session is active can be either IS-IS or OSPF.</p></div></div>                    |                               |
| <div><div></div><div><p><b>NOTE:</b> For information about how to configure BFD, see the <i>Junos Routing Protocols Configuration Guide</i>.</p></div></div> |                               |

## clear bfd adaptation

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear bfd adaptation</code><br><code>&lt;address session-address&gt;</code><br><code>&lt;discriminator discr-number&gt;</code>                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Clear adaptation for Bidirectional Forwarding Detection (BFD) sessions. BFD is a simple hello mechanism that detects failures in a network. Configured BFD interval timers can change, adapting to network situations. Use this command to return BFD interval timers to their configured values.</p> <p>The <b>clear bfd adaptation</b> command is hitless, meaning that the command does not affect traffic flow on the routing device.</p> |
| <b>Options</b>                  | <p><b>none</b>—Clear adaptation for all BFD sessions.</p> <p><b>address session-address</b>—(Optional) Clear adaptation for all BFD sessions matching the specified address.</p> <p><b>discriminator discr-number</b>—(Optional) Clear adaptation for the local BFD session matching the specified discriminator.</p>                                                                                                                            |
| <b>Additional Information</b>   | For more information, see the description of the <b>bfd-liveness-detection</b> configuration statement in the <i>Junos Routing Protocols Configuration Guide</i> .                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>    | <a href="#">clear bfd adaptation on page 3234</a>                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                            |

## Sample Output

### clear bfd adaptation

```
user@host> clear bfd adaptation
```



## clear bfd session

|                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                   | clear bfd session<br><address <i>session-address</i> ><br><discriminator <i>discr-number</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                          |
| <b>Syntax (EX Series Switch and QFX Series)</b> | clear bfd session<br><address <i>session-address</i> ><br><discriminator <i>discr-number</i> >                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>                      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>                              | Drop one or more Bidirectional Forwarding Detection (BFD) sessions.                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                  | <p><b>none</b>—Drop all BFD sessions.</p> <p><b>address <i>session-address</i></b>—(Optional) Drop all BFD sessions matching the specified address.</p> <p><b>discriminator <i>discr-number</i></b>—(Optional) Drop the local BFD session matching the specified discriminator.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                 | clear                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>                    | <ul style="list-style-type: none"> <li>• <a href="#">show bfd session on page 3236</a></li> </ul>                                                                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>                    | <a href="#">clear bfd session on page 3235</a>                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>                            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                           |

## Sample Output

### clear bfd session

```
user@host> clear bfd session
```

```
show bfd session
```

|                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                                   | show bfd session<br><brief   detail   extensive   summary><br><address <i>address</i> ><br><discriminator <i>discriminator</i> ><br><logical-system (all   <i>logical-system-name</i> )><br><prefix <i>address</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Syntax (EX Series Switch and QFX Series) | show bfd session<br><brief   detail   extensive   summary><br><address <i>address</i> ><br><discriminator <i>discriminator</i> ><br><prefix <i>address</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Release Information                      | Command introduced before Junos OS Release 7.4.<br>Options <b>discriminator</b> and <b>address</b> introduced in Junos OS Release 8.2.<br>Option <b>prefix</b> introduced in Junos OS Release 9.0.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Description                              | Display information about active Bidirectional Forwarding Detection (BFD) sessions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Options                                  | <b>none</b> —(Same as <b>brief</b> ) Display information about active BFD sessions.<br><b>brief   detail   extensive   summary</b> —(Optional) Display the specified level of output.<br><b>address <i>address</i></b> —(Optional) Display information about the BFD session for the specified neighbor address.<br><b>discriminator <i>discriminator</i></b> —(Optional) Display information about the BFD session using the specified local discriminator.<br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><b>prefix <i>address</i></b> —(Optional) Display information about all of the BFD sessions for the specified LDP forwarding equivalence class (FEC). |
| Required Privilege Level                 | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Related Documentation                    | <ul style="list-style-type: none"><li>• <a href="#">clear bfd session on page 3235</a></li><li>• <a href="#">Examples: Configuring BFD for Static Routes on page 3194</a></li><li>• <i>Example: Configuring BFD for OSPF</i></li><li>• <i>Example: Configuring BFD for BGP</i></li><li>• <a href="#">Configuring PIM and the Bidirectional Forwarding Detection (BFD) Protocol on page 3725</a></li><li>• <i>Example: Configuring BFD for IS-IS</i></li></ul>                                                                                                                                                                                                                                                                                                                 |

**List of Sample Output** [show bfd session on page 3240](#)  
[show bfd session brief on page 3240](#)  
[show bfd session detail on page 3240](#)  
[show bfd session detail \(with Authentication\) on page 3241](#)  
[show bfd session address extensive on page 3241](#)  
[show bfd session extensive on page 3241](#)  
[show bfd session extensive \(with Authentication\) on page 3242](#)  
[show bfd session summary on page 3242](#)

**Output Fields** [Table 263 on page 3237](#) describes the output fields for the **show bfd session** command. Output fields are listed in the approximate order in which they appear.

**Table 263: show bfd session Output Fields**

| Field Name               | Field Description                                                                                                                                                                                                                                                                               | Level of Output                       |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|
| <b>Address</b>           | Address on which the BFD session is active.                                                                                                                                                                                                                                                     | <b>brief detail extensive</b><br>none |
| <b>State</b>             | State of the BFD session: <b>Up</b> , <b>Down</b> , <b>Init</b> (initializing), or <b>Failing</b> .                                                                                                                                                                                             | <b>brief detail extensive</b><br>none |
| <b>Interface</b>         | Interface on which the BFD session is active.                                                                                                                                                                                                                                                   | <b>brief detail extensive</b><br>none |
| <b>Detect Time</b>       | Negotiated time interval, in seconds, used to detect BFD control packets.                                                                                                                                                                                                                       | <b>brief detail extensive</b><br>none |
| <b>Transmit Interval</b> | Time interval, in seconds, used by the transmitting system to send BFD control packets.                                                                                                                                                                                                         | <b>brief detail extensive</b><br>none |
| <b>Multiplier</b>        | Negotiated multiplier by which the time interval is multiplied to determine the detection time for the transmitting system.                                                                                                                                                                     | <b>detail extensive</b>               |
| <b>Session up time</b>   | How long a BFD session has been established.                                                                                                                                                                                                                                                    | <b>detail extensive</b>               |
| <b>Client</b>            | Protocol for which the BFD session is active: <b>ISIS</b> , <b>OSPF</b> , or <b>Static</b> .                                                                                                                                                                                                    | <b>detail extensive</b>               |
| <b>TX interval</b>       | Time interval, in seconds, used by the host system to transmit BFD control packets.                                                                                                                                                                                                             | <b>brief detail extensive</b><br>none |
| <b>RX interval</b>       | Time interval, in seconds, used by the host system to receive BFD control packets.                                                                                                                                                                                                              | <b>brief detail extensive</b><br>none |
| <b>Authenticate</b>      | Indicates that BFD authentication is configured.                                                                                                                                                                                                                                                | <b>detail extensive</b>               |
| <b>keychain</b>          | Name of the security authentication keychain being used by a specific client.<br><br>BFD authentication information for a client is provided in a single line and includes the <b>keychain</b> , <b>algo</b> , and <b>mode</b> parameters. Multiple clients can be configured on a BFD session. | <b>extensive</b>                      |

Table 263: show bfd session Output Fields (*continued*)

| Field Name                        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output         |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>algo</b>                       | <p>BFD authentication algorithm being used for a specific client: <b>keyed-md5</b>, <b>keyed-sha-1</b>, <b>meticulous-keyed-md5</b>, <b>meticulous-keyed-sha-1</b>, or <b>simple-password</b>.</p> <p>BFD authentication information for a client is provided in a single line and includes the <b>keychain</b>, <b>algo</b>, and <b>mode</b> parameters. Multiple clients can be configured on a BFD session.</p>                                                                                                           | <b>extensive</b>        |
| <b>mode</b>                       | <p>Level of BFD authentication enforcement being used by a specific client: <b>strict</b> or <b>loose</b>. Strict enforcement indicates that authentication is configured at both ends of the session (the default). Loose enforcement indicates that one end of the session might not be authenticated.</p> <p>BFD authentication information for a client is provided in a single line and includes the <b>keychain</b>, <b>algo</b>, and <b>mode</b> parameters. Multiple clients can be configured on a BFD session.</p> | <b>extensive</b>        |
| <b>Local diagnostic</b>           | Local diagnostic information about failing BFD sessions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b> |
| <b>Remote diagnostic</b>          | Remote diagnostic information about failing BFD sessions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b> |
| <b>Remote state</b>               | Reports whether the remote system's BFD packets have been received and whether the remote system is receiving transmitted control packets.                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b> |
| <b>Version</b>                    | BFD version: <b>0</b> or <b>1</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>extensive</b>        |
| <b>Replicated</b>                 | The <b>replicated</b> flag appears when nonstop routing or graceful Routing Engine switchover is configured and the BFD session has been replicated to the backup Routing Engine.                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Min async interval</b>         | Minimum amount of time, in seconds, between asynchronous control packet transmissions across the BFD session.                                                                                                                                                                                                                                                                                                                                                                                                                | <b>extensive</b>        |
| <b>Min slow interval</b>          | Minimum amount of time, in seconds, between synchronous control packet transmissions across the BFD session.                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>extensive</b>        |
| <b>Adaptive async TX interval</b> | Transmission interval being used because of adaptation.                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>extensive</b>        |
| <b>RX interval</b>                | Minimum required receive interval.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>extensive</b>        |
| <b>Local min TX interval</b>      | Minimum amount of time, in seconds, between control packet transmissions on the local system.                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>extensive</b>        |
| <b>Local min RX interval</b>      | Minimum amount of time, in seconds, between control packet detections on the local system.                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>extensive</b>        |
| <b>Remote min TX interval</b>     | Minimum amount of time, in seconds, between control packet transmissions on the remote system.                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>extensive</b>        |
| <b>Remote min RX interval</b>     | Minimum amount of time, in seconds, between control packet detections on the remote system.                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>extensive</b>        |

Table 263: show bfd session Output Fields (*continued*)

| Field Name                          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output  |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Threshold transmission interval     | Threshold for notification if the transmission interval increases.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | extensive        |
| Threshold for detection time        | Threshold for notification if the detection time increases.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | extensive        |
| Local discriminator                 | Authentication code used by the local system to identify that BFD session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | extensive        |
| Remote discriminator                | Authentication code used by the remote system to identify that BFD session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | extensive        |
| Echo mode                           | Information about the state of echo transmissions on the BFD session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | extensive        |
| Prefix                              | LDP FEC address associated with the BFD session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | All levels       |
| Egress, Destination                 | Displays the LDP FEC destination address. This field is displayed only on a router at the egress of an LDP FEC, where the BFD session has an LDP Operation, Administration, and Maintenance (OAM) client.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | All levels       |
| Remote is control-plane independent | <p>The BFD session on the remote peer is running on its Packet Forwarding Engine. In this case, when the remote node undergoes a graceful restart, the local peer can help the remote peer with the graceful restart.</p> <p>The following BFD sessions are not distributed to the Packet Forwarding Engine: multihop sessions, tunnel-encapsulated sessions, and sessions over aggregated Ethernet and integrated routing and bridging (IRB) interfaces.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | extensive        |
| Authentication                      | <p>Summary status of BFD authentication:</p> <ul style="list-style-type: none"> <li>• <b>status</b>—<b>enabled/active</b> indicates authentication is configured and active. <b>enabled/inactive</b> indicates authentication is configured but not active. This only occurs when the remote end of the session does not support authentication and loose checking is configured.</li> <li>• <b>keychain</b>—Name of the security authentication keychain associated with the specified BFD session.</li> <li>• <b>algo</b>—BFD authentication algorithm being used: <b>keyed-md5</b>, <b>keyed-sha-1</b>, <b>meticulous-keyed-md5</b>, <b>meticulous-keyed-sha-1</b>, or <b>simple-password</b>.</li> <li>• <b>mode</b>—Level of BFD authentication enforcement: <b>strict</b> or <b>loose</b>. Strict enforcement indicates authentication is configured at both ends of the session (the default). Loose enforcement indicates that one end of the session might not be authenticated.</li> </ul> <p>This information is only shown if BFD authentication is configured.</p> | extensive        |
| Session ID                          | The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | detail extensive |
| sessions                            | Total number of active BFD sessions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | All levels       |
| clients                             | Total number of clients that are hosting active BFD sessions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | All levels       |

Table 263: show bfd session Output Fields (*continued*)

| Field Name                      | Field Description                                                                  | Level of Output  |
|---------------------------------|------------------------------------------------------------------------------------|------------------|
| <b>Cumulative transmit rate</b> | Total number of BFD control packets transmitted per second on all active sessions. | All levels       |
| <b>Cumulative receive rate</b>  | Total number of BFD control packets received per second on all active sessions.    | All levels       |
| <b>Multi-hop, min-recv-TTL</b>  | Minimum time to live (TTL) accepted if the session is configured for multihop.     | <b>extensive</b> |
| <b>route table</b>              | Route table used if the session is configured for multihop.                        | <b>extensive</b> |
| <b>local address</b>            | Local address of the source used if the session is configured for multihop.        | <b>extensive</b> |

## Sample Output

### show bfd session

```

user@host> show bfd session

Address State Interface Detect Time Transmit
10.9.1.33 Up so-7/1/0.0 0.600 Interval 0.200 Multiplier 3
10.9.1.29 Up ge-4/0/0.0 0.600 Interval 0.200 Multiplier 3

2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

```

### show bfd session brief

The output for the **show bfd session brief** command is identical to that for the **show bfd session** command. For sample output, see [show bfd session on page 3240](#).

### show bfd session detail

```

user@host> show bfd session detail

Address State Interface Detect Time Transmit
10.9.1.33 Up so-7/1/0.0 0.600 Interval 0.200 Multiplier 3
Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3
Session up time 3d 00:34
Local diagnostic None, remote diagnostic None
Remote state Up, version 1
Replicated
10.9.1.29 Up ge-4/0/0.0 0.600 Interval 0.200 Multiplier 3
Client ISIS L2, TX interval 0.200, RX interval 0.200, multiplier 3
Session up time 3d 00:29, previous down time 00:00:01
Local diagnostic NbrSignal, remote diagnostic AdminDown
Remote state Up, version 1

2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

```

**show bfd session detail (with Authentication)**

```

user@host> show bfd session detail

```

| Address                                                                              | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|--------------------------------------------------------------------------------------|-------|------------|-------------|-------------------|------------|
| 10.9.1.33                                                                            | Up    | so-7/1/0.0 | 0.600       | 0.200             | 3          |
| Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3, <b>Authenticate</b> |       |            |             |                   |            |
| Session up time 3d 00:34                                                             |       |            |             |                   |            |
| Local diagnostic None, remote diagnostic None                                        |       |            |             |                   |            |
| Remote state Up, version 1                                                           |       |            |             |                   |            |
| Replicated                                                                           |       |            |             |                   |            |
| 10.9.1.29                                                                            | Up    | ge-4/0/0.0 | 0.600       | 0.200             | 3          |
| Client ISIS L2, TX interval 0.200, RX interval 0.200, multiplier 3                   |       |            |             |                   |            |
| Session up time 3d 00:29, previous down time 00:00:01                                |       |            |             |                   |            |
| Local diagnostic NbrSignal, remote diagnostic AdminDown                              |       |            |             |                   |            |
| Remote state Up, version 1                                                           |       |            |             |                   |            |

2 sessions, 2 clients  
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

**show bfd session address extensive**

```

user@host> show bfd session 10.255.245.212 extensive

```

| Address                                                                   | State | Interface | Detect Time | Transmit Interval | Multiplier |
|---------------------------------------------------------------------------|-------|-----------|-------------|-------------------|------------|
| 10.255.245.212                                                            | Up    |           | 1.200       | 0.400             | 3          |
| Client Static, TX interval 0.400, RX interval 0.400, multiplier 3         |       |           |             |                   |            |
| Session up time 00:17:03, previous down time 00:00:14                     |       |           |             |                   |            |
| Local diagnostic CtlExpire, remote diagnostic NbrSignal                   |       |           |             |                   |            |
| Remote state Up, version 1                                                |       |           |             |                   |            |
| Replicated                                                                |       |           |             |                   |            |
| Min async interval 0.400, min slow interval 1.000                         |       |           |             |                   |            |
| Adaptive async tx interval 0.400, rx interval 0.400                       |       |           |             |                   |            |
| Local min tx interval 0.400, min rx interval 0.400, multiplier 3          |       |           |             |                   |            |
| Remote min tx interval 0.400, min rx interval 0.400, multiplier 3         |       |           |             |                   |            |
| Threshold transmission interval 0.000, Threshold for detection time 0.000 |       |           |             |                   |            |
| Local discriminator 6, remote discriminator 16                            |       |           |             |                   |            |
| Echo mode disabled/inactive                                               |       |           |             |                   |            |
| Multi-hop, min-recv-TTL 255, route-table 0, local-address 10.255.245.205  |       |           |             |                   |            |

1 sessions, 1 clients  
Cumulative transmit rate 2.5 pps, cumulative receive rate 2.5 pps

**show bfd session extensive**

```

user@host> show bfd session extensive

```

| Address                                                                      | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|------------------------------------------------------------------------------|-------|------------|-------------|-------------------|------------|
| 10.31.1.2                                                                    | Up    | ge-2/1/8.0 | 0.030       | 0.010             | 3          |
| Client OSPF realm ospf-v2 Area 0.0.0.0, TX interval 0.010, RX interval 0.010 |       |            |             |                   |            |
| Session up time 00:10:13                                                     |       |            |             |                   |            |
| Local diagnostic None, remote diagnostic None                                |       |            |             |                   |            |
| Remote state Up, version 1                                                   |       |            |             |                   |            |
| Replicated                                                                   |       |            |             |                   |            |
| Min async interval 0.010, min slow interval 1.000                            |       |            |             |                   |            |
| Adaptive async TX interval 0.010, RX interval 0.010                          |       |            |             |                   |            |
| Local min TX interval 0.010, minimum RX interval 0.010, multiplier 3         |       |            |             |                   |            |
| Remote min TX interval 0.010, min RX interval 0.010, multiplier 3            |       |            |             |                   |            |
| Local discriminator 12, remote discriminator 4                               |       |            |             |                   |            |
| Echo mode disabled/inactive                                                  |       |            |             |                   |            |
| Remote is control-plane independent                                          |       |            |             |                   |            |
| Session ID: 0x201                                                            |       |            |             |                   |            |

| Address   | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|-----------|-------|------------|-------------|-------------------|------------|
| 10.31.2.2 | Up    | ge-2/1/4.0 | 0.030       | 0.010             | 3          |

Client OSPF realm ospf-v2 Area 0.0.0.0, TX interval 0.010, RX interval 0.010  
Session up time 00:10:14  
Local diagnostic None, remote diagnostic NbrSignal  
Remote state Up, version 1  
Replicated  
Min async interval 0.010, min slow interval 1.000  
Adaptive async TX interval 0.010, RX interval 0.010  
Local min TX interval 0.010, minimum RX interval 0.010, multiplier 3  
Remote min TX interval 0.010, min RX interval 0.010, multiplier 3  
Local discriminator 13, remote discriminator 5  
Echo mode disabled/inactive  
Remote is control-plane independent  
Session ID: 0x202

2 sessions, 2 clients  
Cumulative transmit rate 200.0 pps, cumulative receive rate 200.0 pps

#### show bfd session extensive (with Authentication)

```
user@host> show bfd session extensive
```

| Address        | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|----------------|-------|------------|-------------|-------------------|------------|
| 192.168.208.26 | Up    | so-1/0/0.0 | 2.400       | 0.800             | 10         |

Client Static, TX interval 0.600, RX interval 0.600, **Authenticate**  
**keychain bfd, algo keyed-md5, mode loose**  
Session up time 00:18:07  
Local diagnostic None, remote diagnostic NbrSignal  
Remote state Up, version 1  
Replicated  
Min async interval 0.600, min slow interval 1.000  
Adaptive async TX interval 0.600, RX interval 0.600  
Local min TX interval 0.600, minimum RX interval 0.600, multiplier 10  
Remote min TX interval 0.800, min RX interval 0.800, multiplier 3  
Local discriminator 2, remote discriminator 3  
Echo mode disabled/inactive  
**Authentication enabled/active, keychain bfd, algo keyed-md5, mode loose**

1 sessions, 1 clients  
Cumulative transmit rate 1.2 pps, cumulative receive rate 1.2 pps

#### show bfd session summary

```
user@host> show bfd session summary
```

2 sessions, 2 clients  
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps



## CHAPTER 12

# LLDP

- [Overview on page 3243](#)
- [Configuration on page 3244](#)
- [Administration on page 3257](#)

## Overview

---

- [LLDP on page 3243](#)

## LLDP

- [LLDP Overview on page 3243](#)

### LLDP Overview

---

The Link Layer Discovery Protocol (LLDP) is an industry-standard, vendor-neutral method to allow networked devices to advertise capabilities, identity, and other information onto a LAN. The Layer 2 protocol, detailed in IEEE 802.1AB-2005, replaces several proprietary protocols implemented by individual vendors for their equipment.

LLDP allows network devices that operate at the lower layers of a protocol stack (such as Layer 2 bridges and switches) to learn some of the capabilities and characteristics of LAN devices available to higher layer protocols, such as IP addresses. The information gathered through LLDP operation is stored in a network device and is queried with SNMP. Topology information can also be gathered from this database.

Some of the information that can be gathered by LLDP (only minimal information is mandatory) is:

- System name and description
- Port name and description
- VLAN name and identifier
- IP network management address
- Capabilities of the device (for example, switch, router, or server)
- MAC address and physical layer information

- Power information
- Link aggregation information



**NOTE:** LLDP media endpoint discovery (LLDP-MED) is not supported on T Series routers.

LLDP frames are sent at fixed intervals on each port that runs LLDP. LLDP protocol data units (LLDP PDUs) are sent inside Ethernet frames and identified by their destination Media Access Control (MAC) address (**01:80:C2:00:00:0E**) and Ethertype (**0x88CC**). Mandatory information supplied by LLDP is chassis ID, port ID, and a time-to-live value for this information.

LLDP is a powerful way to allow Layer 2 devices to gather details about other network-attached devices.

#### Related Documentation

- [Configuring LLDP on page 3246](#)
- [Tracing LLDP Operations on page 3257](#)
- [Example: Configuring LLDP on page 3244](#)
- *LLDP Operational Mode Commands*

---

## Configuration

- [Configuration Examples on page 3244](#)
- [Configuration Tasks on page 3245](#)
- [Configuration Statements on page 3249](#)

### Configuration Examples

- [Example: Configuring LLDP on page 3244](#)

---

#### Example: Configuring LLDP

The following example configures LLDP on interface **ge-1/1/1** but disables LLDP on all other interfaces, explicitly configures the default values for all automatically enabled features, and configures a value of 30 seconds for the LLDP configuration notification interval and a value of 30 seconds for the physical topology trap interval.

```
[edit]
protocols {
 lldp {
 advertisement-interval 30;
 hold-multiplier 4;
 interface all {
 disable;
 }
 interface ge-1/1/1;
 lldp-configuration-notification-interval 30;
 ptopo-configuration-maximum-hold-time 300;
```

```

 ptopo-configuration-trap-interval 30;
 transmit-delay 2;
 }
}

```

You verify operation of LLDP with several show commands:

- **show lldp <detail>**
- **show lldp neighbors *interface-name***
- **show lldp statistics *interface-name***
- **show lldp local-information**
- **show lldp remote-global-statistics**

You can clear LLDP neighbor information or statistics globally or on an interface:

- **clear lldp neighbors *interface-name***
- **clear lldp statistics *interface-name***

You can display basic information about LLDP with the **show lldp detail** command:

```

user@host> show lldp detail
LLDP : Enabled
Advertisement interval : 30 Second(s)
Transmit delay : 2 Second(s)
Hold timer : 4 Second(s)
Notification interval : 30 Second(s)
Config Trap Interval : 300 Second(s)
Connection Hold timer : 60 Second(s)

```

```

Interface LLDP Neighbor count
ge-1/1/1 Enabled 0

```

LLDP basic TLVs supported:

Chassis identifier, Port identifier, Port description, System name, System description, System capabilities, Management address.

LLDP 802 TLVs supported:

Link aggregation, Maximum frame size, MAC/PHY Configuration/Status, Port VLAN ID, Port VLAN name.

For more details about the output of these commands, see the *Junos OS Operational Mode Commands*.

#### Related Documentation

- [LLDP Overview on page 3243](#)
- [Configuring LLDP on page 3246](#)
- [Tracing LLDP Operations on page 3257](#)
- [LLDP Operational Mode Commands](#)

## Configuration Tasks

- [Configuring LLDP on page 3246](#)

## Configuring LLDP

---

You configure LLDP by including the **lldp** statement and associated parameters at the **[edit protocols]** hierarchy level. The complete set of LLDP statements follows:

```
lldp {
 advertisement-interval seconds;
 disable;
 hold-multiplier number;
 interface (all | interface-name) {
 disable;
 }
 lldp-configuration-notification-interval seconds;
 port-id-subtype {
 interface-name;
 locally-assigned;
 }
 ptopo-configuration-maximum-hold-time seconds;
 ptopo-configuration-trap-interval seconds;
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
 transmit-delay seconds
}
```

The following statements have default values:

- **advertisement-interval**—The default value is 30 seconds. The allowable range is from 5 through 32768 seconds.
- **hold-multiplier**—The default values is 4. The allowable range is from 2 through 10.
- **ptopo-configuration-maximum-hold-time**—The default value is 300 seconds. The allowable range is from 1 through 2147483647 seconds.
- **transmit-delay**—The default values is 2 seconds. The allowable range is from 1 through 8192 seconds.

The following statements must be explicitly configured:

- **lldp-configuration-notification-interval**—The allowable range is from 0 through 3600 seconds. There is no default value.
- **ptopo-configuration-trap-interval**—The allowable range is from 1 through 2147483647 seconds. There is no default value.

To disable LLDP on all or a particular interface, include the **interfaces** statement at the **[edit protocols lldp]** hierarchy level:

```
interface (all | interface-name) {
 disable;
}
```

To disable LLDP on all interfaces, use the **all** option. To disable LLDP on a particular interface, include the **disable** statement with the interface name.

To configure LLDP on a T Series router within a TX Matrix, you must specify the interface name in the LLDP configuration for the TX Matrix. For information about interface names for TX Matrix routers, see *TX Matrix Router Chassis and Interface Names*. For information about FPC numbering, see *Routing Matrix with a TX Matrix Router FPC Numbering*.



**NOTE:** The interface-name must be the physical interface (for example, **ge-1/0/0**) and not a logical interface (unit).

The advertisement interval determines the frequency that an LLDP interface sends LLDP advertisement frames. The default value is 30 seconds. The allowable range is from 5 through 32768 seconds. You adjust this parameter by including the **advertisement-interval** statement at the **[edit protocols lldp]** hierarchy level.

The hold multiplier determines the multiplier to apply to the advertisement interval. The resulting value in seconds is used to cache learned LLDP information before discard. The default value is 4. When used with the default advertisement interval value of 30 seconds, this makes the default cache lifetime 120 seconds. The allowable range of the hold multiplier is from 2 through 10. You adjust this parameter by including the **hold-multiplier** statement at the **[edit protocols lldp]** hierarchy level.

The transmit delay determines the delay between any two consecutive LLDP advertisement frames. The default value is 2 seconds. The allowable range is from 1 through 8192 seconds. You adjust this parameter by including the **transmit-delay** statement at the **[edit protocols lldp]** hierarchy level.

The physical topology configuration maximum hold time determines the time interval for which an agent device maintains physical topology database entries. The default value is 300 seconds. The allowable range is from 1 through 2147483647 seconds. You adjust this parameter by including the **ptopo-configuration-maximum-hold-time** statement at the **[edit protocols lldp]** hierarchy level.

The LLDP configuration notification interval determines the period for which trap notifications are sent to the SNMP Master Agent when changes occur in the database of LLDP information. This capability is disabled by default. The allowable range is from 0 (disabled) through 3600 seconds. You adjust this parameter by including the **lldp-configuration-notification-interval** statement at the **[edit protocols lldp]** hierarchy level.

The physical topology configuration trap interval determines the period for which trap notifications are sent to the SNMP Master Agent when changes occur in the global physical topology statistics. This capability is disabled by default. The allowable range is from 0 (disabled) through 3600 seconds. The LLDP agent sends traps to the SNMP Master Agent if this interval has a value greater than 0 and there is any change during the **lldp-configuration-notification-interval** trap interval. You adjust this parameter by including the **ptopo-configuration-trap-interval** statement at the **[edit protocols lldp]** hierarchy level.

By default, LLDP generates the SNMP index of the interface for the port ID Type, Length, and Value (TLV). Starting with Junos OS Release 12.3R1, you can generate the interface

name as the port ID TLV by including the **interface-name** statement at the **[edit protocols lldp port-id-subtype]** hierarchy level. When configure the **interface-name** statement on the remote LLDP neighbor, the **show lldp neighbors** command displays the interface name in the **Port ID** field rather than the SNMP index of the interface, which is displayed by default. If you change the default behavior of generating the SNMP index of the interface as the Port ID TLV, you can reenable the default behavior by including the **locally-assigned** statement at the **[edit protocols lldp port-id-subtype]** hierarchy level.

**Related  
Documentation**

- [LLDP Overview on page 3243](#)
- [Tracing LLDP Operations on page 3257](#)
- [Example: Configuring LLDP on page 3244](#)
- *LLDP Operational Mode Commands*
- *TX Matrix Router Chassis and Interface Names*
- *Miscellaneous Commands for a Routing Matrix with a TX Matrix Router*


## Configuration Statements

### lldp

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>lldp {   advertisement-interval <i>seconds</i>;   disable;   hold-multiplier <i>number</i>;   interface (all   <i>interface-name</i>) {     disable;   }   lldp-configuration-notification-interval <i>seconds</i>;   port-id-subtype {     interface-name;     locally-assigned;   }   ptopo-configuration-maximum-hold-time <i>seconds</i>;   ptopo-configuration-trap-interval <i>seconds</i>;   traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>maximum-file-size</i>&gt; &lt;world-readable         no-world-readable&gt;;     flag <i>flag</i> &lt;disable&gt;;   } }</pre> |
| <b>Hierarchy Level</b>          | [edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | (MX Series and T Series routers and EX Series switches only) Specify LLDP configuration parameters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring LLDP on page 3246</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## advertisement-interval

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>advertisement-interval seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">lldp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">lldp</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for MX Series and T Series routers.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>For MX Series and T Series routers and EX Series switches, configure an interval for LLDP advertisement.</p> <p>For switches configured for Link Layer Discovery Protocol, configure the frequency at which LLDP advertisements are sent.</p> <p>The <b>advertisement-interval</b> value must be greater than or equal to four times the <b>transmit-delay</b> value, or an error will be returned when you attempt to commit the configuration.</p> <div><p><b>NOTE:</b> The default value of <b>transmit-delay</b> is 2 seconds. If you configure the <b>advertisement-interval</b> as less than 8 seconds and you do not configure a value for <b>transmit-delay</b>, the default value of <b>transmit-delay</b> is automatically changed to 1 second in order to satisfy the requirement that the <b>advertisement-interval</b> value must be greater than or equal to four times the <b>transmit-delay</b> value.</p></div> |
| <b>Default</b>                  | Disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b>seconds</b> —Interval between LLDP advertisement.<br><b>Default:</b> 30<br><b>Range:</b> 5 through 32768                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring LLDP on page 3246</a></li><li>• <code>show lldp</code></li><li>• <a href="#">Configuring LLDP (CLI Procedure)</a></li><li>• <a href="#">Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches</a></li><li>• <code>transmit-delay</code></li><li>• <a href="#">Understanding LLDP</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |



## disable


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">lldp</a> ],<br>[edit protocols <a href="#">lldp interface</a> (all   <i>interface-name</i> )],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">lldp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">lldp interface</a> (all   <i>interface-name</i> )]                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | (MX Series and T Series routers and EX Series switches) Disable LLDP globally or on an interface.<br><br>For information about interface names, see <i>Interface Naming Overview</i> . For information about interface names for TX Matrix routers, see <i>TX Matrix Router Chassis and Interface Names</i> . For information about FPC numbering on TX Matrix routers, see <i>Routing Matrix with a TX Matrix Router FPC Numbering</i> . |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring LLDP on page 3246</a></li> </ul>                                                                                                                                                                                                                                                                                                                                         |

## hold-multiplier

|                                 |                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | hold-multiplier <i>number</i> ;                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">lldp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">lldp</a> ]                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.                                                                                                                                                           |
| <b>Description</b>              | (MX Series and T Series routers and EX series switches only) Configure a value for the LLDP hold multiplier.<br><br>Hold timer interval in seconds to cache learned LLDP information before discarding. |
| <b>Options</b>                  | <p><b>number</b>—Advertisement interval multiplier for LLDP cache discard.</p> <p><b>Default:</b> 4 (giving 120 second LLDP cache lifetime with other defaults)</p> <p><b>Range:</b> 2 through 10</p>   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring LLDP on page 3246</a></li> </ul>                                                                                                       |

## interface

---

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |                                                                                                                                                 |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <code>interface (all   <i>interface-name</i>) {<br/>    <b>disable</b>;<br/>}</code>                                                            |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | [edit protocols <b>lldp</b> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <b>lldp</b> ]                                  |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Statement introduced in Junos OS Release 9.6.                                                                                                   |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | (MX Series and T Series routers and EX Series switches only) Specify an LLDP interface.                                                         |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <i>interface-name</i> —A valid physical interface name.                                                                                         |
| <div><p><b>NOTE:</b> On MX Series and T Series routers, you run LLDP on a physical interface, such as <code>ge-1/0/0</code>, and not at the logical interface (unit) level.</p><p>For information about interface names, see <i>Interface Naming Overview</i>. For information about interface names for TX Matrix routers, see <i>TX Matrix Router Chassis and Interface Names</i>. For information about FPC numbering on TX Matrix routers, see <i>Routing Matrix with a TX Matrix Router FPC Numbering</i>.</p></div> |                                                                                                                                                 |
| <div><p><b>all</b>—Run LLDP on all interfaces.</p><p><b>disable</b>—Disable LLDP on the specified interface</p></div>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                                                                                                                                                 |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration. |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <ul style="list-style-type: none"><li>• <a href="#">Configuring LLDP on page 3246</a></li></ul>                                                 |

## lldp-configuration-notification-interval

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | lldp-configuration-notification-interval <i>seconds</i> ;                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">lldp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">lldp</a> ]                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.                                                                                                                                                  |
| <b>Description</b>              | (MX Series and T Series routers and EX Series switches only) Configure a time for the period of SNMP trap notifications to the Master Agent to wait regarding changes in database information. |
| <b>Options</b>                  | <b>seconds</b> —Time for the period of SNMP trap notifications about the LLDP database. This feature is disabled by default.<br><b>Range:</b> 0 through 3600                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring LLDP on page 3246</a></li> </ul>                                                                                              |

## ptopo-configuration-maximum-hold-time

|                                 |                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ptopo-configuration-maximum-hold-time <i>seconds</i> ;                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">lldp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">lldp</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.                                                                                    |
| <b>Description</b>              | (MX Series and T Series routers and EX Series switches only) Configure a time to maintain dynamic topology entries.              |
| <b>Options</b>                  | <b>seconds</b> —Time to maintain interval dynamic topology entries.<br><b>Default:</b> 300<br><b>Range:</b> 1 through 2147483647 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring LLDP on page 3246</a></li> </ul>                                |

## ptopo-configuration-trap-interval

---

|                                 |                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ptopo-configuration-trap-interval seconds;</code>                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">lldp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">lldp</a> ]                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.                                                                                                                                                        |
| <b>Description</b>              | (MX Series and T Series routers and EX Series switches only) Configure a time for the period of SNMP trap notifications to the Master Agent to wait regarding changes in topology global statistics. |
| <b>Options</b>                  | <b>seconds</b> —Time for the period of SNMP trap notifications about global statistics. This feature is disabled by default.<br><b>Range:</b> 0 through 3600                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring LLDP on page 3246</a></li></ul>                                                                                                      |

## traceoptions

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>maximum-file-size</i>&gt; &lt;world-readable       no-world-readable&gt;;     flag <i>flag</i> &lt;disable&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>     | [edit protocols <b>lldp</b> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <b>lldp</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.6.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>         | Set LLDP protocol-level tracing options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Default</b>             | The default LLDP protocol-level trace options are inherited from the global <b>traceoptions</b> statement.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place spanning-tree protocol tracing output in the file <code>/var/log/stp-log</code>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 1 trace file only</p> <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. The following are the LLDP-specific tracing options:</p> <ul style="list-style-type: none"> <li>• <b>all</b>—Trace all operations.</li> <li>• <b>config</b>—Log configuration events.</li> <li>• <b>interface</b>—Trace interface update events.</li> <li>• <b>protocol</b>—Trace protocol information.</li> <li>• <b>rtsock</b>—Trace socket events.</li> <li>• <b>vlan</b>—Trace vlan update events.</li> </ul> |

The following are the global tracing options:

- **all**—All tracing operations.
- **config-internal**—Trace configuration internals.
- **general**—Trace general events.
- **normal**—All normal events. This is the default. If you do not specify this option, only unusual or abnormal operations are traced.
- **parse**—Trace configuration parsing.
- **policy**—Trace policy operations and actions.
- **regex-parse**—Trace regular-expression parsing.
- **route**—Trace routing table changes.
- **state**—Trace state transitions.
- **task**—Trace protocol task processing.
- **timer**—Trace protocol task timer processing.

**no-world-readable**—(Optional) Prevent any user from reading the log file. This is the default. If you do not include this option, tracing output is appended to an existing trace file.

**size maximum-file-size**—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

**Required Privilege  
Level**

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

**Related  
Documentation**

- [Tracing LLDP Operations on page 3257](#)

## transmit-delay

|                                 |                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>transmit-delay <i>seconds</i>;</code>                                                                                      |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">lldp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">lldp</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.                                                                                    |
| <b>Description</b>              | (MX Series and T Series routers and EX Series switches only) Configure a delay between two successive LLDP advertisements.       |
| <b>Options</b>                  | <b><i>seconds</i></b> —Delay between two successive LLDP advertisements.<br><b>Default:</b> 2<br><b>Range:</b> 1 through 8192    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring LLDP on page 3246</a></li> </ul>                                |

## Administration

- [Routine Monitoring on page 3257](#)

## Routine Monitoring

- [Tracing LLDP Operations on page 3257](#)

### Tracing LLDP Operations

To trace LLDP operational traffic, you can specify options in the global **traceoptions** statement included at the [edit **routing-options**] hierarchy level, and you can specify LLDP-specific options by including the **traceoptions** statement:

```
traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols [lldp](#)]
- [edit routing-instances *routing-instance-name* protocols [lldp](#)]

You can specify the following LLDP-specific options in the LLDP **traceoptions** statement:

- **all**—Trace all operations.
- **config**—Log configuration events.
- **interface**—Trace interface update events.

- **protocol**—Trace protocol information.
- **rtsock**—Trace real-time socket events.
- **vlan**—Trace VLAN update events.



**NOTE:** Use the trace flag all with caution. This flag may cause the CPU to become very busy.

For general information about tracing and global tracing options, see the statement summary for the global **traceoptions** statement in the *Junos OS Routing Protocols Configuration Guide*.

**Related  
Documentation**

- [LLDP Overview on page 3243](#)
- [Configuring LLDP on page 3246](#)
- [Example: Configuring LLDP on page 3244](#)
- *LLDP Operational Mode Commands*



## CHAPTER 13

# Logical Systems

- [Overview on page 3259](#)
- [Configuration on page 3268](#)
- [Administration on page 3523](#)

### Overview

---

- [Logical Systems on page 3259](#)

### Logical Systems

- [Introduction to Logical Systems on page 3259](#)
- [Junos OS Features That Are Supported on Logical Systems on page 3262](#)
- [Logical Systems Operations and Restrictions on page 3263](#)
- [Comparing Junos OS Device Virtualization Technologies on page 3265](#)
- [Logical Systems Applications on page 3266](#)
- [Logical Systems Requirements on page 3267](#)
- [Logical Systems Terms and Acronyms on page 3267](#)

#### Introduction to Logical Systems

---

For many years, engineers have combined power supplies, routing hardware and software, forwarding hardware and software, and physical interfaces into a networking device known as a router. Networking vendors have created large routers and small routers, but all routers have been placed into service as individual devices. As a result, the router has been considered a single physical device for most of its history.

The concept of logical systems breaks with this tradition. With the Junos<sup>®</sup> operating system (Junos OS), you can partition a single router into multiple logical devices that perform independent routing tasks. Because logical systems perform a subset of the tasks once handled by the main router, logical systems offer an effective way to maximize the use of a single routing or switching platform.



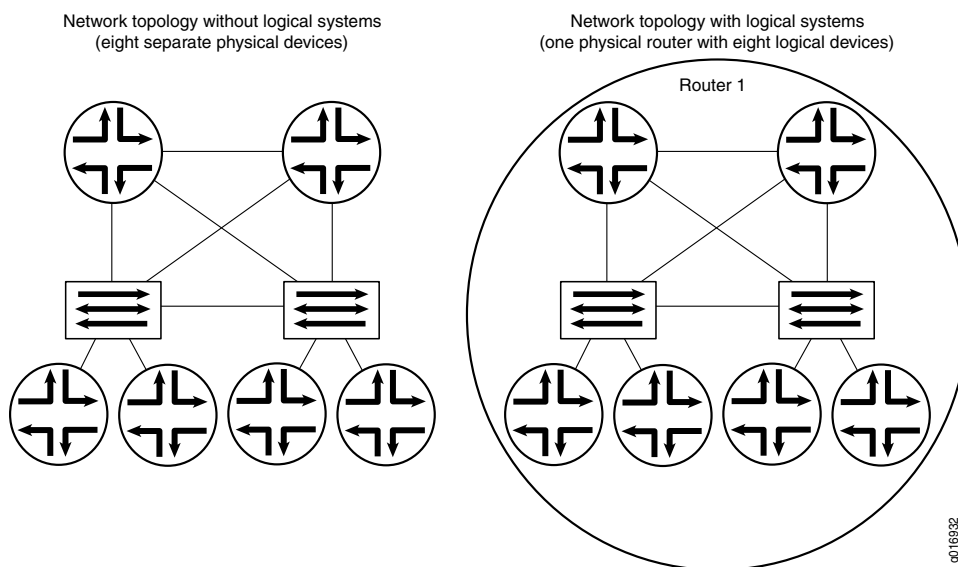
**NOTE:** Beginning with Junos OS Release 9.3, the logical router feature has been renamed logical system.

All configuration statements, operational commands, show command output, error messages, log messages, and SNMP MIB objects that contain the string `logical-router` have been changed to `logical-system`.

Traditionally, service provider network design requires multiple layers of switches and routers. These devices transport packet traffic between customers. As seen on the left side of [Figure 49 on page 3260](#), access devices are connected to edge devices, which are in turn connected to core devices.

However, this complexity can lead to challenges in maintenance, configuration, and operation. To reduce such complexity, Juniper Networks supports logical systems. Logical systems perform a subset of the actions of the main router and have their own unique routing tables, interfaces, policies, and routing instances. As shown on the right side of [Figure 49 on page 3260](#), a set of logical systems within a single router can handle the functions previously performed by several small routers.

**Figure 49: Logical Systems Concepts**



[Figure 50 on page 3261](#) shows the Junos OS architecture without logical systems configured. [Figure 51 on page 3261](#) shows the Junos OS architecture when logical systems are configured. Note that each logical system runs its own routing protocol process (`rpd`).

Figure 50: Junos OS Without Logical Systems

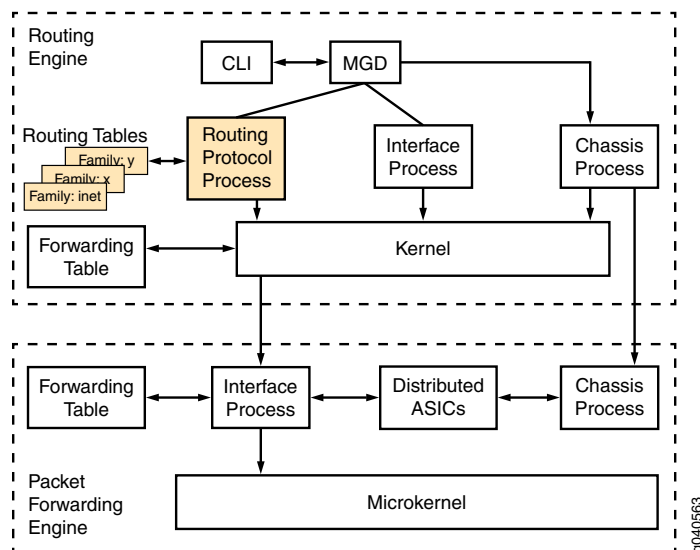
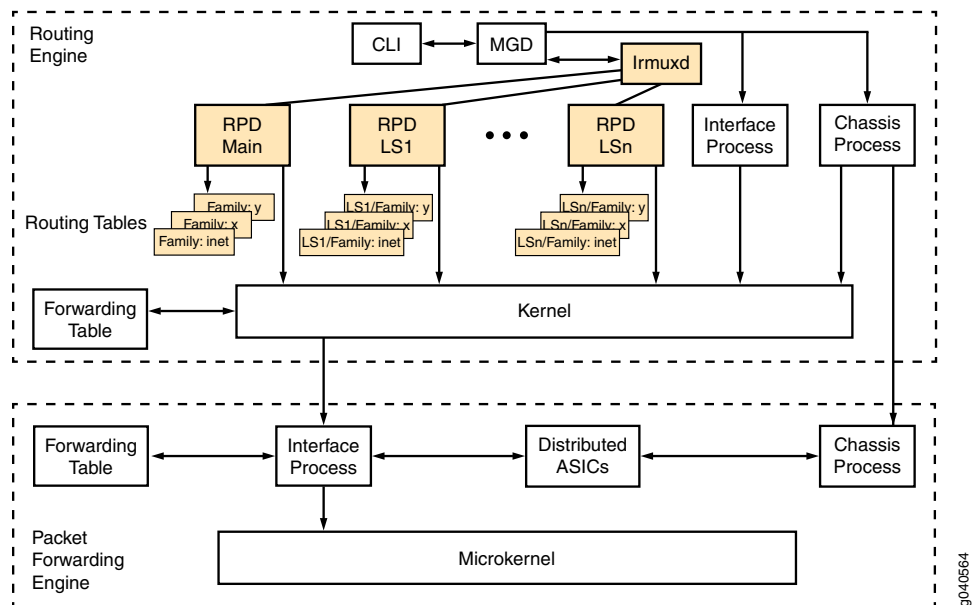


Figure 51: Junos OS With Logical Systems



**Related Documentation**

- [Logical Systems Operations and Restrictions on page 3263](#)
- [Junos OS Features That Are Supported on Logical Systems on page 3262](#)

### Junos OS Features That Are Supported on Logical Systems

---

The following protocols and functions are supported on logical systems:

- Open Shortest Path First (OSPF), Intermediate System-to-Intermediate System (IS-IS), Routing Information Protocol (RIP), RIP next generation (RIPng), Border Gateway Protocol (BGP), Resource Reservation Protocol (RSVP), Label Distribution Protocol (LDP), static routes, and Internet Protocol version 4 (IPv4) and version 6 (IPv6).
- Multiprotocol Label Switching (MPLS) provider edge (PE) and core provider router functions, such as Layer 2 virtual private networks (VPNs), Layer 3 VPNs, circuit cross-connect (CCC), Layer 2 circuits, and virtual private LAN service (VPLS).
- Resource Reservation Protocol (RSVP) point-to-multipoint label-switched paths (LSPs).
- Multicast protocols, such as Protocol Independent Multicast (PIM), Distance Vector Multicast Routing Protocol (DVMRP), rendezvous point (RP), and source designated router (DR).
- All policy-related statements available at the **[edit policy-options]** hierarchy level.
- Most routing options statements available at the **[edit routing-options]** hierarchy level.
- Graceful Routing Engine switchover (GRES). Configure graceful Routing Engine switchover on the main router with the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level.
- Graceful restart. Include the **graceful-restart** statement at the **[edit logical-systems logical-system-name routing-options]** hierarchy level.
- You can assign most interface types to a logical system. For a list of unsupported PICs, see [“Logical Systems Operations and Restrictions” on page 3263](#).
- Flow aggregation in logical systems is supported starting in Junos OS Release 11.4. In the logical system, sampling based on the Routing Engine is not supported. Only PIC-based sampling is supported. Logical systems support only cflowd version 9. Currently, cflowd version 5 and cflowd version 8 are not supported in logical systems. Flow aggregation in logical systems is slightly different from flow aggregation on the main router in that when you configure flow aggregation in logical systems, the **route-record** statement is not required.
- Port mirroring, source class usage, destination class usage, unicast reverse-path forwarding, class of service, firewall filters, class-based forwarding, and policy-based accounting work with logical systems when you configure these features on the main router.
- The Simple Network Management Protocol (SNMP) has been extended to support logical systems and routing instances. A network management system receives instance-aware information in the following format:

***logical-system-name/routing-instance@community***

As a result, a network manager can gather statistics for a specific community within a routing instance within a logical system. The SNMP manager for a routing instance can request and manage SNMP data only for that routing instance and other routing

instances in the same logical system. By default, the SNMP manager for the default routing instance in the main router (**inet.0**) can access SNMP data from all routing instances. To restrict that manager's access to the default routing instance only, include the **routing-instance-access** statement at the **[edit snmp]** hierarchy level.

- For SRX platforms only, support for system logging at the **[edit logical-system logical-system-name system syslog]** hierarchy level is introduced in Junos OS Release 11.4.

**Related  
Documentation**

- [Introduction to Logical Systems on page 3259](#)
- [Logical Systems Operations and Restrictions on page 3263](#)

---

### Logical Systems Operations and Restrictions

Logical systems have the following operations and restrictions:

- You can configure a maximum of 15 logical systems plus the master logical system on a routing device. When a configuration session is in use, users who are tied to the same logical system cannot commit configuration changes.
- The routing device has only one running configuration database, which contains configuration information for the main routing device and all associated logical systems. When configuring a logical system, a user has his own candidate configuration database, which does not become part of the running configuration database until the user issues the **commit** command.
- Configuring the out-of-band management interface, such as **em0** or **fxp0**, in a logical system is not supported.
- Some high availability features are not supported on logical systems. These features include non-stop routing (NSR), non-stop bridging (NSB), and unified in-service software upgrade (unified ISSU). Graceful restart is supported. For a logical system, include the graceful-restart statement at the **[edit logical-systems logical-system-name routing-options]** hierarchy level.
- The following guidelines describe how firewall filters affect the main routing device, logical systems, and virtual routers. The "default loopback interface" refers to **lo0.0** (associated with the default routing table), the "loopback interface in a logical system" refers to **lo0.n** configured in the logical system, and the "loopback interface in the virtual router" refers to **lo0.n** configured in the virtual router.

If you configure Filter A on the default loopback interface in the main routing device but do not configure a filter on the loopback interface in a logical system, the logical system does not use a filter.

If you configure Filter A on the default loopback interface in the main routing device but do not configure a loopback interface in a logical system, the logical system uses Filter A.

If you configure Filter A on the default loopback interface on the main routing device and Filter B on the loopback interface in a logical system, the logical system uses Filter

B. In a special case of this rule, when you also configure a routing instance of type **virtual-router** on the logical system, the following rules apply:

- If you configure Filter C on the loopback interface in the virtual router, traffic belonging to the virtual router uses Filter C.
- If you do not configure a filter on the loopback interface in the virtual router, traffic belonging to the virtual router does not use a filter.
- If you do not configure a loopback interface in the virtual router, traffic belonging to the virtual router uses Filter A.
- If a logical system experiences an interruption of its routing protocol process (**rpd**), the core dump output is placed in **/var/tmp/** in a file called **rpd\_logical-system-name.core-tarball.number.tgz**. Likewise, if you issue the **restart routing** command in a logical system, only the routing protocol process (**rpd**) for the logical system is restarted.
- If you configure trace options for a logical system, the output log file is stored in the following location: **/var/log/logical-system-name**. To monitor a log file within a logical system, issue the **monitor start logical-system-name/filename** command.
- The following PICs are not supported with logical systems: Adaptive Services, Multiservices, ES, Monitoring Services, and Monitoring Services II.
- The Multiservices Dense Port Concentrator (MS-DPC) is not supported with logical systems.
- Generalized MPLS (GMPLS), IP Security (IPsec), and sampling are not supported.
- Class of service (CoS) on a logical tunnel (**lt**) or virtual loopback tunnel (**vt**) interface in a logical system is not supported.
- You cannot include the **vrf-table-label** statement on multiple logical systems if the core-facing interfaces are channelized or configured with multiple logical interfaces (Frame Relay DLCIs or Ethernet VLANs).
- The master administrator must configure global interface properties and physical interface properties at the **[edit interfaces]** hierarchy level. Logical system administrators can only configure and verify configurations for the logical systems to which they are assigned.
- You can configure only Frame Relay interface encapsulation on a logical tunnel interface (**lt-**) when it is configured with an IPv6 address.
- IPv6 tunneling is not supported with point-to-multipoint label-switched paths (LSPs) configured on logical systems.
- IGMP snooping is not supported.
- In a VPLS multihoming scenario in which a logical tunnel interface (**lt-**) is used for connecting the dual-home VPLS, Junos OS creates a unique static MAC address for every logical tunnel interface configured. This MAC address is not flushed when a CCC down event occurs on the link and when traffic is switched from the primary link to the backup link (or the reverse). As a result, any traffic that is destined for hosts behind the logical tunnel MAC address does not take the new path.

- Related Documentation
- [Introduction to Logical Systems on page 3259](#)
  - [Junos OS Features That Are Supported on Logical Systems on page 3262](#)

Comparing Junos OS Device Virtualization Technologies

The Junos OS supports multiple device virtualization technologies. The technologies have similar names, which can lead to confusion.

The Junos OS device virtualization technologies are:

- Logical systems—Offer routing and management separation. Management separation means multiple user access. Each logical system has its own routing tables.  
  
Logical routers is the old name for logical systems. Beginning with Junos OS Release 9.3, the logical router feature has been renamed logical system. All configuration statements, operational commands, **show** command output, error messages, log messages, and SNMP MIB objects that contain the string logical-router have been changed to logical-system.
- Virtual routers—Offer scalable routing separation. A virtual router does not have the same capabilities as a logical system. A virtual router is a type of simplified routing instance that has a single routing table. By contrast, a logical system is a partition of the main routing device and can contain multiple routing instances.
- VRF-Lite—Offers routing separation. The functionality of VRF-Lite is similar to virtual routers, but VRF-Lite is for smaller environments.
- Virtual switches—Offer scalable switching separation.

[Table 264 on page 3265](#) summarizes the benefits of virtual routers, VRF-Lite, and logical systems.

Table 264: Benefits of Virtual Routers, VRF-Lite, and Logical Systems

| Benefits                                     | Virtual Router | VRF-Lite | Logical Systems |
|----------------------------------------------|----------------|----------|-----------------|
| Logical platform partitioning                | Yes            | Yes      | Yes             |
| Fault isolation on the routing plane         | No             | No       | Yes             |
| Multiple user access (management separation) | No             | No       | Yes             |
| Scalable routing separation                  | Yes            | No       | Yes             |

- Related Documentation
- [Logical Systems Applications on page 3266](#)

## Logical Systems Applications

Logical systems are discrete contexts that virtually divide a supported device into multiple devices, isolating one from another and protecting them from faulty conditions outside their own contexts.


The logical systems functionality enables you to partition the device and assign private logical systems to groups or organizations. Logical systems are defined largely by the resources allocated to them, features enabled for the logical context, their routing configurations, and their logical interface assignments. Logical systems segment a physical routing device to be configured and operated as multiple independent routing devices within a platform. This isolates routing protocols and interfaces among up to 16 logical systems (including the master logical system). User permissions and access are defined separately for each logical system, enabling different groups to manage the same physical device. Logical systems enable the use of large routing devices in small routing device roles and provide flexible segmentation of routing by service type. Multiple service capabilities bring improved asset optimization by consolidating services into one device.

For example, logical systems enable the following services on a single routing device platform:

- Internet BGP peering
- Core transit
- Edge aggregation and dedicated access
- MPLS provider edge (PE) and provider (P) VPN label-switched routing routers (LSRs)

Figure 52 on page 3266 shows how logical systems can be used for horizontal consolidation, vertical consolidation, and managed services. Horizontal consolidation occurs when you combine routing device functions of the same layer into a single routing device. Vertical consolidation occurs when you collapse routing device functions of different layers into a single routing device. With managed services, each logical system is a customer routing device.

**Figure 52: Applications of Logical Systems**



### Related Documentation

- [Comparing Junos OS Device Virtualization Technologies on page 3265](#)



### Logical Systems Requirements

---

To implement logical systems, your system must meet these minimum requirements:

- Junos OS Release 8.5 or later for logical system administrator support
- Junos OS Release 8.4 or later for SNMP enhancements and limits
- Junos OS Release 8.3 or later for Bidirectional Forwarding Detection (BFD) on logical systems
- Junos OS Release 8.2 or later for support on MX Series routers
- Junos OS Release 7.5 or later for SNMP support within a logical system
- Junos OS Release 7.4 or later for multicast protocol RP and source designated router functionality within a logical system
- Junos OS Release 7.0 or later to implement a logical tunnel (**lt**) interface on an integrated Adaptive Services Module in an M7i router
- Junos OS Release 6.1 or later, a Tunnel Services PIC, and an Enhanced FPC on M Series or T Series routers to implement a logical tunnel (**lt**) interface
- Junos OS Release 6.0 or later for basic logical system functionality
- One or more M Series, MX Series, or T Series routers
- On M Series and T Series routers, a variety of PICs to assign interfaces to each logical system
- One or more EX Series switches

#### Related Documentation

- *Junos OS Logical Systems Configuration Guide for Security Devices*

### Logical Systems Terms and Acronyms

---

#### A

|                                     |                                                                                                                            |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| <b>logical system administrator</b> | A user account with configuration and verification privileges for only the logical systems to which that user is assigned. |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------|

|                             |                                                                          |
|-----------------------------|--------------------------------------------------------------------------|
| <b>master administrator</b> | A user account with superuser configuration and verification privileges. |
|-----------------------------|--------------------------------------------------------------------------|

#### L

|                       |                                                                                                                                                                 |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>logical system</b> | Segmentation of a system into multiple logical devices. Logical system configuration statements are found at the <b>[edit logical-systems]</b> hierarchy level. |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|

#### M

|                    |                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------|
| <b>main router</b> | The standard concept of a routing device. Main routing configuration statements are found at the <b>[edit]</b> hierarchy level. |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------|

## Configuration

---

- [Configuration Examples on page 3268](#)
- [Configuration Statements on page 3521](#)

### Configuration Examples

- [Example: Configuring User Access for Logical Systems on page 3268](#)
- [Examples: Configuring Logical System Interfaces on page 3274](#)
- [Example: Configuring Static Routing on Logical Systems on page 3287](#)
- [Examples: Configuring Standard Firewall Filters on Logical Systems on page 3293](#)
- [Example: Configuring OSPF on Logical Systems on page 3307](#)
- [Examples: Configuring OSPF Routing Policy on Logical Systems on page 3320](#)
- [Examples: Configuring IS-IS on Logical Systems on page 3343](#)
- [Examples: Configuring BGP on Logical Systems on page 3362](#)
- [Example: Configuring RSVP-Signaled Point-to-Multipoint LSPs on Logical Systems on page 3407](#)
- [Example: Configuring VPNs and VPLS on Logical Systems on page 3431](#)
- [Example: Configuring a Virtualized Data Center on page 3478](#)

#### **Example: Configuring User Access for Logical Systems**

---

- [Understanding Junos OS Access Privilege Levels on page 3268](#)
- [Example: Configuring Logical System Administrators on page 3272](#)

##### ***Understanding Junos OS Access Privilege Levels***

Each top-level command-line interface (CLI) command and each configuration statement have an access privilege level associated with them. Users can execute only those commands and configure and view only those statements for which they have access privileges. The access privileges for each login class are defined by one or more *permission flags*.

For each login class, you can explicitly deny or allow the use of operational and configuration mode commands that would otherwise be permitted or not allowed by a privilege level specified in the **permissions** statement.

The following sections provide additional information about permissions:

- [Junos OS Login Class Permission Flags on page 3268](#)
- [Allowing or Denying Individual Commands for Junos OS Login Classes on page 3272](#)

##### ***Junos OS Login Class Permission Flags***

The **permissions** statement specifies one or more of the permission flags listed in [Table 40 on page 124](#). Permission flags are not cumulative, so for each class you must list all the permission flags needed, including **view** to display information and **configure**

to enter configuration mode. Two forms of permissions control for individual parts of the configuration are:

- "Plain" form—Provides read-only capability for that permission type. An example is **interface**.
- Form that ends in **-control**—Provides read and write capability for that permission type. An example is **interface-control**.

Table 40 on page 124 lists the Junos<sup>®</sup> operating system (Junos OS) login class permission flags that you can configure by including the **permissions** statement at the **[edit system login class class-name]** hierarchy level.

**Table 265: Login Class Permission Flags**

| Permission Flag         | Description                                                                                                                                   |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>access</b>           | Can view the access configuration in configuration mode and with the <b>show configuration</b> operational mode command.                      |
| <b>access-control</b>   | Can view and configure access information at the <b>[edit access]</b> hierarchy level.                                                        |
| <b>admin</b>            | Can view user account information in configuration mode and with the <b>show configuration</b> operational mode command.                      |
| <b>admin-control</b>    | Can view user accounts and configure them at the <b>[edit system login]</b> hierarchy level.                                                  |
| <b>all-control</b>      | Can access all operational mode commands and configuration mode commands. Can modify configuration in all the configuration hierarchy levels. |
| <b>clear</b>            | Can clear (delete) information learned from the network that is stored in various network databases by using the <b>clear</b> commands.       |
| <b>configure</b>        | Can enter configuration mode by using the <b>configure</b> command.                                                                           |
| <b>control</b>          | Can perform all control-level operations—all operations configured with the <b>-control</b> permission flags.                                 |
| <b>field</b>            | Can view field debug commands. Reserved for debugging support.                                                                                |
| <b>firewall</b>         | Can view the firewall filter configuration in configuration mode.                                                                             |
| <b>firewall-control</b> | Can view and configure firewall filter information at the <b>[edit firewall]</b> hierarchy level.                                             |
| <b>floppy</b>           | Can read from and write to the removable media.                                                                                               |
| <b>flow-tap</b>         | Can view the flow-tap configuration in configuration mode.                                                                                    |

Table 265: Login Class Permission Flags (*continued*)

| Permission Flag                       | Description                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>flow-tap-control</b>               | Can view the flow-tap configuration in configuration mode and can configure flow-tap configuration information at the <b>[edit services flow-tap]</b> hierarchy level.                                                                                                                                                                                                                                        |
| <b>flow-tap-operation</b>             | Can make flow-tap requests to the router or switch. For example, a Dynamic Tasking Control Protocol (DTCP) client must authenticate itself to the Junos OS as an administrative user. That account must have <b>flow-tap-operation</b> permission.<br><br><b>NOTE:</b> The <b>flow-tap-operation</b> option is not included in the <b>all-control</b> permissions flag.                                       |
| <b>idp-profiler-operation</b>         | Can view profiler data.                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>interface</b>                      | Can view the interface configuration in configuration mode and with the <b>show configuration</b> operational mode command.                                                                                                                                                                                                                                                                                   |
| <b>interface-control</b>              | Can view chassis, class of service (CoS), groups, forwarding options, and interfaces configuration information. Can edit configuration at the following hierarchy levels: <ul style="list-style-type: none"> <li>• <b>[edit chassis]</b></li> <li>• <b>[edit class-of-service]</b></li> <li>• <b>[edit groups]</b></li> <li>• <b>[edit forwarding-options]</b></li> <li>• <b>[edit interfaces]</b></li> </ul> |
| <b>maintenance</b>                    | Can perform system maintenance, including starting a local shell on the router or switch and becoming the superuser in the shell by using the <b>su root</b> command, and can halt and reboot the router or switch by using the <b>request system</b> commands.                                                                                                                                               |
| <b>network</b>                        | Can access the network by using the <b>ping</b> , <b>ssh</b> , <b>telnet</b> , and <b>traceroute</b> commands.                                                                                                                                                                                                                                                                                                |
| <b>pgcp-session-mirroring</b>         | Can view the <b>pgcp</b> session mirroring configuration.                                                                                                                                                                                                                                                                                                                                                     |
| <b>pgcp-session-mirroring-control</b> | Can modify the <b>pgcp</b> session mirroring configuration.                                                                                                                                                                                                                                                                                                                                                   |
| <b>reset</b>                          | Can restart software processes by using the <b>restart</b> command and can configure whether software processes are enabled or disabled at the <b>[edit system processes]</b> hierarchy level.                                                                                                                                                                                                                |
| <b>rollback</b>                       | Can use the <b>rollback</b> command to return to a previously committed configuration other than the most recently committed one.                                                                                                                                                                                                                                                                             |
| <b>routing</b>                        | Can view general routing, routing protocol, and routing policy configuration information in configuration and operational modes.                                                                                                                                                                                                                                                                              |

Table 265: Login Class Permission Flags (*continued*)

| Permission Flag           | Description                                                                                                                                                                                                                                                                                                                |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>routing-control</b>    | Can view general routing, routing protocol, and routing policy configuration information and can configure general routing at the <b>[edit routing-options]</b> hierarchy level, routing protocols at the <b>[edit protocols]</b> hierarchy level, and routing policy at the <b>[edit policy-options]</b> hierarchy level. |
| <b>secret</b>             | Can view passwords and other authentication keys in the configuration.                                                                                                                                                                                                                                                     |
| <b>secret-control</b>     | Can view passwords and other authentication keys in the configuration and can modify them in configuration mode.                                                                                                                                                                                                           |
| <b>security</b>           | Can view security configuration in configuration mode and with the <b>show configuration</b> operational mode command.                                                                                                                                                                                                     |
| <b>security-control</b>   | Can view and configure security information at the <b>[edit security]</b> hierarchy level.                                                                                                                                                                                                                                 |
| <b>shell</b>              | Can start a local shell on the router or switch by using the <b>start shell</b> command.                                                                                                                                                                                                                                   |
| <b>snmp</b>               | Can view Simple Network Management Protocol (SNMP) configuration information in configuration and operational modes.                                                                                                                                                                                                       |
| <b>snmp-control</b>       | Can view SNMP configuration information and can modify SNMP configuration at the <b>[edit snmp]</b> hierarchy level.                                                                                                                                                                                                       |
| <b>system</b>             | Can view system-level information in configuration and operational modes.                                                                                                                                                                                                                                                  |
| <b>system-control</b>     | Can view system-level configuration information and configure it at the <b>[edit system]</b> hierarchy level.                                                                                                                                                                                                              |
| <b>trace</b>              | Can view trace file settings and configure trace file properties.                                                                                                                                                                                                                                                          |
| <b>trace-control</b>      | Can modify trace file settings and configure trace file properties.                                                                                                                                                                                                                                                        |
| <b>view</b>               | Can use various commands to display current system-wide, routing table, and protocol-specific values and statistics. Cannot view the secret configuration.                                                                                                                                                                 |
| <b>view-configuration</b> | Can view all of the configuration excluding secrets, system scripts, and event options.<br><br><b>NOTE:</b> Only users with the <b>maintenance</b> permission can view commit script, op script, or event script configuration.                                                                                            |

### *Allowing or Denying Individual Commands for Junos OS Login Classes*

By default, all top-level CLI commands have associated access privilege levels. Users can execute only those commands and view only those statements for which they have access privileges. For each login class, you can explicitly deny or allow the use of operational and configuration mode commands that would otherwise be permitted or not allowed by a privilege level specified in the **permissions** statement.

- The **all** login class permission bits take precedence over extended regular expressions when a user with **rollback** permission issues the **rollback** command.
- Expressions used to allow and deny commands for users on RADIUS and TACACS+ servers have been simplified. Instead of a single, long expression with multiple commands (**allow-commands=cmd1 cmd2 ... cmdn**), you can specify each command as a separate expression. This new syntax is valid for **allow-configuration** and **deny-configuration**, **allow-commands** and **deny-commands**, and all user permission bits.
- Users cannot issue the **load override** command when specifying an extended regular expression. Users can only issue the **merge**, **replace**, and **patch** configuration commands.
- If you allow and deny the same commands, the **allow-commands** permissions take precedence over the permissions specified by the **deny-commands**. For example, if you include **allow-commands "request system software add"** and **deny-commands "request system software add"**, the login class user is allowed to install software using the **request system software add** command.
- Regular expressions for **allow-commands** and **deny-commands** can also include the **commit**, **load**, **rollback**, **save**, **status**, and **update** commands.
- If you specify a regular expression for **allow-commands** and **deny-commands** with two different variants of a command, the longest match is always executed.

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

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

### *Example: Configuring Logical System Administrators*

This example shows how to configure logical system administrators.

- [Requirements on page 3273](#)
- [Overview on page 3273](#)

- [Configuration on page 3273](#)
- [Verification on page 3274](#)

### Requirements

You must be the master administrator to assign system administrators to logical systems.

### Overview

The master administrator can assign one or more system administrators to each logical system. Logical system administrators are confined to the context of the logical system to which they are assigned. This means that logical system administrators cannot access any global configuration statements. This also means that command output is restricted to the context to which the logical system administrators are assigned.

Configuring a user account for each logical system helps in navigating the CLI. This enables you to log in to each logical system and be positioned within the root of that logical system as if you were in the root of a physical routing device.

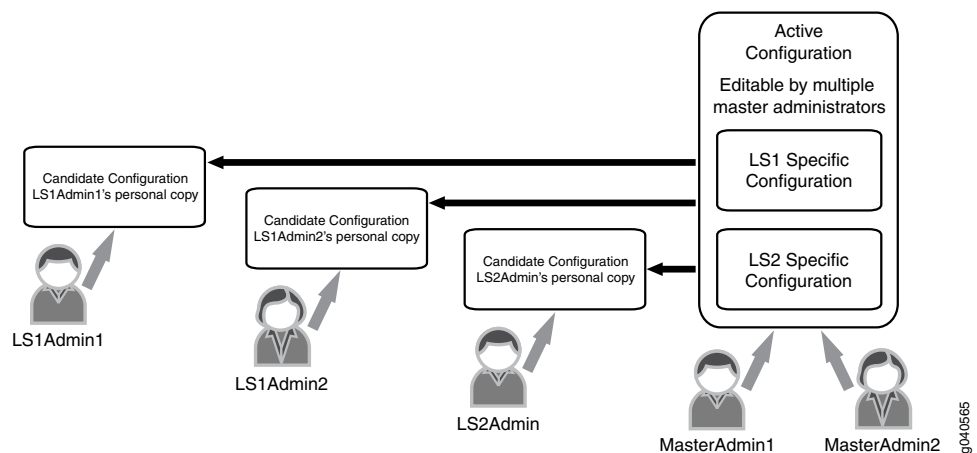
In this example, **LS1Admin** has full permissions on Logical System LS1.

In this example, **LS2Admin** has the ability to view Logical System LS2 but not to change the configuration.

### Diagram

[Figure 53 on page 3273](#) shows how logical system administration works.

**Figure 53: Logical System Administrators**



### Configuration

#### CLI Quick Configuration

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

```

set logical-systems LS1
set logical-systems LS2
set system login class admin1 logical-system LS1

```

```
set system login class admin2 logical-system LS2
set system login class admin1 permissions all
set system login class admin2 permissions view
set system login user LS1Admin class admin1
set system login user LS2Admin class admin2
```

**Step-by-Step  
Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To assign logical system administrators to a logical systems:

1. Configure the logical systems.

```
[edit]
user@host# set logical-systems LS1
user@host# set logical-systems LS2
```

2. Create the login classes and assign logical systems to the classes.

```
[edit]
user@host# set system login class admin1 logical-system LS1
user@host# set system login class admin2 logical-system LS2
```

3. Assign permissions to the login classes.

```
[edit]
user@host# set system login class admin1 permissions all
user@host# set system login class admin2 permissions view
```

4. Assign users to the login classes.

```
[edit]
user@host# set system login user LS1Admin class admin1
user@host# set system login user LS2Admin class admin2
```

5. If you are done configuring the device, commit the configuration.

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

**Verification**

To verify that the configuration is working properly, issue the [show cli authorization](#) command to view permissions for the current user.

**Related  
Documentation**

- [Introduction to Logical Systems on page 3259](#)

---

**Examples: Configuring Logical System Interfaces**

- [Understanding Interfaces on page 3275](#)
- [Example: Creating an Interface on a Logical System on page 3275](#)
- [Example: Connecting a Logical System to a Physical Router on page 3277](#)



- [Example: Connecting Logical Systems Within the Same Device Using Logical Tunnel Interfaces on MX Series Routers and EX Series Switches on page 3279](#)
- [Example: Connecting Logical Systems Within the Same Router Using Logical Tunnel Interfaces on MX Series Routers on page 3283](#)

### **Understanding Interfaces**

Interfaces act as a doorway through which traffic enters and exits a device. Juniper Networks devices support a variety of interface types:

- Network interfaces—Networking interfaces primarily provide traffic connectivity.
- Services interfaces—Services interfaces manipulate traffic before it is delivered to its destination.
- Special interfaces—Special interfaces include management interfaces, the loopback interface, and the discard interface.

Each type of interface uses a particular medium to transmit data. The physical wires and Data Link Layer protocols used by a medium determine how traffic is sent. To configure and monitor interfaces, you need to understand their media characteristics, as well as physical and logical properties such as IP addressing, link-layer protocols, and link encapsulation.



**NOTE:** Most interfaces are configurable, but some internally generated interfaces are not configurable.

### **Example: Creating an Interface on a Logical System**

This example shows how to create an interface on a logical system.

- [Requirements on page 3275](#)
- [Overview on page 3275](#)
- [Configuration on page 3276](#)
- [Verification on page 3277](#)

#### **Requirements**

For the interface on the logical system to have connectivity, the corresponding physical interface must be administratively up, and the physical link must be up. You can verify the status of the physical interface by running the **show interfaces terse** command.

#### **Overview**

In logical systems, you must treat each interface like a point-to-point connection because you can only connect one logical tunnel interface to another at any given time. Also, you must select an interface encapsulation type, specify a DLCI number or VLAN identifier, configure a corresponding protocol family, and set the logical interface unit number of the peering **lt** interface.

To configure the interface encapsulation type, include the **dlci**, **encapsulation**, **family**, **peer-unit**, and **vlan-id** statements at the following hierarchy levels:

- M Series, MX Series, or T Series router (master administrator only)—[edit interfaces *lt-fpc/pic/0* unit *unit-number*]
- Logical system—[edit logical-systems *logical-system-name* interfaces *lt-fpc/pic/0* unit *unit-number*]

```
[edit]
logical-systems logical-system-name {
 interfaces {
 lt-fpc/pic/0 {
 unit unit-number {
 encapsulation (ethernet | ethernet-ccc | ethernet-vpls | frame-relay |
 frame-relay-ccc | vlan | vlan-ccc | vlan-vpls);
 peer-unit number; # The logical unit number of the peering lt interface.
 dlci dlci-number;
 vlan-id vlan-number;
 family (ccc | inet | inet6 | iso | mpls | tcc);
 }
 }
 }
}
```



**NOTE:** When you configure IPv6 addresses on a logical tunnel interface, you must configure unique IPv6 link-local addresses for any logical interfaces that peer with one another. To configure a link-local address, you must be the master administrator. Include a second IPv6 address with the address statement at the [edit interfaces *lt-fpc/pic/port* unit *unit-number* family inet6] hierarchy level. Link-local addresses typically begin with the numbers fe80 (such as fe80::1111:1/64).

In this example, you create the **fe-1/1/3** physical interface on the main router. You can also add values for properties that you need to configure on the physical interface, such as physical encapsulation, VLAN tagging (enabling), and link speed.

The example then shows how to assign logical interfaces to a logical system. Once you do this, the logical interfaces are considered part of the logical system.

Any logical interface unit can only be assigned to one system, including the main router. For example, if you configure logical unit 3 in the main router, you cannot configure logical unit 3 in a logical system.

In this example, you create logical unit 0 on Logical System LS1. You can also add values for properties that you need to configure on the logical interface, such as logical interface encapsulation, VLAN ID number, and protocol family.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces fe-1/1/3 description "main router interface"
set logical-systems LS1 interfaces fe-1/1/3 unit 0 description "LS1 interface"
set logical-systems LS1 interfaces fe-1/1/3 unit 0 family inet address 10.11.2.2/24
```

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure an interface on a logical system:

1. As the master administrator, configure the physical interface on the main router.  
  

```
[edit]
user@host# set interfaces fe-1/1/3 description "main router interface"
```
2. Create the logical system interface on the logical unit.  
  

```
[edit]
user@host# set logical-systems LS1 interfaces fe-1/1/3 unit 0 description "LS1 interface"
user@host# set logical-systems LS1 interfaces fe-1/1/3 unit 0 family inet address 10.11.2.2/24
```
3. If you are done configuring the device, commit the configuration.  
  

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

#### Verification

To verify that the configuration is working properly, issue the **show interfaces** command.

#### Example: Connecting a Logical System to a Physical Router

This example shows how to configure an interface on a logical system to connect to a separate router. The separate router can be a physical router or a logical system on a physical router.

- [Requirements on page 3277](#)
- [Overview on page 3277](#)
- [Configuration on page 3278](#)
- [Verification on page 3279](#)

#### Requirements


PICs must be installed on the two routers.

#### Overview

In this example, Logical System LS1 is configured on Router R1. The Logical System LS1 has a direct connection to Router R2.

[Figure 54 on page 3278](#) shows the topology used in this example.

Figure 54: Logical System Connected to a Physical Router

**Configuration****CLI Quick Configuration**

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

**Router R1**      **set interfaces so-0/0/2 description "main router interface to R2"**  
**set logical-systems LS1 interfaces so-0/0/2 unit 0 description LS1->R2**  
**set logical-systems LS1 interfaces so-0/0/2 unit 0 family inet address 10.0.45.2/30**

**Device R2**      **set interfaces so-0/0/2 description R2->LS1**  
**set interfaces so-0/0/2 unit 0 family inet address 10.0.45.1/30**

**Step-by-Step Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To connect a logical system to a physical router:

1. On Router R1, configure the interface.  

```
[edit]
user@R1# set interfaces so-0/0/2 description "main router interface to R2"
```
2. On Router R1, configure the Logical System LS1 interface.  

```
[edit]
user@R1# set logical-systems LS1 interfaces so-0/0/2 unit 0 description LS1->R2
user@R1# set logical-systems LS1 interfaces so-0/0/2 unit 0 family inet address 10.0.45.2/30
```
3. On Device R2, configure the interface to Logical System LS1.  

```
[edit]
user@R2# set interfaces so-0/0/2 description R2->LS1
user@R2# set interfaces so-0/0/2 unit 0 family inet address 10.0.45.1/30
```

- If you are done configuring the devices, commit the configurations.

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

### Verification

Confirm that the configuration is working properly.

### Verifying Connectivity

**Purpose** Make sure that the devices can ping each other.

**Action** user@R2> ping 10.0.45.2  
 PING 10.0.45.2 (10.0.45.2): 56 data bytes  
 64 bytes from 10.0.45.2: icmp\_seq=0 ttl=64 time=3.910 ms  
 64 bytes from 10.0.45.2: icmp\_seq=1 ttl=64 time=3.559 ms  
 64 bytes from 10.0.45.2: icmp\_seq=2 ttl=64 time=3.503 ms

```
user@R1> set cli logical-system LS1
Logical system: LS1
```

```
user@R1:LS1> ping 10.0.45.1
PING 10.0.45.1 (10.0.45.1): 56 data bytes
64 bytes from 10.0.45.1: icmp_seq=0 ttl=64 time=1.217 ms
64 bytes from 10.0.45.1: icmp_seq=1 ttl=64 time=1.183 ms
64 bytes from 10.0.45.1: icmp_seq=2 ttl=64 time=1.121 ms
```

### Example: Connecting Logical Systems Within the Same Device Using Logical Tunnel Interfaces on MX Series Routers and EX Series Switches

This example shows how to configure logical tunnel interfaces to connect two logical systems that are configured in a single router.

- [Requirements on page 3279](#)
- [Overview on page 3280](#)
- [Configuration on page 3281](#)
- [Verification on page 3282](#)

### Requirements

On M Series and T Series routers, you can create a logical tunnel interface if you have a Tunnel Services PIC installed on an Enhanced FPC in your routing platform.

On M40e routers, you can create a logical tunnel interface if you have a Tunnel Services PIC. (An Enhanced FPC is not required.)

On an M7i router, logical tunnel interfaces can be created by using the integrated Adaptive Services Module.

On an MX Series router, the master administrator can configure logical tunnel interfaces by including the **tunnel-services** statement at the **[edit chassis fpc slot-number pic number]** hierarchy level.

### Overview

To connect two logical systems, you configure a logical tunnel interface on both logical systems. Then you configure a peer relationship between the logical tunnel interfaces, thus creating a point-to-point connection. Logical tunnel interfaces behave like regular interfaces. You can configure them with Ethernet, Frame Relay, or another encapsulation type. You can also configure routing protocols across them. In effect, the logical tunnel (lt) interfaces connect two logical systems within the same router. The two logical systems do not share routing tables. This means that you can run dynamic routing protocols between different logical systems within the same router.

You must treat each interface like a point-to-point connection because you can only connect one logical tunnel interface to another at any given time. Also, you must select an interface encapsulation type, configure a corresponding protocol family, and set the logical interface unit number of the peering lt interface.


In this example, the logical tunnel interfaces are configured to behave as Ethernet interfaces with the **encapsulation ethernet** statement. The IS-IS Protocol is enabled on the logical tunnel interfaces with the **family iso** statement.

When configuring logical tunnel interfaces, note the following:

- The peering logical interfaces must have the same physical lt interface name. For example, a logical unit on lt-0/1/0 cannot peer with a logical unit on lt-0/0/10. The FPC, PIC, and port numbers must match.
- The peering logical interfaces must be derived from the same PIC or module.
- You can configure only one peer unit for each logical interface. For example, unit 0 cannot peer with both unit 1 and unit 2.
- Logical tunnels are not supported with Adaptive Services, MultiServices, or Link Services PICs, but they are supported on the Adaptive Services Module on M7i routers.

Figure 55 on page 3280 shows the topology used in this example.

**Figure 55: Connecting Two Logical Systems**



### Configuration

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

```
set logical-systems LS1 interfaces lt-0/1/0 unit 0 description LS1->LS2
set logical-systems LS1 interfaces lt-0/1/0 unit 0 encapsulation ethernet
set logical-systems LS1 interfaces lt-0/1/0 unit 0 peer-unit 1
set logical-systems LS1 interfaces lt-0/1/0 unit 0 family inet address 10.0.8.13/30
set logical-systems LS1 interfaces lt-0/1/0 unit 0 family iso
set logical-systems LS2 interfaces lt-0/1/0 unit 1 description LS2->LS1
set logical-systems LS2 interfaces lt-0/1/0 unit 1 encapsulation ethernet
set logical-systems LS2 interfaces lt-0/1/0 unit 1 peer-unit 0
set logical-systems LS2 interfaces lt-0/1/0 unit 1 family inet address 10.0.8.14/30
set logical-systems LS2 interfaces lt-0/1/0 unit 1 family iso
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To connect logical system interfaces:

1. Run the **show interfaces terse** command to verify that the router has a logical tunnel (lt) interface.

```
user@host> show interfaces terse
Interface Admin Link Proto Local Remote
so-0/0/0 up down
so-0/0/1 up down
so-0/0/2 up down
so-0/0/3 up down
gr-0/1/0 up up
ip-0/1/0 up up
lt-0/1/0 up up
...
```

2. Configure the logical tunnel interface on Logical System LS1.

```
[edit]
user@host# set logical-systems LS1 interfaces lt-0/1/0 unit 0 description LS1->LS2
user@host# set logical-systems LS1 interfaces lt-0/1/0 unit 0 encapsulation ethernet
user@host# set logical-systems LS1 interfaces lt-0/1/0 unit 0 peer-unit 1
user@host# set logical-systems LS1 interfaces lt-0/1/0 unit 0 family inet address
10.0.8.13/30
user@host# set logical-systems LS1 interfaces lt-0/1/0 unit 0 family iso
```

3. Configure the logical tunnel interface on Logical System LS2.

```
[edit]
user@host# set logical-systems LS2 interfaces lt-0/1/0 unit 1 description LS2->LS1
user@host# set logical-systems LS2 interfaces lt-0/1/0 unit 1 encapsulation ethernet
user@host# set logical-systems LS2 interfaces lt-0/1/0 unit 1 peer-unit 0
user@host# set logical-systems LS2 interfaces lt-0/1/0 unit 1 family inet address
10.0.8.14/30
user@host# set logical-systems LS2 interfaces lt-0/1/0 unit 1 family iso
```

4. If you are done configuring the device, commit the configuration.

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

### ***Verification***

Confirm that the configuration is working properly.

- [Verifying That the Logical Systems Are Up on page 3282](#)
- [Verifying Connectivity Between the Logical Systems on page 3282](#)

### ***Verifying That the Logical Systems Are Up***

**Purpose** Make sure that the interfaces are properly configured.

```
Action user@host> show interfaces terse
```

| Interface  | Admin | Link | Proto | Local        | Remote |
|------------|-------|------|-------|--------------|--------|
| so-0/0/0   | up    | down |       |              |        |
| so-0/0/1   | up    | down |       |              |        |
| so-0/0/2   | up    | down |       |              |        |
| so-0/0/3   | up    | down |       |              |        |
| gr-0/1/0   | up    | up   |       |              |        |
| ip-0/1/0   | up    | up   |       |              |        |
| lt-0/1/0   | up    | up   |       |              |        |
| lt-0/1/0.0 | up    | up   | inet  | 10.0.8.13/30 |        |
|            |       |      | iso   |              |        |
| lt-0/1/0.1 | up    | up   | inet  | 10.0.8.14/30 |        |
|            |       |      | iso   |              |        |
| ...        |       |      |       |              |        |

### ***Verifying Connectivity Between the Logical Systems***

**Purpose** Make sure that the network address appears as directly connected.



```

Action user@host> show route logical-system all
logical-system: LS1

inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.8.12/30 *[Direct/0] 00:00:34
 > via lt-0/1/0.0
10.0.8.13/32 *[Local/0] 00:00:34
 Local via lt-0/1/0.0

logical-system: LS2

inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.8.12/30 *[Direct/0] 00:00:34
 > via lt-0/1/0.1
10.0.8.14/32 *[Local/0] 00:00:34
 Local via lt-0/1/0.1
...

```

***Example: Connecting Logical Systems Within the Same Router Using Logical Tunnel Interfaces on MX Series Routers***

This example shows how to configure logical tunnel interfaces to connect two logical systems that are configured in a single MX Series 3D Universal Edge Router.

- [Requirements on page 3283](#)
- [Overview on page 3284](#)
- [Configuration on page 3284](#)
- [Verification on page 3286](#)

***Requirements***

The MX Series router chassis must have a DPC, MPC, or MIC installed and in the online state.

### Overview

To connect two logical systems, you configure a logical tunnel interface on both logical systems. Then you configure a peer relationship between the logical tunnel interfaces, thus creating a point-to-point connection. Logical tunnel interfaces behave like regular interfaces. You can configure them with Ethernet, Frame Relay, or another encapsulation type. You can also configure routing protocols across them. In effect, the logical tunnel (lt) interfaces connect two logical systems within the same router. The two logical systems do not share routing tables. This means that you can run dynamic routing protocols between different logical systems within the same router.

You must treat each interface like a point-to-point connection because you can only connect one logical tunnel interface to another at any given time. Also, you must select an interface encapsulation type, configure a corresponding protocol family, and set the logical interface unit number of the peering lt interface.


In this example, the logical tunnel interfaces are configured to behave as Ethernet interfaces with the **encapsulation ethernet** statement. The IS-IS Protocol is enabled on the logical tunnel interfaces with the **family iso** statement.

When configuring logical tunnel interfaces, note the following:

- The peering logical interfaces must have the same lt interface name. For example, a logical unit on lt-0/1/0.0 cannot peer with a logical unit on lt-0/0/10. The FPC (DPC, MPC, or MIC), PIC, and port numbers must match.
- The peering logical interfaces must be derived from the same module.
- You can configure only one peer unit for each logical interface. For example, unit 0 cannot peer with both unit 1 and unit 2.

Figure 56 on page 3284 shows the topology used in this example.

**Figure 56: Connecting Two Logical Systems**



### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set chassis fpc 0 pic 1 tunnel-services bandwidth 1g
set logical-systems LS1 interfaces lt-0/1/0 unit 0 description LS1->LS2
set logical-systems LS1 interfaces lt-0/1/0 unit 0 encapsulation ethernet
set logical-systems LS1 interfaces lt-0/1/0 unit 0 peer-unit 1
set logical-systems LS1 interfaces lt-0/1/0 unit 0 family inet address 10.0.8.13/30
set logical-systems LS1 interfaces lt-0/1/0 unit 0 family iso
set logical-systems LS2 interfaces lt-0/1/0 unit 1 description LS2->LS1
set logical-systems LS2 interfaces lt-0/1/0 unit 1 encapsulation ethernet
set logical-systems LS2 interfaces lt-0/1/0 unit 1 peer-unit 0
set logical-systems LS2 interfaces lt-0/1/0 unit 1 family inet address 10.0.8.14/30
set logical-systems LS2 interfaces lt-0/1/0 unit 1 family iso
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To connect logical system interfaces:

1. Run the **show chassis fpc** command to verify that the router has a DPC, MPC, or MIC installed and is in the online state.

```
user@host> show chassis fpc
```

| Slot | State  | Temp (C) | CPU Total | Utilization (%) Interrupt | Memory DRAM (MB) | Utilization (%) Heap | Buffer |
|------|--------|----------|-----------|---------------------------|------------------|----------------------|--------|
| 0    | Online | 31       | 4         | 0                         | 1024             | 14                   | 21     |
| 1    | Empty  |          |           |                           |                  |                      |        |
| 2    | Empty  |          |           |                           |                  |                      |        |

This output shows that slot 1 and slot 2 are empty. Slot 0 is online.

2. Configure FPC slot 0 to support logical tunnel (lt) interfaces.

```
[edit]
user@host# set chassis fpc 0 pic 1 tunnel-services bandwidth 1g
```

This command creates several tunnel interface types, including **gr**, **ip**, and **lt**. For this example, the important one is the logical tunnel (lt) interface.

3. Commit the configuration.

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

4. Run the **show interfaces terse** command to verify that the router has a logical tunnel (lt) interface.

```
user@host> show interfaces terse
```

| Interface | Admin | Link | Proto | Local | Remote |
|-----------|-------|------|-------|-------|--------|
| ...       |       |      |       |       |        |
| gr-0/1/0  | up    | up   |       |       |        |
| ip-0/1/0  | up    | up   |       |       |        |
| lt-0/1/0  | up    | up   |       |       |        |
| ...       |       |      |       |       |        |

- Configure the logical tunnel interface on Logical System LS1.

```
[edit]
user@host# set logical-systems LS1 interfaces lt-0/1/0 unit 0 description LS1->LS2
user@host# set logical-systems LS1 interfaces lt-0/1/0 unit 0 encapsulation ethernet
user@host# set logical-systems LS1 interfaces lt-0/1/0 unit 0 peer-unit 1
user@host# set logical-systems LS1 interfaces lt-0/1/0 unit 0 family inet address
10.0.8.13/30
user@host# set logical-systems LS1 interfaces lt-0/1/0 unit 0 family iso
```

- Configure the logical tunnel interface on Logical System LS2.

```
[edit]
user@host# set logical-systems LS2 interfaces lt-0/1/0 unit 1 description LS2->LS1
user@host# set logical-systems LS2 interfaces lt-0/1/0 unit 1 encapsulation ethernet
user@host# set logical-systems LS2 interfaces lt-0/1/0 unit 1 peer-unit 0
user@host# set logical-systems LS2 interfaces lt-0/1/0 unit 1 family inet address
10.0.8.14/30
user@host# set logical-systems LS2 interfaces lt-0/1/0 unit 1 family iso
```

- If you are done configuring the device, commit the configuration.

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

### Verification

Confirm that the configuration is working properly.

- [Verifying That the Logical Systems Are Up on page 3286](#)
- [Verifying Connectivity Between the Logical Systems on page 3286](#)

### Verifying That the Logical Systems Are Up

**Purpose** Make sure that the interfaces are properly configured.

```
Action user@host> show interfaces terse
```

| Interface  | Admin | Link | Proto | Local        | Remote |
|------------|-------|------|-------|--------------|--------|
| so-0/0/0   | up    | down |       |              |        |
| so-0/0/1   | up    | down |       |              |        |
| so-0/0/2   | up    | down |       |              |        |
| so-0/0/3   | up    | down |       |              |        |
| ge-0/1/0   | up    | up   |       |              |        |
| ip-0/1/0   | up    | up   |       |              |        |
| lt-0/1/0   | up    | up   |       |              |        |
| lt-0/1/0.0 | up    | up   | inet  | 10.0.8.13/30 |        |
|            |       |      | iso   |              |        |
| lt-0/1/0.1 | up    | up   | inet  | 10.0.8.14/30 |        |
|            |       |      | iso   |              |        |
| ...        |       |      |       |              |        |

### Verifying Connectivity Between the Logical Systems

**Purpose** Make sure that the network address appears as directly connected.

```

Action user@host> show route logical-system all
logical-system: LS1

inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.8.12/30 *[Direct/0] 00:00:34
 > via lt-0/1/0.0
10.0.8.13/32 *[Local/0] 00:00:34
 Local via lt-0/1/0.0

logical-system: LS2

inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.8.12/30 *[Direct/0] 00:00:34
 > via lt-0/1/0.1
10.0.8.14/32 *[Local/0] 00:00:34
 Local via lt-0/1/0.1
...

```

**Related Documentation**

- [Introduction to Logical Systems on page 3259](#)

### Example: Configuring Static Routing on Logical Systems

- [Understanding Basic Static Routing on page 3287](#)
- [Example: Configuring Static Routes Between Logical Systems Within the Same Router on page 3287](#)

#### *Understanding Basic Static Routing*

Routes that are permanent fixtures in the routing and forwarding tables are often configured as static routes. These routes generally do not change, and often include only one or very few paths to the destination.

To create a static route in the routing table, you must, at minimum, define the route as static and associate a next-hop address with it. The static route in the routing table is inserted into the forwarding table when the next-hop address is reachable. All traffic destined for the static route is transmitted to the next-hop address for transit.

You can specify options that define additional information about static routes that is included with the route when it is installed in the routing table. All static options are optional.

#### *Example: Configuring Static Routes Between Logical Systems Within the Same Router*

This example shows how to configure static routes between logical systems. The logical systems are configured in a single physical router and are connected by logical tunnel interfaces.

- [Requirements on page 3288](#)
- [Overview on page 3288](#)

- [Configuration on page 3288](#)
- [Verification on page 3292](#)

### Requirements


You must connect the logical systems by using logical tunnel (lt) interfaces. See “[Example: Connecting Logical Systems Within the Same Device Using Logical Tunnel Interfaces on MX Series Routers and EX Series Switches](#)” on page 3279.

### Overview

A static route is a hard-coded path in the device that specifies how the route gets to a certain subnet by using a certain path. Routers that are connected to stub networks are often configured to use static routes. A *stub network* is a network with no knowledge of other networks. Stub networks send non-local traffic by way of a single path, with the network aware only of a default route to non-local destinations. In this example, you configure Logical System LS1 with a static route to the 10.10.10.0/30 network and define the next-hop address as 192.168.10.2. You also configure Logical System LS1 with a static route to the 192.168.10.0/30 network and define a next-hop address of 10.10.10.1.

[Figure 57 on page 3288](#) shows the sample network.

**Figure 57: Static Routes Between Logical Systems**



### Configuration

#### CLI Quick Configuration

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

```

set logical-systems LS1 interfaces lt-1/2/0 unit 0 description LS1->LS2
set logical-systems LS1 interfaces lt-1/2/0 unit 0 encapsulation ethernet
set logical-systems LS1 interfaces lt-1/2/0 unit 0 peer-unit 1
set logical-systems LS1 interfaces lt-1/2/0 unit 0 family inet address 192.168.10.1/30
set logical-systems LS2 interfaces lt-1/2/0 unit 1 description LS2->LS1
set logical-systems LS2 interfaces lt-1/2/0 unit 1 encapsulation ethernet
set logical-systems LS2 interfaces lt-1/2/0 unit 1 peer-unit 0
set logical-systems LS2 interfaces lt-1/2/0 unit 1 family inet address 192.168.10.2/30
set logical-systems LS2 interfaces lt-1/2/0 unit 9 description LS2->LS3
set logical-systems LS2 interfaces lt-1/2/0 unit 9 encapsulation ethernet
set logical-systems LS2 interfaces lt-1/2/0 unit 9 peer-unit 10
set logical-systems LS2 interfaces lt-1/2/0 unit 9 family inet address 10.10.10.1/30
set logical-systems LS3 interfaces lt-1/2/0 unit 10 description LS3->LS2
set logical-systems LS3 interfaces lt-1/2/0 unit 10 encapsulation ethernet
set logical-systems LS3 interfaces lt-1/2/0 unit 10 peer-unit 9
set logical-systems LS3 interfaces lt-1/2/0 unit 10 family inet address 10.10.10.2/30
set logical-systems LS1 routing-options static route 10.10.10.0/30 next-hop 192.168.10.2
set logical-systems LS3 routing-options static route 192.168.10.0/30 next-hop 10.10.10.1

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure static routes between logical systems:

1. Run the **show interfaces terse** command to verify that the router has a logical tunnel (lt) interface.

```

user@host> show interfaces terse
Interface Admin Link Proto Local Remote
so-0/0/0 up down
so-0/0/1 up down
so-0/0/2 up down
so-0/0/3 up down
gr-1/2/0 up up
ip-1/2/0 up up
lt-1/2/0 up up
...

```

2. Configure the logical tunnel interface on Logical System LS1 connecting to Logical System LS2.

```

[edit]
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 0 description LS1->LS2
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 0 encapsulation ethernet
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 0 peer-unit 1
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 0 family inet address
192.168.10.1/30

```

3. Configure the logical tunnel interface on Logical System LS2 connecting to Logical System LS1.

```

[edit]
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 1 description LS2->LS1
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 1 encapsulation ethernet
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 1 peer-unit 0
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 1 family inet address
192.168.10.2/30

```

4. Configure the logical tunnel interface on Logical System LS2 connecting to Logical System LS3.

```
[edit]
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 9 description LS2->LS3
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 9 encapsulation ethernet
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 9 peer-unit 10
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 9 family inet address
10.10.10.1/30
```

5. Configure the logical tunnel interface on Logical System LS3 connecting to Logical System LS2.

```
[edit]
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 10 description LS3->LS2
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 10 encapsulation
ethernet
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 10 peer-unit 9
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 10 family inet address
10.10.10.2/30
```

6. Configure the static route on Logical System LS1 connecting to the 10.10.10.0/30 network.

```
[edit]
user@host# set logical-systems LS1 routing-options static route 10.10.10.0/30
next-hop 192.168.10.2
```

7. Configure the static route on Logical System LS3 connecting to the 192.168.10.0/30 network.

```
[edit]
user@host# set logical-systems LS3 routing-options static route 192.168.10.0/30
next-hop 10.10.10.1
```

8. If you are done configuring the device, commit the configuration.

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

### **Results**

Confirm your configuration by issuing the **show logical-systems** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show logical-systems
LS1 {
 interfaces {
 lt-1/2/0 {
 unit 0 {
 description LS1->LS2;
 encapsulation ethernet;
 peer-unit 1;
 family inet {
 address 192.168.10.1/30;
 }
 }
 }
 }
}
```



```
 }
 }
 routing-options {
 static {
 route 10.10.10.0/30 next-hop 192.168.10.2;
 }
 }
}
LS2 {
 interfaces {
 lt-1/2/0 {
 unit 1 {
 description LS2->LS1;
 encapsulation ethernet;
 peer-unit 0;
 family inet {
 address 192.168.10.2/30;
 }
 }
 }
 unit 9 {
 description LS2->LS3;
 encapsulation ethernet;
 peer-unit 10;
 family inet {
 address 10.10.10.1/30;
 }
 }
 }
}
LS3 {
 interfaces {
 lt-1/2/0 {
 unit 10 {
 description LS3->LS2;
 encapsulation ethernet;
 peer-unit 9;
 family inet {
 address 10.10.10.2/30;
 }
 }
 }
 }
 routing-options {
 static {
 route 192.168.10.0/30 next-hop 10.10.10.1;
 }
 }
}
```

**Verification**

Confirm that the configuration is working properly.

- [Verifying That the Logical Systems Are Up on page 3292](#)
- [Verifying Connectivity Between the Logical Systems on page 3292](#)

**Verifying That the Logical Systems Are Up**

**Purpose** Make sure that the interfaces are properly configured.

**Action** user@host> show interfaces terse

| Interface   | Admin | Link | Proto | Local           | Remote |
|-------------|-------|------|-------|-----------------|--------|
| ...         |       |      |       |                 |        |
| lt-1/2/0    | up    | up   |       |                 |        |
| lt-1/2/0.0  | up    | up   | inet  | 192.168.10.1/30 |        |
| lt-1/2/0.1  | up    | up   | inet  | 192.168.10.2/30 |        |
| lt-1/2/0.9  | up    | up   | inet  | 10.10.10.1/30   |        |
| lt-1/2/0.10 | up    | up   | inet  | 10.10.10.2/30   |        |
| ...         |       |      |       |                 |        |

**Verifying Connectivity Between the Logical Systems**

**Purpose** Make sure that the static routes appear in the routing tables of Logical Systems LS1 and LS3. Also, make sure that the logical systems can ping each other.

```

Action user@host> show route logical-system LS1
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.0/30 *[Static/5] 18:43:25
 > to 192.168.10.2 via lt-1/2/0.0
192.168.10.0/30 *[Direct/0] 18:43:25
 > via lt-1/2/0.0
192.168.10.1/32 *[Local/0] 18:43:25
 Local via lt-1/2/0.0

user@host> show route logical-system LS3
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.0/30 *[Direct/0] 23:11:21
 > via lt-1/2/0.10
10.10.10.2/32 *[Local/0] 23:11:21
 Local via lt-1/2/0.10
192.168.10.0/30 *[Static/5] 00:23:31
 > to 10.10.10.1 via lt-1/2/0.10

```

#### From LS1, Ping LS3

```

user@host> set cli logical-system LS1

user@host:LS1> ping 10.10.10.2
PING 10.10.10.2 (10.10.10.2): 56 data bytes
64 bytes from 10.10.10.2: icmp_seq=0 ttl=63 time=1.263 ms
64 bytes from 10.10.10.2: icmp_seq=1 ttl=63 time=1.086 ms
64 bytes from 10.10.10.2: icmp_seq=2 ttl=63 time=1.077 ms

```

#### From LS3, Ping LS1

```

user@host> set cli logical-system LS3

user@host:LS3> ping 192.168.10.1
PING 192.168.10.1 (192.168.10.1): 56 data bytes
64 bytes from 192.168.10.1: icmp_seq=0 ttl=63 time=10.781 ms
64 bytes from 192.168.10.1: icmp_seq=1 ttl=63 time=1.167 ms
64 bytes from 192.168.10.1: icmp_seq=2 ttl=63 time=1.152 ms

```

**Related Documentation**

- [Introduction to Logical Systems on page 3259](#)

#### Examples: Configuring Standard Firewall Filters on Logical Systems

- [Understanding How to Use Standard Firewall Filters on page 3293](#)
- [Example: Configuring a Stateless Firewall Filter to Protect a Logical System Against ICMP Floods on page 3295](#)
- [Example: Configuring Filter-Based Forwarding on Logical Systems on page 3298](#)

#### ***Understanding How to Use Standard Firewall Filters***

This topic covers the following information:

- [Using Standard Firewall Filters to Affect Local Packets on page 3294](#)
- [Using Standard Firewall Filters to Affect Data Packets on page 3294](#)

### ***Using Standard Firewall Filters to Affect Local Packets***

On a router (or switch), you can configure one physical loopback interface, **lo0**, and one or more addresses on the interface. The loopback interface is the interface to the Routing Engine, which runs and monitors all the control protocols. The loopback interface carries local packets only. Standard firewall filters applied to the loopback interface affect the local packets destined for or transmitted from the Routing Engine.



**NOTE:** When you create an additional loopback interface, it is important to apply a filter to it so the Routing Engine is protected. We recommend that when you apply a filter to the loopback interface, you include the **apply-groups** statement. Doing so ensures that the filter is automatically inherited on every loopback interface, including lo0 and other loopback interfaces.

### ***Trusted Sources***

The typical use of a standard stateless firewall filter is to protect the Routing Engine processes and resources from malicious or untrusted packets. To protect the processes and resources owned by the Routing Engine, you can use a standard stateless firewall filter that specifies which protocols and services, or applications, are allowed to reach the Routing Engine. Applying this type of filter to the loopback interface ensures that the local packets are from a trusted source and protects the processes running on the Routing Engine from an external attack.

### ***Flood Prevention***

You can create standard stateless firewall filters that limit certain TCP and ICMP traffic destined for the Routing Engine. A router (or switch) without this kind of protection is vulnerable to TCP and ICMP flood attacks, which are also called denial-of-service (DoS) attacks. For example:

- A TCP flood attack of SYN packets initiating connection requests can overwhelm the device until it can no longer process legitimate connection requests, resulting in denial of service.
- An ICMP flood can overload the device with so many echo requests (ping requests) that it expends all its resources responding and can no longer process valid network traffic, also resulting in denial of service.

Applying the appropriate firewall filters to the Routing Engine protects against these types of attacks.

### ***Using Standard Firewall Filters to Affect Data Packets***

Standard firewall filters that you apply to your router's (or switch's) transit interfaces evaluate only the user data packets that transit the router (or switch) from one interface directly to another as they are being forwarded from a source to a destination. To protect the network as a whole from unauthorized access and other threats at specific interfaces, you can apply firewall filters router (or switch) transit interfaces .

### Example: Configuring a Stateless Firewall Filter to Protect a Logical System Against ICMP Floods

This example shows how to configure a stateless firewall filter that protects against ICMP denial-of-service attacks on a logical system.

- [Requirements on page 3295](#)
- [Overview on page 3295](#)
- [Configuration on page 3295](#)
- [Verification on page 3297](#)

#### Requirements

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

#### Overview


This example shows a stateless firewall filter called **protect-RE** that polices ICMP packets. The **icmp-policer** limits the traffic rate of the ICMP packets to 1,000,000 bps and the burst size to 15,000 bytes. Packets that exceed the traffic rate are discarded.

The policer is incorporated into the action of a filter term called **icmp-term**.

In this example, a ping is sent from a directly connected physical router to the interface configured on the logical system. The logical system accepts the ICMP packets if they are received at a rate of up to 1 Mbps (bandwidth-limit). The logical system drops all ICMP packets when this rate is exceeded. The **burst-size-limit** statement accepts traffic bursts up to 15 Kbps. If bursts exceed this limit, all packets are dropped. When the flow rate subsides, ICMP packets are again accepted.

[Figure 58 on page 3295](#) shows the topology used in this example.

**Figure 58: Logical System with a Stateless Firewall**



#### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set logical-systems LS1 interfaces so-0/0/2 unit 0 family inet policer input icmp-policer
set logical-systems LS1 interfaces so-0/0/2 unit 0 family inet address 10.0.45.2/30
set logical-systems LS1 firewall family inet filter protect-RE term icmp-term from protocol
icmp
set logical-systems LS1 firewall family inet filter protect-RE term icmp-term then policer
icmp-policer
set logical-systems LS1 firewall family inet filter protect-RE term icmp-term then accept
set logical-systems LS1 firewall policer icmp-policer if-exceeding bandwidth-limit 1m
set logical-systems LS1 firewall policer icmp-policer if-exceeding burst-size-limit 15k
set logical-systems LS1 firewall policer icmp-policer then discard
```

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure an ICMP firewall filter on a logical system:

1. Configure the interface on the logical system.

```
[edit]
user@host# set logical-systems LS1 interfaces so-0/0/2 unit 0 family inet address
10.0.45.2/30
```

2. Explicitly enable ICMP packets to be received on the interface.

```
[edit]
user@host# set logical-systems LS1 firewall family inet filter protect-RE term
icmp-term from protocol icmp
user@host# set logical-systems LS1 firewall family inet filter protect-RE term
icmp-term then accept
```

3. Create the policer.

```
[edit]
user@host# set logical-systems LS1 firewall policer icmp-policer if-exceeding
bandwidth-limit 1m
user@host# set logical-systems LS1 firewall policer icmp-policer if-exceeding
burst-size-limit 15k
user@host# set logical-systems LS1 firewall policer icmp-policer then discard
```

4. Apply the policer to a filter term.

```
[edit]
user@host# set logical-systems LS1 firewall family inet filter protect-RE term
icmp-term then policer icmp-policer
```

5. Apply the policer to the logical system interface.

```
[edit]
user@host# set logical-systems LS1 interfaces so-0/0/2 unit 0 family inet policer
input icmp-policer
```

6. If you are done configuring the device, commit the configuration.

```
[edit]
```

```
user@host# commit
```

### Results

Confirm your configuration by issuing the **show logical-systems LS1** command.

```
user@host# show logical-systems LS1
interfaces {
 so-0/0/2 {
 unit 0 {
 family inet {
 policer {
 input icmp-policer;
 }
 address 10.0.45.2/30;
 }
 }
 }
}
firewall {
 family inet {
 filter protect-RE {
 term icmp-term {
 from {
 protocol icmp;
 }
 then {
 policer icmp-policer;
 accept;
 }
 }
 }
 }
}
policer icmp-policer {
 if-exceeding {
 bandwidth-limit 1m;
 burst-size-limit 15k;
 }
 then discard;
}
}
```

### Verification

Confirm that the configuration is working properly.

#### Verifying That Ping Works Unless the Limits Are Exceeded

**Purpose** Make sure that the logical system interface is protected against ICMP-based DoS attacks.

**Action** Log in to a system that has connectivity to the logical system and run the **ping** command.

```
user@R2> ping 10.0.45.2
PING 10.0.45.2 (10.0.45.2): 56 data bytes
64 bytes from 10.0.45.2: icmp_seq=0 ttl=64 time=1.316 ms
64 bytes from 10.0.45.2: icmp_seq=1 ttl=64 time=1.277 ms
64 bytes from 10.0.45.2: icmp_seq=2 ttl=64 time=1.269 ms
```

```
user@R2> ping 10.0.45.2 size 20000
PING 10.0.45.2 (10.0.45.2): 20000 data bytes
^C
--- 10.0.45.2 ping statistics ---
4 packets transmitted, 0 packets received, 100% packet loss
```

**Meaning** When you send a normal ping, the packet is accepted. When you send a ping packet that exceeds the filter limit, the packet is discarded.

### ***Example: Configuring Filter-Based Forwarding on Logical Systems***

This example shows how to configure filter-based forwarding within a logical system. The filter classifies packets to determine their forwarding path within the ingress routing device.

- [Requirements on page 3298](#)
- [Overview on page 3298](#)
- [Configuration on page 3300](#)
- [Verification on page 3306](#)

### ***Requirements***

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

### ***Overview***

Filter-based forwarding is supported for IP version 4 (IPv4) and IP version 6 (IPv6).

Use filter-based forwarding for service provider selection when customers have Internet connectivity provided by different ISPs yet share a common access layer. When a shared media (such as a cable modem) is used, a mechanism on the common access layer looks at Layer 2 or Layer 3 addresses and distinguishes between customers. You can use filter-based forwarding when the common access layer is implemented using a combination of Layer 2 switches and a single router.

With filter-based forwarding, all packets received on an interface are considered. Each packet passes through a filter that has match conditions. If the match conditions are met for a filter and you have created a routing instance, filter-based forwarding is applied to a packet. The packet is forwarded based on the next hop specified in the routing instance. For static routes, the next hop can be a specific LSP.



**NOTE:** Source-class usage filter matching and unicast reverse-path forwarding checks are not supported on an interface configured with filter-based forwarding (FBF).

---

To configure filter-based forwarding, perform the following tasks:

- Create a match filter on an ingress router or switch. To specify a match filter, include the **filter filter-name** statement at the **[edit firewall]** hierarchy level. A packet that passes through the filter is compared against a set of rules to classify it and to determine its



membership in a set. Once classified, the packet is forwarded to a routing table specified in the accept action in the filter description language. The routing table then forwards the packet to the next hop that corresponds to the destination address entry in the table.

- Create routing instances that specify the routing table(s) to which a packet is forwarded, and the destination to which the packet is forwarded at the **[edit routing-instances]** or **[edit logical-systems *logical-system-name* routing-instances]** hierarchy level. For example:

```
[edit]
routing-instances {
 routing-table-name1 {
 instance-type forwarding;
 routing-options {
 static {
 route 0.0.0.0/0 nexthop 10.0.0.1;
 }
 }
 }
 routing-table-name2 {
 instance-type forwarding;
 routing-options {
 static {
 route 0.0.0.0/0 nexthop 10.0.0.2;
 }
 }
 }
}
```

- Create a routing table group that adds interface routes to the forwarding routing instances used in filter-based forwarding (FBF), as well as to the default routing instance **inet.0**. This part of the configuration resolves the routes installed in the routing instances to directly connected next hops on that interface. Create the routing table group at the **[edit routing-options]** or **[edit logical-systems *logical-system-name* routing-options]** hierarchy level.



**NOTE:** Specify **inet.0** as one of the routing instances that the interface routes are imported into. If the default instance **inet.0** is not specified, interface routes are not imported into the default routing instance.

This example shows a packet filter that directs customer traffic to a next-hop router in the domains, SP 1 or SP 2, based on the packet's source address.

If the packet has a source address assigned to an SP 1 customer, destination-based forwarding occurs using the **sp1-route-table.inet.0** routing table. If the packet has a source address assigned to an SP 2 customer, destination-based forwarding occurs using the **sp2-route-table.inet.0** routing table. If a packet does not match either of these conditions, the filter accepts the packet, and destination-based forwarding occurs using the standard **inet.0** routing table.

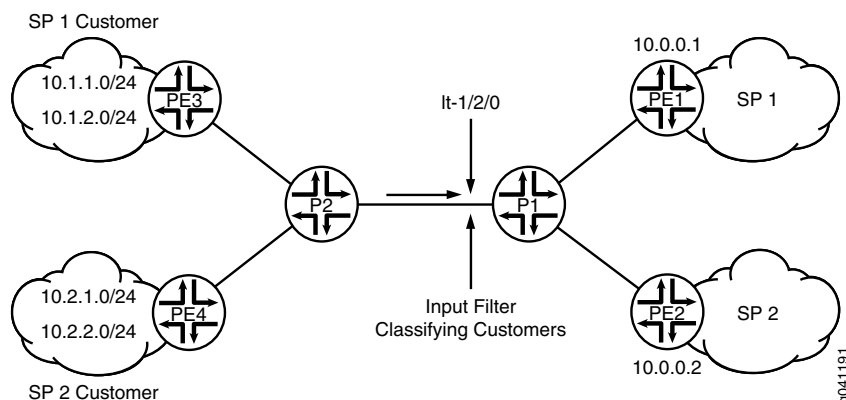
One way to make filter-based forwarding work within a logical system is to configure the firewall filter on the logical system that receives the packets. Another way is to configure the firewall filter on the main router or switch and then reference the logical system in the firewall filter. This example uses the second approach. The specific routing instances are configured within the logical system. Because each routing instance has its own routing table, you have to reference the routing instances in the firewall filter, as well. The syntax looks as follows:

```
[edit firewall filter filter-name term term-name]
user@host# set then logical-system logical-system-name routing-instance
routing-instance-name
```

Figure 59 on page 3300 shows the topology used in this example.

On Logical System P1, an input filter classifies packets received from Logical System PE3 and Logical System PE4. The packets are routed based on the source addresses. Packets with source addresses in the 10.1.1.0/24 and 10.1.2.0/24 networks are routed to Logical System PE1. Packets with source addresses in the 10.2.1.0/24 and 10.2.2.0/24 networks are routed to Logical System PE2.

**Figure 59: Logical Systems with Filter-Based Forwarding**



To establish connectivity, OSPF is configured on all of the interfaces. For demonstration purposes, loopback interface addresses are configured on the routing devices to represent networks in the clouds.

The “CLI Quick Configuration” on page 3300 section shows the entire configuration for all of the devices in the topology. The “Configuring the Routing Instances on the Logical System P1” on page 3303 and “Configuring the Firewall Filter on the Main Router” on page 3302 sections shows the step-by-step configuration of the ingress routing device, Logical System P1.

### Configuration

#### CLI Quick Configuration

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

```
set firewall filter classify-customers term sp1-customers from source-address 10.1.1.0/24
```

```

set firewall filter classify-customers term sp1-customers from source-address 10.1.2.0/24
set firewall filter classify-customers term sp1-customers then log
set firewall filter classify-customers term sp1-customers then logical-system P1
 routing-instance sp1-route-table
set firewall filter classify-customers term sp2-customers from source-address 10.2.1.0/24
set firewall filter classify-customers term sp2-customers from source-address 10.2.2.0/24
set firewall filter classify-customers term sp2-customers then log
set firewall filter classify-customers term sp2-customers then logical-system P1
 routing-instance sp2-route-table
set firewall filter classify-customers term default then accept
set logical-systems P1 interfaces lt-1/2/0 unit 10 encapsulation ethernet
set logical-systems P1 interfaces lt-1/2/0 unit 10 peer-unit 9
set logical-systems P1 interfaces lt-1/2/0 unit 10 family inet filter input classify-customers
set logical-systems P1 interfaces lt-1/2/0 unit 10 family inet address 172.16.0.10/30
set logical-systems P1 interfaces lt-1/2/0 unit 13 encapsulation ethernet
set logical-systems P1 interfaces lt-1/2/0 unit 13 peer-unit 14
set logical-systems P1 interfaces lt-1/2/0 unit 13 family inet address 172.16.0.13/30
set logical-systems P1 interfaces lt-1/2/0 unit 17 encapsulation ethernet
set logical-systems P1 interfaces lt-1/2/0 unit 17 peer-unit 18
set logical-systems P1 interfaces lt-1/2/0 unit 17 family inet address 172.16.0.17/30
set logical-systems P1 protocols ospf rib-group fbf-group
set logical-systems P1 protocols ospf area 0.0.0.0 interface all
set logical-systems P1 protocols ospf area 0.0.0.0 interface fxp0.0 disable
set logical-systems P1 routing-instances sp1-route-table instance-type forwarding
set logical-systems P1 routing-instances sp1-route-table routing-options static route
 0.0.0.0/0 next-hop 172.16.0.13
set logical-systems P1 routing-instances sp2-route-table instance-type forwarding
set logical-systems P1 routing-instances sp2-route-table routing-options static route
 0.0.0.0/0 next-hop 172.16.0.17
set logical-systems P1 routing-options rib-groups fbf-group import-rib inet.0
set logical-systems P1 routing-options rib-groups fbf-group import-rib
 sp1-route-table.inet.0
set logical-systems P1 routing-options rib-groups fbf-group import-rib
 sp2-route-table.inet.0
set logical-systems P2 interfaces lt-1/2/0 unit 2 encapsulation ethernet
set logical-systems P2 interfaces lt-1/2/0 unit 2 peer-unit 1
set logical-systems P2 interfaces lt-1/2/0 unit 2 family inet address 172.16.0.2/30
set logical-systems P2 interfaces lt-1/2/0 unit 6 encapsulation ethernet
set logical-systems P2 interfaces lt-1/2/0 unit 6 peer-unit 5
set logical-systems P2 interfaces lt-1/2/0 unit 6 family inet address 172.16.0.6/30
set logical-systems P2 interfaces lt-1/2/0 unit 9 encapsulation ethernet
set logical-systems P2 interfaces lt-1/2/0 unit 9 peer-unit 10
set logical-systems P2 interfaces lt-1/2/0 unit 9 family inet address 172.16.0.9/30
set logical-systems P2 protocols ospf area 0.0.0.0 interface all
set logical-systems P2 protocols ospf area 0.0.0.0 interface fxp0.0 disable
set logical-systems PE1 interfaces lt-1/2/0 unit 14 encapsulation ethernet
set logical-systems PE1 interfaces lt-1/2/0 unit 14 peer-unit 13
set logical-systems PE1 interfaces lt-1/2/0 unit 14 family inet address 172.16.0.14/30
set logical-systems PE1 interfaces lo0 unit 3 family inet address 1.1.1.1/32
set logical-systems PE1 protocols ospf area 0.0.0.0 interface all
set logical-systems PE1 protocols ospf area 0.0.0.0 interface fxp0.0 disable
set logical-systems PE2 interfaces lt-1/2/0 unit 18 encapsulation ethernet
set logical-systems PE2 interfaces lt-1/2/0 unit 18 peer-unit 17
set logical-systems PE2 interfaces lt-1/2/0 unit 18 family inet address 172.16.0.18/30
set logical-systems PE2 interfaces lo0 unit 4 family inet address 2.2.2.2/32
set logical-systems PE2 protocols ospf area 0.0.0.0 interface all

```

```

set logical-systems PE2 protocols ospf area 0.0.0.0 interface fxp0.0 disable
set logical-systems PE3 interfaces lt-1/2/0 unit 1 encapsulation ethernet
set logical-systems PE3 interfaces lt-1/2/0 unit 1 peer-unit 2
set logical-systems PE3 interfaces lt-1/2/0 unit 1 family inet address 172.16.0.1/30
set logical-systems PE3 interfaces lo0 unit 1 family inet address 10.1.1.1/32
set logical-systems PE3 interfaces lo0 unit 1 family inet address 10.1.2.1/32
set logical-systems PE3 protocols ospf area 0.0.0.0 interface all
set logical-systems PE3 protocols ospf area 0.0.0.0 interface fxp0.0 disable
set logical-systems PE4 interfaces lt-1/2/0 unit 5 encapsulation ethernet
set logical-systems PE4 interfaces lt-1/2/0 unit 5 peer-unit 6
set logical-systems PE4 interfaces lt-1/2/0 unit 5 family inet address 172.16.0.5/30
set logical-systems PE4 interfaces lo0 unit 2 family inet address 10.2.1.1/32
set logical-systems PE4 interfaces lo0 unit 2 family inet address 10.2.2.1/32
set logical-systems PE4 protocols ospf area 0.0.0.0 interface all
set logical-systems PE4 protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

### *Configuring the Firewall Filter on the Main Router*

#### **Step-by-Step Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure the firewall filter on the main router or switch:

1. Configure the source addresses for SP1 customers.  

```

[edit firewall filter classify-customers term sp1-customers]
user@host# set from source-address 10.1.1.0/24
user@host# set from source-address 10.1.2.0/24

```
2. Configure the actions that are taken when packets are received with the specified source addresses.

To track the action of the firewall filter, a log action is configured. The sp1-route-table.inet.0 routing table on Logical System P1 routes the packets.

```

[edit firewall filter classify-customers term sp1-customers]
user@host# set then log
user@host# set then logical-system P1 routing-instance sp1-route-table

```

3. Configure the source addresses for SP2 customers.  

```

[edit firewall filter classify-customers term sp2-customers]
user@host# set from source-address 10.2.1.0/24
user@host# set from source-address 10.2.2.0/24

```
4. Configure the actions that are taken when packets are received with the specified source addresses.

To track the action of the firewall filter, a log action is configured. The sp2-route-table.inet.0 routing table on Logical System P1 routes the packet.

```

[edit firewall filter classify-customers term sp2-customers]
user@host# set then log
user@host# set then logical-system P1 routing-instance sp2-route-table

```

5. Configure the action to take when packets are received from any other source address.

All of these packets are simply accepted and routed using the default IPv4 unicast routing table, inet.0.

```
[edit firewall filter classify-customers term default]
user@host# set then accept
```

### *Configuring the Routing Instances on the Logical System P1*

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure the routing instances on a logical system:

1. Configure the interfaces on the logical system.

```
[edit logical-systems P1 interfaces lt-1/2/0]
user@host# set unit 10 encapsulation ethernet
user@host# set unit 10 peer-unit 9
user@host# set unit 10 family inet address 172.16.0.10/30
```

```
user@host# set unit 13 encapsulation ethernet
user@host# set unit 13 peer-unit 14
user@host# set unit 13 family inet address 172.16.0.13/30
```

```
user@host# set unit 17 encapsulation ethernet
user@host# set unit 17 peer-unit 18
user@host# set unit 17 family inet address 172.16.0.17/30
```

2. Assign the **classify-customers** firewall filter to router interface lt-1/2/0.10 as an input packet filter.

```
[edit logical-systems P1 interfaces lt-1/2/0]
user@host# set unit 10 family inet filter input classify-customers
```

3. Configure connectivity, using either a routing protocol or static routing.

As a best practice, disable routing on the management interface.

```
[edit logical-systems P1 protocols ospf area 0.0.0.0]
user@host# set interface all
user@host# set interface fxp0.0 disable
```

4. Create the routing instances.

These routing instances are referenced in the **classify-customers** firewall filter.

The forwarding instance type provides support for filter-based forwarding, where interfaces are not associated with instances. All interfaces belong to the default instance, in this case Logical System P1.

```
[edit logical-systems P1 routing-instances]
user@host# set sp1-route-table instance-type forwarding

user@host# set sp2-route-table instance-type forwarding
```

5. Resolve the routes installed in the routing instances to directly connected next hops.

```
[edit logical-systems P1 routing-instances]
user@host# set sp1-route-table routing-options static route 0.0.0.0/0 next-hop
172.16.0.13
```

```
user@host# set sp2-route-table routing-options static route 0.0.0.0/0 next-hop
172.16.0.17
```

6. Group together the routing tables to form a routing table group.

The first routing table, inet.0, is the primary routing table, and the additional routing tables are the secondary routing tables.

The primary routing table determines the address family of the routing table group, in this case IPv4.

```
[edit logical-systems P1 routing-options]
user@host# set rib-groups fbf-group import-rib inet.0
user@host# set rib-groups fbf-group import-rib sp1-route-table.inet.0
user@host# set rib-groups fbf-group import-rib sp2-route-table.inet.0
```

7. Apply the routing table group to OSPF.

This causes the OSPF routes to be installed into all the routing tables in the group.

```
[edit logical-systems P1 protocols ospf]
user@host# set rib-group fbf-group
```

8. If you are done configuring the device, commit the configuration.

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

## Results

Confirm your configuration by issuing the **show firewall** and **show logical-systems P1** commands.

```
user@host# show firewall
filter classify-customers {
 term sp1-customers {
 from {
 source-address {
 10.1.1.0/24;
 10.1.2.0/24;
 }
 }
 then {
 log;
 logical-system P1 routing-instance sp1-route-table;
 }
 }
}
term sp2-customers {
 from {
 source-address {
 10.2.1.0/24;
 10.2.2.0/24;
 }
 }
}
```

```

 }
 }
 then {
 log;
 logical-system P1 routing-instance sp2-route-table;
 }
}
term default {
 then accept;
}
}

user@host# show logical-systems P1
interfaces {
 lt-1/2/0 {
 unit 10 {
 encapsulation ethernet;
 peer-unit 9;
 family inet {
 filter {
 input classify-customers;
 }
 address 172.16.0.10/30;
 }
 }
 unit 13 {
 encapsulation ethernet;
 peer-unit 14;
 family inet {
 address 172.16.0.13/30;
 }
 }
 unit 17 {
 encapsulation ethernet;
 peer-unit 18;
 family inet {
 address 172.16.0.17/30;
 }
 }
 }
}
protocols {
 ospf {
 rib-group fbf-group;
 area 0.0.0.0 {
 interface all;
 interface fxp0.0 {
 disable;
 }
 }
 }
}
routing-instances {
 sp1-route-table {
 instance-type forwarding;
 routing-options {

```

```
 static {
 route 0.0.0.0/0 next-hop 172.16.0.13;
 }
 }
 sp2-route-table {
 instance-type forwarding;
 routing-options {
 static {
 route 0.0.0.0/0 next-hop 172.16.0.17;
 }
 }
 }
}
routing-options {
 rib-groups {
 fbf-group {
 import-rib [inet.0 sp1-route-table.inet.0 sp2-route-table.inet.0];
 }
 }
}
```

### **Verification**

Confirm that the configuration is working properly.

### **Pinging with Specified Source Addresses**

**Purpose** Send some ICMP packets across the network to test the firewall filter.

**Action** 1. Log in to Logical System PE3.

```
user@host> set cli logical-system PE3
Logical system: PE3
```

2. Run the **ping** command, pinging the lo0.3 interface on Logical System PE1.

The address configured on this interface is 1.1.1.1.

Specify the source address 10.1.2.1, which is the address configured on the lo0.1 interface on Logical System PE3.

```
user@host:PE3> ping 1.1.1.1 source 10.1.2.1
PING 1.1.1.1 (1.1.1.1): 56 data bytes
64 bytes from 1.1.1.1: icmp_seq=0 ttl=62 time=1.444 ms
64 bytes from 1.1.1.1: icmp_seq=1 ttl=62 time=2.094 ms
^C
--- 1.1.1.1 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.444/1.769/2.094/0.325 ms
```

3. Log in to Logical System PE4.

```
user@host:PE3> set cli logical-system PE4
Logical system: PE4
```

4. Run the **ping** command, pinging the lo0.4 interface on Logical System PE2.

The address configured on this interface is 2.2.2.2.



Specify the source address 10.2.1.1, which is the address configured on the lo0.2 interface on Logical System PE4.

```
user@host:PE4> ping 2.2.2.2 source 10.2.1.1
PING 2.2.2.2 (2.2.2.2): 56 data bytes
64 bytes from 2.2.2.2: icmp_seq=0 ttl=62 time=1.473 ms
64 bytes from 2.2.2.2: icmp_seq=1 ttl=62 time=1.407 ms
^C
--- 2.2.2.2 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.407/1.440/1.473/0.033 ms
```

**Meaning** Sending these pings activates the firewall filter actions.

### *Verifying the Firewall Filter*

**Purpose** Make sure the firewall filter actions take effect.

**Action** 1. Log in to Logical System P1.

```
user@host> set cli logical-system P1
Logical system: P1
```

2. Run the **show firewall log** command on Logical System P1.

```
user@host:P1> show firewall log
Log :
Time Filter Action Interface Protocol Src Addr
Dest Addr
13:52:20 pfe A 1t-1/2/0.10 ICMP 10.2.1.1
2.2.2.2
13:52:19 pfe A 1t-1/2/0.10 ICMP 10.2.1.1
2.2.2.2
13:51:53 pfe A 1t-1/2/0.10 ICMP 10.1.2.1
1.1.1.1
13:51:52 pfe A 1t-1/2/0.10 ICMP 10.1.2.1
1.1.1.1
```

**Related Documentation**

- [Introduction to Logical Systems on page 3259](#)
- [Example: Configuring Multitopology Routing Based on Applications](#)

### Example: Configuring OSPF on Logical Systems

- [OSPF Overview on page 3308](#)
- [Example: Configuring OSPF on Logical Systems Within the Same Router on page 3312](#)

## *OSPF Overview*

OSPF is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). OSPF uses link-state information to make routing decisions, making route calculations using the shortest-path-first (SPF) algorithm (also referred to as the Dijkstra algorithm). Each router running OSPF floods link-state advertisements throughout the AS or area that contain information about that router's attached interfaces and routing metrics. Each router uses the information in these link-state advertisements to calculate the least cost path to each network and create a routing table for the protocol.

Junos OS supports OSPF version 2 (OSPFv2) and OSPF version 3 (OSPFv3), including virtual links, stub areas, and for OSPFv2, authentication. Junos OS does not support type-of-service (ToS) routing.

OSPF was designed for the Transmission Control Protocol/Internet Protocol (TCP/IP) environment and as a result explicitly supports IP subnetting and the tagging of externally derived routing information. OSPF also provides for the authentication of routing updates.

OSPF routes IP packets based solely on the destination IP address contained in the IP packet header. OSPF quickly detects topological changes, such as when router interfaces become unavailable, and calculates new loop-free routes quickly and with a minimum of routing overhead traffic.



**NOTE:** On SRX Series devices, when only one link-protection is configured under the OSPF interface, the device does not install an alternative route in the forwarding table. When the per-packet load-balancing is enabled as a workaround, the device does not observe both the OSPF metric and sending the traffic through both the interfaces.

An OSPF AS can consist of a single area, or it can be subdivided into multiple areas. In a single-area OSPF network topology, each router maintains a database that describes the topology of the AS. Link-state information for each router is flooded throughout the AS. In a multiarea OSPF topology, each router maintains a database that describes the topology of its area, and link-state information for each router is flooded throughout that area. All routers maintain summarized topologies of other areas within an AS. Within each area, OSPF routers have identical topological databases. When the AS or area topology changes, OSPF ensures that the contents of all routers' topological databases converge quickly.

All OSPFv2 protocol exchanges can be authenticated. OSPFv3 relies on IPsec to provide this functionality. This means that only trusted routers can participate in the AS's routing. A variety of authentication schemes can be used. A single authentication scheme is configured for each area, which enables some areas to use stricter authentication than others.

Externally derived routing data (for example, routes learned from BGP) is passed transparently throughout the AS. This externally derived data is kept separate from the OSPF link-state data. Each external route can be tagged by the advertising router, enabling the passing of additional information between routers on the boundaries of the AS.



**NOTE:** By default, Junos OS is compatible with RFC 1583, *OSPF Version 2*. In Junos OS Release 8.5 and later, you can disable compatibility with RFC 1583 by including the `no-rfc-1583` statement. For more information, see *Example: Disabling OSPFv2 Compatibility with RFC 1583*.

This topic describes the following information:

- [OSPF Default Route Preference Values on page 3310](#)
- [OSPF Routing Algorithm on page 3310](#)
- [OSPF Three-Way Handshake on page 3311](#)
- [OSPF Version 3 on page 3312](#)

#### ***OSPF Default Route Preference Values***

The Junos OS routing protocol process assigns a default preference value to each route that the routing table receives. The default value depends on the source of the route. The preference value is from 0 through 4,294,967,295 (232 – 1), with a lower value indicating a more preferred route. [Table 266 on page 3310](#) lists the default preference values for OSPF.

**Table 266: Default Route Preference Values for OSPF**

| How Route Is Learned    | Default Preference | Statement to Modify Default Preference   |
|-------------------------|--------------------|------------------------------------------|
| OSPF internal route     | 10                 | OSPF <a href="#">preference</a>          |
| OSPF AS external routes | 150                | OSPF <a href="#">external-preference</a> |

#### ***OSPF Routing Algorithm***

OSPF uses the shortest-path-first (SPF) algorithm, also referred to as the Dijkstra algorithm, to determine the route to each destination. All routing devices in an area run this algorithm in parallel, storing the results in their individual topological databases. Routing devices with interfaces to multiple areas run multiple copies of the algorithm. This section provides a brief summary of how the SPF algorithm works.

When a routing device starts, it initializes OSPF and waits for indications from lower-level protocols that the router interfaces are functional. The routing device then uses the OSPF hello protocol to acquire neighbors, by sending hello packets to its neighbors and receiving their hello packets.

On broadcast or nonbroadcast multiaccess networks (physical networks that support the attachment of more than two routing devices), the OSPF hello protocol elects a designated router for the network. This routing device is responsible for sending *link-state advertisements* (LSAs) that describe the network, which reduces the amount of network traffic and the size of the routing devices' topological databases.

The routing device then attempts to form *adjacencies* with some of its newly acquired neighbors. (On multiaccess networks, only the designated router and backup designated router form adjacencies with other routing devices.) Adjacencies determine the distribution

of routing protocol packets. Routing protocol packets are sent and received only on adjacencies, and topological database updates are sent only along adjacencies. When adjacencies have been established, pairs of adjacent routers synchronize their topological databases.

A routing device sends LSA packets to advertise its state periodically and when its state changes. These packets include information about the routing device's adjacencies, which allows detection of nonoperational routing devices.

Using a reliable algorithm, the routing device floods LSAs throughout the area, which ensures that all routing devices in an area have exactly the same topological database. Each routing device uses the information in its topological database to calculate a shortest-path tree, with itself as the root. The routing device then uses this tree to route network traffic.


The description of the SPF algorithm up to this point has explained how the algorithm works within a single area (*intra-area routing*). For internal routers to be able to route to destinations outside the area (*interarea routing*), the area border routers must inject additional routing information into the area. Because the area border routers are connected to the backbone, they have access to complete topological data about the backbone. The area border routers use this information to calculate paths to all destinations outside its area and then advertise these paths to the area's internal routers.

Autonomous system (AS) boundary routers flood information about external autonomous systems throughout the AS, except to stub areas. Area border routers are responsible for advertising the paths to all AS boundary routers.

### OSPF Three-Way Handshake

OSPF creates a topology map by flooding LSAs across OSPF-enabled links. LSAs announce the presence of OSPF-enabled interfaces to adjacent OSPF interfaces. The exchange of LSAs establishes bidirectional connectivity between all adjacent OSPF interfaces (neighbors) using a three-way handshake, as shown in [Figure 60 on page 3311](#).

**Figure 60: OSPF Three-Way Handshake**



In [Figure 60 on page 3311](#), Router A sends hello packets out all its OSPF-enabled interfaces when it comes online. Router B receives the packet, which establishes that Router B can receive traffic from Router A. Router B generates a response to Router A to acknowledge receipt of the hello packet. When Router A receives the response, it establishes that Router B can receive traffic from Router A. Router A then generates a final response packet to inform Router B that Router A can receive traffic from Router B. This three-way handshake ensures bidirectional connectivity.

As new neighbors are added to the network or existing neighbors lose connectivity, the adjacencies in the topology map are modified accordingly through the exchange (or absence) of LSAs. These LSAs advertise only the incremental changes in the network, which helps minimize the amount of OSPF traffic on the network. The adjacencies are shared and used to create the network topology in the topological database.

### **OSPF Version 3**

OSPFv3 is a modified version of OSPF that supports IP version 6 (IPv6) addressing. OSPFv3 differs from OSPFv2 in the following ways:

- All neighbor ID information is based on a 32-bit router ID.
- The protocol runs per link rather than per subnet.
- Router and network link-state advertisements (LSAs) do not carry prefix information.
- Two new LSA types are included: link-LSA and intra-area-prefix-LSA.
- Flooding scopes are as follows:
  - Link-local
  - Area
  - AS
- Link-local addresses are used for all neighbor exchanges except virtual links.
- Authentication is removed. The IPv6 authentication header relies on the IP layer.
- The packet format has changed as follows:
  - Version number 2 is now version number 3.
  - The **db** option field has been expanded to 24 bits.
  - Authentication information has been removed.
  - Hello messages do not have address information.
  - Two new option bits are included: **R** and **V6**.
- Type 3 summary LSAs have been renamed *inter-area-prefix-LSAs*.
- Type 4 summary LSAs have been renamed *inter-area-router-LSAs*.

### **Example: Configuring OSPF on Logical Systems Within the Same Router**

This example shows how to configure an OSPF network using multiple logical systems that are running on a single physical router. The logical systems are connected by logical tunnel interfaces.

- [Requirements on page 3313](#)
- [Overview on page 3313](#)
- [Configuration on page 3313](#)
- [Verification on page 3317](#)

### Requirements


You must connect the logical systems by using logical tunnel (lt) interfaces. See “[Example: Connecting Logical Systems Within the Same Device Using Logical Tunnel Interfaces on MX Series Routers and EX Series Switches](#)” on page 3279.

### Overview

This example shows the configuration of a single OSPF area with three logical systems running on one physical router. Each logical system has its own routing table. The configuration enables the protocol on all logical system interfaces that participate in the OSPF domain and specifies the area that the interfaces are in.

[Figure 61 on page 3313](#) shows the sample network.

**Figure 61: OSPF on Logical Systems**



g040567

### Configuration

#### CLI Quick Configuration

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

```
set logical-systems LS1 interfaces lt-1/2/0 unit 0 description LS1->LS3
set logical-systems LS1 interfaces lt-1/2/0 unit 0 encapsulation ethernet
set logical-systems LS1 interfaces lt-1/2/0 unit 0 peer-unit 5
set logical-systems LS1 interfaces lt-1/2/0 unit 0 family inet address 10.0.1.2/30
set logical-systems LS1 interfaces lt-1/2/0 unit 2 description LS1->LS2
set logical-systems LS1 interfaces lt-1/2/0 unit 2 encapsulation ethernet
```

```
set logical-systems LS1 interfaces lt-1/2/0 unit 2 peer-unit 1
set logical-systems LS1 interfaces lt-1/2/0 unit 2 family inet address 10.0.0.1/30
set logical-systems LS1 protocols ospf area 0.0.0.0 interface lt-1/2/0.0
set logical-systems LS1 protocols ospf area 0.0.0.0 interface lt-1/2/0.2
set logical-systems LS2 interfaces lt-1/2/0 unit 1 description LS2->LS1
set logical-systems LS2 interfaces lt-1/2/0 unit 1 encapsulation ethernet
set logical-systems LS2 interfaces lt-1/2/0 unit 1 peer-unit 2
set logical-systems LS2 interfaces lt-1/2/0 unit 1 family inet address 10.0.0.2/30
set logical-systems LS2 interfaces lt-1/2/0 unit 4 description LS2->LS3
set logical-systems LS2 interfaces lt-1/2/0 unit 4 encapsulation ethernet
set logical-systems LS2 interfaces lt-1/2/0 unit 4 peer-unit 3
set logical-systems LS2 interfaces lt-1/2/0 unit 4 family inet address 10.0.2.2/30
set logical-systems LS2 protocols ospf area 0.0.0.0 interface lt-1/2/0.1
set logical-systems LS2 protocols ospf area 0.0.0.0 interface lt-1/2/0.4
set logical-systems LS3 interfaces lt-1/2/0 unit 3 description LS3->LS2
set logical-systems LS3 interfaces lt-1/2/0 unit 3 encapsulation ethernet
set logical-systems LS3 interfaces lt-1/2/0 unit 3 peer-unit 4
set logical-systems LS3 interfaces lt-1/2/0 unit 3 family inet address 10.0.2.1/30
set logical-systems LS3 interfaces lt-1/2/0 unit 5 description LS3->LS1
set logical-systems LS3 interfaces lt-1/2/0 unit 5 encapsulation ethernet
set logical-systems LS3 interfaces lt-1/2/0 unit 5 peer-unit 0
set logical-systems LS3 interfaces lt-1/2/0 unit 5 family inet address 10.0.1.1/30
set logical-systems LS3 protocols ospf area 0.0.0.0 interface lt-1/2/0.5
set logical-systems LS3 protocols ospf area 0.0.0.0 interface lt-1/2/0.3
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure OSPF on logical systems:

1. Configure the logical tunnel interface on Logical System LS1 connecting to Logical System LS2.

```
[edit]
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 2 description LS1->LS2
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 2 encapsulation ethernet
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 2 peer-unit 1
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 2 family inet address
10.0.0.1/30
```

2. Configure the logical tunnel interface on Logical System LS1 connecting to Logical System LS3.

```
[edit]
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 0 description LS1->LS3
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 0 encapsulation ethernet
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 0 peer-unit 5
user@host# set logical-systems LS1 interfaces lt-1/2/0 unit 0 family inet address
10.0.1.2/30
```

3. Configure the logical tunnel interface on Logical System LS2 connecting to Logical System LS1.

```
[edit]
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 1 description LS2->LS1
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 1 encapsulation ethernet
```



```

user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 1 peer-unit 2
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 1 family inet address
10.0.0.2/30

```

4. Configure the logical tunnel interface on Logical System LS2 connecting to Logical System LS3.

```

[edit]
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 4 description LS2->LS3
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 4 encapsulation ethernet
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 4 peer-unit 3
user@host# set logical-systems LS2 interfaces lt-1/2/0 unit 4 family inet address
10.0.2.2/30

```

5. Configure the logical tunnel interface on Logical System LS3 connecting to Logical System LS2.

```

[edit]
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 3 description LS3->LS2
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 3 encapsulation ethernet
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 3 peer-unit 4
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 3 family inet address
10.0.2.1/30

```

6. Configure the logical tunnel interface on Logical System LS3 connecting to Logical System LS1.

```

[edit]
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 5 description LS3->LS1
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 5 encapsulation ethernet
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 5 peer-unit 0
user@host# set logical-systems LS3 interfaces lt-1/2/0 unit 5 family inet address
10.0.1.1/30

```

7. Configure OSPF on all the interfaces.

```

[edit]
user@host# set logical-systems LS1 protocols ospf area 0.0.0.0 interface lt-1/2/0.0
user@host# set logical-systems LS1 protocols ospf area 0.0.0.0 interface lt-1/2/0.2
user@host# set logical-systems LS2 protocols ospf area 0.0.0.0 interface lt-1/2/0.1
user@host# set logical-systems LS2 protocols ospf area 0.0.0.0 interface lt-1/2/0.4
user@host# set logical-systems LS3 protocols ospf area 0.0.0.0 interface lt-1/2/0.5
user@host# set logical-systems LS3 protocols ospf area 0.0.0.0 interface lt-1/2/0.3

```

8. If you are done configuring the device, commit the configuration.

```

[edit]
user@host# commit

```

### Results

Confirm your configuration by issuing the **show logical-systems** command.

```

show logical-systems
LS1 {
 interfaces {
 lt-1/2/0 {
 unit 0 {

```

```
 description LS1->LS3;
 encapsulation ethernet;
 peer-unit 5;
 family inet {
 address 10.0.1.2/30;
 }
 }
 unit 2 {
 description LS1->LS2;
 encapsulation ethernet;
 peer-unit 1;
 family inet {
 address 10.0.0.1/30;
 }
 }
}
protocols {
 ospf {
 area 0.0.0.0 {
 interface lt-1/2/0.0;
 interface lt-1/2/0.2;
 }
 }
}
}
LS2 {
 interfaces {
 lt-1/2/0 {
 unit 1 {
 description LS2->LS1;
 encapsulation ethernet;
 peer-unit 2;
 family inet {
 address 10.0.0.2/30;
 }
 }
 unit 4 {
 description LS2->LS3;
 encapsulation ethernet;
 peer-unit 3;
 family inet {
 address 10.0.2.2/30;
 }
 }
 }
 }
 protocols {
 ospf {
 area 0.0.0.0 {
 interface lt-1/2/0.1;
 interface lt-1/2/0.4;
 }
 }
 }
}
```

```
LS3 {
 interfaces {
 lt-1/2/0 {
 unit 3 {
 description LS3->LS2;
 encapsulation ethernet;
 peer-unit 4;
 family inet {
 address 10.0.2.1/30;
 }
 }
 unit 5 {
 description LS3->LS1;
 encapsulation ethernet;
 peer-unit 0;
 family inet {
 address 10.0.1.1/30;
 }
 }
 }
 }
 protocols {
 ospf {
 area 0.0.0.0 {
 interface lt-1/2/0.5;
 interface lt-1/2/0.3;
 }
 }
 }
}
```

**Verification**

Confirm that the configuration is working properly.

- [Verifying That the Logical Systems Are Up on page 3317](#)
- [Verifying Connectivity Between the Logical Systems on page 3318](#)

**Verifying That the Logical Systems Are Up**

**Purpose** Make sure that the interfaces are properly configured.

**Action** user@host> show interfaces terse

| Interface  | Admin | Link | Proto | Local       | Remote |
|------------|-------|------|-------|-------------|--------|
| ...        |       |      |       |             |        |
| lt-1/2/0   | up    | up   |       |             |        |
| lt-1/2/0.0 | up    | up   | inet  | 10.0.1.2/30 |        |
| lt-1/2/0.1 | up    | up   | inet  | 10.0.0.2/30 |        |
| lt-1/2/0.2 | up    | up   | inet  | 10.0.0.1/30 |        |
| lt-1/2/0.3 | up    | up   | inet  | 10.0.2.1/30 |        |
| lt-1/2/0.4 | up    | up   | inet  | 10.0.2.2/30 |        |
| lt-1/2/0.5 | up    | up   | inet  | 10.0.1.1/30 |        |
| ...        |       |      |       |             |        |

#### ***Verifying Connectivity Between the Logical Systems***

**Purpose** Make sure that the OSPF adjacencies are established by checking the OSPF neighbor tables, checking the routing tables, and pinging the logical systems.

```

Action user@host> show ospf neighbor logical-system LS1
Address Interface State ID Pri Dead
10.0.1.1 lt-1/2/0.0 Full 10.0.1.1 128 37
10.0.0.2 lt-1/2/0.2 Full 10.0.0.2 128 33

user@host> show ospf neighbor logical-system LS2
Address Interface State ID Pri Dead
10.0.0.1 lt-1/2/0.1 Full 10.0.0.1 128 32
10.0.2.1 lt-1/2/0.4 Full 10.0.1.1 128 36

user@host> show ospf neighbor logical-system LS3
Address Interface State ID Pri Dead
10.0.2.2 lt-1/2/0.3 Full 10.0.0.2 128 36
10.0.1.2 lt-1/2/0.5 Full 10.0.0.1 128 37

user@host> show route logical-system LS1
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30 *[Direct/0] 00:28:00
 > via lt-1/2/0.2
10.0.0.1/32 *[Local/0] 00:28:00
 Local via lt-1/2/0.2
10.0.1.0/30 *[Direct/0] 00:28:00
 > via lt-1/2/0.0
10.0.1.2/32 *[Local/0] 00:28:00
 Local via lt-1/2/0.0
10.0.2.0/30 *[OSPF/10] 00:27:05, metric 2
 > to 10.0.1.1 via lt-1/2/0.0
 to 10.0.0.2 via lt-1/2/0.2
224.0.0.5/32 *[OSPF/10] 00:28:03, metric 1
 MultiRecv

user@host> show route logical-system LS2
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30 *[Direct/0] 00:28:31
 > via lt-1/2/0.1
10.0.0.2/32 *[Local/0] 00:28:32
 Local via lt-1/2/0.1
10.0.1.0/30 *[OSPF/10] 00:27:38, metric 2
 > to 10.0.0.1 via lt-1/2/0.1
 to 10.0.2.1 via lt-1/2/0.4
10.0.2.0/30 *[Direct/0] 00:28:32
 > via lt-1/2/0.4
10.0.2.2/32 *[Local/0] 00:28:32
 Local via lt-1/2/0.4
224.0.0.5/32 *[OSPF/10] 00:28:34, metric 1
 MultiRecv

user@host> show route logical-system LS3
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30 *[OSPF/10] 00:28:23, metric 2
 > to 10.0.2.2 via lt-1/2/0.3
 to 10.0.1.2 via lt-1/2/0.5
10.0.1.0/30 *[Direct/0] 00:29:13

```

```
10.0.1.1/32 > via lt-1/2/0.5
 *[Local/0] 00:29:15
 Local via lt-1/2/0.5
10.0.2.0/30 *[Direct/0] 00:29:14
 > via lt-1/2/0.3
10.0.2.1/32 *[Local/0] 00:29:15
 Local via lt-1/2/0.3
224.0.0.5/32 *[OSPF/10] 00:29:16, metric 1
 MultiRecv
```

### From LS1, Ping LS3

```
user@host> set cli logical-system LS1

user@host:LS1> ping 10.0.2.1
PING 10.0.2.1 (10.0.2.1): 56 data bytes
64 bytes from 10.0.2.1: icmp_seq=0 ttl=64 time=1.215 ms
64 bytes from 10.0.2.1: icmp_seq=1 ttl=64 time=1.150 ms
64 bytes from 10.0.2.1: icmp_seq=2 ttl=64 time=1.134 ms
```

### From LS3, Ping LS1

```
user@host> set cli logical-system LS3

user@host:LS3> ping 10.0.0.1
PING 10.0.0.1 (10.0.0.1): 56 data bytes
64 bytes from 10.0.0.1: icmp_seq=0 ttl=64 time=1.193 ms
64 bytes from 10.0.0.1: icmp_seq=1 ttl=64 time=1.114 ms
64 bytes from 10.0.0.1: icmp_seq=2 ttl=64 time=1.190 ms
```

### Related Documentation

- [Introduction to Logical Systems on page 3259](#)

---

### Examples: Configuring OSPF Routing Policy on Logical Systems

---

- [Understanding OSPF Routing Policy on page 3320](#)
- [Example: Configuring an OSPF Default Route Policy on Logical Systems on page 3322](#)
- [Example: Configuring a Conditional OSPF Default Route Policy on Logical Systems on page 3327](#)
- [Example: Configuring an OSPF Import Policy on Logical Systems on page 3334](#)

#### ***Understanding OSPF Routing Policy***

Each routing policy is identified by a policy name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks. Each routing policy name must be unique within a configuration. Once a policy is created and named, it must be applied before it is active.

In the **import** statement, you list the name of the routing policy used to filter OSPF external routes from being installed into the routing tables of OSPF neighbors. You can filter the routes, but not link-state address (LSA) flooding. An external route is a route that is outside the OSPF Autonomous System (AS). The import policy does not impact the OSPF database. This means that the import policy has no impact on the link-state advertisements.

In the **export** statement, you list the name of the routing policy to be evaluated when routes are being exported from the routing table into OSPF.

By default, if a routing device has multiple OSPF areas, learned routes from other areas are automatically installed into area 0 of the routing table.

To specify more than one policy and create a policy chain, you list the policies using a space as a separator. If multiple policies are specified, the policies are evaluated in the order in which they are specified. As soon as an accept or reject action is executed, the policy chain evaluation ends.

This topic describes the following information:

- [Routing Policy Terms on page 3321](#)
- [Routing Policy Match Conditions on page 3321](#)
- [Routing Policy Actions on page 3322](#)

### ***Routing Policy Terms***

Routing policies are made up of one or more terms. A term is a named structure in which match conditions and actions are defined. You can define one or more terms. The name can contain letters, numbers, and hyphens ( - ) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

Each term contains a set of match conditions and a set of actions:

- Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route.
- Actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

### ***Routing Policy Match Conditions***

A match condition defines the criteria that a route must match for an action to take place. You can define one or more match conditions for each term. If a route matches all of the match conditions for a particular term, the actions defined for that term are processed.

Each term can include two statements, **from** and **to**, that define the match conditions:

- In the **from** statement, you define the criteria that an incoming route must match. You can specify one or more match conditions. If you specify more than one, they all must match the route for a match to occur.

The **from** statement is optional. If you omit the **from** and the **to** statements, all routes are considered to match.



**NOTE:** In export policies, omitting the **from** statement from a routing policy term might lead to unexpected results. For more information, see *Applying Routing Policies and Policy Chains to Routing Protocols* in the *Routing Policy Configuration Guide*.

- In the **to** statement, you define the criteria that an outgoing route must match. You can specify one or more match conditions. If you specify more than one, they all must match the route for a match to occur.

The order of the match conditions in a term is not important because a route must match all match conditions in a term for an action to be taken.

For a complete list of match conditions, see *Configuring Match Conditions in Routing Policy Terms* in the *Routing Policy Configuration Guide*.

### ***Routing Policy Actions***

An action defines what the routing device does with the route when the route matches all the match conditions in the **from** and **to** statements for a particular term. If a term does not have **from** and **to** statements, all routes are considered to match and the actions apply to all routes.

Each term can have one or more of the following types of actions. The actions are configured under the **then** statement.

- Flow control actions, which affect whether to accept or reject the route and whether to evaluate the next term or routing policy.
- Actions that manipulate route characteristics.
- Trace action, which logs route matches.

The **then** statement is optional. If you omit it, one of the following occurs:

- The next term in the routing policy, if one exists, is evaluated.
- If the routing policy has no more terms, the next routing policy, if one exists, is evaluated.
- If there are no more terms or routing policies, the **accept** or **reject** action specified by the default policy is executed.

For a complete list of routing policy actions, see *Configuring Actions in Routing Policy Terms* in the *Routing Policy Configuration Guide*.

### ***Example: Configuring an OSPF Default Route Policy on Logical Systems***

This example shows how to configure a default route on one logical system and inject the default route into OSPF area 0. In this example, OSPF area 0 contains three logical systems that are configured on a single physical router.

- [Requirements on page 3323](#)
- [Overview on page 3323](#)
- [Configuration on page 3323](#)
- [Verification on page 3325](#)



### Requirements

Before you begin:

- Connect the logical systems by using logical tunnel (lt) interfaces. See [“Example: Connecting Logical Systems Within the Same Device Using Logical Tunnel Interfaces on MX Series Routers and EX Series Switches”](#) on page 3279.
- Enable OSPF on the interfaces. See [“Example: Configuring OSPF on Logical Systems Within the Same Router”](#) on page 3312.


### Overview

This example shows a logical system redistributing a default route to other logical systems. All logical systems are running OSPF. A common reason for a default route is to provide a path for sending traffic destined outside the OSPF domain.

In this example, the default route is not used for forwarding traffic. The **no-install** statement prevents the route from being installed in the forwarding table of Logical System LS3. If you configure a route so it is not installed in the forwarding table, the route is still eligible to be exported from the routing table to other protocols. The **discard** statement silently drops packets without notice.

[Figure 62 on page 3323](#) shows the sample network.

**Figure 62: OSPF with a Default Route to an ISP**



g040569

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set logical-systems LS3 routing-options static route 0.0.0.0/0 discard
set logical-systems LS3 routing-options static route 0.0.0.0/0 no-install
set logical-systems LS3 policy-options policy-statement ospf-default from protocol
 static
set logical-systems LS3 policy-options policy-statement ospf-default from route-filter
 0.0.0.0/0 exact
set logical-systems LS3 policy-options policy-statement ospf-default then accept
set logical-systems LS3 protocols ospf export ospf-default
```

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure an OSPF default route policy on logical systems:

1. Change the context to Logical System LS3.

```
[edit]
user@host> set cli logical-system LS3
```

2. Configure the default route on Logical System LS3.

```
[edit]
user@host:LS3# set routing-options static route 0.0.0.0/0 discard
user@host:LS3# set routing-options static route 0.0.0.0/0 no-install
```

3. Configure the policy on Logical System LS3.

```
[edit]
user@host:LS3# set policy-options policy-statement ospf-default from protocol
 static
user@host:LS3# set policy-options policy-statement ospf-default from route-filter
 0.0.0.0/0 exact
user@host:LS3# set policy-options policy-statement ospf-default then accept
```

4. Apply the export policy to OSPF on Logical System LS3.

```
[edit]
user@host:LS3# set protocols ospf export ospf-default
```

5. If you are done configuring the device, commit the configuration.

```
[edit]
user@host:LS3# commit
```

#### Results

Confirm your configuration by issuing the **show logical-systems LS3** command.

```
show logical-systems LS3
interfaces {
 lt-1/2/0 {
 unit 3 {
 description LS3->LS2;
 encapsulation ethernet;
```

```
 peer-unit 4;
 family inet {
 address 10.0.2.1/30;
 }
 }
 unit 5 {
 description LS3->LS1;
 encapsulation ethernet;
 peer-unit 0;
 family inet {
 address 10.0.1.1/30;
 }
 }
}
protocols {
 ospf {
 export ospf-default;
 area 0.0.0.0 {
 interface lt-1/2/0.5;
 interface lt-1/2/0.3;
 }
 }
}
policy-options {
 policy-statement ospf-default {
 from {
 protocol static;
 route-filter 0.0.0.0/0 exact;
 }
 then accept;
 }
}
routing-options {
 static {
 route 0.0.0.0/0 {
 discard;
 no-install;
 }
 }
}
```

**Verification**

Confirm that the configuration is working properly.

**Verifying That the Static Route Is Redistributed**

**Purpose** Make sure that the OSPF policy is working by checking the routing tables.

**Action** user@host> show route logical-system LS3  
inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

```
0.0.0.0/0 *[Static/5] 01:04:38
 Discard
10.0.0.0/30 *[OSPF/10] 11:53:55, metric 2
 to 10.0.2.2 via lt-1/2/0.3
 > to 10.0.1.2 via lt-1/2/0.5
10.0.1.0/30 *[Direct/0] 11:54:50
 > via lt-1/2/0.5
10.0.1.1/32 *[Local/0] 11:54:54
 Local via lt-1/2/0.5
10.0.2.0/30 *[Direct/0] 11:54:50
 > via lt-1/2/0.3
10.0.2.1/32 *[Local/0] 11:54:54
 Local via lt-1/2/0.3
224.0.0.5/32 *[OSPF/10] 11:56:55, metric 1
 MultiRecv
```

user@host> show route logical-system LS1  
inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

```
0.0.0.0/0 *[OSPF/150] 01:02:34, metric 0, tag 0
 > to 10.0.1.1 via lt-1/2/0.0
10.0.0.0/30 *[Direct/0] 11:52:46
 > via lt-1/2/0.2
10.0.0.1/32 *[Local/0] 11:52:50
 Local via lt-1/2/0.2
10.0.1.0/30 *[Direct/0] 11:52:46
 > via lt-1/2/0.0
10.0.1.2/32 *[Local/0] 11:52:50
 Local via lt-1/2/0.0
10.0.2.0/30 *[OSPF/10] 11:51:56, metric 2
 > to 10.0.1.1 via lt-1/2/0.0
 to 10.0.0.2 via lt-1/2/0.2
224.0.0.5/32 *[OSPF/10] 11:54:50, metric 1
 MultiRecv
```

user@host> show route logical-system LS2  
inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

```
0.0.0.0/0 *[OSPF/150] 01:05:20, metric 0, tag 0
 > to 10.0.2.1 via lt-1/2/0.4
10.0.0.0/30 *[Direct/0] 11:55:32
 > via lt-1/2/0.1
10.0.0.2/32 *[Local/0] 11:55:36
 Local via lt-1/2/0.1
10.0.1.0/30 *[OSPF/10] 11:54:37, metric 2
 > to 10.0.0.1 via lt-1/2/0.1
 to 10.0.2.1 via lt-1/2/0.4
10.0.2.0/30 *[Direct/0] 11:55:32
 > via lt-1/2/0.4
10.0.2.2/32 *[Local/0] 11:55:36
 Local via lt-1/2/0.4
224.0.0.5/32 *[OSPF/10] 11:57:36, metric 1
 MultiRecv
```

**Meaning** The routing table on Logical System LS3 contains the default 0.0.0.0/0 route from protocol **Static**. The routing tables on Logical System LS1 and Logical System LS2 contain the default 0.0.0.0/0 route from protocol **OSPF**. If Logical System LS1 and Logical System LS2 receive packets destined for networks not specified in their routing tables, those packets will be sent to Logical System LS3 for further processing. This configuration assumes that Logical System LS3 has a connection to an ISP or another external network.

***Example: Configuring a Conditional OSPF Default Route Policy on Logical Systems***

This example shows how to configure a conditional default route on one logical system and inject the default route into OSPF area 0.

- [Requirements on page 3327](#)
- [Overview on page 3327](#)
- [Configuration on page 3328](#)
- [Verification on page 3332](#)

***Requirements***

Before you begin:

- Connect the logical systems by using logical tunnel (lt) interfaces. See “[Example: Connecting Logical Systems Within the Same Device Using Logical Tunnel Interfaces on MX Series Routers and EX Series Switches](#)” on page 3279.
- Enable OSPF on the interfaces. See “[Example: Configuring OSPF on Logical Systems Within the Same Router](#)” on page 3312.


***Overview***

In this example, OSPF area 0 contains three logical systems that are configured on a single physical router. Logical System LS3 has a BGP session with an external peer, for example, an ISP.

The ISP injects a default static route into BGP, which provides the customer network with a default static route to reach external networks. Logical System LS3 exports the default route into OSPF. The route policy on Logical System LS3 is conditional such that if the connection to the external peer goes down, the default route is no longer active in the routing tables of the logical systems in area 0. This policy prevents blackholing of traffic. Blackholing occurs when packets are dropped without notification.

[Figure 63 on page 3328](#) shows the sample network.

Figure 63: OSPF with a Conditional Default Route to an ISP



g040570

**Configuration****CLI Quick  
Configuration**

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

**Device LS1**

```

set logical-systems LS1 interfaces lt-1/2/0 unit 0 description LS1->LS3
set logical-systems LS1 interfaces lt-1/2/0 unit 0 encapsulation ethernet
set logical-systems LS1 interfaces lt-1/2/0 unit 0 peer-unit 5
set logical-systems LS1 interfaces lt-1/2/0 unit 0 family inet address 10.0.1.2/30
set logical-systems LS1 interfaces lt-1/2/0 unit 2 description LS1->LS2
set logical-systems LS1 interfaces lt-1/2/0 unit 2 encapsulation ethernet
set logical-systems LS1 interfaces lt-1/2/0 unit 2 peer-unit 1
set logical-systems LS1 interfaces lt-1/2/0 unit 2 family inet address 10.0.0.1/30
set logical-systems LS1 protocols ospf area 0.0.0.0 interface lt-1/2/0.0
set logical-systems LS1 protocols ospf area 0.0.0.0 interface lt-1/2/0.2

```

**Device LS2**

```

set logical-systems LS2 interfaces lt-1/2/0 unit 1 description LS2->LS1
set logical-systems LS2 interfaces lt-1/2/0 unit 1 encapsulation ethernet
set logical-systems LS2 interfaces lt-1/2/0 unit 1 peer-unit 2
set logical-systems LS2 interfaces lt-1/2/0 unit 1 family inet address 10.0.0.2/30
set logical-systems LS2 interfaces lt-1/2/0 unit 4 description LS2->LS3
set logical-systems LS2 interfaces lt-1/2/0 unit 4 encapsulation ethernet
set logical-systems LS2 interfaces lt-1/2/0 unit 4 peer-unit 3
set logical-systems LS2 interfaces lt-1/2/0 unit 4 family inet address 10.0.2.2/30
set logical-systems LS2 protocols ospf area 0.0.0.0 interface lt-1/2/0.1
set logical-systems LS2 protocols ospf area 0.0.0.0 interface lt-1/2/0.4

```

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Device LS3</b> | <pre> set logical-systems LS3 interfaces lt-1/2/0 unit 3 description LS3-&gt;LS2 set logical-systems LS3 interfaces lt-1/2/0 unit 3 encapsulation ethernet set logical-systems LS3 interfaces lt-1/2/0 unit 3 peer-unit 4 set logical-systems LS3 interfaces lt-1/2/0 unit 3 family inet address 10.0.2.1/30 set logical-systems LS3 interfaces lt-1/2/0 unit 5 description LS3-&gt;LS1 set logical-systems LS3 interfaces lt-1/2/0 unit 5 encapsulation ethernet set logical-systems LS3 interfaces lt-1/2/0 unit 5 peer-unit 0 set logical-systems LS3 interfaces lt-1/2/0 unit 5 family inet address 10.0.1.1/30 set logical-systems LS3 interfaces so-0/0/2 unit 0 description LS3-&gt;ISP set logical-systems LS3 interfaces so-0/0/2 unit 0 family inet address 10.0.45.2/30 set logical-systems LS3 protocols bgp group ext type external set logical-systems LS3 protocols bgp group ext peer-as 65000 set logical-systems LS3 protocols bgp group ext neighbor 10.0.45.1 set logical-systems LS3 protocols ospf export gendefault set logical-systems LS3 protocols ospf area 0.0.0.0 interface lt-1/2/0.5 set logical-systems LS3 protocols ospf area 0.0.0.0 interface lt-1/2/0.3 set logical-systems LS3 policy-options policy-statement gendefault term upstreamroutes   from protocol bgp set logical-systems LS3 policy-options policy-statement gendefault term upstreamroutes   from as-path upstream set logical-systems LS3 policy-options policy-statement gendefault term upstreamroutes   from route-filter 0.0.0.0/0 upto /16 set logical-systems LS3 policy-options policy-statement gendefault term upstreamroutes   then next-hop 10.0.45.1 set logical-systems LS3 policy-options policy-statement gendefault term upstreamroutes   then accept set logical-systems LS3 policy-options policy-statement gendefault term end then reject set logical-systems LS3 policy-options as-path upstream "^65000 " set logical-systems LS3 routing-options generate route 0.0.0.0/0 policy gendefault set logical-systems LS3 routing-options autonomous-system 65001 </pre> |
| <b>Device ISP</b> | <pre> set interfaces so-0/0/2 unit 0 family inet address 10.0.45.1/30 set protocols bgp group ext type external set protocols bgp group ext export advertise-default set protocols bgp group ext peer-as 65001 set protocols bgp group ext neighbor 10.0.45.2 set policy-options policy-statement advertise-default term 1 from route-filter 0.0.0.0/0   exact set policy-options policy-statement advertise-default term 1 then accept set routing-options static route 0.0.0.0/0 discard set routing-options autonomous-system 65000 </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode”](#) on page 4704 in the *CLI User Guide*.

To configure a conditional default route:

1. Configure the interfaces.

```

[edit logical-systems LS3 interfaces]
user@R3# set lt-1/2/0 unit 3 description LS3->LS2
user@R3# set lt-1/2/0 unit 3 encapsulation ethernet
user@R3# set lt-1/2/0 unit 3 peer-unit 4
user@R3# set lt-1/2/0 unit 3 family inet address 10.0.2.1/30

```

```
user@R3# set lt-1/2/0 unit 5 description LS3->LS1
user@R3# set lt-1/2/0 unit 5 encapsulation ethernet
user@R3# set lt-1/2/0 unit 5 peer-unit 0
user@R3# set lt-1/2/0 unit 5 family inet address 10.0.1.1/30
user@R3# set so-0/0/2 unit 0 description LS3->ISP
user@R3# set so-0/0/2 unit 0 encapsulation ethernet
user@R3# set so-0/0/2 unit 0 peer-unit 7
user@R3# set so-0/0/2 unit 0 family inet address 10.0.45.2/30
```

2. Configure the autonomous system (AS) number.

```
[edit logical-systems LS3 routing-options]
user@R3# set autonomous-system 65001
```

3. Configure the BGP session with the ISP device.

```
[edit logical-systems LS3 protocols bgp group ext]
user@R3# set type external
user@R3# set peer-as 65000
user@R3# set neighbor 10.0.45.1
```

4. Configure OSPF.

```
[edit logical-systems LS3 protocols ospf area 0.0.0.0]
user@R3# set interface lt-1/2/0.5
user@R3# set interface lt-1/2/0.3
```

5. Configure the routing policy.

```
[edit logical-systems LS3 policy-options policy-statement gendefault]
user@R3# set term upstreamroutes from protocol bgp
user@R3# set term upstreamroutes from as-path upstream
user@R3# set term upstreamroutes from route-filter 0.0.0.0/0 upto /16
user@R3# set term upstreamroutes then next-hop 10.0.45.1
user@R3# set term upstreamroutes then accept
```

```
user@R3# set term end then reject
```

```
[edit logical-systems LS3 policy-options]
user@R3# set as-path upstream "^65000 "
```

6. Configure the generated route.

```
[edit logical-systems LS3 routing-options]
user@R3# set generate route 0.0.0.0/0 policy gendefault
```

7. Apply the export policy to OSPF.

```
[edit logical-systems LS3 protocols ospf]
user@R3# set export gendefault
```

8. If you are done configuring the device, commit the configuration.

```
[edit]
user@R3# commit
```

### **Results**

Confirm your configuration by issuing the **show logical-systems LS3** command.



```
show logical-systems LS3
interfaces {
 lt-1/2/0 {
 unit 3 {
 description LS3->LS2;
 encapsulation ethernet;
 peer-unit 4;
 family inet {
 address 10.0.2.1/30;
 }
 }
 unit 5 {
 description LS3->LS1;
 encapsulation ethernet;
 peer-unit 0;
 family inet {
 address 10.0.1.1/30;
 }
 }
 unit 6 {
 description LS3->ISP;
 encapsulation ethernet;
 peer-unit 7;
 family inet {
 address 10.0.45.2/30;
 }
 }
 }
}
protocols {
 bgp {
 group ext {
 type external;
 peer-as 65000;
 neighbor 10.0.45.1;
 }
 }
 ospf {
 export gendefault;
 area 0.0.0.0 {
 interface lt-1/2/0.5;
 interface lt-1/2/0.3;
 }
 }
}
policy-options {
 policy-statement gendefault {
 term upstreamroutes {
 from {
 protocol bgp;
 as-path upstream;
 route-filter 0.0.0.0/0 upto /16;
 }
 then {
 next-hop 10.0.45.1;
 accept;
 }
 }
 }
}
```

```
 }
 }
 term end {
 then reject;
 }
}
as-path upstream "~65000";
}
routing-options {
 generate {
 route 0.0.0.0/0 policy gendefault;
 }
 autonomous-system 65001;
}
```

### **Verification**

Confirm that the configuration is working properly.

- [Verifying that the Route to the ISP Is Working on page 3332](#)
- [Verifying That the Static Route Is Redistributed on page 3332](#)
- [Testing the Policy Condition on page 3333](#)

### **Verifying that the Route to the ISP Is Working**

**Purpose** Make sure connectivity is established between Logical System LS3 and the ISP's router.

**Action**

```
user@host>set cli logical-system LS3
Logical system: LS3

user@host:LS3>ping 10.0.45.1
PING 10.0.45.1 (10.0.45.1): 56 data bytes
64 bytes from 10.0.45.1: icmp_seq=0 ttl=64 time=1.185 ms
64 bytes from 10.0.45.1: icmp_seq=1 ttl=64 time=1.199 ms
64 bytes from 10.0.45.1: icmp_seq=2 ttl=64 time=1.186 ms
```

**Meaning** The ping command confirms reachability.

### **Verifying That the Static Route Is Redistributed**

**Purpose** Make sure that the BGP policy is redistributing the static route into Logical System LS3's routing table. Also make sure that the OSPF policy is redistributing the static route into the routing tables of Logical System LS1 and Logical System LS2.

**Action** user@host> show route logical-system LS3 protocol bgp

```
inet.0: 9 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
0.0.0.0/0 *[BGP/170] 00:00:25, localpref 100
 AS path: 65000 I
 > to 10.0.45.1 via so-0/0/2.0
```

user@host> show route logical-system LS1 protocol ospf

```
inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
0.0.0.0/0 *[OSPF/150] 00:03:58, metric 0, tag 0
 > to 10.0.1.1 via lt-1/2/0.0
10.0.2.0/30 *[OSPF/10] 03:37:45, metric 2
 to 10.0.1.1 via lt-1/2/0.0
 > to 10.0.0.2 via lt-1/2/0.2
224.0.0.5/32 *[OSPF/10] 03:38:41, metric 1
 MultiRecv
```

user@host> show route logical-system LS2 protocol ospf

```
inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
0.0.0.0/0 *[OSPF/150] 00:04:04, metric 0, tag 0
 > to 10.0.2.1 via lt-1/2/0.4
10.0.1.0/30 *[OSPF/10] 03:37:46, metric 2
 to 10.0.0.1 via lt-1/2/0.1
 > to 10.0.2.1 via lt-1/2/0.4
224.0.0.5/32 *[OSPF/10] 03:38:47, metric 1
 MultiRecv
```

**Meaning** The routing tables contain the default 0.0.0.0/0 route. If Logical System LS1 and Logical System LS2 receive packets destined for networks not specified in their routing tables, those packets will be sent to Logical System LS3 for further processing. If Logical System LS3 receives packets destined for networks not specified in its routing table, those packets will be sent to the ISP for further processing.

#### *Testing the Policy Condition*

**Purpose** Deactivate the interface to make sure that the route is removed from the routing tables if the external network becomes unreachable.

**Action** user@host> deactivate logical-systems LS3 interfaces so-0/0/2 unit 0 family inet address 10.0.45.2/30

user@host> commit

user@host> show route logical-system LS1 protocol ospf

inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

```
10.0.2.0/30 *[OSPF/10] 03:41:48, metric 2
 to 10.0.1.1 via lt-1/2/0.0
 > to 10.0.0.2 via lt-1/2/0.2
224.0.0.5/32 *[OSPF/10] 03:42:44, metric 1
 MultiRecv
```

user@host> show route logical-system LS2 protocol ospf

inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

```
10.0.1.0/30 *[OSPF/10] 03:42:10, metric 2
 to 10.0.0.1 via lt-1/2/0.1
 > to 10.0.2.1 via lt-1/2/0.4
224.0.0.5/32 *[OSPF/10] 03:43:11, metric 1
 MultiRecv
```

**Meaning** The routing tables on Logical System LS1 and Logical System LS2 do not contain the default 0.0.0.0/0. This verifies that the default route is no longer present in the OSPF domain. To reactivate the **so-0/0/2.0** interface, issue the **activate logical-systems LS3 interfaces so-0/0/2 unit 0 family inet address 10.0.45.2/30** configuration-mode command.

### *Example: Configuring an OSPF Import Policy on Logical Systems*

This example shows how to configure an OSPF import policy on logical systems. OSPF import policies apply to external routes only. An external route is a route that is outside the OSPF AS.

- [Requirements on page 3334](#)
- [Overview on page 3334](#)
- [Configuration on page 3336](#)
- [Verification on page 3340](#)

#### **Requirements**

This example shows logical systems that are configured within a single physical router. The logical systems connect to each other by using logical tunnel (lt) interfaces. See [“Example: Connecting Logical Systems Within the Same Device Using Logical Tunnel Interfaces on MX Series Routers and EX Series Switches” on page 3279](#). Alternatively, you can use multiple physical routers.

#### **Overview**

External routes are learned by Autonomous System Border Routers (ASBRs). External routes can be advertised throughout the OSPF domain if you configure the ASBR to redistribute the route into OSPF. An external route might be learned by the ASBR from


a routing protocol other than OSPF, or the external route might be a static route that you configure on the ASBR.

OSPF import policy allows you to prevent external routes from being added to the routing tables of OSPF neighbors. The import policy does not impact the OSPF database. This means that the import policy has no impact on the link-state advertisements.

OSPF import policies have practical applications. Suppose, for example, that you are using OSPF to advertise a static route to the devices in your datacenter because you want some of the devices in the datacenter to use the static route. However, you want other devices in the datacenter to ignore the static route. So, you apply the OSPF import policy on the devices that you want to ignore the static route. The filtering is done only on external routes in OSPF. The intra-area and inter-area routes are not considered for filtering. The default action is to accept the route when the route does not match the policy.

Figure 64 on page 3335 shows the sample network.

**Figure 64: OSPF Import Policy on Logical Systems**



In this example, the logical systems operate as follows:

1. **LS3**—Logical System LS3 has a static route to the 10.0.16.0/30 network. The next hop for the static route is 10.0.60.1. LS3 has an OSPF export policy configured. The export policy redistributes static routes from LS3's routing table into LS3's OSPF database. Because the static route is in LS3's OSPF database, the route is advertised in a link state advertisement (LSA) to LS3's OSPF neighbor. LS3's OSPF neighbor is Logical System LS2.
2. **LS2**—Logical System LS2 receives the route advertisement from LS3. LS2 then installs the route into LS2's OSPF database. LS2 has an OSPF import policy configured that matches the static route to the 10.0.16.0/30 network and prevents the static route from being installed in LS2's routing table. However, because the route is in LS2's OSPF database, LS2 advertises the route to its OSPF neighbor, Logical System LS1.

3. LS1—Logical System LS1 receives the route advertisement from LS2. LS1 then installs the route into LS1's OSPF database. LS1 does not have an OSPF import policy configured that matches the static route to the 10.0.16.0/30 network. Therefore, the route gets installed in LS1's routing table.

### Configuration

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

```

LS3 set logical-systems LS3 interfaces so-0/0/0 unit 0 family inet address 10.0.60.2/30
 set logical-systems LS3 interfaces lt-1/2/0 unit 3 description LS3->LS2
 set logical-systems LS3 interfaces lt-1/2/0 unit 3 encapsulation ethernet
 set logical-systems LS3 interfaces lt-1/2/0 unit 3 peer-unit 4
 set logical-systems LS3 interfaces lt-1/2/0 unit 3 family inet address 10.0.2.1/30
 set logical-systems LS3 protocols ospf export export_static
 set logical-systems LS3 protocols ospf area 0.0.0.0 interface lt-1/2/0.3
 set logical-systems LS3 policy-options policy-statement export_static from protocol
 static
 set logical-systems LS3 policy-options policy-statement export_static then accept
 set logical-systems LS3 routing-options static route 10.0.16.0/30 next-hop 10.0.60.1

LS2 set logical-systems LS2 interfaces lt-1/2/0 unit 1 description LS2->LS1
 set logical-systems LS2 interfaces lt-1/2/0 unit 1 encapsulation ethernet
 set logical-systems LS2 interfaces lt-1/2/0 unit 1 peer-unit 2
 set logical-systems LS2 interfaces lt-1/2/0 unit 1 family inet address 10.0.0.2/30
 set logical-systems LS2 interfaces lt-1/2/0 unit 4 description LS2->LS3
 set logical-systems LS2 interfaces lt-1/2/0 unit 4 encapsulation ethernet
 set logical-systems LS2 interfaces lt-1/2/0 unit 4 peer-unit 3
 set logical-systems LS2 interfaces lt-1/2/0 unit 4 family inet address 10.0.2.2/30
 set logical-systems LS2 protocols ospf import filter_routes
 set logical-systems LS2 protocols ospf area 0.0.0.0 interface lt-1/2/0.1
 set logical-systems LS2 protocols ospf area 0.0.0.0 interface lt-1/2/0.4
 set logical-systems LS2 policy-options policy-statement filter_routes from route-filter
 10.0.16.0/30 exact
 set logical-systems LS2 policy-options policy-statement filter_routes then reject

LS1 set logical-systems LS1 interfaces lt-1/2/0 unit 2 description LS1->LS2
 set logical-systems LS1 interfaces lt-1/2/0 unit 2 encapsulation ethernet
 set logical-systems LS1 interfaces lt-1/2/0 unit 2 peer-unit 1
 set logical-systems LS1 interfaces lt-1/2/0 unit 2 family inet address 10.0.0.1/30
 set logical-systems LS1 protocols ospf area 0.0.0.0 interface lt-1/2/0.2

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure an OSPF import policy on logical systems:

1. Configure the interfaces.

[edit]

```

user@R1# set logical-systems LS3 interfaces so-0/0/0 unit 0 family inet address
10.0.60.2/30
user@R1# set logical-systems LS3 interfaces lt-1/2/0 unit 3 description LS3->LS2
user@R1# set logical-systems LS3 interfaces lt-1/2/0 unit 3 encapsulation ethernet
user@R1# set logical-systems LS3 interfaces lt-1/2/0 unit 3 peer-unit 4
user@R1# set logical-systems LS3 interfaces lt-1/2/0 unit 3 family inet address
10.0.2.1/30
user@R1# set logical-systems LS2 interfaces lt-1/2/0 unit 1 description LS2->LS1
user@R1# set logical-systems LS2 interfaces lt-1/2/0 unit 1 encapsulation ethernet
user@R1# set logical-systems LS2 interfaces lt-1/2/0 unit 1 peer-unit 2
user@R1# set logical-systems LS2 interfaces lt-1/2/0 unit 1 family inet address
10.0.0.2/30
user@R1# set logical-systems LS2 interfaces lt-1/2/0 unit 4 description LS2->LS3
user@R1# set logical-systems LS2 interfaces lt-1/2/0 unit 4 encapsulation ethernet
user@R1# set logical-systems LS2 interfaces lt-1/2/0 unit 4 peer-unit 3
user@R1# set logical-systems LS2 interfaces lt-1/2/0 unit 4 family inet address
10.0.2.2/30
user@R1# set logical-systems LS1 interfaces lt-1/2/0 unit 2 description LS1->LS2
user@R1# set logical-systems LS1 interfaces lt-1/2/0 unit 2 encapsulation ethernet
user@R1# set logical-systems LS1 interfaces lt-1/2/0 unit 2 peer-unit 1
user@R1# set logical-systems LS1 interfaces lt-1/2/0 unit 2 family inet address
10.0.0.1/30

```

2. Enable OSPF on the interfaces.

```

[edit]
user@R1# set logical-systems LS3 protocols ospf area 0.0.0.0 interface lt-1/2/0.3
user@R1# set logical-systems LS2 protocols ospf area 0.0.0.0 interface lt-1/2/0.1
user@R1# set logical-systems LS2 protocols ospf area 0.0.0.0 interface lt-1/2/0.4
user@R1# set logical-systems LS1 protocols ospf area 0.0.0.0 interface lt-1/2/0.2

```

3. Configure the static route on Logical System LS3.

```

[edit]
user@R1# set logical-systems LS3 routing-options static route 10.0.16.0/30 next-hop
10.0.60.1

```

4. On Logical System LS3, redistribute the static route into OSPF.

```

[edit]
user@R1# set logical-systems LS3 protocols ospf export export_static
user@R1# set logical-systems LS3 policy-options policy-statement export_static
from protocol static
user@R1# set logical-systems LS3 policy-options policy-statement export_static
then accept

```

5. On Logical System LS2, configure the OSPF import policy.

```

[edit]
user@R1# set logical-systems LS2 protocols ospf import filter_routes
user@R1# set logical-systems LS2 policy-options policy-statement filter_routes
from route-filter 10.0.16.0/30 exact
user@R1# set logical-systems LS2 policy-options policy-statement filter_routes
then reject

```

6. If you are done configuring the device, commit the configuration.

```

[edit]
user@R1# commit

```

## Results

Confirm your configuration by issuing the **show logical-systems** command.

```
user@R1# show logical-systems
LS1 {
 interfaces {
 lt-1/2/0 {
 unit 2 {
 description LS1->LS2;
 encapsulation ethernet;
 peer-unit 1;
 family inet {
 address 10.0.0.1/30;
 }
 }
 }
 }
 protocols {
 ospf {
 area 0.0.0.0 {
 interface lt-1/2/0.2;
 }
 }
 }
}
LS2 {
 interfaces {
 lt-1/2/0 {
 unit 1 {
 description LS2->LS1;
 encapsulation ethernet;
 peer-unit 2;
 family inet {
 address 10.0.0.2/30;
 }
 }
 unit 4 {
 description LS2->LS3;
 encapsulation ethernet;
 peer-unit 3;
 family inet {
 address 10.0.2.2/30;
 }
 }
 }
 }
 protocols {
 ospf {
 import filter_routes;
 area 0.0.0.0 {
 interface lt-1/2/0.1;
 interface lt-1/2/0.4;
 }
 }
 }
}
```



```

 }
 policy-options {
 policy-statement filter_routes {
 from {
 route-filter 10.0.16.0/30 exact;
 }
 then reject;
 }
 }
 }
LS3 {
 interfaces {
 so-0/0/0 {
 unit 0 {
 family inet {
 address 10.0.60.2/30;
 }
 }
 }
 lt-1/2/0 {
 unit 3 {
 description LS3->LS2;
 encapsulation ethernet;
 peer-unit 4;
 family inet {
 address 10.0.2.1/30;
 }
 }
 }
 }
 protocols {
 ospf {
 export export_static;
 area 0.0.0.0 {
 interface lt-1/2/0.3;
 }
 }
 }
 policy-options {
 policy-statement export_static {
 from protocol static;
 then accept;
 }
 }
 routing-options {
 static {
 route 10.0.16.0/30 next-hop 10.0.60.1;
 }
 }
}

```

**Verification**

Confirm that the configuration is working properly.

- [Viewing the OSPF Databases of the Logical Systems on page 3340](#)
- [Viewing the Routing Tables of the Logical Systems on page 3341](#)

**Viewing the OSPF Databases of the Logical Systems**

**Purpose** Verify that OSPF is advertising the static route.

**Action** user@R1> show ospf database logical-system all  
logical-system: LS2

```

 OSPF database, Area 0.0.0.0
 Type ID Adv Rtr Seq Age Opt Cksum Len
Router 10.0.0.1 10.0.0.1 0x8000001f 107 0x22 0x8f59 36
Router *10.0.0.2 10.0.0.2 0x80000025 101 0x22 0x4074 48
Router 10.0.2.1 10.0.2.1 0x80000018 107 0x22 0xab3a 36
Network 10.0.0.1 10.0.0.1 0x80000001 107 0x22 0x7b94 32
Network 10.0.2.1 10.0.2.1 0x8000000c 190 0x22 0x53ab 32
 OSPF AS SCOPE link state database
 Type ID Adv Rtr Seq Age Opt Cksum Len
Extern 10.0.16.0 10.0.2.1 0x80000007 1785 0x22 0x4147 36

```

logical-system: LS1

```

 OSPF database, Area 0.0.0.0
 Type ID Adv Rtr Seq Age Opt Cksum Len
Router *10.0.0.1 10.0.0.1 0x8000001f 107 0x22 0x8f59 36
Router 10.0.0.2 10.0.0.2 0x80000025 103 0x22 0x4074 48
Router 10.0.2.1 10.0.2.1 0x80000018 109 0x22 0xab3a 36
Network *10.0.0.1 10.0.0.1 0x80000001 107 0x22 0x7b94 32
Network 10.0.2.1 10.0.2.1 0x8000000c 192 0x22 0x53ab 32
 OSPF AS SCOPE link state database
 Type ID Adv Rtr Seq Age Opt Cksum Len
Extern 10.0.16.0 10.0.2.1 0x80000007 1787 0x22 0x4147 36

```

logical-system: LS3

```

 OSPF database, Area 0.0.0.0
 Type ID Adv Rtr Seq Age Opt Cksum Len
Router 10.0.0.1 10.0.0.1 0x8000001f 109 0x22 0x8f59 36
Router 10.0.0.2 10.0.0.2 0x80000025 103 0x22 0x4074 48
Router *10.0.2.1 10.0.2.1 0x80000018 107 0x22 0xab3a 36
Network 10.0.0.1 10.0.0.1 0x80000001 109 0x22 0x7b94 32
Network *10.0.2.1 10.0.2.1 0x8000000c 190 0x22 0x53ab 32
 OSPF AS SCOPE link state database
 Type ID Adv Rtr Seq Age Opt Cksum Len
Extern *10.0.16.0 10.0.2.1 0x80000007 1785 0x22 0x4147 36
...

```

**Meaning** The Extern \*10.0.16.0 output shows that OSPF is advertising the external route.

***Viewing the Routing Tables of the Logical Systems***

**Purpose** Make sure that Logical System LS3 and Logical System LS1 have the route to the 10.0.16.0/30 network installed in their respective routing tables. Make sure that Logical System LS2 does not have the route installed in its routing table.

**Action** user@R1> show route logical-system all  
logical-system: LS2

inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

```
10.0.0.0/30 *[Direct/0] 04:22:19
 > via lt-1/2/0.1
10.0.0.2/32 *[Local/0] 04:22:19
 Local via lt-1/2/0.1
10.0.2.0/30 *[Direct/0] 04:22:19
 > via lt-1/2/0.4
10.0.2.2/32 *[Local/0] 04:22:19
 Local via lt-1/2/0.4
224.0.0.5/32 *[OSPF/10] 04:22:23, metric 1
 MultiRecv
```

-----

logical-system: LS1

inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

```
10.0.0.0/30 *[Direct/0] 04:22:19
 > via lt-1/2/0.2
10.0.0.1/32 *[Local/0] 04:22:19
 Local via lt-1/2/0.2
10.0.2.0/30 *[OSPF/10] 00:07:52, metric 2
 > to 10.0.0.2 via lt-1/2/0.2
10.0.16.0/30 *[OSPF/150] 00:07:52, metric 0, tag 0
 > to 10.0.0.2 via lt-1/2/0.2
224.0.0.5/32 *[OSPF/10] 04:22:23, metric 1
 MultiRecv
```

-----

logical-system: LS3

inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

```
10.0.0.0/30 *[OSPF/10] 00:07:57, metric 2
 > to 10.0.2.2 via lt-1/2/0.3
10.0.2.0/30 *[Direct/0] 04:22:19
 > via lt-1/2/0.3
10.0.2.1/32 *[Local/0] 04:22:19
 Local via lt-1/2/0.3
10.0.16.0/30 *[Static/5] 03:51:18
 > to 10.0.60.1 via so-0/0/0.0
10.0.60.0/30 *[Direct/0] 03:53:52
 > via so-0/0/0.0
10.0.60.2/32 *[Local/0] 03:53:58
 Local via so-0/0/0.0
224.0.0.5/32 *[OSPF/10] 04:22:23, metric 1
 MultiRecv
```

**Meaning** The route to 10.0.16.0/30 is not installed in Logical System LS2's routing table. The route to 10.0.16.0/30 is installed in Logical System LS1's routing table as a route learned from OSPF. Because it is an OSPF external route, it has a preference value of 150 (instead of 10). By default, routes resulting from OSPF external LSAs are installed with a preference

value of 150. The route to 10.0.16.0/30 is installed in Logical System LS3's routing table as a static route.

#### Related Documentation

- [Introduction to Logical Systems on page 3259](#)

#### Examples: Configuring IS-IS on Logical Systems

- [IS-IS Overview on page 3343](#)
- [Example: Configuring IS-IS on Logical Systems Within the Same Router on page 3347](#)
- [Example: Configuring an IS-IS Default Route Policy on Logical Systems on page 3356](#)

#### IS-IS Overview

The IS-IS protocol is an interior gateway protocol (IGP) that uses link-state information to make routing decisions.

IS-IS is a link-state IGP that uses the shortest-path-first (SPF) algorithm to determine routes. IS-IS evaluates the topology changes and determines whether to perform a full SPF recalculation or a partial route calculation (PRC). This protocol originally was developed for routing International Organization for Standardization (ISO) Connectionless Network Protocol (CLNP) packets.

Like OSPF routing, IS-IS uses hello packets that allow network convergence to occur quickly when network changes are detected. IS-IS uses the SPF algorithm to determine routes. Using SPF, IS-IS evaluates network topology changes and determines if a full or partial route calculation is required.



**NOTE:** Because IS-IS uses ISO addresses, the configuration of IP version 6 (IPv6) and IP version 4 (IPv4) implementations of IS-IS is identical.

This section discusses the following topics:

- [IS-IS Terminology on page 3343](#)
- [ISO Network Addresses on page 3344](#)
- [IS-IS Packets on page 3345](#)
- [Persistent Route Reachability on page 3346](#)
- [IS-IS Support for Multipoint Network Clouds on page 3346](#)
- [Installing a Default Route to the Nearest Routing Device That Operates at Both IS-IS Levels on page 3347](#)

#### IS-IS Terminology

An IS-IS network is a single autonomous system (AS), also called a *routing domain*, that consists of *end systems* and *intermediate systems*. End systems are network entities that send and receive packets. Intermediate systems send and receive packets and relay (forward) packets. (Intermediate system is the Open System Interconnection [OSI] term for a router.) ISO packets are called network PDUs.

In IS-IS, a single AS can be divided into smaller groups called *areas*. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring *Level 1* and *Level 2* intermediate systems. Level 1 systems route within an area; when the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs. No IS-IS area functions strictly as a backbone.

Level 1 routers share intra-area routing information, and Level 2 routers share interarea information about IP addresses available within each area. Uniquely, IS-IS routers can act as both Level 1 and Level 2 routers, sharing intra-area routes with other Level 1 routers and interarea routes with other Level 2 routers.

The propagation of link-state updates is determined by the level boundaries. All routers within a level maintain a complete link-state database of all other routers in the same level. Each router then uses the Dijkstra algorithm to determine the shortest path from the local router to other routers in the link-state database.

### **ISO Network Addresses**

IS-IS uses ISO network addresses. Each address identifies a point of connection to the network, such as a router interface, and is called a *network service access point (NSAP)*.

IS-IS supports multiple NSAP addresses on the loopback lo0 interface.

An end system can have multiple NSAP addresses, in which case the addresses differ only by the last byte (called the *n-selector*). Each NSAP represents a service that is available at that node. In addition to having multiple services, a single node can belong to multiple areas.

Each network entity also has a special network address called a *network entity title (NET)*. Structurally, an NET is identical to an NSAP address but has an n-selector of 00. Most end systems and intermediate systems have one NET. Intermediate systems that participate in multiple areas can have multiple NETs.

The following ISO addresses illustrate the IS-IS address format:

```
49.0001.00a0.c96b.c490.00
49.0001.2081.9716.9018.00
```

NETs take several forms, depending on your network requirements. NET addresses are hexadecimal and range from 8 octets to 20 octets in length. Generally, the format consists of an authority and format Identifier (AFI), a domain ID, an area ID, a system identifier, and a selector. The simplest format omits the domain ID and is 10 octets long. For example, the NET address 49.0001.1921.6800.1001.00 consists of the following parts:

- 49—AFI
- 0001—Area ID
- 1921.6800.1001—System identifier
- 00—Selector

The system identifier must be unique within the network. For an IP-only network, we recommend using the IP address of an interface on the router. Configuring a loopback

NET address with the IP address is helpful when troubleshooting is required on the network.

The first portion of the address is the area number, which is a variable number from 1 through 13 bytes. The first byte of the area number (49) is the authority and format indicator (AFI). The next bytes are the assigned domain (area) identifier, which can be from 0 through 12 bytes. In the examples above, the area identifier is 0001.

The next six bytes form the system identifier. The system identifier can be any six bytes that are unique throughout the entire domain. The system identifier commonly is the media access control (MAC) address (as in the first example, 00a0.c96b.c490) or the IP address expressed in binary-coded decimal (BCD) (as in the second example, 2081.9716.9018, which corresponds to IP address 208.197.169.18). The last byte (00) is the n-selector.



**NOTE:** The system identifier cannot be 0000.0000.0000. All 0s is an illegal setting, and the adjacency is not formed with this setting.

To provide help with IS-IS debugging, the Junos<sup>®</sup> operating system (Junos OS) supports dynamic mapping of ISO system identifiers to the hostname. Each system can be configured with a hostname, which allows the system identifier-to-hostname mapping to be carried in a dynamic hostname type, length, and value (TLV) tuple in IS-IS link-state PDUs. This enables intermediate systems in the routing domain to learn about the ISO system identifier of a particular intermediate system.

### ***IS-IS Packets***

Each IS-IS PDU shares a common header. IS-IS uses the following PDUs to exchange protocol information:

- IS-IS hello (IIH) PDUs—Broadcast to discover the identity of neighboring IS-IS systems and to determine whether the neighbors are Level 1 or Level 2 intermediate systems.

IS-IS hello PDUs establish adjacencies with other routers and have three different formats: one for point-to-point hello packets, one for Level 1 broadcast links, and one for Level 2 broadcast links. Level 1 routers must share the same area address to form an adjacency, while Level 2 routers do not have this limitation. The request for adjacency is encoded in the Circuit type field of the PDU.

Hello PDUs have a preset length assigned to them. The IS-IS router does not resize any PDU to match the maximum transmission unit (MTU) on a router interface. Each interface supports the maximum IS-IS PDU of 1492 bytes, and hello PDUs are padded to meet the maximum value. When the hello is sent to a neighboring router, the connecting interface supports the maximum PDU size.

- Link-state PDUs—Contain information about the state of adjacencies to neighboring IS-IS systems. Link-state PDUs are flooded periodically throughout an area.

Also included is metric and IS-IS neighbor information. Each link-state PDU must be refreshed periodically on the network and is acknowledged by information within a sequence number PDU.

On point-to-point links, each link-state PDU is acknowledged by a partial sequence number PDU (PSNP), but on broadcast links, a complete sequence number PDU (CSNP) is sent out over the network. Any router that finds newer link-state PDU information in the CSNP then purges the out-of-date entry and updates the link-state database.

Link-state PDUs support variable-length subnet mask addressing.

- Complete sequence number PDUs (CSNPs)—Contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases. The designated router multicasts CSNPs on broadcast links in place of sending explicit acknowledgments for each link-state PDU.

Contained within the CSNP is a link-state PDU identifier, a lifetime, a sequence number, and a checksum for each entry in the database. Periodically, a CSNP is sent on both broadcast and point-to-point links to maintain a correct database. Also, the advertisement of CSNPs occurs when an adjacency is formed with another router. Like IS-IS hello PDUs, CSNPs come in two types: Level 1 and Level 2.

When a device receives a CSNP, it checks the database entries against its own local link-state database. If it detects missing information, the device requests specific link-state PDU details using a partial sequence number PDU (PSNP).

- Partial sequence number PDUs (PSNPs)—Sent multicast by a receiver when it detects that it is missing a link-state PDU (when its link-state PDU database is out of date). The receiver sends a PSNP to the system that transmitted the CSNP, effectively requesting that the missing link-state PDU be transmitted. That routing device, in turn, forwards the missing link-state PDU to the requesting routing device.

A PSNP is used by an IS-IS router to request link-state PDU information from a neighboring router. A PSNP can also explicitly acknowledge the receipt of a link-state PDU on a point-to-point link. On a broadcast link, a CSNP is used as implicit knowledge. Like hello PDUs and CSNPs, the PSNP also has two types: Level 1 and Level 2.

When a device compares a CSNP to its local database and determines that a link-state PDU is missing, the router issues a PSNP for the missing link-state PDU, which is returned in a link-state PDU from the router sending the CSNP. The received link-state PDU is then stored in the local database, and an acknowledgment is sent back to the originating router.

### ***Persistent Route Reachability***

IPv4 and IPv6 route reachability information in IS-IS link-state PDUs is preserved when you commit a configuration. IP prefixes are preserved with their original packet fragment upon link-state PDU regeneration.

### ***IS-IS Support for Multipoint Network Clouds***


IS-IS does not support multipoint configurations. Therefore, when configuring Frame Relay or Asynchronous Transfer Mode (ATM) networks, you must configure them as collections of point-to-point links, not as multipoint clouds.



### ***Installing a Default Route to the Nearest Routing Device That Operates at Both IS-IS Levels***

When a routing device that operates as both a Level 1 and Level 2 router (Router B) determines that it can reach at least one area other than its own (for example, in Area Y), it sets the ATTACHED bit in its Level 1 link-state PDU. Thereafter, the Level 1 router (Router A) introduces a default route pointing to the nearest attached routing device that operates as both a Level 1 and Level 2 router (Router B). See [Figure 65 on page 3347](#).

**Figure 65: Install Default Route to Nearest Routing Device That Operates at Both Level 1 and Level 2**



### ***Example: Configuring IS-IS on Logical Systems Within the Same Router***

This example shows how to configure an IS-IS network by using multiple logical systems that are running on a single physical router. The logical systems are connected by logical tunnel interfaces.

- [Requirements on page 3347](#)
- [Overview on page 3347](#)
- [Configuration on page 3348](#)
- [Verification on page 3353](#)

#### ***Requirements***


You must connect the logical systems by using logical tunnel (lt) interfaces. See [“Example: Connecting Logical Systems Within the Same Device Using Logical Tunnel Interfaces on MX Series Routers and EX Series Switches” on page 3279](#).

#### ***Overview***

This example shows an IS-IS configuration with three logical systems running on one physical router. Each logical system has its own routing table. The configuration enables the protocol on all logical tunnel interfaces that participate in the IS-IS domain.

[Figure 66 on page 3348](#) shows the sample network.

Figure 66: IS-IS on Logical Systems



### Configuration

#### CLI Quick Configuration

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

```
set logical-systems LS1 interfaces lt-0/1/0 unit 2 description LS1->LS2
set logical-systems LS1 interfaces lt-0/1/0 unit 2 encapsulation ethernet
set logical-systems LS1 interfaces lt-0/1/0 unit 2 peer-unit 1
set logical-systems LS1 interfaces lt-0/1/0 unit 2 family inet address 10.0.0.1/30
set logical-systems LS1 interfaces lt-0/1/0 unit 2 family iso
set logical-systems LS1 interfaces lt-0/1/0 unit 0 description LS1->LS3
set logical-systems LS1 interfaces lt-0/1/0 unit 0 encapsulation ethernet
set logical-systems LS1 interfaces lt-0/1/0 unit 0 peer-unit 5
set logical-systems LS1 interfaces lt-0/1/0 unit 0 family inet address 10.0.1.2/30
set logical-systems LS1 interfaces lt-0/1/0 unit 0 family iso
set logical-systems LS1 interfaces lo0 unit 1 family iso address 49.0001.1720.1600.1001.00
set logical-systems LS1 protocols isis interface lt-0/1/0.0
set logical-systems LS1 protocols isis interface lt-0/1/0.2
set logical-systems LS1 protocols isis interface lo0.1 passive
set logical-systems LS2 interfaces lt-0/1/0 unit 1 description LS2->LS1
set logical-systems LS2 interfaces lt-0/1/0 unit 1 encapsulation ethernet
set logical-systems LS2 interfaces lt-0/1/0 unit 1 peer-unit 2
set logical-systems LS2 interfaces lt-0/1/0 unit 1 family inet address 10.0.0.2/30
set logical-systems LS2 interfaces lt-0/1/0 unit 1 family iso
set logical-systems LS2 interfaces lt-0/1/0 unit 4 description LS2->LS3
set logical-systems LS2 interfaces lt-0/1/0 unit 4 encapsulation ethernet
set logical-systems LS2 interfaces lt-0/1/0 unit 4 peer-unit 3
set logical-systems LS2 interfaces lt-0/1/0 unit 4 family inet address 10.0.2.2/30
set logical-systems LS2 interfaces lt-0/1/0 unit 4 family iso
```

```

set logical-systems LS2 interfaces lo0 unit 2 family iso address
 49.0001.1720.1600.2002.00
set logical-systems LS2 protocols isis interface lt-0/1/0.1
set logical-systems LS2 protocols isis interface lt-0/1/0.4
set logical-systems LS2 protocols isis interface lo0.2 passive
set logical-systems LS3 interfaces lt-0/1/0 unit 3 description LS3->LS2
set logical-systems LS3 interfaces lt-0/1/0 unit 3 encapsulation ethernet
set logical-systems LS3 interfaces lt-0/1/0 unit 3 peer-unit 4
set logical-systems LS3 interfaces lt-0/1/0 unit 3 family inet address 10.0.2.1/30
set logical-systems LS3 interfaces lt-0/1/0 unit 3 family iso
set logical-systems LS3 interfaces lt-0/1/0 unit 5 description LS3->LS1
set logical-systems LS3 interfaces lt-0/1/0 unit 5 encapsulation ethernet
set logical-systems LS3 interfaces lt-0/1/0 unit 5 peer-unit 0
set logical-systems LS3 interfaces lt-0/1/0 unit 5 family inet address 10.0.1.1/30
set logical-systems LS3 interfaces lt-0/1/0 unit 5 family iso
set logical-systems LS3 interfaces lo0 unit 3 family iso address 49.0001.1234.1600.2231.00
set logical-systems LS3 protocols isis interface lt-0/1/0.5
set logical-systems LS3 protocols isis interface lt-0/1/0.3
set logical-systems LS3 protocols isis interface lo0.3 passive

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure IS-IS on logical systems:

1. Configure the logical tunnel interface on Logical System LS1 connecting to Logical System LS2.

```

[edit logical-systems LS1]
user@host# set interfaces lt-0/1/0 unit 2 description LS1->LS2
user@host# set interfaces lt-0/1/0 unit 2 encapsulation ethernet
user@host# set interfaces lt-0/1/0 unit 2 peer-unit 1
user@host# set interfaces lt-0/1/0 unit 2 family inet address 10.0.0.1/30
user@host# set interfaces lt-0/1/0 unit 2 family iso

```

2. Configure the logical tunnel interface on Logical System LS1 connecting to Logical System LS3.

```

[edit logical-systems LS1]
user@host# set interfaces lt-0/1/0 unit 0 description LS1->LS3
user@host# set interfaces lt-0/1/0 unit 0 encapsulation ethernet
user@host# set interfaces lt-0/1/0 unit 0 peer-unit 5
user@host# set interfaces lt-0/1/0 unit 0 family inet address 10.0.1.2/30
user@host# set interfaces lt-0/1/0 unit 0 family iso

```

3. Configure the logical tunnel interface on Logical System LS2 connecting to Logical System LS1.

```

[edit logical-systems LS2]
user@host# set interfaces lt-0/1/0 unit 1 description LS2->LS1
user@host# set interfaces lt-0/1/0 unit 1 encapsulation ethernet
user@host# set interfaces lt-0/1/0 unit 1 peer-unit 2
user@host# set interfaces lt-0/1/0 unit 1 family inet address 10.0.0.2/30
user@host# set interfaces lt-0/1/0 unit 1 family iso

```

4. Configure the logical tunnel interface on Logical System LS2 connecting to Logical System LS3.

```
[edit logical-systems LS2]
user@host# set interfaces lt-0/1/0 unit 4 description LS2->LS3
user@host# set interfaces lt-0/1/0 unit 4 encapsulation ethernet
user@host# set interfaces lt-0/1/0 unit 4 peer-unit 3
user@host# set interfaces lt-0/1/0 unit 4 family inet address 10.0.2.2/30
user@host# set interfaces lt-0/1/0 unit 4 family iso
```

5. Configure the logical tunnel interface on Logical System LS3 connecting to Logical System LS2.

```
[edit logical-systems LS3]
user@host# set interfaces lt-0/1/0 unit 3 description LS3->LS2
user@host# set interfaces lt-0/1/0 unit 3 encapsulation ethernet
user@host# set interfaces lt-0/1/0 unit 3 peer-unit 4
user@host# set interfaces lt-0/1/0 unit 3 family inet address 10.0.2.1/30
user@host# set interfaces lt-0/1/0 unit 3 family iso
```

6. Configure the logical tunnel interface on Logical System LS3 connecting to Logical System LS1.

```
[edit logical-systems LS3]
user@host# set interfaces lt-0/1/0 unit 5 description LS3->LS1
user@host# set interfaces lt-0/1/0 unit 5 encapsulation ethernet
user@host# set interfaces lt-0/1/0 unit 5 peer-unit 0
user@host# set interfaces lt-0/1/0 unit 5 family inet address 10.0.1.1/30
user@host# set interfaces lt-0/1/0 unit 5 family iso
```

7. Configure the ISO address on the loopback interface for the three logical systems.

```
[edit logical-systems LS1]
user@host# set interfaces lo0 unit 1 family iso address 49.0001.1720.1600.1001.00
user@host# set protocols isis interface lo0.1 passive
```

```
[edit logical-systems LS2]
user@host# set interfaces lo0 unit 2 family iso address 49.0001.1720.1600.2002.00
user@host# set protocols isis interface lo0.2 passive
```

```
[edit logical-systems LS3]
user@host# set interfaces lo0 unit 3 family iso address 49.0001.1234.1600.2231.00
user@host# set protocols isis interface lo0.3 passive
```

8. Configure IS-IS on all the interfaces.

```
[edit logical-systems LS1 protocols isis]
user@host# set interface lt-0/1/0.0
user@host# set interface lt-0/1/0.2
```

```
[edit logical-systems LS2 protocols isis]
user@host# set interface lt-0/1/0.1
user@host# set interface lt-0/1/0.4
```

```
[edit logical-systems LS3 protocols isis]
user@host# set interface lt-0/1/0.5
user@host# set interface lt-0/1/0.3
```

9. If you are done configuring the device, commit the configuration.

```
[edit]
```

```
user@host# commit
```

### Results

From configuration mode, confirm your configuration by issuing the **show logical-systems** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show logical-systems
LS1 {
 interfaces {
 lt-0/1/0 {
 unit 0 {
 description LS1->LS3;
 encapsulation ethernet;
 peer-unit 5;
 family inet {
 address 10.0.1.2/30;
 }
 family iso;
 }
 unit 2 {
 description LS1->LS2;
 encapsulation ethernet;
 peer-unit 1;
 family inet {
 address 10.0.0.1/30;
 }
 family iso;
 }
 }
 lo0 {
 unit 1 {
 family iso {
 address 49.0001.1720.1600.1001.00;
 }
 }
 }
 }
 protocols {
 isis {
 interface lt-0/1/0.0;
 interface lt-0/1/0.2;
 interface lo0.1 {
 passive;
 }
 }
 }
}
LS2 {
 interfaces {
 lt-0/1/0 {
 unit 1 {
 description LS2->LS1;
 encapsulation ethernet;
 peer-unit 2;
 family inet {
 address 10.0.0.2/30;
 }
 family iso;
 }
 }
 }
}
```

```
 }
 unit 4 {
 description LS2->LS3;
 encapsulation ethernet;
 peer-unit 3;
 family inet {
 address 10.0.2.2/30;
 }
 family iso;
 }
}
lo0 {
 unit 2 {
 family iso {
 address 49.0001.1720.1600.2002.00;
 }
 }
}
}
protocols {
 isis {
 interface lt-0/1/0.1;
 interface lt-0/1/0.4;
 interface lo0.2 {
 passive;
 }
 }
}
}
LS3 {
 interfaces {
 lt-0/1/0 {
 unit 3 {
 description LS3->LS2;
 encapsulation ethernet;
 peer-unit 4;
 family inet {
 address 10.0.2.1/30;
 }
 family iso;
 }
 unit 5 {
 description LS3->LS1;
 encapsulation ethernet;
 peer-unit 0;
 family inet {
 address 10.0.1.1/30;
 }
 family iso;
 }
 }
 lo0 {
 unit 3 {
 family iso {
 address 49.0001.1234.1600.2231.00;
 }
 }
 }
 }
}
protocols {
```

```

isis {
 interface lt-0/1/0.3;
 interface lt-0/1/0.5;
 interface lo0.3 {
 passive;
 }
}
}

```

### Verification

Confirm that the configuration is working properly.

- [Verifying That the Logical Systems Are Up on page 3353](#)
- [Verifying Connectivity Between the Logical Systems on page 3353](#)

### Verifying That the Logical Systems Are Up

**Purpose** Make sure that the interfaces are properly configured.

|               |                                  |       |      |       |             |        |
|---------------|----------------------------------|-------|------|-------|-------------|--------|
| <b>Action</b> | user@host> show interfaces terse |       |      |       |             |        |
|               | Interface                        | Admin | Link | Proto | Local       | Remote |
|               | ...                              |       |      |       |             |        |
|               | lt-0/1/0                         | up    | up   |       |             |        |
|               | lt-0/1/0.0                       | up    | up   | inet  | 10.0.1.2/30 |        |
|               |                                  |       |      | iso   |             |        |
|               | lt-0/1/0.1                       | up    | up   | inet  | 10.0.0.2/30 |        |
|               |                                  |       |      | iso   |             |        |
|               | lt-0/1/0.2                       | up    | up   | inet  | 10.0.0.1/30 |        |
|               |                                  |       |      | iso   |             |        |
|               | lt-0/1/0.3                       | up    | up   | inet  | 10.0.2.1/30 |        |
|               |                                  |       |      | iso   |             |        |
|               | lt-0/1/0.4                       | up    | up   | inet  | 10.0.2.2/30 |        |
|               |                                  |       |      | iso   |             |        |
|               | lt-0/1/0.5                       | up    | up   | inet  | 10.0.1.1/30 |        |
|               |                                  |       |      | iso   |             |        |
|               | ...                              |       |      |       |             |        |

### Verifying Connectivity Between the Logical Systems

**Purpose** Make sure that the IS-IS adjacencies are established by checking the logical system routing entries and by pinging the logical systems.

```

Action user@host> show route logical-system LS1
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30 *[Direct/0] 3w0d 01:37:52
 > via lt-0/1/0.2
10.0.0.1/32 *[Local/0] 3w0d 01:37:52
 Local via lt-0/1/0.2
10.0.1.0/30 *[Direct/0] 3w0d 01:37:52
 > via lt-0/1/0.0
10.0.1.2/32 *[Local/0] 3w0d 01:37:52
 Local via lt-0/1/0.0
10.0.2.0/30 *[IS-IS/15] 3w0d 01:37:13, metric 20
 > to 10.0.1.1 via lt-0/1/0.0
 to 10.0.0.2 via lt-0/1/0.2

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0001.1720.1600.1001/72
 *[Direct/0] 3w0d 01:37:52
 > via lo0.1

user@host> show route logical-system LS2
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30 *[Direct/0] 3w0d 01:38:01
 > via lt-0/1/0.1
10.0.0.2/32 *[Local/0] 3w0d 01:38:01
 Local via lt-0/1/0.1
10.0.1.0/30 *[IS-IS/15] 3w0d 01:37:01, metric 20
 to 10.0.0.1 via lt-0/1/0.1
 > to 10.0.2.1 via lt-0/1/0.4
10.0.2.0/30 *[Direct/0] 3w0d 01:38:01
 > via lt-0/1/0.4
10.0.2.2/32 *[Local/0] 3w0d 01:38:01
 Local via lt-0/1/0.4

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0001.1720.1600.2002/72
 *[Direct/0] 3w0d 01:38:01
 > via lo0.2

user@host> show route logical-system LS3
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30 *[IS-IS/15] 3w0d 01:37:10, metric 20
 to 10.0.2.2 via lt-0/1/0.3
 > to 10.0.1.2 via lt-0/1/0.5
10.0.1.0/30 *[Direct/0] 3w0d 01:38:10
 > via lt-0/1/0.5
10.0.1.1/32 *[Local/0] 3w0d 01:38:11
 Local via lt-0/1/0.5
10.0.2.0/30 *[Direct/0] 3w0d 01:38:11
 > via lt-0/1/0.3

```



```

10.0.2.1/32 *[Local/0] 3w0d 01:38:11
 Local via lt-0/1/0.3

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0001.1234.1600.2231/72
 *[Direct/0] 3w0d 01:38:11
 > via lo0.3

```

#### From LS1, Ping LS3

```

user@host> set cli logical-system LS1

user@host:LS1> ping 10.0.2.1
PING 10.0.2.1 (10.0.2.1): 56 data bytes
64 bytes from 10.0.2.1: icmp_seq=0 ttl=63 time=1.264 ms
64 bytes from 10.0.2.1: icmp_seq=1 ttl=63 time=1.189 ms
64 bytes from 10.0.2.1: icmp_seq=2 ttl=63 time=1.165 ms
^C
--- 10.0.2.1 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.165/1.206/1.264/0.042 ms

```

#### From LS3, Ping LS1

```

user@host> set cli logical-system LS3

user@host:LS3> ping 10.0.0.1
PING 10.0.0.1 (10.0.0.1): 56 data bytes
64 bytes from 10.0.0.1: icmp_seq=0 ttl=63 time=1.254 ms
64 bytes from 10.0.0.1: icmp_seq=1 ttl=63 time=1.210 ms
^C
--- 10.0.0.1 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.210/1.232/1.254/0.022 ms

```

#### From LS1, Ping LS2

```

user@host> set cli logical-system LS1

user@host:LS1> ping 10.0.2.2
PING 10.0.2.2 (10.0.2.2): 56 data bytes
64 bytes from 10.0.2.2: icmp_seq=0 ttl=64 time=1.240 ms
64 bytes from 10.0.2.2: icmp_seq=1 ttl=64 time=1.204 ms
64 bytes from 10.0.2.2: icmp_seq=2 ttl=64 time=1.217 ms
^C
--- 10.0.2.2 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.204/1.220/1.240/0.015 ms

```

#### From LS2, Ping LS1

```

user@host> set cli logical-system LS2

user@host:LS2> ping 10.0.1.2
PING 10.0.1.2 (10.0.1.2): 56 data bytes
64 bytes from 10.0.1.2: icmp_seq=0 ttl=64 time=1.308 ms
64 bytes from 10.0.1.2: icmp_seq=1 ttl=64 time=1.235 ms
^C
--- 10.0.1.2 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.235/1.272/1.308/0.037 ms

```

### From LS2, Ping LS3

```
user@host> set cli logical-system LS2

user@host:LS2> ping 10.0.1.1
PING 10.0.1.1 (10.0.1.1): 56 data bytes
64 bytes from 10.0.1.1: icmp_seq=0 ttl=64 time=1.253 ms
64 bytes from 10.0.1.1: icmp_seq=1 ttl=64 time=1.194 ms
64 bytes from 10.0.1.1: icmp_seq=2 ttl=64 time=1.212 ms
64 bytes from 10.0.1.1: icmp_seq=3 ttl=64 time=1.221 ms
64 bytes from 10.0.1.1: icmp_seq=4 ttl=64 time=1.195 ms
^C
--- 10.0.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.194/1.215/1.253/0.022 ms
```

### From LS3, Ping LS2

```
user@host> set cli logical-system LS3

user@host:LS3> ping 10.0.0.2
PING 10.0.0.2 (10.0.0.2): 56 data bytes
64 bytes from 10.0.0.2: icmp_seq=0 ttl=64 time=1.240 ms
64 bytes from 10.0.0.2: icmp_seq=1 ttl=64 time=1.217 ms
^C
--- 10.0.0.2 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.217/1.228/1.240/0.012 ms
```

#### *Example: Configuring an IS-IS Default Route Policy on Logical Systems*

This example shows logical systems configured on a single physical router and explains how to configure a default route on one logical system.

- [Requirements on page 3356](#)
- [Overview on page 3356](#)
- [Configuration on page 3357](#)
- [Verification on page 3360](#)

#### **Requirements**

No special configuration beyond device initialization is required before configuring this example.


#### **Overview**

This example shows a logical system redistributing a default route to other logical systems. All logical systems are running IS-IS. A common reason for a default route is to provide a path for sending traffic destined outside the IS-IS domain.

In this example, the default route is not used for forwarding traffic. The **no-install** statement prevents the route from being installed in the forwarding table of Logical System LS3. If you configure a route so it is not installed in the forwarding table, the route is still eligible to be exported from the routing table to other protocols. The **discard** statement silently drops packets without notice.

[Figure 67 on page 3357](#) shows the sample network.

Figure 67: IS-IS with a Default Route to an ISP



g040918

**Configuration****CLI Quick Configuration**

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

```

set logical-systems LS3 interfaces lt-1/2/0 unit 3 description LS3->LS2
set logical-systems LS3 interfaces lt-1/2/0 unit 3 encapsulation ethernet
set logical-systems LS3 interfaces lt-1/2/0 unit 3 peer-unit 4
set logical-systems LS3 interfaces lt-1/2/0 unit 3 family inet address 10.0.2.1/30
set logical-systems LS3 interfaces lt-1/2/0 unit 3 family iso
set logical-systems LS3 interfaces lt-1/2/0 unit 5 description LS3->LS1
set logical-systems LS3 interfaces lt-1/2/0 unit 5 encapsulation ethernet
set logical-systems LS3 interfaces lt-1/2/0 unit 5 peer-unit 0
set logical-systems LS3 interfaces lt-1/2/0 unit 5 family inet address 10.0.1.1/30
set logical-systems LS3 interfaces lt-1/2/0 unit 5 family iso
set logical-systems LS3 interfaces lo0 unit 3 family iso address 49.0001.1234.1600.2231.00
set logical-systems LS3 protocols isis export isis-default
set logical-systems LS3 protocols isis interface lt-1/2/0.3
set logical-systems LS3 protocols isis interface lt-1/2/0.5
set logical-systems LS3 protocols isis interface lo0.3 passive
set logical-systems LS3 routing-options static route 0.0.0.0/0 discard
set logical-systems LS3 routing-options static route 0.0.0.0/0 no-install
set logical-systems LS3 policy-options policy-statement isis-default from protocol static
set logical-systems LS3 policy-options policy-statement isis-default from route-filter
 0.0.0.0/0 exact
set logical-systems LS3 policy-options policy-statement isis-default then accept

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure an IS-IS default route policy on logical systems:

1. Configure the logical tunnel interfaces.  

```
[edit logical-systems LS3 interfaces lt-1/2/0]
user@R1# set unit 3 description LS3->LS2
user@R1# set unit 3 encapsulation ethernet
user@R1# set unit 3 peer-unit 4
user@R1# set unit 3 family inet address 10.0.2.1/30
user@R1# set unit 3 family iso
user@R1# set unit 5 description LS3->LS1
user@R1# set unit 5 encapsulation ethernet
user@R1# set unit 5 peer-unit 0
user@R1# set unit 5 family inet address 10.0.1.1/30
user@R1# set unit 5 family iso
[edit logical-systems LS3 interfaces lo0 unit 3]
user@R1# set family iso address 49.0001.1234.1600.2231.00
```
2. Enable IS-IS on the interfaces.  

```
[edit logical-systems LS3 protocols isis]
user@R1# set interface lt-1/2/0.3
user@R1# set interface lt-1/2/0.5
user@R1# set interface lo0.3 passive
```
3. Configure the default route on Logical System LS3.  

```
[edit logical-systems LS3 routing-options]
user@R1# set static route 0.0.0.0/0 discard
user@R1# set static route 0.0.0.0/0 no-install
```
4. Configure the default route policy on Logical System LS3.  

```
[edit logical-systems LS3 policy-options]
user@R1# set policy-statement isis-default from protocol static
user@R1# set policy-statement isis-default from route-filter 0.0.0.0/0 exact
user@R1# set policy-statement isis-default then accept
```
5. Apply the export policy to IS-IS on Logical System LS3.  

```
[edit logical-systems LS3 protocols isis]
user@R1# set export isis-default
```
6. If you are done configuring the device, commit the configuration.  

```
[edit]
user@R1# commit
```

### Results

From configuration mode, confirm your configuration by issuing the **show logical-systems LS3** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show logical-systems LS3
interfaces {
 lt-1/2/0 {
 unit 3 {
 description LS3->LS2;
 encapsulation ethernet;
 peer-unit 4;
 family inet {
 address 10.0.2.1/30;
 }
 family iso;
 }
 unit 5 {
 description LS3->LS1;
 encapsulation ethernet;
 peer-unit 0;
 family inet {
 address 10.0.1.1/30;
 }
 family iso;
 }
 }
 lo0 {
 unit 3 {
 family iso {
 address 49.0001.1234.1600.2231.00;
 }
 }
 }
}
protocols {
 isis {
 export isis-default;
 interface lt-1/2/0.3;
 interface lt-1/2/0.5;
 interface lo0.3 {
 passive;
 }
 }
}
policy-options {
 policy-statement isis-default {
 from {
 protocol static;
 route-filter 0.0.0.0/0 exact;
 }
 then accept;
 }
}
routing-options {
 static {
 route 0.0.0.0/0 {
 discard;
 no-install;
 }
 }
}
```

```
}
}
```

### ***Verification***

Confirm that the configuration is working properly.

### ***Verifying That the Static Route Is Redistributed***

**Purpose** Make sure that the IS-IS policy is working by checking the routing tables.

```

Action user@R1> show route logical-system LS3
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0 *[Static/5] 00:00:45
 Discard
10.0.0.0/30 *[IS-IS/15] 1w0d 10:14:14, metric 20
 to 10.0.2.2 via lt-1/2/0.3
 > to 10.0.1.2 via lt-1/2/0.5
10.0.1.0/30 *[Direct/0] 1w0d 10:15:18
 > via lt-1/2/0.5
10.0.1.1/32 *[Local/0] 1w0d 10:15:18
 Local via lt-1/2/0.5
10.0.2.0/30 *[Direct/0] 1w0d 10:15:18
 > via lt-1/2/0.3
10.0.2.1/32 *[Local/0] 1w0d 10:15:18
 Local via lt-1/2/0.3

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0001.1234.1600.2231/72
 *[Direct/0] 1w0d 10:17:19
 > via lo0.3

user@R1> show route logical-system LS2
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0 *[IS-IS/160] 00:01:38, metric 10
 > to 10.0.2.1 via lt-1/2/0.4
10.0.0.0/30 *[Direct/0] 1w0d 10:16:11
 > via lt-1/2/0.1
10.0.0.2/32 *[Local/0] 1w0d 10:16:11
 Local via lt-1/2/0.1
10.0.1.0/30 *[IS-IS/15] 1w0d 10:15:07, metric 20
 > to 10.0.0.1 via lt-1/2/0.1
 to 10.0.2.1 via lt-1/2/0.4
10.0.2.0/30 *[Direct/0] 1w0d 10:16:11
 > via lt-1/2/0.4
10.0.2.2/32 *[Local/0] 1w0d 10:16:11
 Local via lt-1/2/0.4

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0001.1720.1600.2002/72
 *[Direct/0] 1w0d 10:18:12
 > via lo0.2

user@R1> show route logical-system LS1
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0 *[IS-IS/160] 00:02:01, metric 10
 > to 10.0.1.1 via lt-1/2/0.0
10.0.0.0/30 *[Direct/0] 1w0d 10:16:34
 > via lt-1/2/0.2
10.0.0.1/32 *[Local/0] 1w0d 10:16:34
 Local via lt-1/2/0.2

```

```
10.0.1.0/30 *[Direct/0] 1w0d 10:16:34
 > via lt-1/2/0.0
10.0.1.2/32 *[Local/0] 1w0d 10:16:34
 Local via lt-1/2/0.0
10.0.2.0/30 *[IS-IS/15] 1w0d 10:15:55, metric 20
 to 10.0.1.1 via lt-1/2/0.0
 > to 10.0.0.2 via lt-1/2/0.2

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0001.1720.1600.1001/72
 *[Direct/0] 1w0d 10:18:35
 > via lo0.1
```

**Meaning** The routing table on Logical System LS3 contains the default 0.0.0.0/0 route from protocol **IS-IS**. The routing tables on Logical System LS1 and Logical System LS2 contain the default 0.0.0.0/0 route from protocol **IS-IS**. If Logical System LS1 and Logical System LS2 receive packets destined for networks not specified in their routing tables, those packets will be sent to Logical System LS3 for further processing. This configuration assumes that Logical System LS3 has a connection to an ISP or another external network.

---

#### Examples: Configuring BGP on Logical Systems

- [Understanding BGP on page 3363](#)
- [Example: Configuring Internal BGP Peering Sessions on Logical Systems on page 3365](#)
- [Example: Configuring External BGP on Logical Systems with IPv6 Interfaces on page 3375](#)
- [Example: Configuring BFD on Internal BGP Peer Sessions on page 3390](#)
- [Example: Configuring EBGP Multihop Sessions on Logical Systems on page 3398](#)



### ***Understanding BGP***

BGP is an exterior gateway protocol (EGP) that is used to exchange routing information among routing devices in different autonomous systems (ASs). BGP routing information includes the complete route to each destination. BGP uses the routing information to maintain a database of network reachability information, which it exchanges with other BGP systems. BGP uses the network reachability information to construct a graph of AS connectivity, which enables BGP to remove routing loops and enforce policy decisions at the AS level.

Multiprotocol BGP (MBGP) extensions enable BGP to support IP version 6 (IPv6). MBGP defines the attributes `MP_REACH_NLRI` and `MP_UNREACH_NLRI`, which are used to carry IPv6 reachability information. Network layer reachability information (NLRI) update messages carry IPv6 address prefixes of feasible routes.

BGP allows for policy-based routing. You can use routing policies to choose among multiple paths to a destination and to control the redistribution of routing information.

BGP uses TCP as its transport protocol, using port 179 for establishing connections. Running over a reliable transport protocol eliminates the need for BGP to implement update fragmentation, retransmission, acknowledgment, and sequencing.

The Junos OS routing protocol software supports BGP version 4. This version of BGP adds support for Classless Interdomain Routing (CIDR), which eliminates the concept of network classes. Instead of assuming which bits of an address represent the network by looking at the first octet, CIDR allows you to explicitly specify the number of bits in the network address, thus providing a means to decrease the size of the routing tables. BGP version 4 also supports aggregation of routes, including the aggregation of AS paths.

This section discusses the following topics:

- [Autonomous Systems on page 3363](#)
- [AS Paths and Attributes on page 3363](#)
- [External and Internal BGP on page 3364](#)
- [Multiple Instances of BGP on page 3364](#)

### ***Autonomous Systems***

An *autonomous system* (AS) is a set of routing devices that are under a single technical administration and normally use a single interior gateway protocol and a common set of metrics to propagate routing information within the set of routing devices. To other ASs, an AS appears to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it.

### ***AS Paths and Attributes***

The routing information that BGP systems exchange includes the complete route to each destination, as well as additional information about the route. The route to each destination is called the *AS path*, and the additional route information is included in *path attributes*. BGP uses the AS path and the path attributes to completely determine the network topology. Once BGP understands the topology, it can detect and eliminate


routing loops and select among groups of routes to enforce administrative preferences and routing policy decisions.

### External and Internal BGP

BGP supports two types of exchanges of routing information: exchanges among different ASs and exchanges within a single AS. When used among ASs, BGP is called *external BGP* (EBGP) and BGP sessions perform *inter-AS routing*. When used within an AS, BGP is called *internal BGP* (IBGP) and BGP sessions perform *intra-AS routing*.

Figure 68 on page 3364 illustrates ASs, IBGP, and EBGP.

Figure 68: ASs, EBGP, and IBGP



A BGP system shares network reachability information with adjacent BGP systems, which are referred to as *neighbors* or *peers*.

BGP systems are arranged into *groups*. In an IBGP group, all peers in the group—called *internal peers*—are in the same AS. Internal peers can be anywhere in the local AS and do not have to be directly connected to one another. Internal groups use routes from an IGP to resolve forwarding addresses. They also propagate external routes among all other internal routing devices running IBGP, computing the next hop by taking the BGP next hop received with the route and resolving it using information from one of the interior gateway protocols.

In an EBGP group, the peers in the group—called *external peers*—are in different ASs and normally share a subnet. In an external group, the next hop is computed with respect to the interface that is shared between the external peer and the local routing device.

### Multiple Instances of BGP

You can configure multiple instances of BGP at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]

Multiple instances of BGP are primarily used for Layer 3 VPN support.

IGP peers and external BGP (EBGP) peers (both nonmultihop and multihop) are all supported for routing instances. BGP peering is established over one of the interfaces configured under the **routing-instances** hierarchy.



**NOTE:** When a BGP neighbor sends BGP messages to the local routing device, the incoming interface on which these messages are received must be configured in the same routing instance that the BGP neighbor configuration exists in. This is true for neighbors that are a single hop away or multiple hops away.

Routes learned from the BGP peer are added to the **instance-name.inet.0** table by default. You can configure import and export policies to control the flow of information into and out of the instance routing table.

For Layer 3 VPN support, configure BGP on the provider edge (PE) router to receive routes from the customer edge (CE) router and to send the instances' routes to the CE router if necessary. You can use multiple instances of BGP to maintain separate per-site forwarding tables for keeping VPN traffic separate on the PE router.

You can configure import and export policies that allow the service provider to control and rate-limit traffic to and from the customer.

You can configure an EBGP multihop session for a VRF routing instance. Also, you can set up the EBGP peer between the PE and CE routers by using the loopback address of the CE router instead of the interface addresses.

### ***Example: Configuring Internal BGP Peering Sessions on Logical Systems***

This example shows how to configure internal BGP peer sessions on logical systems.

- [Requirements on page 3365](#)
- [Overview on page 3365](#)
- [Configuration on page 3366](#)
- [Verification on page 3372](#)

#### ***Requirements***

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


#### ***Overview***

In this example, you configure internal BGP (IBGP) peering sessions.

In the sample network, the devices in AS 17 are fully meshed in the group **internal-peers**. The devices have loopback addresses 192.168.6.5, 192.163.6.4, and 192.168.40.4.

[Figure 69 on page 3366](#) shows a typical network with internal peer sessions.

Figure 69: Typical Network with IBGP Sessions

**Configuration****CLI Quick Configuration**

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

```

set logical-systems A interfaces lt-0/1/0 unit 1 description to-B
set logical-systems A interfaces lt-0/1/0 unit 1 encapsulation ethernet
set logical-systems A interfaces lt-0/1/0 unit 1 peer-unit 2
set logical-systems A interfaces lt-0/1/0 unit 1 family inet address 10.10.10.1/30
set logical-systems A interfaces lo0 unit 1 family inet address 192.168.6.5/32
set logical-systems A protocols bgp group internal-peers type internal
set logical-systems A protocols bgp group internal-peers local-address 192.168.6.5
set logical-systems A protocols bgp group internal-peers export send-direct
set logical-systems A protocols bgp group internal-peers neighbor 192.163.6.4
set logical-systems A protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems A protocols ospf area 0.0.0.0 interface lo0.1 passive
set logical-systems A protocols ospf area 0.0.0.0 interface lt-0/1/0.1
set logical-systems A policy-options policy-statement send-direct term 2 from protocol
 direct
set logical-systems A policy-options policy-statement send-direct term 2 then accept
set logical-systems A routing-options router-id 192.168.6.5
set logical-systems A routing-options autonomous-system 17
set logical-systems B interfaces lt-0/1/0 unit 2 description to-A
set logical-systems B interfaces lt-0/1/0 unit 2 encapsulation ethernet
set logical-systems B interfaces lt-0/1/0 unit 2 peer-unit 1
set logical-systems B interfaces lt-0/1/0 unit 2 family inet address 10.10.10.2/30
set logical-systems B interfaces lt-0/1/0 unit 5 description to-C
set logical-systems B interfaces lt-0/1/0 unit 5 encapsulation ethernet
set logical-systems B interfaces lt-0/1/0 unit 5 peer-unit 6
set logical-systems B interfaces lt-0/1/0 unit 5 family inet address 10.10.10.5/30
set logical-systems B interfaces lo0 unit 2 family inet address 192.163.6.4/32
set logical-systems B protocols bgp group internal-peers type internal
set logical-systems B protocols bgp group internal-peers local-address 192.163.6.4
set logical-systems B protocols bgp group internal-peers export send-direct
set logical-systems B protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems B protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems B protocols ospf area 0.0.0.0 interface lo0.2 passive

```

```

set logical-systems B protocols ospf area 0.0.0.0 interface lt-0/1/0.2
set logical-systems B protocols ospf area 0.0.0.0 interface lt-0/1/0.5
set logical-systems B policy-options policy-statement send-direct term 2 from protocol
 direct
set logical-systems B policy-options policy-statement send-direct term 2 then accept
set logical-systems B routing-options router-id 192.163.6.4
set logical-systems B routing-options autonomous-system 17
set logical-systems C interfaces lt-0/1/0 unit 6 description to-B
set logical-systems C interfaces lt-0/1/0 unit 6 encapsulation ethernet
set logical-systems C interfaces lt-0/1/0 unit 6 peer-unit 5
set logical-systems C interfaces lt-0/1/0 unit 6 family inet address 10.10.10.6/30
set logical-systems C interfaces lo0 unit 3 family inet address 192.168.40.4/32
set logical-systems C protocols bgp group internal-peers type internal
set logical-systems C protocols bgp group internal-peers local-address 192.168.40.4
set logical-systems C protocols bgp group internal-peers export send-direct
set logical-systems C protocols bgp group internal-peers neighbor 192.163.6.4
set logical-systems C protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems C protocols ospf area 0.0.0.0 interface lo0.3 passive
set logical-systems C protocols ospf area 0.0.0.0 interface lt-0/1/0.6
set logical-systems C policy-options policy-statement send-direct term 2 from protocol
 direct
set logical-systems C policy-options policy-statement send-direct term 2 then accept
set logical-systems C routing-options router-id 192.168.40.4
set logical-systems C routing-options autonomous-system 17

```

#### Device A

##### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure internal BGP peer sessions on Device A:

1. Configure the interfaces.

```

[edit logical-systems A interfaces lt-0/1/0 unit 1]
user@R1# set description to-B
user@R1# set encapsulation ethernet
user@R1# set peer-unit 2
user@R1# set family inet address 10.10.10.1/30
user@R1# set family inet address 192.168.6.5/32
user@R1# up
user@R1# up
[edit logical-systems A interfaces]
user@R1# set lo0 unit 1 family inet address 192.168.6.5/32
user@R1# exit
[edit]
user@R1# edit logical-systems B interfaces lt-0/1/0
[edit logical-systems B interfaces lt-0/1/0]
user@R1# set unit 2 description to-A
user@R1# set unit 2 encapsulation ethernet
user@R1# set unit 2 peer-unit 1
user@R1# set unit 2 family inet address 10.10.10.2/30
user@R1# set unit 5 description to-C
user@R1# set unit 5 encapsulation ethernet
user@R1# set unit 5 peer-unit 6

```

```
user@R1# set family inet address 10.10.10.5/30
user@R1# up
[edit logical-systems B interfaces]
user@R1# set lo0 unit 2 family inet address 192.163.6.4/32
user@R1# exit
[edit]
user@R1# edit logical-systems C interfaces lt-0/1/0 unit 6
[edit logical-systems C interfaces lt-0/1/0 unit 6]
set description to-B
set encapsulation ethernet
set peer-unit 5
set family inet address 10.10.10.6/30
user@R1# up
user@R1# up
[edit logical-systems C interfaces]
set lo0 unit 3 family inet address 192.168.40.4/32
```

2. Configure BGP.

On Logical System A, the **neighbor** statements are included for both Device B and Device C, even though Logical System A is not directly connected to Device C.

```
[edit logical-systems A protocols bgp group internal-peers]
user@R1# set type internal
user@R1# set local-address 192.168.6.5
user@R1# set export send-direct
user@R1# set neighbor 192.163.6.4
user@R1# set neighbor 192.168.40.4
```

```
[edit logical-systems B protocols bgp group internal-peers]
user@R1# set type internal
user@R1# set local-address 192.163.6.4
user@R1# set export send-direct
user@R1# set neighbor 192.168.40.4
user@R1# set neighbor 192.168.6.5
```

```
[edit logical-systems C protocols bgp group internal-peers]
user@R1# set type internal
user@R1# set local-address 192.168.40.4
user@R1# set export send-direct
user@R1# set neighbor 192.163.6.4
user@R1# set neighbor 192.168.6.5
```

3. Configure OSPF.

```
[edit logical-systems A protocols ospf area 0.0.0.0]
user@R1# set interface lo0.1 passive
user@R1# set interface lt-0/1/0.1
```

```
[edit logical-systems A protocols ospf area 0.0.0.0]
user@R1# set interface lo0.2 passive
user@R1# set interface lt-0/1/0.2
user@R1# set interface lt-0/1/0.5
```

```
[edit logical-systems A protocols ospf area 0.0.0.0]
user@R1# set interface lo0.3 passive
```

```
user@R1# set interface lt-0/1/0.6
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit logical-systems A policy-options policy-statement send-direct term 2]
user@R1# set from protocol direct
user@R1# set then accept
```

```
[edit logical-systems B policy-options policy-statement send-direct term 2]
user@R1# set from protocol direct
user@R1# set then accept
```

```
[edit logical-systems C policy-options policy-statement send-direct term 2]
user@R1# set from protocol direct
user@R1# set then accept
```

5. Configure the router ID and the autonomous system (AS) number.

```
[edit logical-systems A routing-options]
user@R1# set router-id 192.168.6.5
user@R1# set autonomous-system 17
```

```
[edit logical-systems B routing-options]
user@R1# set router-id 192.163.6.4
user@R1# set autonomous-system 17
```

```
[edit logical-systems C routing-options]
user@R1# set router-id 192.168.40.4
user@R1# set autonomous-system 17
```

**Results** From configuration mode, confirm your configuration by entering the **show logical-systems** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show logical-systems
A {
 interfaces {
 lt-0/1/0 {
 unit 1 {
 description to-B;
 encapsulation ethernet;
 peer-unit 2;
 family inet {
 address 10.10.10.1/30;
 }
 }
 }
 }
 lo0 {
 unit 1 {
 family inet {
 address 192.168.6.5/32;
 }
 }
 }
}
```

```
 }
 }
}
protocols {
 bgp {
 group internal-peers {
 type internal;
 local-address 192.168.6.5;
 export send-direct;
 neighbor 192.163.6.4;
 neighbor 192.168.40.4;
 }
 }
 ospf {
 area 0.0.0.0 {
 interface lo0.1 {
 passive;
 }
 interface lt-0/1/0.1;
 }
 }
}
policy-options {
 policy-statement send-direct {
 term 2 {
 from protocol direct;
 then accept;
 }
 }
}
routing-options {
 router-id 192.168.6.5;
 autonomous-system 17;
}
}
B {
 interfaces {
 lt-0/1/0 {
 unit 2 {
 description to-A;
 encapsulation ethernet;
 peer-unit 1;
 family inet {
 address 10.10.10.2/30;
 }
 }
 unit 5 {
 description to-C;
 encapsulation ethernet;
 peer-unit 6;
 family inet {
 address 10.10.10.5/30;
 }
 }
 }
 }
 lo0 {
```



```

 unit 2 {
 family inet {
 address 192.163.6.4/32;
 }
 }
 }
}
protocols {
 bgp {
 group internal-peers {
 type internal;
 local-address 192.163.6.4;
 export send-direct;
 neighbor 192.168.40.4;
 neighbor 192.168.6.5;
 }
 }
 ospf {
 area 0.0.0.0 {
 interface lo0.2 {
 passive;
 }
 interface lt-0/1/0.2;
 interface lt-0/1/0.5;
 }
 }
}
policy-options {
 policy-statement send-direct {
 term 2 {
 from protocol direct;
 then accept;
 }
 }
}
routing-options {
 router-id 192.163.6.4;
 autonomous-system 17;
}
}
C {
 interfaces {
 lt-0/1/0 {
 unit 6 {
 description to-B;
 encapsulation ethernet;
 peer-unit 5;
 family inet {
 address 10.10.10.6/30;
 }
 }
 }
 }
 lo0 {
 unit 3 {
 family inet {
 address 192.168.40.4/32;
 }
 }
 }
}

```

```
 }
 }
}
protocols {
 bgp {
 group internal-peers {
 type internal;
 local-address 192.168.40.4;
 export send-direct;
 neighbor 192.163.6.4;
 neighbor 192.168.6.5;
 }
 }
 ospf {
 area 0.0.0.0 {
 interface lo0.3 {
 passive;
 }
 interface lt-0/1/0.6;
 }
 }
}
policy-options {
 policy-statement send-direct {
 term 2 {
 from protocol direct;
 then accept;
 }
 }
}
routing-options {
 router-id 192.168.40.4;
 autonomous-system 17;
}
}
```

If you are done configuring the device, enter **commit** from configuration mode.

### **Verification**

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3372](#)
- [Verifying BGP Groups on page 3374](#)
- [Verifying BGP Summary Information on page 3374](#)
- [Verifying That BGP Routes Are Installed in the Routing Table on page 3374](#)

### **Verifying BGP Neighbors**

**Purpose** Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

**Action** From the operational mode, enter the **show bgp neighbor** command.

```

user@R1> show bgp neighbor logical-system A
Peer: 192.163.6.4+179 AS 17 Local: 192.168.6.5+58852 AS 17
 Type: Internal State: Established Flags: <Sync>
 Last State: OpenConfirm Last Event: RecvKeepAlive
 Last Error: None
 Export: [send-direct]
 Options: <Preference LocalAddress Refresh>
 Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
 Number of flaps: 0
 Peer ID: 192.163.6.4 Local ID: 192.168.6.5 Active Holdtime: 90
 Keepalive Interval: 30 Peer index: 0
 BFD: disabled, down
 NLRI for restart configured on peer: inet-unicast
 NLRI advertised by peer: inet-unicast
 NLRI for this session: inet-unicast
 Peer supports Refresh capability (2)
 Restart time configured on the peer: 120
 Stale routes from peer are kept for: 300
 Restart time requested by this peer: 120
 NLRI that peer supports restart for: inet-unicast
 NLRI that restart is negotiated for: inet-unicast
 NLRI of received end-of-rib markers: inet-unicast
 NLRI of all end-of-rib markers sent: inet-unicast
 Peer supports 4 byte AS extension (peer-as 17)
 Peer does not support Addpath
 Table inet.0 Bit: 10000
 RIB State: BGP restart is complete
 Send state: in sync
 Active prefixes: 0
 Received prefixes: 3
 Accepted prefixes: 3
 Suppressed due to damping: 0
 Advertised prefixes: 2
 Last traffic (seconds): Received 16 Sent 1 Checked 63
 Input messages: Total 15713 Updates 4 Refreshes 0 Octets 298622
 Output messages: Total 15690 Updates 2 Refreshes 0 Octets 298222
 Output Queue[0]: 0

Peer: 192.168.40.4+179 AS 17 Local: 192.168.6.5+56466 AS 17
 Type: Internal State: Established Flags: <Sync>
 Last State: OpenConfirm Last Event: RecvKeepAlive
 Last Error: None
 Export: [send-direct]
 Options: <Preference LocalAddress Refresh>
 Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
 Number of flaps: 0
 Peer ID: 192.168.40.4 Local ID: 192.168.6.5 Active Holdtime: 90
 Keepalive Interval: 30 Peer index: 1
 BFD: disabled, down
 NLRI for restart configured on peer: inet-unicast
 NLRI advertised by peer: inet-unicast
 NLRI for this session: inet-unicast
 Peer supports Refresh capability (2)
 Restart time configured on the peer: 120
 Stale routes from peer are kept for: 300
 Restart time requested by this peer: 120
 NLRI that peer supports restart for: inet-unicast
 NLRI that restart is negotiated for: inet-unicast
 NLRI of received end-of-rib markers: inet-unicast

```

```

NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 17)
Peer does not support Addpath
Table inet.0 Bit: 10000
 RIB State: BGP restart is complete
 Send state: in sync
 Active prefixes: 0
 Received prefixes: 2
 Accepted prefixes: 2
 Suppressed due to damping: 0
 Advertised prefixes: 2
Last traffic (seconds): Received 15 Sent 22 Checked 68
Input messages: Total 15688 Updates 2 Refreshes 0 Octets 298111
Output messages: Total 15688 Updates 2 Refreshes 0 Octets 298184
Output Queue[0]: 0

```

### Verifying BGP Groups

**Purpose** Verify that the BGP groups are configured correctly.

**Action** From the operational mode, enter the **show bgp group** command.

```

user@A> show bgp group logical-system A
Group Type: Internal AS: 17 Local AS: 17
Name: internal-peers Index: 0 Flags: <Export Eval>
Export: [send-direct]
Holdtime: 0
Total peers: 2 Established: 2
192.163.6.4+179
192.168.40.4+179
inet.0: 0/5/5/0

Groups: 1 Peers: 2 External: 0 Internal: 2 Down peers: 0 Flaps: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 5 0 0 0 0 0 0

```

### Verifying BGP Summary Information

**Purpose** Verify that the BGP configuration is correct.

**Action** From the operational mode, enter the **show bgp summary** command.

```

user@A> show bgp summary logical-system A
Groups: 1 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 5 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.163.6.4 17 15723 15700 0 0 4d 22:13:15
0/3/3/0 0/0/0/0
192.168.40.4 17 15698 15699 0 0 4d 22:13:11
0/2/2/0 0/0/0/0

```

### Verifying That BGP Routes Are Installed in the Routing Table

**Purpose** Verify that the export policy configuration is working.

**Action** From the operational mode, enter the **show route protocol bgp** command.

```
user@A> show route protocol bgp logical-system A
inet.0: 7 destinations, 12 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.0/30 [BGP/170] 4d 11:05:55, localpref 100, from 192.163.6.4
 AS path: I
 > to 10.10.10.2 via lt-0/1/0.1
10.10.10.4/30 [BGP/170] 4d 11:05:55, localpref 100, from 192.163.6.4
 AS path: I
 > to 10.10.10.2 via lt-0/1/0.1
 [BGP/170] 4d 11:03:10, localpref 100, from 192.168.40.4
 AS path: I
 > to 10.10.10.2 via lt-0/1/0.1
192.163.6.4/32 [BGP/170] 4d 11:05:55, localpref 100, from 192.163.6.4
 AS path: I
 > to 10.10.10.2 via lt-0/1/0.1
192.168.40.4/32 [BGP/170] 4d 11:03:10, localpref 100, from 192.168.40.4
 AS path: I
 > to 10.10.10.2 via lt-0/1/0.1
```

#### *Example: Configuring External BGP on Logical Systems with IPv6 Interfaces*

This example shows how to configure external BGP (EBGP) point-to-point peer sessions on logical systems with IPv6 interfaces.

- [Requirements on page 3375](#)
- [Overview on page 3375](#)
- [Configuration on page 3376](#)
- [Verification on page 3385](#)

#### **Requirements**

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

#### **Overview**

Junos OS supports EBGP peer sessions by means of IPv6 addresses. An IPv6 peer session can be configured when an IPv6 address is specified in the **neighbor** statement. This example uses EUI-64 to generate IPv6 addresses that are automatically applied to the interfaces. An EUI-64 address is an IPv6 address that uses the IEEE EUI-64 format for the interface identifier portion of the address (the last 64 bits).



**NOTE:** Alternatively, you can configure EBGP sessions using manually assigned 128-bit IPv6 addresses.

If you use 128-bit link-local addresses for the interfaces, you must include the **local-interface** statement. This statement is valid only for 128-bit IPv6 link-local addresses and is mandatory for configuring an IPv6 EBGP link-local peer session.


Configuring EBGP peering using link-local addresses is only applicable for directly connected interfaces. There is no support for multihop peering.

After your interfaces are up, you can use the **show interfaces terse** command to view the EUI-64-generated IPv6 addresses on the interfaces. You must use these generated addresses in the BGP **neighbor** statements. This example demonstrates the full end-to-end procedure.

In this example, Frame Relay interface encapsulation is applied to the logical tunnel (**lt**) interfaces. This is a requirement because only Frame Relay encapsulation is supported when IPv6 addresses are configured on the **lt** interfaces.

Figure 70 on page 3376 shows a network with BGP peer sessions. In the sample network, Router R1 has five logical systems configured. Device E in autonomous system (AS) 17 has BGP peer sessions to a group of peers called **external-peers**. Peers A, B, and C reside in AS 22. This example shows the step-by-step configuration on Logical System A and Logical System E.

Figure 70: Typical Network with BGP Peer Sessions



### Configuration

#### CLI Quick Configuration

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

Device A      **set logical-systems A interfaces lt-0/1/0 unit 1 description to-E**

```

set logical-systems A interfaces lt-0/1/0 unit 1 encapsulation frame-relay
set logical-systems A interfaces lt-0/1/0 unit 1 dlci 1
set logical-systems A interfaces lt-0/1/0 unit 1 peer-unit 25
set logical-systems A interfaces lt-0/1/0 unit 1 family inet6 address 2001:db8:0:1::/64
 eui-64
set logical-systems A interfaces lo0 unit 1 family inet6 address 2001:db8::1/128
set logical-systems A protocols bgp group external-peers type external
set logical-systems A protocols bgp group external-peers peer-as 17
set logical-systems A protocols bgp group external-peers neighbor
 2001:db8:0:1:2a0:a502:0:19da
set logical-systems A routing-options router-id 1.1.1.1
set logical-systems A routing-options autonomous-system 22

```

**Device B**

```

set logical-systems B interfaces lt-0/1/0 unit 6 description to-E
set logical-systems B interfaces lt-0/1/0 unit 6 encapsulation frame-relay
set logical-systems B interfaces lt-0/1/0 unit 6 dlci 6
set logical-systems B interfaces lt-0/1/0 unit 6 peer-unit 5
set logical-systems B interfaces lt-0/1/0 unit 6 family inet6 address 2001:db8:0:2::/64
 eui-64
set logical-systems B interfaces lo0 unit 2 family inet6 address 2001:db8::2/128
set logical-systems B protocols bgp group external-peers type external
set logical-systems B protocols bgp group external-peers peer-as 17
set logical-systems B protocols bgp group external-peers neighbor
 2001:db8:0:2:2a0:a502:0:5da
set logical-systems B routing-options router-id 2.2.2.2
set logical-systems B routing-options autonomous-system 22

```

**Device C**

```

set logical-systems C interfaces lt-0/1/0 unit 10 description to-E
set logical-systems C interfaces lt-0/1/0 unit 10 encapsulation frame-relay
set logical-systems C interfaces lt-0/1/0 unit 10 dlci 10
set logical-systems C interfaces lt-0/1/0 unit 10 peer-unit 9
set logical-systems C interfaces lt-0/1/0 unit 10 family inet6 address 2001:db8:0:3::/64
 eui-64
set logical-systems C interfaces lo0 unit 3 family inet6 address 2001:db8::3/128
set logical-systems C protocols bgp group external-peers type external
set logical-systems C protocols bgp group external-peers peer-as 17
set logical-systems C protocols bgp group external-peers neighbor
 2001:db8:0:3:2a0:a502:0:9da
set logical-systems C routing-options router-id 3.3.3.3
set logical-systems C routing-options autonomous-system 22

```

**Device D**

```

set logical-systems D interfaces lt-0/1/0 unit 7 description to-E
set logical-systems D interfaces lt-0/1/0 unit 7 encapsulation frame-relay
set logical-systems D interfaces lt-0/1/0 unit 7 dlci 7
set logical-systems D interfaces lt-0/1/0 unit 7 peer-unit 21
set logical-systems D interfaces lt-0/1/0 unit 7 family inet6 address 2001:db8:0:4::/64
 eui-64
set logical-systems D interfaces lo0 unit 4 family inet6 address 2001:db8::4/128
set logical-systems D protocols bgp group external-peers type external
set logical-systems D protocols bgp group external-peers peer-as 17
set logical-systems D protocols bgp group external-peers neighbor
 2001:db8:0:4:2a0:a502:0:15da
set logical-systems D routing-options router-id 4.4.4.4
set logical-systems D routing-options autonomous-system 79

```

```

Device E set logical-systems E interfaces lt-0/1/0 unit 5 description to-B
 set logical-systems E interfaces lt-0/1/0 unit 5 encapsulation frame-relay
 set logical-systems E interfaces lt-0/1/0 unit 5 dlci 6
 set logical-systems E interfaces lt-0/1/0 unit 5 peer-unit 6
 set logical-systems E interfaces lt-0/1/0 unit 5 family inet6 address 2001:db8:0:2::/64
 eui-64
 set logical-systems E interfaces lt-0/1/0 unit 9 description to-C
 set logical-systems E interfaces lt-0/1/0 unit 9 encapsulation frame-relay
 set logical-systems E interfaces lt-0/1/0 unit 9 dlci 10
 set logical-systems E interfaces lt-0/1/0 unit 9 peer-unit 10
 set logical-systems E interfaces lt-0/1/0 unit 9 family inet6 address 2001:db8:0:3::/64
 eui-64
 set logical-systems E interfaces lt-0/1/0 unit 21 description to-D
 set logical-systems E interfaces lt-0/1/0 unit 21 encapsulation frame-relay
 set logical-systems E interfaces lt-0/1/0 unit 21 dlci 7
 set logical-systems E interfaces lt-0/1/0 unit 21 peer-unit 7
 set logical-systems E interfaces lt-0/1/0 unit 21 family inet6 address 2001:db8:0:4::/64
 eui-64
 set logical-systems E interfaces lt-0/1/0 unit 25 description to-A
 set logical-systems E interfaces lt-0/1/0 unit 25 encapsulation frame-relay
 set logical-systems E interfaces lt-0/1/0 unit 25 dlci 1
 set logical-systems E interfaces lt-0/1/0 unit 25 peer-unit 1
 set logical-systems E interfaces lt-0/1/0 unit 25 family inet6 address 2001:db8:0:1::/64
 eui-64
 set logical-systems E interfaces lo0 unit 5 family inet6 address 2001:db8::5/128
 set logical-systems E protocols bgp group external-peers type external
 set logical-systems E protocols bgp group external-peers peer-as 22
 set logical-systems E protocols bgp group external-peers neighbor
 2001:db8:0:1:2a0:a502:0:1da
 set logical-systems E protocols bgp group external-peers neighbor
 2001:db8:0:2:2a0:a502:0:6da
 set logical-systems E protocols bgp group external-peers neighbor
 2001:db8:0:3:2a0:a502:0:ada
 set logical-systems E protocols bgp group external-peers neighbor
 2001:db8:0:4:2a0:a502:0:7da peer-as 79
 set logical-systems E routing-options router-id 5.5.5.5
 set logical-systems E routing-options autonomous-system 17

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure the BGP peer sessions:

1. Run the **show interfaces terse** command to verify that the physical router has a logical tunnel (lt) interface.
 

```

user@R1> show interfaces terse
Interface Admin Link Proto Local Remote
...
lt-0/1/0 up up
...

```
2. On Logical System A, configure the interface encapsulation, peer-unit number, and DLCI to reach Logical System E.



```

user@R1> set cli logical-system A
Logical system: A
[edit]
user@R1:A> edit
Entering configuration mode
[edit]
user@R1:A# edit interfaces
[edit interfaces]
user@R1:A# set lt-0/1/0 unit 1 encapsulation frame-relay
user@R1:A# set lt-0/1/0 unit 1 dlci 1
user@R1:A# set lt-0/1/0 unit 1 peer-unit 25

```

3. On Logical System A, configure the network address for the link to Peer E, and configure a loopback interface.

```

[edit interfaces]
user@R1:A# set lt-0/1/0 unit 1 description to-E
user@R1:A# set lt-0/1/0 unit 1 family inet6 address 2001:db8:0:1::/64 eui-64
user@R1:A# set lo0 unit 1 family inet6 address 2001:db8::1/128

```

4. On Logical System E, configure the interface encapsulation, peer-unit number, and DLCI to reach Logical System A.

```

user@R1> set cli logical-system E
Logical system: E
[edit]
user@R1:E> edit
Entering configuration mode
[edit]
user@R1:E# edit interfaces
[edit interfaces]
user@R1:E# set lt-0/1/0 unit 25 encapsulation frame-relay
user@R1:E# set lt-0/1/0 unit 25 dlci 1
user@R1:E# set lt-0/1/0 unit 25 peer-unit 1

```

5. On Logical System E, configure the network address for the link to Peer A, and configure a loopback interface.

```

[edit interfaces]
user@R1:E# set lt-0/1/0 unit 25 description to-A
user@R1:E# set lt-0/1/0 unit 25 family inet6 address 2001:db8:0:1::/64 eui-64
user@R1:E# set lo0 unit 5 family inet6 address 2001:db8::5/128

```

6. Run the **show interfaces terse** command to see the IPv6 addresses that are generated by EUI-64.

The 2001 addresses are used in this example in the BGP **neighbor** statements.



**NOTE:** The fe80 addresses are link-local addresses and are not used in this example.

```

user@R1:A> show interfaces terse
Interface Admin Link Proto Local Remote
Logical system: A

betsy@tp8:A> show interfaces terse

```

```

Interface Admin Link Proto Local Remote
1t-0/1/0
1t-0/1/0.1 up up inet6 2001:db8:0:1:2a0:a502:0:1da/64
 fe80::2a0:a502:0:1da/64

1o0
1o0.1 up up inet6 2001:db8::1
 fe80::2a0:a50f:fc56:1da

user@R1:E> show interfaces terse
Interface Admin Link Proto Local Remote
1t-0/1/0
1t-0/1/0.25 up up inet6 2001:db8:0:1:2a0:a502:0:19da/64
 fe80::2a0:a502:0:19da/64

1o0
1o0.5 up up inet6 2001:db8::5
 fe80::2a0:a50f:fc56:1da

```

7. Repeat the interface configuration on the other logical systems.

### Configuring the External BGP Sessions

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure the BGP peer sessions:

1. On Logical System A, create the BGP group, and add the external neighbor address.
 

```
[edit protocols bgp group external-peers]
user@R1:A# set neighbor 2001:db8:0:1:2a0:a502:0:19da
```
2. On Logical System E, create the BGP group, and add the external neighbor address.
 

```
[edit protocols bgp group external-peers]
user@R1:E# set neighbor 2001:db8:0:1:2a0:a502:0:1da
```
3. On Logical System A, specify the autonomous system (AS) number of the external AS.
 

```
[edit protocols bgp group external-peers]
user@R1:A# set peer-as 17
```
4. On Logical System E, specify the autonomous system (AS) number of the external AS.
 

```
[edit protocols bgp group external-peers]
user@R1:E# set peer-as 22
```
5. On Logical System A, set the peer type to EBGp.
 

```
[edit protocols bgp group external-peers]
user@R1:A# set type external
```
6. On Logical System E, set the peer type to EBGp.
 

```
[edit protocols bgp group external-peers]
user@R1:E# set type external
```
7. On Logical System A, set the autonomous system (AS) number and router ID.

```
[edit routing-options]
user@R1:A# set router-id 1.1.1.1
user@R1:A# set autonomous-system 22
```

8. On Logical System E, set the AS number and router ID.

```
[edit routing-options]
user@R1:E# set router-id 5.5.5.5
user@R1:E# set autonomous-system 17
```

9. Repeat these steps for Peers A, B, C, and D.

**Results** From configuration mode, confirm your configuration by entering the **show logical-systems** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@R1# show logical-systems
A {
 interfaces {
 lt-0/1/0 {
 unit 1 {
 description to-E;
 encapsulation frame-relay;
 dlci 1;
 peer-unit 25;
 family inet6 {
 address 2001:db8:0:1::/64 {
 eui-64;
 }
 }
 }
 }
 }
 lo0 {
 unit 1 {
 family inet6 {
 address 2001:db8::1/128;
 }
 }
 }
}
protocols {
 bgp {
 group external-peers {
 type external;
 peer-as 17;
 neighbor 2001:db8:0:1:2a0:a502:0:19da;
 }
 }
 routing-options {
 router-id 1.1.1.1;
 autonomous-system 22;
 }
}
B {
 interfaces {
```

```
lt-0/1/0 {
 unit 6 {
 description to-E;
 encapsulation frame-relay;
 dlci 6;
 peer-unit 5;
 family inet6 {
 address 2001:db8:0:2::/64 {
 eui-64;
 }
 }
 }
}
lo0 {
 unit 2 {
 family inet6 {
 address 2001:db8::2/128;
 }
 }
}
}
protocols {
 bgp {
 group external-peers {
 type external;
 peer-as 17;
 neighbor 2001:db8:0:2:2a0:a502:0:5da;
 }
 }
 routing-options {
 router-id 2.2.2.2;
 autonomous-system 22;
 }
}
}
C {
 interfaces {
 lt-0/1/0 {
 unit 10 {
 description to-E;
 encapsulation frame-relay;
 dlci 10;
 peer-unit 9;
 family inet6 {
 address 2001:db8:0:3::/64 {
 eui-64;
 }
 }
 }
 }
 }
 lo0 {
 unit 3 {
 family inet6 {
 address 2001:db8::3/128;
 }
 }
 }
}
```

```

}
protocols {
 bgp {
 group external-peers {
 type external;
 peer-as 17;
 neighbor 2001:db8:0:3:2a0:a502:0:9da;
 }
 }
}
routing-options {
 router-id 3.3.3.3;
 autonomous-system 22;
}
}
D {
 interfaces {
 lt-0/1/0 {
 unit 7 {
 description to-E;
 encapsulation frame-relay;
 dlci 7;
 peer-unit 21;
 family inet6 {
 address 2001:db8:0:4::/64 {
 eui-64;
 }
 }
 }
 }
 }
 lo0 {
 unit 4 {
 family inet6 {
 address 2001:db8::4/128;
 }
 }
 }
}
protocols {
 bgp {
 group external-peers {
 type external;
 peer-as 17;
 neighbor 2001:db8:0:4:2a0:a502:0:15da;
 }
 }
 routing-options {
 router-id 4.4.4.4;
 autonomous-system 79;
 }
}
}
E {
 interfaces {
 lt-0/1/0 {
 unit 5 {
 description to-B;

```

```
 encapsulation frame-relay;
 dlci 6;
 peer-unit 6;
 family inet6 {
 address 2001:db8:0:2::/64 {
 eui-64;
 }
 }
 }
 unit 9 {
 description to-C;
 encapsulation frame-relay;
 dlci 10;
 peer-unit 10;
 family inet6 {
 address 2001:db8:0:3::/64 {
 eui-64;
 }
 }
 }
 unit 21 {
 description to-D;
 encapsulation frame-relay;
 dlci 7;
 peer-unit 7;
 family inet6 {
 address 2001:db8:0:4::/64 {
 eui-64;
 }
 }
 }
 unit 25 {
 description to-A;
 encapsulation frame-relay;
 dlci 1;
 peer-unit 1;
 family inet6 {
 address 2001:db8:0:1::/64 {
 eui-64;
 }
 }
 }
}
lo0 {
 unit 5 {
 family inet6 {
 address 2001:db8::5/128;
 }
 }
}
}
protocols {
 bgp {
 group external-peers {
 type external;
 peer-as 22;
 }
 }
}
```

```

 neighbor 2001:db8:0:1:2a0:a502:0:1da;
 neighbor 2001:db8:0:2:2a0:a502:0:6da;
 neighbor 2001:db8:0:3:2a0:a502:0:ada;
 neighbor 2001:db8:0:4:2a0:a502:0:7da {
 peer-as 79;
 }
 }
}
routing-options {
 router-id 5.5.5.5;
 autonomous-system 17;
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3385](#)
- [Verifying BGP Groups on page 3388](#)
- [Verifying BGP Summary Information on page 3388](#)
- [Checking the Routing Table on page 3389](#)

### Verifying BGP Neighbors

**Purpose** Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

**Action** From operational mode, run the **show bgp neighbor** command.

```

user@R1:E> show bgp neighbor
Peer: 2001:db8:0:1:2a0:a502:0:1da+54987 AS 22 Local:
2001:db8:0:1:2a0:a502:0:19da+179 AS 17
 Type: External State: Established Flags: <Sync>
 Last State: OpenConfirm Last Event: RecvKeepAlive
 Last Error: Open Message Error
 Options: <Preference PeerAS Refresh>
 Holdtime: 90 Preference: 170
 Number of flaps: 0
 Error: 'Open Message Error' Sent: 20 Recv: 0
 Peer ID: 1.1.1.1 Local ID: 5.5.5.5 Active Holdtime: 90
 Keepalive Interval: 30 Peer index: 0
 BFD: disabled, down
 Local Interface: lt-0/1/0.25
 NLRI for restart configured on peer: inet6-unicast
 NLRI advertised by peer: inet6-unicast
 NLRI for this session: inet6-unicast
 Peer supports Refresh capability (2)
 Stale routes from peer are kept for: 300
 Peer does not support Restarter functionality
 NLRI that restart is negotiated for: inet6-unicast
 NLRI of received end-of-rib markers: inet6-unicast
 NLRI of all end-of-rib markers sent: inet6-unicast
 Peer supports 4 byte AS extension (peer-as 22)

```

```

Peer does not support Addpath
Table inet6.0 Bit: 10000
 RIB State: BGP restart is complete
 Send state: in sync
 Active prefixes: 0
 Received prefixes: 0
 Accepted prefixes: 0
 Suppressed due to damping: 0
 Advertised prefixes: 0
Last traffic (seconds): Received 7 Sent 18 Checked 81
Input messages: Total 1611 Updates 1 Refreshes 0 Octets 30660
Output messages: Total 1594 Updates 0 Refreshes 0 Octets 30356
Output Queue[0]: 0

Peer: 2001:db8:0:2:2a0:a502:0:6da+179 AS 22 Local:
2001:db8:0:2:2a0:a502:0:5da+55502 AS 17
 Type: External State: Established Flags: <Sync>
 Last State: OpenConfirm Last Event: RecvKeepAlive
 Last Error: Open Message Error
 Options: <Preference PeerAS Refresh>
 Holdtime: 90 Preference: 170
 Number of flaps: 0
 Error: 'Open Message Error' Sent: 26 Recv: 0
 Peer ID: 2.2.2.2 Local ID: 5.5.5.5 Active Holdtime: 90
 Keepalive Interval: 30 Peer index: 2
 BFD: disabled, down
 Local Interface: lt-0/1/0.5
 NLRI for restart configured on peer: inet6-unicast
 NLRI advertised by peer: inet6-unicast
 NLRI for this session: inet6-unicast
 Peer supports Refresh capability (2)
 Stale routes from peer are kept for: 300
 Peer does not support Restarter functionality
 NLRI that restart is negotiated for: inet6-unicast
 NLRI of received end-of-rib markers: inet6-unicast
 NLRI of all end-of-rib markers sent: inet6-unicast
 Peer supports 4 byte AS extension (peer-as 22)
 Peer does not support Addpath
Table inet6.0 Bit: 10000
 RIB State: BGP restart is complete
 Send state: in sync
 Active prefixes: 0
 Received prefixes: 0
 Accepted prefixes: 0
 Suppressed due to damping: 0
 Advertised prefixes: 0
Last traffic (seconds): Received 15 Sent 8 Checked 8
Input messages: Total 1610 Updates 1 Refreshes 0 Octets 30601
Output messages: Total 1645 Updates 0 Refreshes 0 Octets 32417
Output Queue[0]: 0

Peer: 2001:db8:0:3:2a0:a502:0:ada+55983 AS 22 Local:
2001:db8:0:3:2a0:a502:0:9da+179 AS 17
 Type: External State: Established Flags: <Sync>
 Last State: OpenConfirm Last Event: RecvKeepAlive
 Last Error: None
 Options: <Preference PeerAS Refresh>
 Holdtime: 90 Preference: 170
 Number of flaps: 0
 Peer ID: 3.3.3.3 Local ID: 5.5.5.5 Active Holdtime: 90
 Keepalive Interval: 30 Peer index: 3

```



```

BFD: disabled, down
Local Interface: lt-0/1/0.9
NLRI for restart configured on peer: inet6-unicast
NLRI advertised by peer: inet6-unicast
NLRI for this session: inet6-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet6-unicast
NLRI of received end-of-rib markers: inet6-unicast
NLRI of all end-of-rib markers sent: inet6-unicast
Peer supports 4 byte AS extension (peer-as 22)
Peer does not support Addpath
Table inet6.0 Bit: 10000
 RIB State: BGP restart is complete
 Send state: in sync
 Active prefixes: 0
 Received prefixes: 0
 Accepted prefixes: 0
 Suppressed due to damping: 0
 Advertised prefixes: 0
Last traffic (seconds): Received 21 Sent 21 Checked 67
Input messages: Total 1610 Updates 1 Refreshes 0 Octets 30641
Output messages: Total 1587 Updates 0 Refreshes 0 Octets 30223
Output Queue[0]: 0

Peer: 2001:db8:0:4:2a0:a502:0:7da+49255 AS 79 Local:
2001:db8:0:4:2a0:a502:0:15da+179 AS 17
 Type: External State: Established Flags: <Sync>
 Last State: OpenConfirm Last Event: RecvKeepAlive
 Last Error: None
 Options: <Preference PeerAS Refresh>
 Holdtime: 90 Preference: 170
 Number of flaps: 0
 Peer ID: 4.4.4.4 Local ID: 5.5.5.5 Active Holdtime: 90
 Keepalive Interval: 30 Peer index: 1
 BFD: disabled, down
 Local Interface: lt-0/1/0.21
 NLRI for restart configured on peer: inet6-unicast
 NLRI advertised by peer: inet6-unicast
 NLRI for this session: inet6-unicast
 Peer supports Refresh capability (2)
 Stale routes from peer are kept for: 300
 Peer does not support Restarter functionality
 NLRI that restart is negotiated for: inet6-unicast
 NLRI of received end-of-rib markers: inet6-unicast
 NLRI of all end-of-rib markers sent: inet6-unicast
 Peer supports 4 byte AS extension (peer-as 79)
 Peer does not support Addpath
Table inet6.0 Bit: 10000
 RIB State: BGP restart is complete
 Send state: in sync
 Active prefixes: 0
 Received prefixes: 0
 Accepted prefixes: 0
 Suppressed due to damping: 0
 Advertised prefixes: 0
Last traffic (seconds): Received 6 Sent 17 Checked 25
Input messages: Total 1615 Updates 1 Refreshes 0 Octets 30736
Output messages: Total 1593 Updates 0 Refreshes 0 Octets 30337
Output Queue[0]: 0

```

**Meaning** IPv6 unicast network layer reachability information (NLRI) is being exchanged between the neighbors.

#### *Verifying BGP Groups*

**Purpose** Verify that the BGP groups are configured correctly.

**Action** From operational mode, run the **show bgp group** command.

```
user@R1:~> show bgp group
Group Type: External Local AS: 17
 Name: external-peers Index: 0 Flags: <>
 Holdtime: 0
 Total peers: 4 Established: 4
 2001:db8:0:1:2a0:a502:0:1da+54987
 2001:db8:0:2:2a0:a502:0:6da+179
 2001:db8:0:3:2a0:a502:0:ada+55983
 2001:db8:0:4:2a0:a502:0:7da+49255
 inet6.0: 0/0/0/0

Groups: 1 Peers: 4 External: 4 Internal: 0 Down peers: 0 Flaps: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet6.0 0 0 0 0 0 0 0
inet6.2 0 0 0 0 0 0 0
```

**Meaning** The group type is external, and the group has four peers.

#### *Verifying BGP Summary Information*

**Purpose** Verify that the BGP that the peer relationships are established.

**Action** From operational mode, run the **show bgp summary** command.

```
user@R1:~> show bgp summary
Groups: 1 Peers: 4 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet6.0 0 0 0 0 0 0 0
inet6.2 0 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
2001:db8:0:1:2a0:a502:0:1da 22 1617 1600 0 0
 12:07:00 Establ
 inet6.0: 0/0/0/0
2001:db8:0:2:2a0:a502:0:6da 22 1616 1651 0 0
 12:06:56 Establ
 inet6.0: 0/0/0/0
2001:db8:0:3:2a0:a502:0:ada 22 1617 1594 0 0
 12:04:32 Establ
 inet6.0: 0/0/0/0
2001:db8:0:4:2a0:a502:0:7da 79 1621 1599 0 0
 12:07:00 Establ
 inet6.0: 0/0/0/0
```

**Meaning** The Down peers: 0 output shows that the BGP peers are in the established state.

### Checking the Routing Table

**Purpose** Verify that the inet6.0 routing table is populated with local and direct routes.

**Action** From operational mode, run the **show route** command.

```
user@R1:E> show route
inet6.0: 15 destinations, 18 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::5/128 *[Direct/0] 12:41:18
 > via lo0.5
2001:db8:0:1::/64 *[Direct/0] 14:40:01
 > via lt-0/1/0.25
2001:db8:0:1:2a0:a502:0:19da/128
 *[Local/0] 14:40:01
 Local via lt-0/1/0.25
2001:db8:0:2::/64 *[Direct/0] 14:40:02
 > via lt-0/1/0.5
2001:db8:0:2:2a0:a502:0:5da/128
 *[Local/0] 14:40:02
 Local via lt-0/1/0.5
2001:db8:0:3::/64 *[Direct/0] 14:40:02
 > via lt-0/1/0.9
2001:db8:0:3:2a0:a502:0:9da/128
 *[Local/0] 14:40:02
 Local via lt-0/1/0.9
2001:db8:0:4::/64 *[Direct/0] 14:40:01
 > via lt-0/1/0.21
2001:db8:0:4:2a0:a502:0:15da/128
 *[Local/0] 14:40:01
 Local via lt-0/1/0.21
fe80::/64 *[Direct/0] 14:40:02
 > via lt-0/1/0.5
 [Direct/0] 14:40:02
 > via lt-0/1/0.9
 [Direct/0] 14:40:01
 > via lt-0/1/0.21
 [Direct/0] 14:40:01
 > via lt-0/1/0.25
fe80::2a0:a502:0:5da/128
 *[Local/0] 14:40:02
 Local via lt-0/1/0.5
fe80::2a0:a502:0:9da/128
 *[Local/0] 14:40:02
 Local via lt-0/1/0.9
fe80::2a0:a502:0:15da/128
 *[Local/0] 14:40:01
 Local via lt-0/1/0.21
fe80::2a0:a502:0:19da/128
 *[Local/0] 14:40:01
 Local via lt-0/1/0.25
fe80::2a0:a50f:fc56:1da/128
 *[Direct/0] 12:41:18
 > via lo0.5
```

**Meaning** The inet6.0 routing table contains local and direct routes. To populate the routing table with other types of routes, you must configure routing policies.

### **Example: Configuring BFD on Internal BGP Peer Sessions**

This example shows how to configure internal BGP (IBGP) peer sessions with the Bidirectional Forwarding Detection (BFD) protocol to detect failures in a network.

- [Requirements on page 3390](#)
- [Overview on page 3390](#)
- [Configuration on page 3391](#)
- [Verification on page 3395](#)

#### **Requirements**

No special configuration beyond device initialization is required before you configure this example.

#### **Overview**

The minimum configuration to enable BFD on IBGP sessions is to include the [bfd-liveness-detection minimum-interval](#) statement in the BGP configuration of all neighbors participating in the BFD session. The **minimum-interval** statement specifies the minimum transmit and receive intervals for failure detection. Specifically, this value represents the minimum interval after which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value from 1 through 255,000 milliseconds.

Optionally, you can specify the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements. For information about these and other optional BFD configuration statements, see [bfd-liveness-detection](#).



.....  
**NOTE:** BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD less than 100 ms for Routing Engine-based sessions and less than 10 ms for distributed BFD sessions can cause undesired BFD flapping.


Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
  - For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
  - For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.
- .....

BFD is supported on the default routing instance (the main router), routing instances, and logical systems. This example shows BFD on logical systems.

Figure 71 on page 3391 shows a typical network with internal peer sessions.

**Figure 71: Typical Network with IBGP Sessions**



#### Configuration

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

**Device A**

```

set logical-systems A interfaces lt-1/2/0 unit 1 description to-B
set logical-systems A interfaces lt-1/2/0 unit 1 encapsulation ethernet
set logical-systems A interfaces lt-1/2/0 unit 1 peer-unit 2
set logical-systems A interfaces lt-1/2/0 unit 1 family inet address 10.10.10.1/30
set logical-systems A interfaces lo0 unit 1 family inet address 192.168.6.5/32
set logical-systems A protocols bgp group internal-peers type internal
set logical-systems A protocols bgp group internal-peers traceoptions file bgp-bfd
set logical-systems A protocols bgp group internal-peers traceoptions flag bfd detail
set logical-systems A protocols bgp group internal-peers local-address 192.168.6.5
set logical-systems A protocols bgp group internal-peers export send-direct
set logical-systems A protocols bgp group internal-peers bfd-liveness-detection
 minimum-interval 1000
set logical-systems A protocols bgp group internal-peers neighbor 192.163.6.4
set logical-systems A protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems A protocols ospf area 0.0.0.0 interface lo0.1 passive
set logical-systems A protocols ospf area 0.0.0.0 interface lt-1/2/0.1
set logical-systems A policy-options policy-statement send-direct term 2 from protocol
 direct
set logical-systems A policy-options policy-statement send-direct term 2 then accept
set logical-systems A routing-options router-id 192.168.6.5
set logical-systems A routing-options autonomous-system 17

```

**Device B**

```

set logical-systems B interfaces lt-1/2/0 unit 2 description to-A
set logical-systems B interfaces lt-1/2/0 unit 2 encapsulation ethernet
set logical-systems B interfaces lt-1/2/0 unit 2 peer-unit 1
set logical-systems B interfaces lt-1/2/0 unit 2 family inet address 10.10.10.2/30
set logical-systems B interfaces lt-1/2/0 unit 5 description to-C

```

```

set logical-systems B interfaces lt-1/2/0 unit 5 encapsulation ethernet
set logical-systems B interfaces lt-1/2/0 unit 5 peer-unit 6
set logical-systems B interfaces lt-1/2/0 unit 5 family inet address 10.10.10.5/30
set logical-systems B interfaces lo0 unit 2 family inet address 192.163.6.4/32
set logical-systems B protocols bgp group internal-peers type internal
set logical-systems B protocols bgp group internal-peers local-address 192.163.6.4
set logical-systems B protocols bgp group internal-peers export send-direct
set logical-systems B protocols bgp group internal-peers bfd-liveness-detection
 minimum-interval 1000
set logical-systems B protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems B protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems B protocols ospf area 0.0.0.0 interface lo0.2 passive
set logical-systems B protocols ospf area 0.0.0.0 interface lt-1/2/0.2
set logical-systems B protocols ospf area 0.0.0.0 interface lt-1/2/0.5
set logical-systems B policy-options policy-statement send-direct term 2 from protocol
 direct
set logical-systems B policy-options policy-statement send-direct term 2 then accept
set logical-systems B routing-options router-id 192.163.6.4
set logical-systems B routing-options autonomous-system 17

```

|          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Device C | <pre> set logical-systems C interfaces lt-1/2/0 unit 6 description to-B set logical-systems C interfaces lt-1/2/0 unit 6 encapsulation ethernet set logical-systems C interfaces lt-1/2/0 unit 6 peer-unit 5 set logical-systems C interfaces lt-1/2/0 unit 6 family inet address 10.10.10.6/30 set logical-systems C interfaces lo0 unit 3 family inet address 192.168.40.4/32 set logical-systems C protocols bgp group internal-peers type internal set logical-systems C protocols bgp group internal-peers local-address 192.168.40.4 set logical-systems C protocols bgp group internal-peers export send-direct set logical-systems C protocols bgp group internal-peers bfd-liveness-detection     minimum-interval 1000 set logical-systems C protocols bgp group internal-peers neighbor 192.163.6.4 set logical-systems C protocols bgp group internal-peers neighbor 192.168.6.5 set logical-systems C protocols ospf area 0.0.0.0 interface lo0.3 passive set logical-systems C protocols ospf area 0.0.0.0 interface lt-1/2/0.6 set logical-systems C policy-options policy-statement send-direct term 2 from protocol     direct set logical-systems C policy-options policy-statement send-direct term 2 then accept set logical-systems C routing-options router-id 192.168.40.4 set logical-systems C routing-options autonomous-system 17 </pre> |
|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

### Configuring Device A

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure Device A:

1. Set the CLI to Logical System A.  

```
user@host> set cli logical-system A
```
2. Configure the interfaces.  

```
[edit interfaces lt-1/2/0 unit 1]
user@host:A# set description to-B
user@host:A# set encapsulation ethernet
```

```
user@host:A# set peer-unit 2
user@host:A# set family inet address 10.10.10.1/30
```

```
[edit interfaces lo0 unit 1]
user@host:A# set family inet address 192.168.6.5/32
```

3. Configure BGP.

The **neighbor** statements are included for both Device B and Device C, even though Device A is not directly connected to Device C.

```
[edit protocols bgp group internal-peers]
user@host:A# set type internal
user@host:A# set local-address 192.168.6.5
user@host:A# set export send-direct
user@host:A# set neighbor 192.163.6.4
user@host:A# set neighbor 192.168.40.4
```

4. Configure BFD.

```
[edit protocols bgp group internal-peers]
user@host:A# set bfd-liveness-detection minimum-interval 1000
```

You must configure the same minimum interval on the connecting peer.

5. (Optional) Configure BFD tracing.

```
[edit protocols bgp group internal-peers]
user@host:A# set traceoptions file bgp-bfd
user@host:A# set traceoptions flag bfd detail
```

6. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@host:A# set interface lo0.1 passive
user@host:A# set interface lt-1/2/0.1
```

7. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 2]
user@host:A# set from protocol direct
user@host:A# set then accept
```

8. Configure the router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@host:A# set router-id 192.168.6.5
user@host:A# set autonomous-system 17
```

9. If you are done configuring the device, enter **commit** from configuration mode. Repeat these steps to configure Device B and Device C.

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host:A# show interfaces
lt-1/2/0 {
 unit 1 {
 description to-B;
 encapsulation ethernet;
 peer-unit 2;
 family inet {
 address 10.10.10.1/30;
 }
 }
}
lo0 {
 unit 1 {
 family inet {
 address 192.168.6.5/32;
 }
 }
}

user@host:A# show policy-options
policy-statement send-direct {
 term 2 {
 from protocol direct;
 then accept;
 }
}

user@host:A# show protocols
bgp {
 group internal-peers {
 type internal;
 traceoptions {
 file bgp-bfd;
 flag bfd detail;
 }
 local-address 192.168.6.5;
 export send-direct;
 bfd-liveness-detection {
 minimum-interval 1000;
 }
 neighbor 192.163.6.4;
 neighbor 192.168.40.4;
 }
}
ospf {
 area 0.0.0.0 {
 interface lo0.1 {
 passive;
 }
 interface lt-1/2/0.1;
 }
}

user@host:A# show routing-options
router-id 192.168.6.5;
autonomous-system 17;
```



**Verification**

Confirm that the configuration is working properly.

- [Verifying That BFD Is Enabled on page 3395](#)
- [Verifying That BFD Sessions Are Up on page 3395](#)
- [Viewing Detailed BFD Events on page 3396](#)
- [Viewing Detailed BFD Events After Deactivating and Reactivating a Loopback Interface on page 3397](#)

**Verifying That BFD Is Enabled**

**Purpose** Verify that BFD is enabled between the IBGP peers.

**Action** From operational mode, enter the **show bgp neighbor** command. You can use the **| match bfd** filter to narrow the output.

```
user@host:A> show bgp neighbor | match bfd
Options: <BfdEnabled>
 BFD: enabled, up
 Trace file: /var/log/A/bgp-bfd size 131072 files 10
Options: <BfdEnabled>
 BFD: enabled, up
 Trace file: /var/log/A/bgp-bfd size 131072 files 10
```

**Meaning** The output shows that Logical System A has two neighbors with BFD enabled. When BFD is not enabled, the output displays **BFD: disabled, down**, and the **<BfdEnabled>** option is absent. If BFD is enabled and the session is down, the output displays **BFD: enabled, down**. The output also shows that BFD-related events are being written to a log file because trace operations are configured.

**Verifying That BFD Sessions Are Up**

**Purpose** Verify that the BFD sessions are up, and view details about the BFD sessions.

**Action** From operational mode, enter the **show bfd session extensive** command.

```
user@host:A> show bfd session extensive
```

| Address     | State | Interface | Detect Time | Transmit Interval | Multiplier |
|-------------|-------|-----------|-------------|-------------------|------------|
| 192.163.6.4 | Up    |           | 3.000       | 1.000             | 3          |

```
Client BGP, TX interval 1.000, RX interval 1.000
Session up time 00:54:40
Local diagnostic None, remote diagnostic None
Remote state Up, version 1
Logical system 12, routing table index 25
Min async interval 1.000, min slow interval 1.000
Adaptive async TX interval 1.000, RX interval 1.000
Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
Local discriminator 10, remote discriminator 9
Echo mode disabled/inactive
Multi-hop route table 25, local-address 192.168.6.5
```

|  | Detect | Transmit |
|--|--------|----------|
|--|--------|----------|

```

Address State Interface Time Interval Multiplier
192.168.40.4 Up Up 3.000 1.000 3
Client BGP, TX interval 1.000, RX interval 1.000
Session up time 00:48:03
Local diagnostic None, remote diagnostic None
Remote state Up, version 1
Logical system 12, routing table index 25
Min async interval 1.000, min slow interval 1.000
Adaptive async TX interval 1.000, RX interval 1.000
Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
Local discriminator 14, remote discriminator 13
Echo mode disabled/inactive
Multi-hop route table 25, local-address 192.168.6.5

2 sessions, 2 clients
Cumulative transmit rate 2.0 pps, cumulative receive rate 2.0 pps

```

**Meaning** The TX interval 1.000, RX interval 1.000 output represents the setting configured with the **minimum-interval** statement. All of the other output represents the default settings for BFD. To modify the default settings, include the optional statements under the **bfd-liveness-detection** statement.

#### *Viewing Detailed BFD Events*

**Purpose** View the contents of the BFD trace file to assist in troubleshooting, if needed.

**Action** From operational mode, enter the **file show /var/log/A/bgp-bfd** command.

```

user@host:A> file show /var/log/A/bgp-bfd
Aug 15 17:07:25 trace_on: Tracing to "/var/log/A/bgp-bfd" started
Aug 15 17:07:26.492190 bgp_peer_init: BGP peer 192.163.6.4 (Internal AS 17) local
address 192.168.6.5 not found. Leaving peer idled
Aug 15 17:07:26.493176 bgp_peer_init: BGP peer 192.168.40.4 (Internal AS 17) local
address 192.168.6.5 not found. Leaving peer idled
Aug 15 17:07:32.597979 task_connect: task BGP_17.192.163.6.4+179 addr
192.163.6.4+179: No route to host
Aug 15 17:07:32.599623 bgp_connect_start: connect 192.163.6.4 (Internal AS 17):
No route to host
Aug 15 17:07:36.869394 task_connect: task BGP_17.192.168.40.4+179 addr
192.168.40.4+179: No route to host
Aug 15 17:07:36.870624 bgp_connect_start: connect 192.168.40.4 (Internal AS 17):
No route to host
Aug 15 17:08:04.599220 task_connect: task BGP_17.192.163.6.4+179 addr
192.163.6.4+179: No route to host
Aug 15 17:08:04.601135 bgp_connect_start: connect 192.163.6.4 (Internal AS 17):
No route to host
Aug 15 17:08:08.869717 task_connect: task BGP_17.192.168.40.4+179 addr
192.168.40.4+179: No route to host
Aug 15 17:08:08.869934 bgp_connect_start: connect 192.168.40.4 (Internal AS 17):
No route to host
Aug 15 17:08:36.603544 advertising receiving-speaker only capability to neighbor
192.163.6.4 (Internal AS 17)
Aug 15 17:08:36.606726 bgp_read_message: 192.163.6.4 (Internal AS 17): 0 bytes
buffered
Aug 15 17:08:36.609119 Initiated BFD session to peer 192.163.6.4 (Internal AS
17): address=192.163.6.4 ifindex=0 ifname=(none) txivl=1000 rxivl=1000 mult=3
ver=255
Aug 15 17:08:36.734033 advertising receiving-speaker only capability to neighbor

```

```

192.168.40.4 (Internal AS 17)
Aug 15 17:08:36.738436 Initiated BFD session to peer 192.168.40.4 (Internal AS
17): address=192.168.40.4 ifindex=0 ifname=(none) txivl=1000 rxivl=1000 mult=3
ver=255
Aug 15 17:08:40.537552 BFD session to peer 192.163.6.4 (Internal AS 17) up
Aug 15 17:08:40.694410 BFD session to peer 192.168.40.4 (Internal AS 17) up

```

**Meaning** Before the routes are established, the **No route to host** message appears in the output. After the routes are established, the last two lines show that both BFD sessions come up.

### *Viewing Detailed BFD Events After Deactivating and Reactivating a Loopback Interface*

**Purpose** Check to see what happens after bringing down a router or switch and then bringing it back up. To simulate bringing down a router or switch, deactivate the loopback interface on Logical System B.

**Action** 1. From configuration mode, enter the **deactivate logical-systems B interfaces lo0 unit 2 family inet** command.

```

user@host:A# deactivate logical-systems B interfaces lo0 unit 2 family inet
user@host:A# commit

```

2. From operational mode, enter the **file show /var/log/A/bgp-bfd** command.

```

user@host:A> file show /var/log/A/bgp-bfd
...
Aug 15 17:20:55.995648 bgp_read_v4_message:9747: NOTIFICATION received from
192.163.6.4 (Internal AS 17): code 6 (Cease) subcode 6 (Other Configuration
Change)
Aug 15 17:20:56.004508 Terminated BFD session to peer 192.163.6.4 (Internal
AS 17)
Aug 15 17:21:28.007755 task_connect: task BGP_17.192.163.6.4+179 addr
192.163.6.4+179: No route to host
Aug 15 17:21:28.008597 bgp_connect_start: connect 192.163.6.4 (Internal AS
17): No route to host

```

3. From configuration mode, enter the **activate logical-systems B interfaces lo0 unit 2 family inet** command.

```

user@host:A# activate logical-systems B interfaces lo0 unit 2 family inet
user@host:A# commit

```

4. From operational mode, enter the **file show /var/log/A/bgp-bfd** command.

```

user@host:A> file show /var/log/A/bgp-bfd
...
Aug 15 17:25:53.623743 advertising receiving-speaker only capability to neighbor
192.163.6.4 (Internal AS 17)
Aug 15 17:25:53.631314 Initiated BFD session to peer 192.163.6.4 (Internal AS
17): address=192.163.6.4 ifindex=0 ifname=(none) txivl=1000 rxivl=1000 mult=3
ver=255
Aug 15 17:25:57.570932 BFD session to peer 192.163.6.4 (Internal AS 17) up

```

### **Example: Configuring EBGMP Multihop Sessions on Logical Systems**

This example shows how to configure an external BGP (EBGP) peer that is more than one hop away from the local router. This type of session is called a *multihop* EBGMP session.

- [Requirements on page 3398](#)
- [Overview on page 3398](#)
- [Configuration on page 3399](#)
- [Verification on page 3406](#)

#### **Requirements**

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

#### **Overview**

When EBGMP peers are not directly connected to each other, they must cross one or more non-BGP routing devices to reach each other. Configuring multihop EBGMP enables the peers to pass through the other routing devices to form peer relationships and exchange update messages. This type of configuration is typically used when a Juniper Networks routing device needs to run EBGMP with a third-party routing device that does not allow direct connection of the two EBGMP peers. EBGMP multihop enables a neighbor connection between two EBGMP peers that do not have a direct connection.

The configuration to enable multihop EBGMP sessions requires connectivity between the two EBGMP peers. This example uses static routes to provide connectivity between the devices.


For directly connected EBGMP sessions, physical addresses are typically used in the **neighbor** statements. For multihop EBGMP, you must use loopback interface addresses, and specify the loopback interface address of the indirectly connected peer. In the use of loopback interfaces addresses, EBGMP multihop is similar to internal BGP (IBGP).

Finally, you must add the **multihop** statement. Optionally, you can set a maximum time-to-live (TTL) value with the **ttl** statement. The TTL is carried in the IP header of BGP packets. If you do not specify a TTL value, the system's default maximum TTL value is used. The default TTL value is 64 for multihop EBGMP sessions. Another option is to retain the BGP next-hop value for route advertisements by including the **no-nexthop-change** statement.

[Figure 72 on page 3399](#) shows a typical EBGMP multihop network.

Device C and Device E have an established EBGMP session. Device D is not a BGP-enabled device. All of the devices have connectivity via static routes.

Figure 72: Typical Network with EBGP Multihop Sessions

**Configuration**

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

**Device C**

```

set logical-systems C interfaces lt-1/2/0 unit 9 description to-D
set logical-systems C interfaces lt-1/2/0 unit 9 encapsulation ethernet
set logical-systems C interfaces lt-1/2/0 unit 9 peer-unit 10
set logical-systems C interfaces lt-1/2/0 unit 9 family inet address 10.10.10.9/30
set logical-systems C interfaces lo0 unit 3 family inet address 192.168.40.4/32
set logical-systems C protocols bgp group external-peers type external
set logical-systems C protocols bgp group external-peers multihop ttl 2
set logical-systems C protocols bgp group external-peers local-address 192.168.40.4
set logical-systems C protocols bgp group external-peers export send-static
set logical-systems C protocols bgp group external-peers peer-as 18
set logical-systems C protocols bgp group external-peers neighbor 192.168.6.7
set logical-systems C policy-options policy-statement send-static term 1 from protocol static
set logical-systems C policy-options policy-statement send-static term 1 then accept
set logical-systems C routing-options static route 10.10.10.14/32 next-hop 10.10.10.10
set logical-systems C routing-options static route 192.168.6.7/32 next-hop 10.10.10.10
set logical-systems C routing-options router-id 192.168.40.4
set logical-systems C routing-options autonomous-system 17

```

**Device D**

```

set logical-systems D interfaces lt-1/2/0 unit 10 description to-C
set logical-systems D interfaces lt-1/2/0 unit 10 encapsulation ethernet
set logical-systems D interfaces lt-1/2/0 unit 10 peer-unit 9
set logical-systems D interfaces lt-1/2/0 unit 10 family inet address 10.10.10.10/30
set logical-systems D interfaces lt-1/2/0 unit 13 description to-E
set logical-systems D interfaces lt-1/2/0 unit 13 encapsulation ethernet
set logical-systems D interfaces lt-1/2/0 unit 13 peer-unit 14
set logical-systems D interfaces lt-1/2/0 unit 13 family inet address 10.10.10.13/30
set logical-systems D interfaces lo0 unit 4 family inet address 192.168.6.6/32
set logical-systems D routing-options static route 192.168.40.4/32 next-hop 10.10.10.9
set logical-systems D routing-options static route 192.168.6.7/32 next-hop 10.10.10.14
set logical-systems D routing-options router-id 192.168.6.6

```

**Device E**

```

set logical-systems E interfaces lt-1/2/0 unit 14 description to-D
set logical-systems E interfaces lt-1/2/0 unit 14 encapsulation ethernet
set logical-systems E interfaces lt-1/2/0 unit 14 peer-unit 13

```

```
set logical-systems E interfaces lt-1/2/0 unit 14 family inet address 10.10.10.14/30
set logical-systems E interfaces lo0 unit 5 family inet address 192.168.6.7/32
set logical-systems E protocols bgp group external-peers multihop ttl 2
set logical-systems E protocols bgp group external-peers local-address 192.168.6.7
set logical-systems E protocols bgp group external-peers export send-static
set logical-systems E protocols bgp group external-peers peer-as 17
set logical-systems E protocols bgp group external-peers neighbor 192.168.40.4
set logical-systems E policy-options policy-statement send-static term 1 from protocol
static
set logical-systems E policy-options policy-statement send-static term 1 then accept
set logical-systems E routing-options static route 10.10.10.8/30 next-hop 10.10.10.13
set logical-systems E routing-options static route 192.168.40.4/32 next-hop 10.10.10.13
set logical-systems E routing-options router-id 192.168.6.7
set logical-systems E routing-options autonomous-system 18
```

### *Device C*

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure Device C:

1. Set the CLI to Logical System C.  

```
user@host> set cli logical-system C
```
2. Configure the interface to the directly-connected device (to-D), and configure the loopback interface.  

```
[edit interfaces lt-1/2/0 unit 9]
user@host:C# set description to-D
user@host:C# set encapsulation ethernet
user@host:C# set peer-unit 10
user@host:C# set family inet address 10.10.10.9/30
```

```
[edit interfaces lo0 unit 3]
user@host:C# set family inet address 192.168.40.4/32
```

3. Configure an EBGp session with Logical System E.

The **neighbor** statement points to the loopback interface on Logical System E.

```
[edit protocols bgp group external-peers]
user@host:C# set type external
user@host:C# set local-address 192.168.40.4
user@host:C# set export send-static
user@host:C# set peer-as 18
user@host:C# set neighbor 192.168.6.7
```

4. Configure the multihop statement to enable Logical System C and Logical System E to become EBGp peers.

Because the peers are two hops away from each other, the example uses the **ttl 2** statement.

```
[edit protocols bgp group external-peers]
user@host:C# set multihop ttl 2
```

5. Configure connectivity to Logical System E, using static routes.

You must configure a route to both the loopback interface address and to the address on the physical interface.

```
[edit logical-systems C routing-options]
user@host:C# set static route 10.10.10.14/32 next-hop 10.10.10.10
user@host:C# set static route 192.168.6.7/32 next-hop 10.10.10.10
```

6. Configure the local router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@host:C# set router-id 192.168.40.4
user@host:C# set autonomous-system 17
```

7. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-static term 1]
user@host:C# set from protocol static
user@host:C# set then accept
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host:C# show interfaces
lt-1/2/0 {
 unit 9 {
 description to-D;
 encapsulation ethernet;
 peer-unit 10;
 family inet {
 address 10.10.10.9/30;
 }
 }
}
lo0 {
 unit 3 {
 family inet {
 address 192.168.40.4/32;
 }
 }
}

user@host:C# show protocols
bgp {
 group external-peers {
 type external;
 multihop {
 ttl 2;
 }
 local-address 192.168.40.4;
 export send-static;
```

```
 peer-as 18;
 neighbor 192.168.6.7;
 }
}

user@host:C# show policy-options
policy-statement send-static {
 term 1 {
 from protocol static;
 then accept;
 }
}

user@host:C# show routing-options
static {
 route 10.10.10.14/32 next-hop 10.10.10.10;
 route 192.168.6.7/32 next-hop 10.10.10.10;
}
router-id 192.168.40.4;
autonomous-system 17;
```

If you are done configuring the device, enter **commit** from configuration mode. Repeat these steps for all BFD sessions in the topology.

#### **Device D**

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure Device D:

1. Set the CLI to Logical System D.  

```
user@host> set cli logical-system D
```
2. Configure the interfaces to the directly-connected devices, and configure a loopback interface.  

```
[edit interfaces lt-1/2/0 unit 10]
user@host:D# set description to-C
user@host:D# set encapsulation ethernet
user@host:D# set peer-unit 9
user@host:D# set family inet address 10.10.10.30

[edit interfaces lt-1/2/0 unit 13]
user@host:D# set description to-E
user@host:D# set encapsulation ethernet
user@host:D# set peer-unit 14
user@host:D# set family inet address 10.10.10.13/30

[edit interfaces lo0 unit 4]
user@host:D# set family inet address 192.168.6.6/32
```
3. Configure connectivity to the other devices using a static routes to the loopback interface addresses.



On Logical System D, you do not need static routes to the physical addresses because Logical System D is directly connected to Logical System C and Logical System E.

```
[edit routing-options]
user@host:D# set static route 192.168.40.4/32 next-hop 10.10.10.9
user@host:D# set static route 192.168.6.7/32 next-hop 10.10.10.14
```

4. Configure the local router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@host:D# set router-id 192.168.6.6
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host:D# show interfaces
lt-1/2/0 {
 unit 10 {
 description to-C;
 encapsulation ethernet;
 peer-unit 9;
 family inet {
 address 10.10.10.10/30;
 }
 }
 unit 13 {
 description to-E;
 encapsulation ethernet;
 peer-unit 14;
 family inet {
 address 10.10.10.13/30;
 }
 }
}
lo0 {
 unit 4 {
 family inet {
 address 192.168.6.6/32;
 }
 }
}

user@host:D# show protocols

user@host:D# show routing-options
static {
 route 192.168.40.4/32 next-hop 10.10.10.9;
 route 192.168.6.7/32 next-hop 10.10.10.14;
}
router-id 192.168.6.6;
```

If you are done configuring the device, enter **commit** from configuration mode. Repeat these steps for all BFD sessions in the topology.

### Device E

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure Device E:

1. Set the CLI to Logical System E.

```
user@host> set cli logical-system E
```

2. Configure the interface to the directly-connected device (to-D), and configure the loopback interface.

```
[edit interfaces lt-1/2/0 unit 14]
user@host:E# set description to-D
user@host:E# set encapsulation ethernet
user@host:E# set peer-unit 13
user@host:E# set family inet address 10.10.10.14/30
```

```
[edit interfaces lo0 unit 5]
user@host:E# set family inet address 192.168.6.7/32
```

3. Configure an EBGp session with Logical System E.

The **neighbor** statement points to the loopback interface on Logical System C.

```
[edit protocols bgp group external-peers]
user@host:E# set local-address 192.168.6.7
user@host:E# set export send-static
user@host:E# set peer-as 17
user@host:E# set neighbor 192.168.40.4
```

4. Configure the **multihop** statement to enable Logical System C and Logical System E to become EBGp peers.

Because the peers are two hops away from each other, the example uses the **ttl 2** statement.

```
[edit protocols bgp group external-peers]
user@host:E# set multihop ttl 2
```

5. Configure connectivity to Logical System E, using static routes.

You must configure a route to both the loopback interface address and to the address on the physical interface.

```
[edit routing-options]
user@host:E# set static route 10.10.10.8/30 next-hop 10.10.10.13
user@host:E# set static route 192.168.40.4/32 next-hop 10.10.10.13
```

6. Configure the local router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@host:E# set router-id 192.168.6.7
user@host:E# set autonomous-system 18
```

7. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-static term 1]
user@host:E# set from protocol static
user@host:E# set send-static then accept
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host:E# show interfaces
lt-1/2/0 {
 unit 14 {
 description to-D;
 encapsulation ethernet;
 peer-unit 13;
 family inet {
 address 10.10.10.14/30;
 }
 }
}
lo0 {
 unit 5 {
 family inet {
 address 192.168.6.7/32;
 }
 }
}

user@host:E# show protocols
bgp {
 group external-peers {
 multihop {
 ttl 2;
 }
 local-address 192.168.6.7;
 export send-static;
 peer-as 17;
 neighbor 192.168.40.4;
 }
}

user@host:E# show policy-options
policy-statement send-static {
 term 1 {
 from protocol static;
 then accept;
 }
}

user@host:E# show routing-options
static {
 route 10.10.10.8/30 next-hop 10.10.10.13;
 route 192.168.40.4/32 next-hop 10.10.10.13;
```

```
}
router-id 192.168.6.7;
autonomous-system 18;
```

If you are done configuring the device, enter **commit** from configuration mode.

### **Verification**

Confirm that the configuration is working properly.

- [Verifying Connectivity on page 3406](#)
- [Verifying the BGP Sessions Are Established on page 3406](#)
- [Viewing Advertised Routes on page 3407](#)

### **Verifying Connectivity**

**Purpose** Make sure that Device C can ping Device E, specifying the loopback interface address as the source of the ping request.

The loopback interface address is the source address that BGP will be using.

**Action** From operational mode, enter the **ping 10.10.10.14 source 192.168.40.4** command from Logical System C, and enter the **ping 10.10.10.9 source 192.168.6.7** command from Logical System E.

```
user@host:C> ping 10.10.10.14 source 192.168.40.4
```

```
PING 10.10.10.14 (10.10.10.14): 56 data bytes
64 bytes from 10.10.10.14: icmp_seq=0 ttl=63 time=1.262 ms
64 bytes from 10.10.10.14: icmp_seq=1 ttl=63 time=1.202 ms
^C
--- 10.10.10.14 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.202/1.232/1.262/0.030 ms
```

```
user@host:E> ping 10.10.10.9 source 192.168.6.7
```

```
PING 10.10.10.9 (10.10.10.9): 56 data bytes
64 bytes from 10.10.10.9: icmp_seq=0 ttl=63 time=1.255 ms
64 bytes from 10.10.10.9: icmp_seq=1 ttl=63 time=1.158 ms
^C
--- 10.10.10.9 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.158/1.206/1.255/0.049 ms
```

**Meaning** The static routes are working if the pings work.

### **Verifying the BGP Sessions Are Established**

**Purpose** Verify that the BGP sessions are up.

**Action** From operational mode, enter the **show bgp summary** command.

```
user@host:C> show bgp summary
```

```
Groups: 1 Peers: 1 Down peers: 0
```

```

Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 2 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.168.6.7 18 147 147 0 1 1:04:27
0/2/2/0 0/0/0/0

```

```
user@host:E> show bgp summary
```

```

Groups: 1 Peers: 1 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 2 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.168.40.4 17 202 202 0 1 1:02:18
0/2/2/0 0/0/0/0

```

**Meaning** The output shows that both devices have one peer each. No peers are down.

### Viewing Advertised Routes

**Purpose** Checking to make sure that routes are being advertised by BGP.

**Action** From operational mode, enter the `show route advertising-protocol bgp neighbor` command.

```
user@host:C> show route advertising-protocol bgp 192.168.6.7
```

```

inet.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
 Prefix Nexthop MED Lclpref AS path
* 10.10.10.14/32 Self 0
* 192.168.6.7/32 Self 0

```

```
user@host:E> show route advertising-protocol bgp 192.168.40.4
```

```

inet.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
 Prefix Nexthop MED Lclpref AS path
* 10.10.10.8/30 Self 0
* 192.168.40.4/32 Self 0

```

**Meaning** The `send-static` routing policy is exporting the static routes from the routing table into BGP. BGP is advertising these routes between the peers because the BGP peer session is established.

**Related Documentation**

- [Introduction to Logical Systems on page 3259](#)

### Example: Configuring RSVP-Signaled Point-to-Multipoint LSPs on Logical Systems

- [Point-to-Multipoint LSPs Overview on page 3407](#)
- [Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP on Logical Systems on page 3409](#)

### Point-to-Multipoint LSPs Overview


A point-to-multipoint MPLS LSP is an LSP with a single source and multiple destinations. By taking advantage of the MPLS packet replication capability of the network,

point-to-multipoint LSPs avoid unnecessary packet replication at the ingress router. Packet replication takes place only when packets are forwarded to two or more different destinations requiring different network paths.

This process is illustrated in [Figure 73 on page 3408](#). Router PE1 is configured with a point-to-multipoint LSP to Routers PE2, PE3, and PE4. When Router PE1 sends a packet on the point-to-multipoint LSP to Routers P1 and P2, Router P1 replicates the packet and forwards it to Routers PE2 and PE3. Router P2 sends the packet to Router PE4.

This feature is described in detail in the Internet drafts [draft-raggarwa-mpls-p2mp-te-02.txt](#) (expired February 2004), *Establishing Point to Multipoint MPLS TE LSPs*, [draft-ietf-mpls-rsvp-te-p2mp-02.txt](#), *Extensions to Resource Reservation Protocol-Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label-Switched Paths (LSPs)*, and [draft-ietf-mpls-ldp-p2mp-10.txt](#), *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*.

**Figure 73: Point-to-Multipoint LSPs**



The following are some of the properties of point-to-multipoint LSPs:

- A point-to-multipoint LSP enables you to use MPLS for point-to-multipoint data distribution. This functionality is similar to that provided by IP multicast.
- You can add and remove branch LSPs from a main point-to-multipoint LSP without disrupting traffic. The unaffected parts of the point-to-multipoint LSP continue to function normally.

- You can configure a node to be both a transit and an egress router for different branch LSPs of the same point-to-multipoint LSP.
- You can enable link protection on a point-to-multipoint LSP. Link protection can provide a bypass LSP for each of the branch LSPs that make up the point-to-multipoint LSP. If any of the primary paths fail, traffic can be quickly switched to the bypass.
- You can configure branch LSPs either statically, dynamically, or as a combination of static and dynamic LSPs.
- You can enable graceful Routing Engine switchover (GRES) and graceful restart for point-to-multipoint LSPs at ingress and egress routers. The point-to-multipoint LSPs must be configured using either static routes or circuit cross-connect (CCC). GRES and graceful restart allow the traffic to be forwarded at the Packet Forwarding Engine based on the old state while the control plane recovers. Feature parity for GRES and graceful restart for MPLS point-to-multipoint LSPs on the Junos Trio chipset is supported in Junos OS Releases 11.1R2, 11.2R2, and 11.4.

***Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP on Logical Systems***

In this example, multiple logical systems in a physical routing device act as a collection of paths for an RSVP-signaled point-to-multipoint LSP. The logical systems are chained together and connected internally over a series of logical tunnel (**lt**) interfaces.

- [Requirements on page 3409](#)
- [Overview on page 3410](#)
- [Configuration on page 3411](#)
- [Verification on page 3429](#)

***Requirements***

This example uses the following hardware and software components:

- One MX Series router running logical systems. You do not need to use an MX Series router for the logical systems. You can use any Juniper Networks router that supports logical systems.
- On the MX Series router, the logical systems are connected using logical tunnel (**lt**) interfaces. For more information, see [“Example: Connecting Logical Systems Within the Same Router Using Logical Tunnel Interfaces on MX Series Routers” on page 3283](#) and [“Example: Connecting Logical Systems Within the Same Device Using Logical Tunnel Interfaces on MX Series Routers and EX Series Switches” on page 3279](#). An alternative to using **lt** interfaces is to create external back-to-back interconnections between ports on the router.
- Four customer-edge (CE) devices running in separate physical devices. You do not need to use routers for the CE devices. For example, the CE devices can be EX Series Ethernet Switches.
- Junos OS Release 12.1 or later running on the MX Series router.

On M Series Multiservice Edge and T Series Core Routers, you can create an **lt** interface if you have a Tunnel Services PIC installed on an Enhanced FPC in your routing platform.

On M40e routers, you can create an **lt** interface if you have a Tunnel Services PIC. (An Enhanced FPC is not required.)

On an M7i router, **lt** interfaces can be created by using the integrated Adaptive Services Module.

On an MX Series router, as is shown in this example, the master administrator can configure **lt** interfaces by including the **tunnel-services** statement at the **[edit chassis fpc slot-number pic number]** hierarchy level.

### **Overview**

In this example, the logical systems serve as the transit, branch, and leaf nodes of a single point-to-multipoint LSP. Logical system LS1 is the ingress node. The branches go from LS1 to LS5, LS1 to LS7, and LS1 to LS4. Static unicast routes on the ingress node (LS1) point to the egress nodes.

The following topologies are supported:


- A single logical system in a physical router. The logical system is one node in an RSVP-signaled point-to-multipoint LSP.
- Multiple logical systems in a physical router, with each logical system acting as a label-switched router (LSR). The multiple logical systems can be unconnected, connected to each other internally with **lt** interfaces, or connected to each other externally with back-to-back connections.
- One RSVP-signaled point-to-multipoint LSP, with some nodes being logical systems and other nodes being physical routers.

### **Topology Diagram**

[Figure 74 on page 3411](#) shows the topology used in this example.



Figure 74: RSVP-Signaled Point-to-Multipoint LSP on Logical Systems



g040705

**Configuration**

- [Configuring the MX Series Router to Support Logical Tunnel Interfaces on page 3413](#)
- [Configuring the Ingress LSR \(Logical System LS1\) on page 3413](#)
- [Configuring the Transit and Egress LSRs \(Logical Systems LS2, LS3, LS4, LS5, LS6, and LS7\) on page 3415](#)
- [Configuring Device CE1 on page 3426](#)
- [Configuring Device CE2 on page 3427](#)
- [Configuring Device CE3 on page 3427](#)
- [Configuring Device CE4 on page 3428](#)

**CLI Quick Configuration**

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

```
Router R1
set logical-systems LS1 interfaces ge-2/0/2 unit 0 description LS1-to-CE1
set logical-systems LS1 interfaces ge-2/0/2 unit 0 family inet address 10.0.244.10/30
set logical-systems LS1 interfaces lt-2/0/10 unit 1 description LS1-to-LS2
set logical-systems LS1 interfaces lt-2/0/10 unit 1 encapsulation ethernet
set logical-systems LS1 interfaces lt-2/0/10 unit 1 peer-unit 2
set logical-systems LS1 interfaces lt-2/0/10 unit 1 family inet address 2.2.2.1/24
set logical-systems LS1 interfaces lt-2/0/10 unit 1 family mpls
set logical-systems LS1 interfaces lt-2/0/10 unit 8 description LS1-to-LS6
set logical-systems LS1 interfaces lt-2/0/10 unit 8 encapsulation ethernet
set logical-systems LS1 interfaces lt-2/0/10 unit 8 peer-unit 6
set logical-systems LS1 interfaces lt-2/0/10 unit 8 family inet address 6.6.6.1/24
set logical-systems LS1 interfaces lt-2/0/10 unit 8 family mpls
set logical-systems LS1 interfaces lt-2/0/10 unit 9 description LS1-to-LS3
```

```
set logical-systems LS1 interfaces lt-2/0/10 unit 9 encapsulation ethernet
set logical-systems LS1 interfaces lt-2/0/10 unit 9 peer-unit 3
set logical-systems LS1 interfaces lt-2/0/10 unit 9 family inet address 3.3.3.1/24
set logical-systems LS1 interfaces lt-2/0/10 unit 9 family mpls
set logical-systems LS1 interfaces lo0 unit 1 family inet address 100.10.10.10/32
set logical-systems LS1 protocols rsvp interface lt-2/0/10.1
set logical-systems LS1 protocols rsvp interface lt-2/0/10.8
set logical-systems LS1 protocols rsvp interface lt-2/0/10.9
set logical-systems LS1 protocols rsvp interface lo0.1
set logical-systems LS1 protocols mpls traffic-engineering bgp-igp
set logical-systems LS1 protocols mpls label-switched-path LS1-LS5 to 100.50.50.50
set logical-systems LS1 protocols mpls label-switched-path LS1-LS5 p2mp p2mp1
set logical-systems LS1 protocols mpls label-switched-path LS1-LS7 to 100.70.70.70
set logical-systems LS1 protocols mpls label-switched-path LS1-LS7 p2mp p2mp1
set logical-systems LS1 protocols mpls label-switched-path LS1-LS4 to 100.40.40.40
set logical-systems LS1 protocols mpls label-switched-path LS1-LS4 p2mp p2mp1
set logical-systems LS1 protocols mpls interface lt-2/0/10.1
set logical-systems LS1 protocols mpls interface lt-2/0/10.8
set logical-systems LS1 protocols mpls interface lt-2/0/10.9
set logical-systems LS1 protocols mpls interface lo0.1
set logical-systems LS1 protocols ospf traffic-engineering
set logical-systems LS1 protocols ospf area 0.0.0.0 interface ge-2/0/2.0
set logical-systems LS1 protocols ospf area 0.0.0.0 interface lt-2/0/10.1
set logical-systems LS1 protocols ospf area 0.0.0.0 interface lt-2/0/10.8
set logical-systems LS1 protocols ospf area 0.0.0.0 interface lt-2/0/10.9
set logical-systems LS1 protocols ospf area 0.0.0.0 interface lo0.1
set logical-systems LS1 routing-options static route 5.5.5.0/24 p2mp-lsp-next-hop p2mp1
set logical-systems LS1 routing-options static route 7.7.7.0/24 p2mp-lsp-next-hop p2mp1
set logical-systems LS1 routing-options static route 4.4.4.0/24 p2mp-lsp-next-hop p2mp1
set logical-systems LS1 routing-options router-id 100.10.10.10
```

|            |                                                                                                                                                                                                                                                                                                                                               |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Device CE1 | <pre>set interfaces ge-1/3/2 unit 0 family inet address 10.0.244.9/30 set interfaces ge-1/3/2 unit 0 description CE1-to-LS1 set routing-options static route 10.0.104.8/30 next-hop 10.0.244.10 set routing-options static route 10.0.134.8/30 next-hop 10.0.244.10 set routing-options static route 10.0.224.8/30 next-hop 10.0.244.10</pre> |
| Device CE2 | <pre>set interfaces ge-1/3/3 unit 0 family inet address 10.0.224.9/30 set interfaces ge-1/3/3 unit 0 description CE2-to-LS5 set routing-options static route 10.0.244.8/30 next-hop 10.0.224.10</pre>                                                                                                                                         |
| Device CE3 | <pre>set interfaces ge-2/0/1 unit 0 family inet address 10.0.134.9/30 set interfaces ge-2/0/1 unit 0 description CE3-to-LS7 set routing-options static route 10.0.244.8/30 next-hop 10.0.134.10</pre>                                                                                                                                         |
| Device CE4 | <pre>set interfaces ge-3/1/3 unit 0 family inet address 10.0.104.10/30 set interfaces ge-3/1/3 unit 0 description CE4-to-LS4 set routing-options static route 10.0.244.8/30 next-hop 10.0.104.9</pre>                                                                                                                                         |

### Configuring the MX Series Router to Support Logical Tunnel Interfaces

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

This procedure is required for MX Series routers only. If you have an M Series or T Series router, skip this procedure.

To enable **lt** interfaces on the MX Series router:

1. Run the **show chassis fpc** command to verify that the router has a DPC, MPC, or MIC installed and is in the online state.

```
user@host> show chassis fpc
```

| Slot | State  | Temp (C) | CPU Total | Utilization (%) Interrupt | Memory DRAM (MB) | Utilization (%) Heap | Utilization (%) Buffer |
|------|--------|----------|-----------|---------------------------|------------------|----------------------|------------------------|
| 0    | Empty  |          |           |                           |                  |                      |                        |
| 1    | Empty  |          |           |                           |                  |                      |                        |
| 2    | Online | 31       | 3         | 0                         | 1024             | 14                   | 21                     |

This output shows that slot 0 and slot 1 are empty. Slot 2 is online.

2. Configure FPC slot 2 to support **lt** interfaces.

[edit]

```
user@host# set chassis fpc 2 pic 0 tunnel-services bandwidth 1g
```

This command creates several tunnel interface types, including **gr**, **ip**, and **lt**. For this example, the important one is the **lt** interface.

3. Commit the configuration.

[edit]

```
user@host# commit
```

```
user@host# exit
```

4. Run the **show interfaces terse** command to verify that the router has an **lt** interface.

```
user@host> show interfaces terse
```

| Interface | Admin | Link | Proto | Local | Remote |
|-----------|-------|------|-------|-------|--------|
| ...       |       |      |       |       |        |
| gr-2/0/10 | up    | up   |       |       |        |
| ip-2/0/10 | up    | up   |       |       |        |
| lt-2/0/10 | up    | up   |       |       |        |
| ...       |       |      |       |       |        |

### Configuring the Ingress LSR (Logical System LS1)

#### Step-by-Step Procedure

To configure Logical System LS1:

1. From the main router, configure the logical system.

[edit]

```
user@R1# set logical-systems LS1
```

2. Commit the configuration.

[edit]

```
user@R1# commit
user@R1# exit
```

3. Set the CLI to view the logical system.

```
user@R1> set cli logical-system LS1
Logical system: LS1
```

```
user@R1:LS1>
```

4. Configure the interfaces, interface encapsulation, and protocol families.

```
[edit]
user@R1:LS1# edit interfaces
[edit interfaces]
user@R1:LS1# set ge-2/0/2 unit 0 description R1-to-CE1
user@R1:LS1# set ge-2/0/2 unit 0 family inet address 10.0.244.10/30
user@R1:LS1# set lt-2/0/10 unit 1 description LS1-to-LS2
user@R1:LS1# set lt-2/0/10 unit 1 encapsulation ethernet
user@R1:LS1# set lt-2/0/10 unit 1 peer-unit 2
user@R1:LS1# set lt-2/0/10 unit 1 family inet address 2.2.2.1/24
user@R1:LS1# set lt-2/0/10 unit 1 family mpls
user@R1:LS1# set lt-2/0/10 unit 8 description LS1-to-LS6
user@R1:LS1# set lt-2/0/10 unit 8 encapsulation ethernet
user@R1:LS1# set lt-2/0/10 unit 8 peer-unit 6
user@R1:LS1# set lt-2/0/10 unit 8 family inet address 6.6.6.1/24
user@R1:LS1# set lt-2/0/10 unit 8 family mpls
user@R1:LS1# set lt-2/0/10 unit 9 description LS1-to-LS3
user@R1:LS1# set lt-2/0/10 unit 9 encapsulation ethernet
user@R1:LS1# set lt-2/0/10 unit 9 peer-unit 3
user@R1:LS1# set lt-2/0/10 unit 9 family inet address 3.3.3.1/24
user@R1:LS1# set lt-2/0/10 unit 9 family mpls
user@R1:LS1# set lo0 unit 1 family inet address 100.10.10.10/32
user@R1:LS1# exit
```

5. Enable RSVP, MPLS, and OSPF on the interfaces.

```
[edit]
user@R1:LS1# edit protocols
[edit protocols]
user@R1:LS1# set rsvp interface lt-2/0/10.1
user@R1:LS1# set rsvp interface lt-2/0/10.8
user@R1:LS1# set rsvp interface lt-2/0/10.9
user@R1:LS1# set rsvp interface lo0.1
user@R1:LS1# set mpls interface lt-2/0/10.1
user@R1:LS1# set mpls interface lt-2/0/10.8
user@R1:LS1# set mpls interface lt-2/0/10.9
user@R1:LS1# set mpls interface lo0.1
user@R1:LS1# set ospf area 0.0.0.0 interface ge-2/0/2.0
user@R1:LS1# set ospf area 0.0.0.0 interface lt-2/0/10.1
user@R1:LS1# set ospf area 0.0.0.0 interface lt-2/0/10.8
user@R1:LS1# set ospf area 0.0.0.0 interface lt-2/0/10.9
user@R1:LS1# set ospf area 0.0.0.0 interface lo0.1
```

6. Configure the MPLS point-to-multipoint LSPs.

```
[edit protocols]
user@R1:LS1# set mpls label-switched-path LS1-LS5 to 100.50.50.50
user@R1:LS1# set mpls label-switched-path LS1-LS5 p2mp p2mp1
```

```

user@R1:LS1# set mpls label-switched-path LS1-LS7 to 100.70.70.70
user@R1:LS1# set mpls label-switched-path LS1-LS7 p2mp p2mp1
user@R1:LS1# set mpls label-switched-path LS1-LS4 to 100.40.40.40
user@R1:LS1# set mpls label-switched-path LS1-LS4 p2mp p2mp1

```

7. Enable MPLS to perform traffic engineering for OSPF.

```

[edit protocols]
user@R1:LS1# set mpls traffic-engineering bgp-igp
user@R1:LS1# exit

```

This causes the ingress routes to be installed in the `inet.0` routing table. By default, MPLS performs traffic engineering for BGP only. You need to enable MPLS traffic engineering on the ingress LSR only.

8. Enable traffic engineering for OSPF.

```

[edit protocols]
user@R1:LS1# set ospf traffic-engineering
user@R1:LS1# exit

```

This causes the shortest-path first (SPF) algorithm to take into account the LSPs configured under MPLS.

9. Configure the router ID.

```

[edit]
user@R1:LS1# edit routing-options
[edit routing-options]
user@R1:LS1# set router-id 100.10.10.10

```

10. Configure static IP unicast routes with the point-to-multipoint LSP name as the next hop for each route.

```

[edit routing-options]
user@R1:LS1# set static route 5.5.5.0/24 p2mp-lsp-next-hop p2mp1
user@R1:LS1# set static route 7.7.7.0/24 p2mp-lsp-next-hop p2mp1
user@R1:LS1# set static route 4.4.4.0/24 p2mp-lsp-next-hop p2mp1
user@R1:LS1# exit

```

11. If you are done configuring the device, commit the configuration.

```

[edit]
user@R1:LS1# commit

```

### *Configuring the Transit and Egress LSRs (Logical Systems LS2, LS3, LS4, LS5, LS6, and LS7)*

#### **Step-by-Step Procedure**

To configure the transit and egress LSRs:

1. Configure the interfaces, interface encapsulation, and protocol families.

```

[edit]
user@R1# edit logical-systems
[edit logical-systems]
user@R1# set LS2 interfaces lt-2/0/10 unit 2 description LS2-to-LS1
user@R1# set LS2 interfaces lt-2/0/10 unit 2 encapsulation ethernet
user@R1# set LS2 interfaces lt-2/0/10 unit 2 peer-unit 1
user@R1# set LS2 interfaces lt-2/0/10 unit 2 family inet address 2.2.2.2/24

```

```
user@R1# set LS2 interfaces lt-2/0/10 unit 2 family mpls
user@R1# set LS2 interfaces lt-2/0/10 unit 10 description LS2-to-LS5
user@R1# set LS2 interfaces lt-2/0/10 unit 10 encapsulation ethernet
user@R1# set LS2 interfaces lt-2/0/10 unit 10 peer-unit 5
user@R1# set LS2 interfaces lt-2/0/10 unit 10 family inet address 5.5.5.1/24
user@R1# set LS2 interfaces lt-2/0/10 unit 10 family mpls
user@R1# set LS2 interfaces lo0 unit 2 family inet address 100.20.20.20/32
user@R1# set LS3 interfaces lt-2/0/10 unit 3 description LS3-to-LS1
user@R1# set LS3 interfaces lt-2/0/10 unit 3 encapsulation ethernet
user@R1# set LS3 interfaces lt-2/0/10 unit 3 peer-unit 9
user@R1# set LS3 interfaces lt-2/0/10 unit 3 family inet address 3.3.3.2/24
user@R1# set LS3 interfaces lt-2/0/10 unit 3 family mpls
user@R1# set LS3 interfaces lt-2/0/10 unit 12 description LS3-to-LS4
user@R1# set LS3 interfaces lt-2/0/10 unit 12 encapsulation ethernet
user@R1# set LS3 interfaces lt-2/0/10 unit 12 peer-unit 4
user@R1# set LS3 interfaces lt-2/0/10 unit 12 family inet address 4.4.4.1/24
user@R1# set LS3 interfaces lt-2/0/10 unit 12 family mpls
user@R1# set LS3 interfaces lo0 unit 3 family inet address 100.30.30.30/32
user@R1# set LS4 interfaces ge-2/0/0 unit 0 description R1-to-CE4
user@R1# set LS4 interfaces ge-2/0/0 unit 0 family inet address 10.0.104.9/30
user@R1# set LS4 interfaces lt-2/0/10 unit 4 description LS4-to-LS3
user@R1# set LS4 interfaces lt-2/0/10 unit 4 encapsulation ethernet
user@R1# set LS4 interfaces lt-2/0/10 unit 4 peer-unit 12
user@R1# set LS4 interfaces lt-2/0/10 unit 4 family inet address 4.4.4.2/24
user@R1# set LS4 interfaces lt-2/0/10 unit 4 family mpls
user@R1# set LS4 interfaces lo0 unit 4 family inet address 100.40.40.40/32
user@R1# set LS5 interfaces ge-2/0/3 unit 0 description LS1-to-CE1
user@R1# set LS5 interfaces ge-2/0/3 unit 0 family inet address 10.0.224.10/30
user@R1# set LS5 interfaces lt-2/0/10 unit 5 description LS5-to-LS2
user@R1# set LS5 interfaces lt-2/0/10 unit 5 encapsulation ethernet
user@R1# set LS5 interfaces lt-2/0/10 unit 5 peer-unit 10
user@R1# set LS5 interfaces lt-2/0/10 unit 5 family inet address 5.5.5.2/24
user@R1# set LS5 interfaces lt-2/0/10 unit 5 family mpls
user@R1# set LS5 interfaces lo0 unit 5 family inet address 100.50.50.50/32
user@R1# set LS6 interfaces lt-2/0/10 unit 6 description LS6-to-LS1
user@R1# set LS6 interfaces lt-2/0/10 unit 6 encapsulation ethernet
user@R1# set LS6 interfaces lt-2/0/10 unit 6 peer-unit 8
user@R1# set LS6 interfaces lt-2/0/10 unit 6 family inet address 6.6.6.2/24
user@R1# set LS6 interfaces lt-2/0/10 unit 6 family mpls
user@R1# set LS6 interfaces lt-2/0/10 unit 11 description LS6-to-LS7
user@R1# set LS6 interfaces lt-2/0/10 unit 11 encapsulation ethernet
user@R1# set LS6 interfaces lt-2/0/10 unit 11 peer-unit 7
user@R1# set LS6 interfaces lt-2/0/10 unit 11 family inet address 7.7.7.1/24
user@R1# set LS6 interfaces lt-2/0/10 unit 11 family mpls
user@R1# set LS6 interfaces lo0 unit 6 family inet address 100.60.60.60/32
user@R1# set LS7 interfaces ge-2/0/1 unit 0 description R1-to-CE3
user@R1# set LS7 interfaces ge-2/0/1 unit 0 family inet address 10.0.134.10/30
user@R1# set LS7 interfaces lt-2/0/10 unit 7 description LS7-to-LS6
user@R1# set LS7 interfaces lt-2/0/10 unit 7 encapsulation ethernet
user@R1# set LS7 interfaces lt-2/0/10 unit 7 peer-unit 11
user@R1# set LS7 interfaces lt-2/0/10 unit 7 family inet address 7.7.7.2/24
user@R1# set LS7 interfaces lt-2/0/10 unit 7 family mpls
user@R1# set LS7 interfaces lo0 unit 7 family inet address 100.70.70.70/32
```

2. Enable RSVP, MPLS, and OSPF on the interfaces.

```
[edit]
user@R1# edit logical-systems
[edit logical-systems]
user@R1# set LS2 protocols rsvp interface lt-2/0/10.2
user@R1# set LS2 protocols rsvp interface lt-2/0/10.10
user@R1# set LS2 protocols rsvp interface lo0.2
user@R1# set LS2 protocols mpls interface lt-2/0/10.2
user@R1# set LS2 protocols mpls interface lt-2/0/10.10
user@R1# set LS2 protocols mpls interface lo0.2
user@R1# set LS2 protocols ospf area 0.0.0.0 interface lt-2/0/10.2
user@R1# set LS2 protocols ospf area 0.0.0.0 interface lt-2/0/10.10
user@R1# set LS2 protocols ospf area 0.0.0.0 interface lo0.2
user@R1# set LS3 protocols rsvp interface lt-2/0/10.3
user@R1# set LS3 protocols rsvp interface lt-2/0/10.12
user@R1# set LS3 protocols rsvp interface lo0.3
user@R1# set LS3 protocols mpls interface lt-2/0/10.3
user@R1# set LS3 protocols mpls interface lt-2/0/10.12
user@R1# set LS3 protocols mpls interface lo0.3
user@R1# set LS3 protocols ospf area 0.0.0.0 interface lt-2/0/10.3
user@R1# set LS3 protocols ospf area 0.0.0.0 interface lt-2/0/10.12
user@R1# set LS3 protocols ospf area 0.0.0.0 interface lo0.3
user@R1# set LS4 protocols rsvp interface lt-2/0/10.4
user@R1# set LS4 protocols rsvp interface lo0.4
user@R1# set LS4 protocols mpls interface lt-2/0/10.4
user@R1# set LS4 protocols mpls interface lo0.4
user@R1# set LS4 protocols ospf area 0.0.0.0 interface ge-2/0/0.0
user@R1# set LS4 protocols ospf area 0.0.0.0 interface lt-2/0/10.4
user@R1# set LS4 protocols ospf area 0.0.0.0 interface lo0.4
user@R1# set LS5 protocols rsvp interface lt-2/0/10.5
user@R1# set LS5 protocols rsvp interface lo0.5
user@R1# set LS5 protocols mpls interface lt-2/0/10.5
user@R1# set LS5 protocols mpls interface lo0.5
user@R1# set LS5 protocols ospf area 0.0.0.0 interface ge-2/0/3.0
user@R1# set LS5 protocols ospf area 0.0.0.0 interface lt-2/0/10.5
user@R1# set LS5 protocols ospf area 0.0.0.0 interface lo0.5
user@R1# set LS6 protocols rsvp interface lt-2/0/10.6
user@R1# set LS6 protocols rsvp interface lt-2/0/10.11
user@R1# set LS6 protocols rsvp interface lo0.6
user@R1# set LS6 protocols mpls interface lt-2/0/10.6
user@R1# set LS6 protocols mpls interface lt-2/0/10.11
user@R1# set LS6 protocols mpls interface lo0.6
user@R1# set LS6 protocols ospf area 0.0.0.0 interface lt-2/0/10.6
user@R1# set LS6 protocols ospf area 0.0.0.0 interface lt-2/0/10.11
user@R1# set LS6 protocols ospf area 0.0.0.0 interface lo0.6
user@R1# set LS7 protocols rsvp interface lt-2/0/10.7
user@R1# set LS7 protocols rsvp interface lo0.7
user@R1# set LS7 protocols mpls interface lt-2/0/10.7
user@R1# set LS7 protocols mpls interface lo0.7
user@R1# set LS7 protocols ospf area 0.0.0.0 interface ge-2/0/1.0
user@R1# set LS7 protocols ospf area 0.0.0.0 interface lt-2/0/10.7
user@R1# set LS7 protocols ospf area 0.0.0.0 interface lo0.7
```

3. Enable traffic engineering for OSPF.

```
[edit logical-systems]
user@R1# set LS2 protocols ospf traffic-engineering
user@R1# set LS3 protocols ospf traffic-engineering
user@R1# set LS4 protocols ospf traffic-engineering
user@R1# set LS5 protocols ospf traffic-engineering
user@R1# set LS6 protocols ospf traffic-engineering
user@R1# set LS7 protocols ospf traffic-engineering
```

This causes the SPF algorithm to take into account the LSPs configured under MPLS.

4. Configure the router IDs.

```
[edit logical-systems]
user@R1# set LS2 routing-options router-id 100.20.20.20
user@R1# set LS3 routing-options router-id 100.30.30.30
user@R1# set LS4 routing-options router-id 100.40.40.40
user@R1# set LS5 routing-options router-id 100.50.50.50
user@R1# set LS6 routing-options router-id 100.60.60.60
user@R1# set LS7 routing-options router-id 100.70.70.70
```

5. If you are done configuring the device, commit the configuration.

```
[edit logical-systems]
user@R1# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show logical-systems** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show logical-systems
LS1 {
 interfaces {
 ge-2/0/2 {
 unit 0 {
 description R1-to-CE1;
 family inet {
 address 10.0.244.10/30;
 }
 }
 }
 }
 lt-2/0/10 {
 unit 1 {
 description LS1-to-LS2;
 encapsulation ethernet;
 peer-unit 2;
 family inet {
 address 2.2.2.1/24;
 }
 family mpls;
 }
 unit 8 {
 description LS1-to-LS6;
 encapsulation ethernet;
```



```

 peer-unit 6;
 family inet {
 address 6.6.6.1/24;
 }
 family mpls;
 }
 unit 9 {
 description LS1-to-LS3;
 encapsulation ethernet;
 peer-unit 3;
 family inet {
 address 3.3.3.1/24;
 }
 family mpls;
 }
}
lo0 {
 unit 1 {
 family inet {
 address 100.10.10.10/32;
 }
 }
}
}
protocols {
 rsvp {
 interface lt-2/0/10.1;
 interface lt-2/0/10.8;
 interface lt-2/0/10.9;
 interface lo0.1;
 }
 mpls {
 traffic-engineering bgp-igp;
 label-switched-path LS1-to-LS5 {
 to 100.50.50.50;
 p2mp p2mpl;
 }
 label-switched-path LS1-to-LS7 {
 to 100.70.70.70;
 p2mp p2mpl;
 }
 label-switched-path LS1-to-LS4 {
 to 100.40.40.40;
 p2mp p2mpl;
 }
 interface lt-2/0/10.1;
 interface lt-2/0/10.8;
 interface lt-2/0/10.9;
 interface lo0.1;
 }
}
ospf {
 traffic-engineering;
 area 0.0.0.0 {
 interface ge-2/0/2.0;
 interface lt-2/0/10.1;
 interface lt-2/0/10.8;

```

```
 interface lt-2/0/10.9;
 interface lo0.1;
 }
}
routing-options {
 static {
 route 5.5.5.0/24 {
 p2mp-lsp-next-hop p2mp1;
 }
 route 7.7.7.0/24 {
 p2mp-lsp-next-hop p2mp1;
 }
 route 4.4.4.0/24 {
 p2mp-lsp-next-hop p2mp1;
 }
 }
 router-id 100.10.10.10;
}
LS2 {
 interfaces {
 lt-2/0/10 {
 unit 2 {
 description LS2-to-LS1;
 encapsulation ethernet;
 peer-unit 1;
 family inet {
 address 2.2.2.2/24;
 }
 family mpls;
 }
 unit 10 {
 description LS2-to-LS5;
 encapsulation ethernet;
 peer-unit 5;
 family inet {
 address 5.5.5.1/24;
 }
 family mpls;
 }
 }
 lo0 {
 unit 2 {
 family inet {
 address 100.20.20.20/32;
 }
 }
 }
 }
 protocols {
 rsvp {
 interface lt-2/0/10.2;
 interface lt-2/0/10.10;
 interface lo0.2;
 }
 }
}
```

```
 mpls {
 interface lt-2/0/10.2;
 interface lt-2/0/10.10;
 interface lo0.2;
 }
 ospf {
 traffic-engineering;
 area 0.0.0.0 {
 interface lt-2/0/10.2;
 interface lt-2/0/10.10;
 interface lo0.2;
 }
 }
 }
 routing-options {
 router-id 100.20.20.20;
 }
}
LS3 {
 interfaces {
 lt-2/0/10 {
 unit 3 {
 description LS3-to-LS1;
 encapsulation ethernet;
 peer-unit 9;
 family inet {
 address 3.3.3.2/24;
 }
 family mpls;
 }
 unit 12 {
 description LS3-to-LS4;
 encapsulation ethernet;
 peer-unit 4;
 family inet {
 address 4.4.4.1/24;
 }
 family mpls;
 }
 }
 }
 lo0 {
 unit 3 {
 family inet {
 address 100.30.30.30/32;
 }
 }
 }
}
protocols {
 rsvp {
 interface lt-2/0/10.3;
 interface lt-2/0/10.12;
 interface lo0.3;
 }
 mpls {
 interface lt-2/0/10.3;
```

```
 interface lt-2/0/10.12;
 interface lo0.3;
 }
 ospf {
 traffic-engineering;
 area 0.0.0.0 {
 interface lt-2/0/10.3;
 interface lt-2/0/10.12;
 interface lo0.3;
 }
 }
}
routing-options {
 router-id 100.30.30.30;
}
}
LS4 {
 interfaces {
 ge-2/0/0 {
 unit 0 {
 description R1-to-CE4;
 family inet {
 address 10.0.104.9/30;
 }
 }
 }
 }
 lt-2/0/10 {
 unit 4 {
 description LS4-to-LS3;
 encapsulation ethernet;
 peer-unit 12;
 family inet {
 address 4.4.4.2/24;
 }
 family mpls;
 }
 }
 lo0 {
 unit 4 {
 family inet {
 address 100.40.40.40/32;
 }
 }
 }
}
protocols {
 rsvp {
 interface lt-2/0/10.4;
 interface lo0.4;
 }
 mpls {
 interface lt-2/0/10.4;
 interface lo0.4;
 }
 ospf {
 traffic-engineering;
 }
}
```

```
 area 0.0.0.0 {
 interface ge-2/0/0.0;
 interface lt-2/0/10.4;
 interface lo0.4;
 }
 }
 routing-options {
 router-id 100.40.40.40;
 }
}
LS5 {
 interfaces {
 ge-2/0/3 {
 unit 0 {
 description LS1-to-CE1;
 family inet {
 address 10.0.224.10/30;
 }
 }
 }
 lt-2/0/10 {
 unit 5 {
 description LS5-to-LS2;
 encapsulation ethernet;
 peer-unit 10;
 family inet {
 address 5.5.5.2/24;
 }
 family mpls;
 }
 }
 lo0 {
 unit 5 {
 family inet {
 address 100.50.50.50/32;
 }
 }
 }
 }
 protocols {
 rsvp {
 interface lt-2/0/10.5;
 interface lo0.5;
 }
 mpls {
 interface lt-2/0/10.5;
 interface lo0.5;
 }
 ospf {
 traffic-engineering;
 area 0.0.0.0 {
 interface ge-2/0/3.0;
 interface lt-2/0/10.5;
 interface lo0.5;
 }
 }
 }
}
```

```
 }
 }
 routing-options {
 router-id 100.50.50.50;
 }
}
LS6 {
 interfaces {
 lt-2/0/10 {
 unit 6 {
 description LS6-to-LS1;
 encapsulation ethernet;
 peer-unit 8;
 family inet {
 address 6.6.6.2/24;
 }
 family mpls;
 }
 unit 11 {
 description LS6-to-LS7;
 encapsulation ethernet;
 peer-unit 7;
 family inet {
 address 7.7.7.1/24;
 }
 family mpls;
 }
 }
 }
 lo0 {
 unit 6 {
 family inet {
 address 100.60.60.60/32;
 }
 }
 }
}
protocols {
 rsvp {
 interface lt-2/0/10.6;
 interface lt-2/0/10.11;
 interface lo0.6;
 }
 mpls {
 interface lt-2/0/10.6;
 interface lt-2/0/10.11;
 interface lo0.6;
 }
 ospf {
 traffic-engineering;
 area 0.0.0.0 {
 interface lt-2/0/10.6;
 interface lt-2/0/10.11;
 interface lo0.6;
 }
 }
}
```

```
 routing-options {
 router-id 100.60.60.60;
 }
 }
LS7 {
 interfaces {
 ge-2/0/1 {
 unit 0 {
 description R1-to-CE3;
 family inet {
 address 10.0.134.10/30;
 }
 }
 }
 lt-2/0/10 {
 unit 7 {
 description LS7-to-LS6;
 encapsulation ethernet;
 peer-unit 11;
 family inet {
 address 7.7.7.2/24;
 }
 family mpls;
 }
 }
 lo0 {
 unit 7 {
 family inet {
 address 100.70.70.70/32;
 }
 }
 }
 }
 protocols {
 rsvp {
 interface lt-2/0/10.7;
 interface lo0.7;
 }
 mpls {
 interface lt-2/0/10.7;
 interface lo0.7;
 }
 ospf {
 traffic-engineering;
 area 0.0.0.0 {
 interface ge-2/0/1.0;
 interface lt-2/0/10.7;
 interface lo0.7;
 }
 }
 }
 routing-options {
 router-id 100.70.70.70;
 }
}
```

### *Configuring Device CE1*

#### **Step-by-Step Procedure**

To configure Device CE1:

1. Configure an interface to Logical System LS1.

```
[edit]
user@CE1# edit interfaces
[edit interfaces]
user@CE1# set ge-1/3/2 unit 0 family inet address 10.0.244.9/30
user@CE1# set ge-1/3/2 unit 0 description CE1-to-LS1
user@CE1# exit
```

2. Configure static routes from Device CE1 to the three other customer networks, with Logical System LS1 as the next hop.

```
[edit]
user@CE1# edit routing-options
[edit routing-options]
set static route 10.0.104.8/30 next-hop 10.0.244.10
set static route 10.0.134.8/30 next-hop 10.0.244.10
set static route 10.0.224.8/30 next-hop 10.0.244.10
user@CE1# exit
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@CE1# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE1# show interfaces
interfaces {
 ge-1/3/2 {
 unit 0 {
 family inet {
 address 10.0.244.9/30;
 description CE1-to-LS1;
 }
 }
 }
}

user@CE1# show routing-options
routing-options {
 static {
 route 10.0.104.8/30 next-hop 10.0.244.10;
 route 10.0.134.8/30 next-hop 10.0.244.10;
 route 10.0.224.8/30 next-hop 10.0.244.10;
 }
}
```



**Configuring Device CE2****Step-by-Step Procedure**

To configure Device CE2:

1. Configure an interface to Logical System LS5.

```
[edit]
user@CE2# edit interfaces
[edit interfaces]
user@CE2# set ge-1/3/3 unit 0 family inet address 10.0.224.9/30
user@CE2# set ge-1/3/3 unit 0 description CE2-to-LS5
user@CE2# exit
```

2. Configure a static route from Device CE2 to CE1, with Logical System LS5 as the next hop.

```
[edit]
user@CE2# edit routing-options
[edit routing-options]
user@CE2# set static route 10.0.244.8/30 next-hop 10.0.224.10
user@CE2# exit
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@CE2# commit
```

**Results**

From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE2# show interfaces
interfaces {
 ge-1/3/3 {
 unit 0 {
 family inet {
 address 10.0.224.9/30;
 description CE2-to-LS5;
 }
 }
 }
}

user@CE2# show routing-options
routing-options {
 static {
 route 10.0.244.8/30 next-hop 10.0.224.10;
 }
}
```

**Configuring Device CE3****Step-by-Step Procedure**

To configure Device CE3:

1. Configure an interface to Logical System LS7.

```
[edit]
```

```
user@CE3# edit interfaces
[edit interfaces]
user@CE3# set ge-2/0/1 unit 0 family inet address 10.0.134.9/30
user@CE3# set ge-2/0/1 unit 0 description CE3-to-LS7
user@CE3# exit
```

2. Configure a static route from Device CE3 to CE1, with Logical System LS7 as the next hop.

```
[edit]
user@CE3# edit routing-options
[edit routing-options]
user@CE3# set static route 10.0.244.8/30 next-hop 10.0.134.10
user@CE3# exit
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@CE3# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE3# show interfaces
interfaces {
 ge-2/0/1 {
 unit 0 {
 family inet {
 address 10.0.134.9/30;
 description CE3-to-LS7;
 }
 }
 }
}

user@CE3# show routing-options
routing-options {
 static {
 route 10.0.244.8/30 next-hop 10.0.134.10;
 }
}
```

#### *Configuring Device CE4*

**Step-by-Step Procedure** To configure Device CE4:

1. Configure an interface to Logical System LS4.

```
[edit]
user@CE4# edit interfaces
[edit interfaces]
user@CE4# set ge-3/1/3 unit 0 family inet address 10.0.104.10/30
user@CE4# set ge-3/1/3 unit 0 description CE4-to-LS4
```

2. Configure a static route from Device CE4 to CE1, with Logical System LS4 as the next hop.

```
[edit]
user@CE4# edit routing-options
[edit routing-options]
user@CE4# set static route 10.0.244.8/30 next-hop 10.0.104.9
user@CE4# exit
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@CE4# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE4# show interfaces
interfaces {
 ge-3/1/3 {
 unit 0 {
 family inet {
 address 10.0.104.10/30;
 description CE4-to-LS4;
 }
 }
 }
}

user@CE4# show routing-options
routing-options {
 static {
 route 10.0.244.8/30 next-hop 10.0.104.9;
 }
}
```

### **Verification**

Confirm that the configuration is working properly.

- [Verifying Connectivity on page 3429](#)
- [Verifying the State of the Point-to-Multipoint LSP on page 3430](#)
- [Checking the Forwarding Table on page 3431](#)

### **Verifying Connectivity**

**Purpose** Make sure that the devices can ping each other.

**Action** Run the **ping** command from CE1 to the interface on CE2 connecting to LS5.

```
user@CE1> ping 10.0.224.9
PING 10.0.224.9 (10.0.224.9): 56 data bytes
64 bytes from 10.0.224.9: icmp_seq=0 ttl=61 time=1.387 ms
64 bytes from 10.0.224.9: icmp_seq=1 ttl=61 time=1.394 ms
64 bytes from 10.0.224.9: icmp_seq=2 ttl=61 time=1.506 ms
^C
--- 10.0.224.9 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.387/1.429/1.506/0.055 ms
```

Run the **ping** command from CE1 to the interface on CE3 connecting to LS7.

```
user@CE1> ping 10.0.134.9
PING 10.0.134.9 (10.0.134.9): 56 data bytes
64 bytes from 10.0.134.9: icmp_seq=0 ttl=61 time=1.068 ms
64 bytes from 10.0.134.9: icmp_seq=1 ttl=61 time=1.062 ms
64 bytes from 10.0.134.9: icmp_seq=2 ttl=61 time=1.053 ms
^C
--- 10.0.134.9 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.053/1.061/1.068/0.006 ms
```

Run the **ping** command from CE1 to the interface on CE4 connecting to LS4.

```
user@CE1> ping 10.0.104.10
PING 10.0.104.10 (10.0.104.10): 56 data bytes
64 bytes from 10.0.104.10: icmp_seq=0 ttl=61 time=1.079 ms
64 bytes from 10.0.104.10: icmp_seq=1 ttl=61 time=1.048 ms
64 bytes from 10.0.104.10: icmp_seq=2 ttl=61 time=1.070 ms
^C
--- 10.0.104.10 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.048/1.066/1.079/0.013 ms
```

#### *Verifying the State of the Point-to-Multipoint LSP*

**Purpose** Make sure that the ingress, transit, and egress LSRs are in the Up state.



.....  
**NOTE:** For this example, the **show rsvp session** command displays the same output as the **show mpls lsp p2mp** command.  
.....

**Action** Run the `show mpls lsp p2mp` command on all of the LSRs. Only the ingress LSR is shown here.

```
user@R1> set cli logical-system LS1
Logical system: LS1

user@R1:LS1> show mpls lsp p2mp
Ingress LSP: 1 sessions
P2MP name: p2mp1, P2MP branch count: 3
To From State Rt P ActivePath LSPname
100.40.40.40 100.10.10.10 Up 0 * LS1-LS4
100.70.70.70 100.10.10.10 Up 0 * LS1-LS7
100.50.50.50 100.10.10.10 Up 0 * LS1-LS5
Total 3 displayed, Up 3, Down 0
...
```

### Checking the Forwarding Table

**Purpose** Make sure that the routes are set up as expected by running the `show route forwarding-table` command. Only the routes to the remote customer networks are shown here.

**Action**

```
user@R1:LS1> show route forwarding-table
Routing table: default.inet
Internet:
Destination Type RtRef Next hop Type Index NhRef Netif
...
10.0.104.8/30 user 0 3.3.3.2 ucst 1006 6 1t-2/0/10.9
10.0.134.8/30 user 0 6.6.6.2 ucst 1010 6 1t-2/0/10.8
10.0.224.8/30 user 0 2.2.2.2 ucst 1008 6 1t-2/0/10.1
...
```

### Example: Configuring VPNs and VPLS on Logical Systems

- [Layer 3 VPN Introduction on page 3431](#)
- [Introduction to VPLS on page 3432](#)
- [Example: Using Logical Systems to Configure Provider Edge and Provider Routers in a Layer 3 VPN and VPLS Scenario on page 3432](#)

#### Layer 3 VPN Introduction

In Junos OS, Layer 3 VPNs are based on RFC 4364. RFC 4364 defines a mechanism by which service providers can use their IP backbones to provide VPN services to their customers. A Layer 3 VPN is a set of sites that share common routing information and whose connectivity is controlled by a collection of policies. The sites that make up a Layer 3 VPN are connected over a provider's existing public Internet backbone.

RFC 4364 VPNs are also known as BGP/MPLS VPNs because BGP is used to distribute VPN routing information across the provider's backbone, and MPLS is used to forward VPN traffic across the backbone to remote VPN sites.

Customer networks, because they are private, can use either public addresses or private addresses, as defined in RFC 1918, *Address Allocation for Private Internets*. When customer networks that use private addresses connect to the public Internet infrastructure, the

private addresses might overlap with the same private addresses used by other network users. MPLS/BGP VPNs solve this problem by adding a VPN identifier prefix to each address from a particular VPN site, thereby creating an address that is unique both within the VPN and within the public Internet. In addition, each VPN has its own VPN-specific routing table that contains the routing information for that VPN only.

### ***Introduction to VPLS***

VPLS is an Ethernet-based point-to-multipoint Layer 2 VPN. It allows you to connect geographically dispersed Ethernet local area networks (LAN) sites to each other across an MPLS backbone. For customers who implement VPLS, all sites appear to be in the same Ethernet LAN even though traffic travels across the service provider's network.

VPLS, in its implementation and configuration, has much in common with a Layer 2 VPN. In VPLS, a packet originating within a service provider customer's network is sent first to a customer edge (CE) device (for example, a router or Ethernet switch). It is then sent to a provider edge (PE) router within the service provider's network. The packet traverses the service provider's network over a MPLS label-switched path (LSP). It arrives at the egress PE router, which then forwards the traffic to the CE device at the destination customer site.



**NOTE:** In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

---

The difference is that for VPLS, packets can traverse the service provider's network in point-to-multipoint fashion, meaning that a packet originating from a CE device can be broadcast to all the PE routers participating in a VPLS routing instance. In contrast, a Layer 2 VPN forwards packets in point-to-point fashion only.

The paths carrying VPLS traffic between each PE router participating in a routing instance are called pseudowires. The pseudowires are signaled using either BGP or LDP.

### ***Example: Using Logical Systems to Configure Provider Edge and Provider Routers in a Layer 3 VPN and VPLS Scenario***

This example provides step-by-step procedures to configure provider edge (PE) and provider (P) routers in a VPN and VPLS scenario using logical systems.

- [Requirements on page 3433](#)
- [Overview on page 3433](#)
- [Configuration on page 3435](#)
- [Verification on page 3462](#)

**Requirements**

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

**Overview**

In this example, VPNs are used to separate customer traffic across a provider backbone.

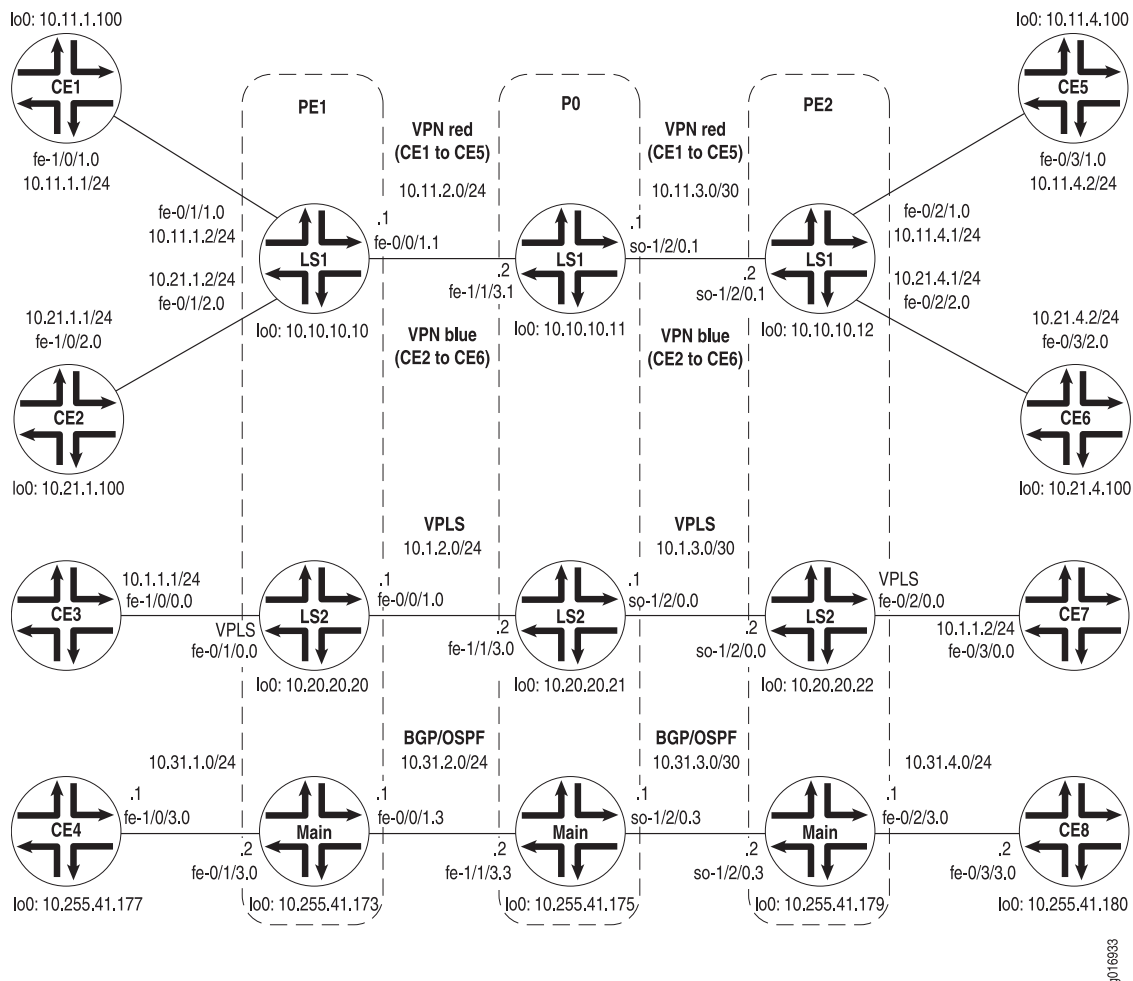
**Topology**

[Figure 75 on page 3434](#) shows four pairs of CE routers that are connected across an MPLS backbone:

- Routers CE1 and CE5 are part of the red VPN.
- Routers CE2 and CE6 are in the blue VPN.
- Routers CE3 and CE7 belong to a VPLS domain.
- Routers CE4 and CE8 are connected with standard protocols.

Two logical systems are configured on PE routers PE1 and PE2 and provider core Router P0. Each of these three routers has two logical systems: LS1 and LS2. To illustrate the concept of a logical system, both VPNs are part of Logical System LS1, the VPLS instance belongs to Logical System LS2, and the remaining routers use the main router portion of routers PE1, P0, and PE2.

### Figure 75: Provider Edge and Provider Logical System Topology Diagram



On Router PE1, two VPN routing and forwarding (VRF) routing instances are created in Logical System LS1. The routing instances are called red and blue. The example configures the customer edge (CE)-facing logical interfaces so that traffic from Router CE1 is placed in the red VPN, and traffic from Router CE2 is placed in the blue VPN. A logical interface at **fe-0/0/1.1** connects to Logical System LS1 on Router P0. A VPLS routing instance is in Logical System LS2. The logical interface is configured so that traffic from Router CE3 is sent into the VPLS domain. This logical interface connects to Logical System LS2 on Router P0. The example also contains an administrator for Logical System LS1. The logical system administrator is responsible for the maintenance of this logical system. Finally, the example shows how to configure a logical interface to interconnect Router CE4 with the main router portion of Router PE1.

Router PE2 has the two VRF routing instances in Logical System LS1: red and blue. The CE-facing logical interfaces enable traffic from Router CE5 to be placed in the red VPN, and traffic from Router CE6 in the blue VPN. One logical interface on **so-1/2/0.1** connects to Logical System LS1 on Router P0. The VPLS routing instance is configured in Logical System LS2. A logical interface enables traffic from Router CE7 to be sent into the VPLS domain and connects to Logical System LS2 on Router P0. The example shows how to



configure a logical interface to interconnect Router CE8 with the main router portion of Router P0. Finally, you can optionally create a logical system administrator that has configuration privileges for Logical System LS1 and viewing privileges for Logical System LS2.

On Router P0, the example shows how to configure Logical Systems LS1, LS2, and the main router. You must configure physical interface properties at the main router **[edit interfaces]** hierarchy level. Next, the example shows how to configure protocols (such as RSVP, MPLS, BGP, and IS-IS), routing options, and policy options for the logical systems. Last, the example shows how to configure the same administrator for Logical System LS1 that is configured on Router PE1. This system administrator for Logical System LS2 has permission to view the LS2 configuration, but not change the configuration for Logical System LS2.

Logical System LS1 transports traffic for the red VPN that exists between routers CE1 and CE5. Logical System LS1 also connects the blue VPN that exists between routers CE2 and CE6. Logical System LS2 transports VPLS traffic between routers CE3 and CE7. For the main router on Router P0, you can configure the router as usual. The main router transports traffic between routers CE4 and CE8. The example shows how to configure the interfaces and routing protocols (OSPF, BGP) to connect to the main router portion of routers PE1 and PE2.

### Configuration

To configure the PE and P routers in logical systems involves performing the following tasks:

- [Configuring Interfaces on the Customer Edge Devices on page 3435](#)
- [Configuring Router PE1 on page 3437](#)
- [Configuring Router PE2 on page 3440](#)
- [Configuring Router P0 on page 3442](#)
- [Results on page 3444](#)

### Configuring Interfaces on the Customer Edge Devices

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

1. On Router CE1, configure OSPF to connect to the red VPN in Logical System LS1 on Router PE1.
 

```

user@CE1# set interfaces fe-1/0/1 vlan-tagging
user@CE1# set interfaces fe-1/0/1 unit 0 description "routing-instance red CE"
user@CE1# set interfaces fe-1/0/1 unit 0 vlan-id 101
user@CE1# set interfaces fe-1/0/1 unit 0 family inet address 10.11.1.24
user@CE1# set interfaces lo0 unit 0 family inet address 10.11.1.100/32
user@CE1# set protocols ospf area 0.0.0.0 interface fe-1/0/1.0
user@CE1# set protocols ospf area 0.0.0.0 interface lo0.0

```
2. On Router CE2, configure BGP to connect to the blue VPN in Logical System LS1 on Router PE1.

```

user@CE2# set interfaces fe-1/0/2 vlan-tagging
user@CE2# set interfaces fe-1/0/2 unit 0 description "routing-instance blue CE"
user@CE2# set interfaces fe-1/0/2 unit 0 vlan-id 102
user@CE2# set interfaces fe-1/0/2 unit 0 family inet address 10.21.1.1/24
user@CE2# set interfaces lo0 unit 0 family inet address 10.21.1.100/32
user@CE2# set policy-options policy-statement export_loopback from route-filter
10.21.1.100/32 exact
user@CE2# set policy-options policy-statement export_loopback then accept
user@CE2# set protocols bgp export export_loopback
user@CE2# set protocols bgp group to_PE type external
user@CE2# set protocols bgp group to_PE local-address 10.21.1.1
user@CE2# set protocols bgp group to_PE peer-as 100
user@CE2# set protocols bgp group to_PE neighbor 10.21.1.2
user@CE2# set routing-options autonomous-system 200

```

3. On Router CE3, configure the Fast Ethernet interface in VLAN 600 to connect with the VPLS routing instance in Logical System LS2 on Router PE1.

```

user@CE3# set interfaces fe-1/0/0 vlan-tagging
user@CE3# set interfaces fe-1/0/0 unit 0 description "vpls interface"
user@CE3# set interfaces fe-1/0/0 unit 0 vlan-id 600
user@CE3# set interfaces fe-1/0/0 unit 0 family inet address 10.1.1.1/24

```

4. On Router CE4, configure the Fast Ethernet interface to connect with the main router at Router PE1.

```

user@CE4# set interfaces fe-1/0/3 vlan-tagging
user@CE4# set interfaces fe-1/0/3 unit 0 description "main router interface"
user@CE4# set interfaces fe-1/0/3 unit 0 vlan-id 103
user@CE4# set interfaces fe-1/0/3 unit 0 family inet address 10.31.1.1/24
user@CE4# set interfaces lo0 unit 0 family inet address 10.255.41.177/32

```

5. On Router CE5, configure OSPF to connect to the red VPN in Logical System LS1 on Router PE2.

```

user@CE5# set interfaces fe-0/3/1 vlan-tagging
user@CE5# set interfaces fe-0/3/1 unit 0 description "routing-instance red CE"
user@CE5# set interfaces fe-0/3/1 unit 0 vlan-id 101
user@CE5# set interfaces fe-0/3/1 unit 0 family inet address 10.11.4.2/24
user@CE5# set interfaces lo0 unit 0 family inet address 10.11.4.100/32
user@CE5# set protocols ospf area 0.0.0.0 interface fe-0/3/1.0
user@CE5# set protocols ospf area 0.0.0.0 interface lo0.0
user@CE5# set system login class LS1admin logical-system LS1
user@CE5# set system login class LS1admin permissions all
user@CE5# set system login class LS1onlooker logical-system LS2
user@CE5# set system login class LS1onlooker permissions view
user@CE5# set system login user LS1admin class LS1admin

```

6. On Router CE6, configure BGP to connect to the blue VPN in Logical System LS1 on Router PE2.

```

user@CE6# set interfaces fe-0/3/2 vlan-tagging
user@CE6# set interfaces fe-0/3/2 unit 0 description "routing-instance blue CE"
user@CE6# set interfaces fe-0/3/2 unit 0 vlan-id 102
user@CE6# set interfaces fe-0/3/2 unit 0 family inet address 10.21.4.2/24
user@CE6# set interfaces lo0 unit 0 family inet address 10.21.4.100/32
user@CE6# set routing-options autonomous-system 300
user@CE6# set protocols bgp export export_loopback

```

```

user@CE6# set protocols bgp group to_PE type external
user@CE6# set protocols bgp group to_PE local-address 10.21.4.2
user@CE6# set protocols bgp group to_PE peer-as 100
user@CE6# set protocols bgp group to_PE neighbor 10.21.4.1
user@CE6# set policy-options policy-statement export_loopback from route-filter
10.21.4.100/32 exact
user@CE6# set policy-options policy-statement export_loopback then accept

```

7. On Router CE7, configure the Fast Ethernet interface in VLAN 600 to connect with the VPLS routing instance in Logical System LS2 on Router PE2.

```

user@CE7# set interfaces fe-0/3/0 vlan-tagging
user@CE7# set interfaces fe-0/3/0 unit 0 description "vpls interface"
user@CE7# set interfaces fe-0/3/0 unit 0 vlan-id 600
user@CE7# set interfaces fe-0/3/0 unit 0 family inet address 10.1.1.2/24

```

8. On Router CE8, configure the Fast Ethernet interface to connect with the main router at Router PE2.

```

user@CE8# set interfaces fe-0/3/3 vlan-tagging
user@CE8# set interfaces fe-0/3/3 unit 0 description "main router interface"
user@CE8# set interfaces fe-0/3/3 unit 0 vlan-id 103
user@CE8# set interfaces fe-0/3/3 unit 0 family inet address 10.31.4.2/24
user@CE8# set interfaces lo0 unit 0 family inet address 10.255.41.180/32

```

### Configuring Router PE1

#### Step-by-Step Procedure

1. Configure the main router on Router PE1.

```

user@PE1# set interfaces fe-0/0/1 vlan-tagging
user@PE1# set interfaces fe-0/0/1 unit 3 description "main router to P0"
user@PE1# set interfaces fe-0/0/1 unit 3 vlan-id 103
user@PE1# set interfaces fe-0/0/1 unit 3 family inet address 10.31.2.1/24
user@PE1# set interfaces fe-0/0/1 unit 3 family iso
user@PE1# set interfaces fe-0/0/1 unit 3 family mpls
user@PE1# set interfaces fe-0/1/0 vlan-tagging
user@PE1# set interfaces fe-0/1/0 encapsulation vlan-vpls
user@PE1# set interfaces fe-0/1/1 vlan-tagging
user@PE1# set interfaces fe-0/1/2 vlan-tagging
user@PE1# set interfaces fe-0/1/3 vlan-tagging
user@PE1# set interfaces fe-0/1/3 unit 0 description "main router to CE4"
user@PE1# set interfaces fe-0/1/3 unit 0 vlan-id 103
user@PE1# set interfaces fe-0/1/3 unit 0 family inet address 10.31.1.2/24
user@PE1# set interfaces lo0 unit 0 description "main router loopback"
user@PE1# set interfaces lo0 unit 0 family inet address 10.255.41.173/32
user@PE1# set protocols bgp group to_main_ls type internal
user@PE1# set protocols bgp group to_main_ls local-address 10.255.41.173
user@PE1# set protocols bgp group to_main_ls export export_address
user@PE1# set protocols bgp group to_main_ls neighbor 10.255.41.179
user@PE1# set protocols bgp group to_main_ls neighbor 10.255.41.175
user@PE1# set protocols ospf area 0.0.0.0 interface lo0.0
user@PE1# set protocols ospf area 0.0.0.0 interface fe-0/0/1.3
user@PE1# set routing-options static route 10.255.41.177/32 next-hop 10.31.1.1
user@PE1# set routing-options autonomous-system 500
user@PE1# set policy-options policy-statement export_address from route-filter
10.255.41.177/32 exact
user@PE1# set policy-options policy-statement export_address then accept

```

```
user@PE1# set system login class LS1-admin logical-system LS1
user@PE1# set system login class LS1-admin permissions all
user@PE1# set system login user LS1-admin class LS1-admin
user@PE1# set system login user LS1-admin authentication plain-text-password
New password:
Retype new password:
```

2. Configure Logical System LS1 on Router PE1.

```
user@PE1# set logical-systems LS1 interfaces fe-0/0/1 unit 1 description "LS1
interface"
user@PE1# set logical-systems LS1 interfaces fe-0/0/1 unit 1 vlan-id 101
user@PE1# set logical-systems LS1 interfaces fe-0/0/1 unit 1 family inet address
10.11.2.1/24
user@PE1# set logical-systems LS1 interfaces fe-0/0/1 unit 1 family iso
user@PE1# set logical-systems LS1 interfaces fe-0/0/1 unit 1 family mpls
user@PE1# set logical-systems LS1 interfaces fe-0/1/1 unit 0 description
"routing-instance red interface"
user@PE1# set logical-systems LS1 interfaces fe-0/1/1 unit 0 vlan-id 101
user@PE1# set logical-systems LS1 interfaces fe-0/1/1 unit 0 family inet address
10.11.1.2/24
user@PE1# set logical-systems LS1 interfaces fe-0/1/2 unit 0 description
"routing-instance blue interface"
user@PE1# set logical-systems LS1 interfaces fe-0/1/2 unit 0 vlan-id 102
user@PE1# set logical-systems LS1 interfaces fe-0/1/2 unit 0 family inet address
10.21.1.2/24
user@PE1# set logical-systems LS1 interfaces lo0 unit 1 description "LS1 loopback"
user@PE1# set logical-systems LS1 interfaces lo0 unit 1 family inet address
10.10.10.10/32
user@PE1# set logical-systems LS1 interfaces lo0 unit 1 family iso address
47.1111.1111.1111.00
user@PE1# set logical-systems LS1 protocols rsvp interface all
user@PE1# set logical-systems LS1 protocols mpls label-switched-path to_10.10.10.12
to 10.10.10.12
user@PE1# set logical-systems LS1 protocols mpls interface all
user@PE1# set logical-systems LS1 protocols bgp group to_other_PE type internal
user@PE1# set logical-systems LS1 protocols bgp group to_other_PE local-address
10.10.10.10
user@PE1# set logical-systems LS1 protocols bgp group to_other_PE family inet-vpn
any
user@PE1# set logical-systems LS1 protocols bgp group to_other_PE neighbor
10.10.10.12
user@PE1# set logical-systems LS1 protocols isis interface all
user@PE1# set logical-systems LS1 policy-options policy-statement
from_bgp_to_ospf then accept
user@PE1# set logical-systems LS1 routing-instances blue instance-type vrf
user@PE1# set logical-systems LS1 routing-instances blue interface fe-0/1/2.0
user@PE1# set logical-systems LS1 routing-instances blue route-distinguisher
10.10.10.10:200
user@PE1# set logical-systems LS1 routing-instances blue vrf-target target:20:20
user@PE1# set logical-systems LS1 routing-instances blue protocols bgp group
to_CE type external
user@PE1# set logical-systems LS1 routing-instances blue protocols bgp group
to_CE local-address 10.21.1.2
user@PE1# set logical-systems LS1 routing-instances blue protocols bgp group
to_CE peer-as 200
```

```

user@PE1# set logical-systems LS1 routing-instances blue protocols bgp group
to_CE neighbor 10.21.1.1
user@PE1# set logical-systems LS1 routing-instances red instance-type vrf
user@PE1# set logical-systems LS1 routing-instances red interface fe-0/1/1.0
user@PE1# set logical-systems LS1 routing-instances red route-distinguisher
10.10.10.10:100
user@PE1# set logical-systems LS1 routing-instances red vrf-target target:10:10
user@PE1# set logical-systems LS1 routing-instances red protocols ospf export
from_bgp_to_ospf
user@PE1# set logical-systems LS1 routing-instances red protocols ospf area 0.0.0.0
interface all
user@PE1# set logical-systems LS1 routing-options autonomous-system 100

```

### 3. Configure Logical System LS2 on Router PE1.

```

user@PE1# set logical-systems LS2 interfaces fe-0/0/1 unit 0 description
"core-facing LS2 interface"
user@PE1# set logical-systems LS2 interfaces fe-0/0/1 unit 0 vlan-id 100
user@PE1# set logical-systems LS2 interfaces fe-0/0/1 unit 0 family inet address
10.1.2.1/24
user@PE1# set logical-systems LS2 interfaces fe-0/0/1 unit 0 family iso
user@PE1# set logical-systems LS2 interfaces fe-0/0/1 unit 0 family mpls
user@PE1# set logical-systems LS2 interfaces fe-0/1/0 unit 0 description "vpls
interface to ce3"
user@PE1# set logical-systems LS2 interfaces fe-0/1/0 unit 0 encapsulation
vlan-vpls
user@PE1# set logical-systems LS2 interfaces fe-0/1/0 unit 0 vlan-id 600
user@PE1# set logical-systems LS2 interfaces fe-0/1/0 unit 0 family vpls
user@PE1# set logical-systems LS2 interfaces lo0 unit 2 description "LS2 loopback"
user@PE1# set logical-systems LS2 interfaces lo0 unit 2 family inet address
10.20.20.20/32
user@PE1# set logical-systems LS2 interfaces lo0 unit 2 family iso address
47.2222.2222.2222.00
user@PE1# set logical-systems LS2 protocols rsvp interface all
user@PE1# set logical-systems LS2 protocols mpls label-switched-path
to_10.20.20.22 to 10.20.20.22
user@PE1# set logical-systems LS2 protocols mpls interface all
user@PE1# set logical-systems LS2 protocols bgp group to_VPLS_PE type internal
user@PE1# set logical-systems LS2 protocols bgp group to_VPLS_PE local-address
10.20.20.20
user@PE1# set logical-systems LS2 protocols bgp group to_VPLS_PE family l2vpn
signaling
user@PE1# set logical-systems LS2 protocols bgp group to_VPLS_PE neighbor
10.20.20.22
user@PE1# set logical-systems LS2 protocols isis interface fe-0/0/1.0
user@PE1# set logical-systems LS2 protocols isis interface lo0.2
user@PE1# set logical-systems LS2 routing-instances new instance-type vpls
user@PE1# set logical-systems LS2 routing-instances new interface fe-0/1/0.0
user@PE1# set logical-systems LS2 routing-instances new route-distinguisher
10.20.20.20:100
user@PE1# set logical-systems LS2 routing-instances new vrf-target target:30:30
user@PE1# set logical-systems LS2 routing-instances new protocols vpls site-range
10
user@PE1# set logical-systems LS2 routing-instances new protocols vpls site newPE
site-identifier 1
user@PE1# set logical-systems LS2 routing-options autonomous-system 400

```

*Configuring Router PE2***Step-by-Step  
Procedure**

1. Configure the main router on Router PE2.

```

user@PE2# set interfaces fe-0/2/0 vlan-tagging
user@PE2# set interfaces fe-0/2/0 encapsulation vlan-vpls
user@PE2# set interfaces fe-0/2/1 vlan-tagging
user@PE2# set interfaces fe-0/2/2 vlan-tagging
user@PE2# set interfaces fe-0/2/3 vlan-tagging
user@PE2# set interfaces fe-0/2/3 unit 0 description "main router to CE8"
user@PE2# set interfaces fe-0/2/3 unit 0 vlan-id 103
user@PE2# set interfaces fe-0/2/3 unit 0 family inet address 10.31.4.1/24
user@PE2# set interfaces so-1/2/0 encapsulation frame-relay
user@PE2# set interfaces so-1/2/0 unit 3 description "main router to P0"
user@PE2# set interfaces so-1/2/0 unit 3 dlci 103
user@PE2# set interfaces so-1/2/0 unit 3 family inet address 10.31.3.2/24
user@PE2# set interfaces so-1/2/0 unit 3 family iso
user@PE2# set interfaces so-1/2/0 unit 3 family mpls
user@PE2# set interfaces lo0 unit 0 description "main router loopback"
user@PE2# set interfaces lo0 unit 0 family inet address 10.155.41.179/32
user@PE2# set protocols bgp group to_main_ls type internal
user@PE2# set protocols bgp group to_main_ls local-address 10.255.41.179
user@PE2# set protocols bgp group to_main_ls export export_address
user@PE2# set protocols bgp group to_main_ls neighbor 10.255.41.173
user@PE2# set protocols bgp group to_main_ls neighbor 10.255.41.175
user@PE2# set protocols ospf area 0.0.0.0 interface so-1/2/0.3
user@PE2# set protocols ospf area 0.0.0.0 interface fe-0/2/3.0
user@PE2# set protocols ospf area 0.0.0.0 interface lo0.0
user@PE2# set routing-options static route 10.255.41.180/32 next-hop 10.31.4.2
user@PE2# set routing-options autonomous-system 500
user@PE2# set policy-options policy-statement export_address from route-filter
10.255.41.180/32 exact
user@PE2# set policy-options policy-statement export_address then accept
user@PE2# set system login class LS1-admin logical-system LS1
user@PE2# set system login class LS1-admin permissions all
user@PE2# set system login class LS1-onlooker logical-system LS2
user@PE2# set system login class LS1-onlooker permissions view
user@PE2# set system login user LS1-admin class LS1-admin

```

2. Configure Logical System LS1 on Router PE2.

```

user@PE2# set logical-systems LS1 interfaces fe-0/2/0 unit 1 description
"routing-instance red interface connects to Router CE5"
user@PE2# set logical-systems LS1 interfaces fe-0/2/0 unit 1 vlan-id 101
user@PE2# set logical-systems LS1 interfaces fe-0/2/0 unit 1 family inet address
10.11.4.1/24
user@PE2# set logical-systems LS1 interfaces fe-0/2/0 unit 2 description
"routing-instance blue interface connects to Router CE6"
user@PE2# set logical-systems LS1 interfaces fe-0/2/0 unit 2 vlan-id 102
user@PE2# set logical-systems LS1 interfaces fe-0/2/0 unit 2 family inet address
10.21.4.1/24
user@PE2# set logical-systems LS1 interfaces so-1/2/0 unit 1 description "core-facing
LS1 interface"
user@PE2# set logical-systems LS1 interfaces so-1/2/0 unit 1 dlci 101
user@PE2# set logical-systems LS1 interfaces so-1/2/0 unit 1 family inet address
10.11.3.2/24

```

```

user@PE2# set logical-systems LS1 interfaces so-1/2/0 unit 1 family iso
user@PE2# set logical-systems LS1 interfaces so-1/2/0 unit 1 family mpls
user@PE2# set logical-systems LS1 interfaces lo0 unit 1 description "LS1 loopback"
user@PE2# set logical-systems LS1 interfaces lo0 unit 1 family inet address
10.10.10.12/32
user@PE2# set logical-systems LS1 interfaces lo0 unit 1 family iso address
47.1111.1111.1111.1113.00
user@PE2# set logical-systems LS1 protocols rsvp interface all
user@PE2# set logical-systems LS1 protocols mpls label-switched-path
to_10.10.10.10 to 10.10.10.10
user@PE2# set logical-systems LS1 protocols mpls interface all
user@PE2# set logical-systems LS1 protocols bgp group to_other_PE type internal
user@PE2# set logical-systems LS1 protocols bgp group to_other_PE local-address
10.10.10.12
user@PE2# set logical-systems LS1 protocols bgp group to_other_PE family inet
any
user@PE2# set logical-systems LS1 protocols bgp group to_other_PE family inet-vpn
any
user@PE2# set logical-systems LS1 protocols bgp group to_other_PE neighbor
10.10.10.10
user@PE2# set logical-systems LS1 protocols isis interface all
user@PE2# set logical-systems LS1 policy-options policy-statement
from_bgp_to_ospf then accept
user@PE2# set logical-systems LS1 routing-instances blue instance-type vrf
user@PE2# set logical-systems LS1 routing-instances blue interface fe-0/2/2.0
user@PE2# set logical-systems LS1 routing-instances blue route-distinguisher
10.10.10.12:200
user@PE2# set logical-systems LS1 routing-instances blue vrf-target target:20:20
user@PE2# set logical-systems LS1 routing-instances blue protocols bgp group
to_CE local-address 10.21.4.1
user@PE2# set logical-systems LS1 routing-instances blue protocols bgp group
to_CE peer-as 300
user@PE2# set logical-systems LS1 routing-instances blue protocols bgp group
to_CE neighbor 10.21.4.2
user@PE2# set logical-systems LS1 routing-instances red instance-type vrf
user@PE2# set logical-systems LS1 routing-instances red interface fe-0/2/1.0
user@PE2# set logical-systems LS1 routing-instances red route-distinguisher
10.10.10.12:100
user@PE2# set logical-systems LS1 routing-instances red vrf-target target:10:10
user@PE2# set logical-systems LS1 routing-instances red protocols ospf export
from_bgp_to_ospf
user@PE2# set logical-systems LS1 routing-instances red protocols ospf area 0.0.0.0
interface all
user@PE2# set logical-systems LS1 routing-options autonomous-system 100

```

3. Configure Logical System LS2 on Router PE2.

```

user@PE2# set logical-systems LS2 interfaces fe-0/2/0 unit 0 description "vpls
interface connects to Router CE7"
user@PE2# set logical-systems LS2 interfaces fe-0/2/0 unit 0 encapsulation
vlan-vpls
user@PE2# set logical-systems LS2 interfaces fe-0/2/0 unit 0 vlan-id 600
user@PE2# set logical-systems LS2 interfaces fe-0/2/0 unit 0 family vpls
user@PE2# set logical-systems LS2 interfaces so-1/2/0 unit 0 description
"core-facing LS2 interface"
user@PE2# set logical-systems LS2 interfaces so-1/2/0 unit 0 dlci 100

```

```

user@PE2# set logical-systems LS2 interfaces so-1/2/0 unit 0 family inet address
10.1.3.2/24
user@PE2# set logical-systems LS2 interfaces so-1/2/0 unit 0 family iso
user@PE2# set logical-systems LS2 interfaces so-1/2/0 unit 0 family mpls
user@PE2# set logical-systems LS2 interfaces lo0 unit 2 description "LS2 loopback"
user@PE2# set logical-systems LS2 interfaces lo0 unit 2 family inet address
10.20.20.22/32
user@PE2# set logical-systems LS2 interfaces lo0 unit 2 family iso address
47.2222.2222.2222.2224.00
user@PE2# set logical-systems LS2 protocols rsvp interface all
user@PE2# set logical-systems LS2 protocols mpls label-switched-path
to_10.20.20.20 to 10.20.20.20
user@PE2# set logical-systems LS2 protocols mpls interface all
user@PE2# set logical-systems LS2 protocols bgp group to_VPLS_PE type internal
user@PE2# set logical-systems LS2 protocols bgp group to_VPLS_PE local-address
10.20.20.22
user@PE2# set logical-systems LS2 protocols bgp group to_VPLS_PE family l2vpn
signaling
user@PE2# set logical-systems LS2 protocols bgp group to_VPLS_PE neighbor
10.20.20.20
user@PE2# set logical-systems LS2 protocols isis interface so-1/2/0.0
user@PE2# set logical-systems LS2 protocols isis interface lo0.2
user@PE2# set logical-systems LS2 routing-instances new instance-type vpls
user@PE2# set logical-systems LS2 routing-instances new interface fe-0/2/0.0
user@PE2# set logical-systems LS2 routing-instances new route-distinguisher
10.20.20.22:100
user@PE2# set logical-systems LS2 routing-instances new vrf-target target:30:30
user@PE2# set logical-systems LS2 routing-instances new protocols vpls site-range
10
user@PE2# set logical-systems LS2 routing-instances new protocols vpls site newPE
site-identifier 2
user@PE2# set logical-systems LS2 routing-options autonomous-system 400

```

### Configuring Router P0

#### Step-by-Step Procedure

1. Configure the main router on Router P0.

```

user@P0# set interfaces fe-1/1/3 vlan-tagging
user@P0# set interfaces fe-1/1/3 unit 3 description "connects to the main router on
pe1"
user@P0# set interfaces fe-1/1/3 unit 3 vlan-id 103
user@P0# set interfaces fe-1/1/3 unit 3 family inet address 10.31.2.2/24
user@P0# set interfaces fe-1/1/3 unit 3 family iso
user@P0# set interfaces fe-1/1/3 unit 3 family mpls
user@P0# set interfaces so-1/2/0 dce
user@P0# set interfaces so-1/2/0 encapsulation frame-relay
user@P0# set interfaces so-1/2/0 unit 3 description "connects to the main router
on pe2"
user@P0# set interfaces so-1/2/0 unit 3 dlci 103
user@P0# set interfaces so-1/2/0 unit 3 family inet address 10.31.3.1/24
user@P0# set interfaces so-1/2/0 unit 3 family iso
user@P0# set interfaces so-1/2/0 unit 3 family mpls
user@P0# set interfaces lo0 unit 0 description "main router loopback"
user@P0# set interfaces lo0 unit 0 family inet address 10.255.41.175/32
user@P0# set routing-options autonomous-system 500
user@P0# set protocols bgp group to_main_ls type internal

```



```

user@P0# set protocols bgp group to_main_ls local-address 10.255.41.175
user@P0# set protocols bgp group to_main_ls neighbor 10.255.41.179
user@P0# set protocols bgp group to_main_ls neighbor 10.255.41.173
user@P0# set protocols ospf area 0.0.0.0 interface lo0.0
user@P0# set protocols ospf area 0.0.0.0 interface fe-1/1/3.3
user@P0# set protocols ospf area 0.0.0.0 interface so-1/2/0.3
user@P0# set system login class LS1-admin logical-system LS1
user@P0# set system login class LS1-admin permissions all
user@P0# set system login class LS1-onlooker logical-system LS2
user@P0# set system login class LS1-onlooker permissions view
user@P0# set system login user LS1-admin class LS1-admin

```

2. Configure Logical System LS1 on Router P0.

```

user@P0# set logical-systems LS1 interfaces fe-1/1/3 unit 1 description "LS1 interface
connects to LS1 on pe1"
user@P0# set logical-systems LS1 interfaces fe-1/1/3 unit 1 vlan-id 101
user@P0# set logical-systems LS1 interfaces fe-1/1/3 unit 1 family inet address
10.11.2.2/24
user@P0# set logical-systems LS1 interfaces fe-1/1/3 unit 1 family iso
user@P0# set logical-systems LS1 interfaces fe-1/1/3 unit 1 family mpls
user@P0# set logical-systems LS1 interfaces so-1/2/0 unit 1 description "LS1
interface connects to LS1 on pe2"
user@P0# set logical-systems LS1 interfaces so-1/2/0 unit 1 dlci 101
user@P0# set logical-systems LS1 interfaces so-1/2/0 unit 1 family inet address
10.11.3.1/24
user@P0# set logical-systems LS1 interfaces so-1/2/0 unit 1 family iso
user@P0# set logical-systems LS1 interfaces so-1/2/0 unit 1 family mpls
user@P0# set logical-systems LS1 interfaces lo0 unit 1 description "LS1 loopback"
user@P0# set logical-systems LS1 interfaces lo0 unit 1 family inet address
10.10.10.11/32
user@P0# set logical-systems LS1 interfaces lo0 unit 1 family iso address
47.1111.1111.1111.1112.00
user@P0# set logical-systems LS1 protocols rsvp interface all
user@P0# set logical-systems LS1 protocols mpls interface all
user@P0# set logical-systems LS1 protocols isis interface all

```

3. Configure Logical System LS2 on Router P0.

```

user@P0# set logical-systems LS2 interfaces fe-1/1/3 unit 0 description "LS2
interface connects to LS2 on pe1"
user@P0# set logical-systems LS2 interfaces fe-1/1/3 unit 0 vlan-id 100
user@P0# set logical-systems LS2 interfaces fe-1/1/3 unit 0 family inet address
10.1.2.2/24
user@P0# set logical-systems LS2 interfaces fe-1/1/3 unit 0 family iso
user@P0# set logical-systems LS2 interfaces fe-1/1/3 unit 0 family mpls
user@P0# set logical-systems LS2 interfaces so-1/2/0 unit 0 description "LS2
interface connects to LS2 on pe2"
user@P0# set logical-systems LS2 interfaces so-1/2/0 unit 0 dlci 100
user@P0# set logical-systems LS2 interfaces so-1/2/0 unit 0 family inet address
10.1.3.1/24
user@P0# set logical-systems LS2 interfaces so-1/2/0 unit 0 family iso
user@P0# set logical-systems LS2 interfaces so-1/2/0 unit 0 family mpls
user@P0# set logical-systems LS2 interfaces lo0 unit 2 description "LS2 loopback"
user@P0# set logical-systems LS2 interfaces lo0 unit 2 family inet address
10.20.20.21/32

```

```
user@P0# set logical-systems LS2 interfaces lo0 unit 2 family iso address
47.2222.2222.2222.00
user@P0# set logical-systems LS2 protocols rsvp interface all
user@P0# set logical-systems LS2 protocols mpls interface all
user@P0# set logical-systems LS2 protocols isis interface fe-1/1/3.0
user@P0# set logical-systems LS2 protocols isis interface so-1/2/0.0
user@P0# set logical-systems LS2 protocols isis interface lo0.2
```

### Results

On Router CE1, configure OSPF to connect to the red VPN in Logical System LS1 on Router PE1:

```
Router CE1 [edit]
interfaces {
 fe-1/0/1 {
 vlan-tagging;
 unit 0 {
 description "routing-instance red CE";
 vlan-id 101;
 family inet {
 address 10.11.1.1/24;
 }
 }
 }
 lo0 {
 unit 0 {
 family inet {
 address 10.11.1.100/32;
 }
 }
 }
}
protocols {
 ospf {
 area 0.0.0.0 {
 interface fe-1/0/1.0;
 interface lo0.0;
 }
 }
}
```

On Router CE2, configure BGP to connect to the blue VPN in Logical System LS1 on Router PE1:

```
Router CE2 [edit]
interfaces {
 fe-1/0/2 {
 vlan-tagging;
 unit 0 {
 description "routing-instance blue CE";
 vlan-id 102;
 family inet {
 address 10.21.1.1/24;
 }
 }
 }
}
```

```

 }
 lo0 {
 unit 0 {
 family inet {
 address 10.21.1.100/32;
 }
 }
 }
}
routing-options {
 autonomous-system 200;
}
protocols {
 bgp {
 export export_loopback;
 group to_PE {
 type external;
 local-address 10.21.1.1;
 peer-as 100;
 neighbor 10.21.1.2;
 }
 }
}
policy-options {
 policy-statement export_loopback {
 from {
 route-filter 10.21.1.100/32 exact;
 }
 then accept;
 }
}

```

On Router CE3, configure the Fast Ethernet interface in VLAN 600 to connect with the VPLS routing instance in Logical System LS2 on Router PE1:

```

Router CE3 [edit]
interfaces {
 fe-1/0/0 {
 vlan-tagging;
 unit 0 {
 description "vpls interface";
 vlan-id 600;
 family inet {
 address 10.1.1.1/24;
 }
 }
 }
}

```

On Router CE4, configure the Fast Ethernet interface to connect with the main router at Router PE1:

```

Router CE4 [edit]
interfaces {
 fe-1/0/3 {
 vlan-tagging;
 }
}

```

```
 unit 0 {
 description "main router interface";
 vlan-id 103;
 family inet {
 address 10.31.1.1/24;
 }
 }
 }
 lo0 {
 unit 0 {
 family inet {
 address 10.255.41.177/32;
 }
 }
 }
}
```

On Router PE1, create two VPN routing and forwarding (VRF) routing instances in Logical System LS1: red and blue. Configure the CE-facing logical interfaces so that traffic from Router CE1 is placed in the red VPN, and traffic from Router CE2 is placed in the blue VPN. Next, create a logical interface at **fe-0/0/1.1** to connect to Logical System LS1 on Router P0.

Also on Router PE1, create a VPLS routing instance in Logical System LS2. Configure a logical interface so that traffic from Router CE3 is sent into the VPLS domain and connects to Logical System LS2 on Router P0.

Create an administrator for Logical System LS1. The logical system administrator can be responsible for the maintenance of this logical system.

Finally, configure a logical interface to interconnect Router CE4 with the main router portion of Router P0.

```
Router PE1 [edit]
logical-systems {
 LS1 { # The configuration for the first logical system begins here.
 interfaces {
 fe-0/0/1 {
 unit 1 { # This is the core-facing interface for Logical System LS1.
 description "LS1 interface";
 vlan-id 101;
 family inet {
 address 10.11.2.1/24;
 }
 family iso;
 family mpls;
 }
 }
 }
 fe-0/1/1 {
 unit 0 { # This logical interface connects to Router CE1.
 description "routing-instance red interface";
 vlan-id 101;
 family inet {
 address 10.11.1.2/24;
 }
 }
 }
 }
}
```

```

 }
 }
 fe-0/1/2 {
 unit 0 { # This logical interface connects to Router CE2.
 description "routing-instance blue interface";
 vlan-id 102;
 family inet {
 address 10.21.1.2/24;
 }
 }
 }
 lo0 {
 unit 1 {
 description "LS1 loopback";
 family inet {
 address 10.10.10.10/32;
 }
 family iso {
 address 47.1111.1111.1111.00;
 }
 }
 }
}
protocols { # You configure RSVP, MPLS, IS-IS, and BGP for Logical System LS1.
 rsvp {
 interface all;
 }
 mpls {
 label-switched-path to_10.10.10.12 {
 to 10.10.10.12;
 }
 interface all;
 }
 bgp {
 group to_other_PE {
 type internal;
 local-address 10.10.10.10;
 family inet-vpn {
 any;
 }
 neighbor 10.10.10.12;
 }
 }
 isis {
 interface all;
 }
}
policy-options {
 policy-statement from_bgp_to_ospf {
 then accept;
 }
}
routing-instances {
 blue {
 instance-type vrf; # You configure instance blue within Logical System LS1.
 interface fe-0/1/2.0;
 }
}

```

```
route-distinguisher 10.10.10.10:200;
vrf-target target:20:20;
protocols {
 bgp { #BGP connects the blue instance with Router CE2.
 group to_CE {
 type external;
 local-address 10.21.1.2;
 peer-as 200;
 neighbor 10.21.1.1;
 }
 }
}
red {
 instance-type vrf; # You configure instance red within Logical System LS1.
 interface fe-0/1/1.0;
 route-distinguisher 10.10.10.10:100;
 vrf-target target:10:10;
 protocols {
 ospf {#OSPF connects the red instance with Router CE1.
 export from_bgp_to_ospf;
 area 0.0.0.0 {
 interface all;
 }
 }
 }
}
routing-options {
 autonomous-system 100;
}
}
LS2 { # The configuration for the second logical system begins here.
 interfaces {
 fe-0/0/1 {
 unit 0 { # This is the core-facing interface for Logical System LS2.
 description "LS2 interface";
 vlan-id 100;
 family inet {
 address 10.1.2.1/24;
 }
 family iso;
 family mpls;
 }
 }
 fe-0/1/0 {
 unit 0 { # This logical interface connects to Router CE3.
 description "vpls interface";
 encapsulation vlan-vpls;
 vlan-id 600;
 family vpls;
 }
 }
 lo0 {
 unit 2 {
 description "LS2 loopback";
```

```

 family inet {
 address 10.20.20.20/32;
 }
 family iso {
 address 47.2222.2222.2222.00;
 }
 }
}
protocols { # You configure RSVP, MPLS, IS-IS, and BGP for Logical System LS2.
 rsvp {
 interface all;
 }
 mpls {
 label-switched-path to_10.20.20.22 {
 to 10.20.20.22;
 }
 interface all;
 }
 bgp {
 group to_VPLS_PE {
 type internal;
 local-address 10.20.20.20;
 family l2vpn {
 signaling;
 }
 neighbor 10.20.20.22;
 }
 }
 isis {
 interface fe-0/0/1.0;
 interface lo0.2;
 }
}
routing-instances {
 new {
 instance-type vpls; # You configure VPLS within Logical System LS2.
 interface fe-0/1/0.0;
 route-distinguisher 10.20.20.20:100;
 vrf-target target:30:30;
 protocols {
 vpls {
 site-range 10;
 site newPE {
 site-identifier 1;
 }
 }
 }
 }
}
routing-options {
 autonomous-system 400;
}
}
interfaces {

```

```
fe-0/0/1 {
 vlan-tagging;
 unit 3 { # This is the core-facing interface for the main router of PE1.
 description "main router to P0";
 vlan-id 103;
 family inet {
 address 10.31.2.1/24;
 }
 family iso;
 family mpls;
 }
}
fe-0/1/3 {
 vlan-tagging;
 unit 0 { # This logical interface in the main router of PE1 connects to CE4.
 description "main router to CE4";
 vlan-id 103;
 family inet {
 address 10.31.1.2/24;
 }
 }
}
fe-0/1/0 { # You must always configure physical interface statements for
 vlan-tagging; # logical system interfaces at the [edit interfaces] hierarchy level.
 encapsulation vlan-vpls;
}
fe-0/1/1 {
 vlan-tagging;
}
fe-0/1/2 {
 vlan-tagging;
}
lo0 {
 unit 0 {
 description "main router loopback";
 family inet {
 address 10.255.41.173/32;
 }
 }
}
}
routing-options {
 static {
 route 10.255.41.177/32 next-hop 10.31.1.1;
 }
 autonomous-system 500;
}
protocols {
 bgp { # The main router uses BGP as the exterior gateway protocol.
 group to_main_ls {
 type internal;
 local-address 10.255.41.173;
 export export_address;
 neighbor 10.255.41.179;
 neighbor 10.255.41.175;
 }
 }
}
```



```

 }
 ospf { # The main router uses OSPF as the interior gateway protocol.
 area 0.0.0.0 {
 interface lo0.0;
 interface fe-0/0/1.3;
 }
 }
}
policy-options {
 policy-statement export_address {
 from {
 route-filter 10.255.41.177/32 exact;
 }
 then accept;
 }
}
system {
 login {
 class LS1-admin {
 permissions all;
 logical-system LS1;
 }
 user LS1-admin {
 class LS1-admin;
 authentication plain-text password;
 New password: password
 Retype new password: password
 }
 }
}
}

```

On Router P0, configure Logical Systems LS1, LS2, and the main router. For the logical system, you must configure physical interface properties at the main router **[edit interfaces]** hierarchy level and assign the logical interfaces to the logical systems. Next, you must configure protocols (such as RSVP, MPLS, BGP, and IS-IS), routing options, and policy options for the logical systems. Last, configure the same administrator for Logical System LS1 that you configured on Router PE1. Configure this same administrator for Logical System LS2 to have permission to view the LS2 configuration, but not change the configuration for LS2.

In this example, Logical System LS1 transports traffic for the red VPN that exists between routers CE1 and CE5. Logical System LS1 also connects the blue VPN that exists between routers CE2 and CE6. Logical System LS2 transports VPLS traffic between routers CE3 and CE7.

For the main router on Router P0, you can configure the router as usual. In this example, the main router transports traffic between routers CE4 and CE8. As a result, configure the interfaces and routing protocols (OSPF, BGP) to connect to the main router portion of routers PE1 and PE2.

```

Router P0 [edit]
logical-systems {
 LS1 { # The configuration for the first logical system begins here.
 interfaces {

```

```
fe-1/1/3 {
 unit 1 { # This logical interface connects to LS1 on Router PE1.
 description "LS1 interface";
 vlan-id 101;
 family inet {
 address 10.11.2.2/24;
 }
 family iso;
 family mpls;
 }
}
so-1/2/0 {
 unit 1 { # This logical interface connects to LS1 on Router PE2.
 description "LS1 interface";
 dlci 101;
 family inet {
 address 10.11.3.1/24;
 }
 family iso;
 family mpls;
 }
}
lo0 {
 unit 1 {
 description "LS1 loopback";
 family inet {
 address 10.10.10.11/32;
 }
 family iso {
 address 47.1111.1111.1111.1112.00;
 }
 }
}
}
protocols { # You configure RSVP, MPLS, and IS-IS for Logical System LS1.
 rsvp {
 interface all;
 }
 mpls {
 interface all;
 }
 isis {
 interface all;
 }
}
}
LS2 { # The configuration for the second logical system begins here.
 interfaces {
 fe-1/1/3 {
 unit 0 { # This logical interface connects to LS2 on Router PE1.
 description "LS2 interface";
 vlan-id 100;
 family inet {
 address 10.1.2.2/24;
 }
 family iso;
 }
 }
 }
}
```

```

 family mpls;
 }
}
so-1/2/0 {
 unit 0 { # This logical interface connects to LS2 on Router PE2.
 description "LS2 interface";
 dlci 100;
 family inet {
 address 10.1.3.1/24;
 }
 family iso;
 family mpls;
 }
}
lo0 {
 unit 2 {
 description "LS2 loopback";
 family inet {
 address 10.20.20.21/32;
 }
 family iso {
 address 47.2222.2222.2223.00;
 }
 }
}
}
protocols { # You configure RSVP, MPLS, and IS-IS for Logical System LS2.
 rsvp {
 interface all;
 }
 mpls {
 interface all;
 }
 isis {
 interface fe-1/1/3.0;
 interface so-1/2/0.0;
 interface lo0.2;
 }
}
}
}
interfaces {
 fe-1/1/3 {
 vlan-tagging;
 unit 3 { # This logical interface connects to the main router on Router PE1.
 description "main router interface";
 vlan-id 103;
 family inet {
 address 10.31.2.2/24;
 }
 family iso;
 family mpls;
 }
 }
}
so-1/2/0 {
 dce; # You must configure all physical interface statements for logical

```

```
encapsulation frame-relay; # routers at the [edit interfaces] hierarchy level.
unit 3 { # This logical interface connects to the main router on Router PE2.
 description "main router interface";
 dlci 103;
 family inet {
 address 10.31.3.1/24;
 }
 family iso;
 family mpls;
}
}
lo0 {
 unit 0 {
 description "main router loopback";
 family inet {
 address 10.255.41.175/32;
 }
 }
}
}
routing-options {
 autonomous-system 500;
}
protocols { # You configure BGP and OSPF for the main router.
 bgp {
 group to_main_ls {
 type internal;
 local-address 10.255.41.175;
 neighbor 10.255.41.179;
 neighbor 10.255.41.173;
 }
 }
 ospf {
 area 0.0.0.0 {
 interface lo0.0;
 interface fe-1/1/3.3;
 interface so-1/2/0.3;
 }
 }
}
system {
 login {
 class LS1-admin {
 permissions all;
 logical-system LS1;
 }
 class LS1-onlooker {
 permissions view;
 logical-system LS2;
 }
 user LS1-admin {
 class LS1-admin;
 }
 }
}
```

On Router PE2, create two VRF routing instances in Logical System LS1: red and blue. Configure the CE-facing logical interfaces so that traffic from Router CE5 is placed in the red VPN and traffic from Router CE6 is placed in the blue VPN. Next, create one logical interface on **so-1/2/0.1** to connect to Logical System LS1 on Router P0.

Also on Router PE2, create a VPLS routing instance in Logical System LS2. Configure a logical interface so that traffic from Router CE7 is sent into the VPLS domain and connects to Logical System LS2 on Router P0.

Configure a logical interface to interconnect Router CE8 with the main router portion of Router P0.

Finally, you can optionally create a logical system administrator that has configuration privileges for Logical System LS1 and viewing privileges for Logical System LS2.

```
Router PE2 [edit]
logical-systems {
 LS1 { # The configuration for the first logical system begins here.
 interfaces {
 fe-0/2/0 {
 unit 1 { # This logical interface connects to Router CE5.
 description "routing-instance red interface";
 vlan-id 101;
 family inet {
 address 10.11.4.1/24;
 }
 }
 unit 2 { # This logical interface connects to Router CE6.
 description "routing-instance blue interface";
 vlan-id 102;
 family inet {
 address 10.21.4.1/24;
 }
 }
 }
 }
 so-1/2/0 {
 unit 1 { # This is the core-facing interface for Logical System LS1.
 description "LS1 interface";
 dlci 101;
 family inet {
 address 10.11.3.2/24;
 }
 family iso;
 family mpls;
 }
 }
 }
 lo0 {
 unit 1 {
 description "LS1 loopback";
 family inet {
 address 10.10.10.12/32;
 }
 family iso {
 address 47.1111.1111.1111.1113.00;
 }
 }
 }
}
```

```
 }
 }
}
protocols {
 rsvp {# You configure RSVP, MPLS, IS-IS, and BGP for Logical System LS1.
 interface all;
 }
 mpls {
 label-switched-path to_10.10.10.10 {
 to 10.10.10.10;
 }
 interface all;
 }
 bgp {
 group to_other_PE {
 type internal;
 local-address 10.10.10.12;
 family inet {
 any;
 }
 family inet-vpn {
 any;
 }
 neighbor 10.10.10.10;
 }
 }
 isis {
 interface all;
 }
}
policy-options {
 policy-statement from_bgp_to_ospf {
 then accept;
 }
}
routing-instances {
 blue {
 instance-type vrf; # You configure instance blue within Logical System LS1.
 interface fe-0/2/2.0;
 route-distinguisher 10.10.10.12:200;
 vrf-target target:20:20;
 protocols {
 bgp { # BGP connects the blue instance with Router CE6.
 group to_CE {
 local-address 10.21.4.1;
 peer-as 300;
 neighbor 10.21.4.2;
 }
 }
 }
 }
 red {
 instance-type vrf; # You configure instance red within Logical System LS1.
 interface fe-0/2/1.0;
 route-distinguisher 10.10.10.12:100;
 vrf-target target:10:10;
```

```

 protocols {
 ospf { # OSPF connects the red instance with Router CE5.
 export from_bgp_to_ospf;
 area 0.0.0.0 {
 interface all;
 }
 }
 }
}
routing-options {
 autonomous-system 100;
}
}
logical-systems {
 LS2 { # The configuration for the second logical system begins here.
 interfaces {
 fe-0/2/0 {
 unit 0 { # This logical interface connects to Router CE7.
 description "vpls interface";
 encapsulation vlan-vpls;
 vlan-id 600;
 family vpls;
 }
 }
 so-1/2/0 {
 unit 0 { # This is the core-facing interface for Logical System LS2.
 description "LS2 interface";
 dlci 100;
 family inet {
 address 10.1.3.2/24;
 }
 family iso;
 family mpls;
 }
 }
 lo0 {
 unit 2 {
 description "LS2 loopback";
 family inet {
 address 10.20.20.22/32;
 }
 family iso {
 address 47.2222.2222.2222.2224.00;
 }
 }
 }
 }
 protocols { # You configure RSVP, MPLS, IS-IS, and BGP for Logical System LS2.
 rsvp {
 interface all;
 }
 mpls {
 label-switched-path to_10.20.20.20 {
 to 10.20.20.20;
 }
 }
 }
 }
}

```

```
 interface all;
 }
 bgp {
 group to_VPLS_PE {
 type internal;
 local-address 10.20.20.22;
 family l2vpn {
 signaling;
 }
 neighbor 10.20.20.20;
 }
 }
 isis {
 interface so-1/2/0.0;
 interface lo0.2;
 }
}
routing-instances {
 new {
 instance-type vpls; # You configure VPLS within Logical System LS2.
 interface fe-0/2/0.0;
 route-distinguisher 10.20.20.22:100;
 vrf-target target:30:30;
 protocols {
 vpls {
 site-range 10;
 site newPE {
 site-identifier 2;
 }
 }
 }
 }
}
routing-options {
 autonomous-system 400;
}
}
interfaces {
 fe-0/2/0 { # You must always configure physical interface statements for the
 vlan-tagging; # logical system interfaces at the [edit interfaces] hierarchy level.
 encapsulation vlan-vpls;
 }
 fe-0/2/1 {
 vlan-tagging;
 }
 fe-0/2/2 {
 vlan-tagging;
 }
 fe-0/2/3 {
 vlan-tagging;
 unit 0 { # This logical interface in the main router of PE2 connects to CE8.
 description "main router to CE8";
 vlan-id 103;
 family inet {
 address 10.31.4.1/24;
 }
 }
 }
}
```



```

 }
 }
 so-1/2/0 {
 encapsulation frame-relay;
 unit 3 { # This is the core-facing interface for the main router of PE2.
 description "main router to P0";
 dlcI 103;
 family inet {
 address 10.31.3.2/24;
 }
 family iso;
 family mpls;
 }
 }
 lo0 {
 unit 0 {
 description "main router loopback";
 family inet {
 address 10.155.41.179/32;
 }
 }
 }
}
routing-options {
 static {
 route 10.255.41.180/32 next-hop 10.31.4.2;
 }
 autonomous-system 500;
}
protocols {
 bgp {# The main router uses BGP as the exterior gateway protocol.
 group to_main_ls {
 type internal;
 local-address 10.255.41.179;
 export export_address;
 neighbor 10.255.41.173;
 neighbor 10.255.41.175;
 }
 }
 ospf {# The main router uses OSPF as the interior gateway protocol.
 area 0.0.0.0 {
 interface so-1/2/0.3;
 interface fe-0/2/3.0;
 interface lo0.0;
 }
 }
}
policy-options {
 policy-statement export_address {
 from {
 route-filter 10.255.41.180/32 exact;
 }
 then accept;
 }
}
}

```

```
system {
 login {
 class LS1-admin {
 permissions all;
 logical-system LS1;
 }
 class LS1-onlooker {
 permissions view;
 logical-system LS2;
 }
 user LS1-admin {
 class LS1-admin;
 }
 }
}
```

On Router CE5, configure OSPF to connect to the red VPN in Logical System LS1 on Router PE2:

```
Router CE5 [edit]
interfaces {
 fe-0/3/1 {
 vlan-tagging;
 unit 0 {
 description "routing-instance red CE";
 vlan-id 101;
 family inet {
 address 10.11.4.2/24;
 }
 }
 }
 lo0 {
 unit 0 {
 family inet {
 address 10.11.4.100/32;
 }
 }
 }
}
protocols {
 ospf {
 area 0.0.0.0 {
 interface fe-0/3/1.0;
 interface lo0.0;
 }
 }
}
system {
 login {
 class LS1-admin {
 permissions all;
 logical-system LS1;
 }
 class LS1-onlooker {
 permissions view;
 logical-system LS2;
 }
 }
}
```

```

 }
 user LS1-admin {
 class LS1-admin;
 }
}

```

On Router CE6, configure BGP to connect to the blue VPN in Logical System LS1 on Router PE2:

```

Router CE6 [edit]
interfaces {
 fe-0/3/2 {
 vlan-tagging;
 unit 0 {
 description "routing-instance blue CE";
 vlan-id 102;
 family inet {
 address 10.21.4.2/24;
 }
 }
 }
 lo0 {
 unit 0 {
 family inet {
 address 10.21.4.100/32;
 }
 }
 }
}
routing-options {
 autonomous-system 300;
}
protocols {
 bgp {
 export export_loopback;
 group to_PE {
 type external;
 local-address 10.21.4.2;
 peer-as 100;
 neighbor 10.21.4.1;
 }
 }
}
policy-options {
 policy-statement export_loopback {
 from {
 route-filter 10.21.4.100/32 exact;
 }
 then accept;
 }
}

```

On Router CE7, configure the Fast Ethernet interface in VLAN 600 to connect with the VPLS routing instance in Logical System LS2 on Router PE2:

```
Router CE7 [edit]
 interfaces {
 fe-0/3/0 {
 vlan-tagging;
 unit 0 {
 description "vpls interface";
 vlan-id 600;
 family inet {
 address 10.1.1.2/24;
 }
 }
 }
 }
```

On Router CE8, configure the Fast Ethernet interface to connect with the main router at Router PE2:

```
Router CE8 [edit]
 interfaces {
 fe-0/3/3 {
 vlan-tagging;
 unit 0 {
 description "main router interface";
 vlan-id 103;
 family inet {
 address 10.31.4.2/24;
 }
 }
 }
 }
 lo0 {
 unit 0 {
 family inet {
 address 10.255.41.180/32;
 }
 }
 }
 }
```

### Verification

Confirm that the configuration is working properly by running these commands:

- `show bgp summary` (logical-system *logical-system-name*)
- `show isis adjacency` (logical-system *logical-system-name*)
- `show mpls lsp` (logical-system *logical-system-name*)
- `show (ospf | ospf3) neighbor` (logical-system *logical-system-name*)
- `show route` (logical-system *logical-system-name*)
- `show route protocol` (logical-system *logical-system-name*)
- `show rsvp session` (logical-system *logical-system-name*)

The following sections show the output of commands used with the configuration example:

- [Router CE1 Status on page 3463](#)
- [Router CE2 Status on page 3463](#)
- [Router CE3 Status on page 3464](#)
- [Router PE1 Status: Main Router on page 3464](#)
- [Router PE1 Status: Logical System LS1 on page 3465](#)
- [Router PE1 Status: Logical System LS2 on page 3467](#)
- [Router P0 Status: Main Router on page 3468](#)
- [Router P0 Status: Main Router on page 3468](#)
- [Router P0 Status: Logical System LS1 on page 3469](#)
- [Router P0 Status: Logical System LS2 on page 3469](#)
- [Router PE2 Status: Main Router on page 3470](#)
- [Router PE2 Status: Logical System LS1 on page 3472](#)
- [Router PE2 Status: Logical System LS2 on page 3474](#)
- [Router CE5 Status on page 3476](#)
- [Router CE6 Status on page 3476](#)
- [Router CE7 Status on page 3476](#)
- [Logical System Administrator Verification Output on page 3477](#)

#### ***Router CE1 Status***

**Purpose** Verify connectivity.

**Action** user@CE1> show route table  
 inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)  
 + = Active Route, - = Last Active, \* = Both

```

10.11.1.0/24 * [Direct/0] 00:20:20
 > via fe-1/0/1.0
10.11.1.1/32 * [Local/0] 00:20:24
 Local via fe-1/0/1.0
10.11.1.100/32 * [Direct/0] 00:21:53
 > via lo0.0
10.11.4.0/24 * [OSPF/150] 00:18:30, metric 0, tag 3489661028
 > to 10.11.1.2 via fe-1/0/1.0
10.11.4.100/32 * [OSPF/10] 00:18:30, metric 2
 > to 10.11.1.2 via fe-1/0/1.0
224.0.0.5/32 * [OSPF/10] 00:21:58, metric 1
 MultiRecv

```

#### ***Router CE2 Status***

**Purpose** Verify connectivity.

**Action** user@CE2> show route table  
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

```
10.21.1.0/24 *[Direct/0] 00:20:30
 > via fe-1/0/2.0
10.21.1.1/32 *[Local/0] 00:20:34
 Local via fe-1/0/2.0
10.21.1.100/32 *[Direct/0] 00:22:03
 > via lo0.0
10.21.4.0/24 *[BGP/170] 00:18:43, localpref 100
 AS path: 100 I
 > to 10.21.1.2 via fe-1/0/2.0
10.21.4.100/32 *[BGP/170] 00:18:43, localpref 100
 AS path: 100 300 I
 > to 10.21.1.2 via fe-1/0/2.0
```

#### ***Router CE3 Status***

**Purpose** Verify connectivity.

**Action** user@CE3> show route table  
inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

```
10.1.1.0/24 *[Direct/0] 00:20:13
 > via fe-1/0/0.0
10.1.1.1/32 *[Local/0] 00:20:17
 Local via fe-1/0/0.0
```

#### ***Router PE1 Status: Main Router***

**Purpose** Verify BGP operation.

**Action** user@PE1> show bgp summary

Groups: 1 Peers: 2 Down peers: 0

| Table  | Tot Paths | Act Paths | Suppressed | History | Damp | State | Pending |
|--------|-----------|-----------|------------|---------|------|-------|---------|
| inet.0 | 1         | 0         | 0          | 0       | 0    | 0     | 0       |

Peer AS InPkt OutPkt OutQ Flaps Last Up/DwnState|#Active/Received/Damped...

|               |     |   |   |   |   |      |       |
|---------------|-----|---|---|---|---|------|-------|
| 10.255.41.175 | 500 | 5 | 8 | 0 | 0 | 2:31 | 0/0/0 |
| 10.255.41.179 | 500 | 6 | 9 | 0 | 0 | 2:35 | 0/1/0 |

user@PE1> show route protocol bgp

inet.0: 20 destinations, 21 routes (20 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, \* = Both

10.255.41.180/32 [BGP/170] 00:02:48, localpref 100, from 10.255.41.179

AS path: I

> to 10.31.2.2 via fe-0/0/1.3

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

user@PE1> show ospf neighbor

| Address   | Interface  | State | ID            | Pri | Dead |
|-----------|------------|-------|---------------|-----|------|
| 10.31.2.2 | fe-0/0/1.3 | Full  | 10.255.41.175 | 128 | 32   |

user@PE1> show isis adjacency

IS-IS instance is not running

**Router PE1 Status: Logical System LS1**

**Purpose** Verify BGP operation.

```

Action user@PE1> show bgp summary logical-system LS1
Groups: 2 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.l3vpn.0 4 4 0 0 0
bgp.l3vpn.2 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last
Up/DwnState|#Active/Received/Damped...
10.10.10.12 100 13 14 0 0 2:50 Establ
 bgp.l3vpn.0: 4/4/0
 bgp.l3vpn.2: 0/0/0
 blue.inet.0: 2/2/0
 red.inet.0: 2/2/0
10.21.1.1 200 13 14 0 0 4:33 Establ
 blue.inet.0: 1/1/0

```

## Red VPN

The master administrator or logical system administrator can issue the following command to view the output for a specific logical system.

```

user@PE1> show route logical-system LS1 table red
red.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.11.1.0/24 *[Direct/0] 00:04:51
 > via fe-0/1/1.0
10.11.1.2/32 *[Local/0] 00:05:45
 Local via fe-0/1/1.0
10.11.1.100/32 *[OSPF/10] 00:04:02, metric 1
 > to 10.11.1.1 via fe-0/1/1.0
10.11.4.0/24 *[BGP/170] 00:03:05, localpref 100, from 10.10.10.12
 AS path: I
 > to 10.11.2.2 via fe-0/0/1.1, label-switched-path
to_10.10.10.12
10.11.4.100/32 *[BGP/170] 00:03:05, MED 1, localpref 100, from 10.10.10.12
 AS path: I
 > to 10.11.2.2 via fe-0/0/1.1, label-switched-path
to_10.10.10.12
224.0.0.5/32 *[OSPF/10] 00:07:02, metric 1
 MultiRecv

```

## Blue VPN

The master administrator or logical system administrator can issue the following command to view the output for a specific logical system.

```

user@PE1> show route logical-system LS1 table blue
blue.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.21.1.0/24 *[Direct/0] 00:05:29
 > via fe-0/1/2.0
10.21.1.2/32 *[Local/0] 00:06:23
 Local via fe-0/1/2.0
10.21.1.100/32 *[BGP/170] 00:05:26, localpref 100
 AS path: 200 I
 > to 10.21.1.1 via fe-0/1/2.0
10.21.4.0/24 *[BGP/170] 00:03:43, localpref 100, from 10.10.10.12
 AS path: I
 > to 10.11.2.2 via fe-0/0/1.1, label-switched-path

```



```

to_10.10.10.12
10.21.4.100/32 *[BGP/170] 00:03:43, localpref 100, from 10.10.10.12
 AS path: 300 I
 > to 10.11.2.2 via fe-0/0/1.1, label-switched-path
to_10.10.10.12

```

```

user@PE1> show route logical-system LS1 table inet.0
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

10.10.10.10/32 *[Direct/0] 00:08:05
 > via lo0.1
10.10.10.11/32 *[IS-IS/15] 00:05:07, metric 10
 > to 10.11.2.2 via fe-0/0/1.1
10.10.10.12/32 *[IS-IS/15] 00:04:58, metric 20
 > to 10.11.2.2 via fe-0/0/1.1
10.11.2.0/24 *[Direct/0] 00:05:38
 > via fe-0/0/1.1
10.11.2.1/32 *[Local/0] 00:06:51
 Local via fe-0/0/1.1
10.11.3.0/24 *[IS-IS/15] 00:05:07, metric 20
 > to 10.11.2.2 via fe-0/0/1.1

```

```

user@PE1> ping logical-system LS1 routing-instance red 10.11.4.100
PING 10.11.4.100 (10.11.4.100): 56 data bytes
64 bytes from 10.11.4.100: icmp_seq=0 ttl=251 time=1.055 ms
^C
--- 10.11.4.100 ping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.055/1.055/1.055/0.000 ms

```

### ***Router PE1 Status: Logical System LS2***

**Purpose** Verify VPLS operation.

**Action** user@PE1> show vpls connections logical-system LS2

Layer-2 VPN Connections:

Legend for connection status (St)

```
OR -- out of range WE -- intf encaps != instance encaps
EI -- encapsulation invalid Dn -- down
EM -- encapsulation mismatch VC-Dn -- Virtual circuit down
CM -- control-word mismatch -> -- only outbound conn is up
CN -- circuit not provisioned <- -- only inbound conn is up
OL -- no outgoing label Up -- operational
NC -- intf encaps not CCC/TCC XX -- unknown
NP -- intf h/w not present
```

Legend for interface status

```
Up -- operational
Dn -- down
```

Instance: new

Local site: newPE (1)

```
connection-site Type St Time last up # Up trans
2 rmt Up Jul 16 14:05:25 2003 1
 Local interface: vt-1/2/0.49152, Status: Up, Encapsulation: VPLS
 Remote PE: 10.20.20.22, Negotiated control-word: No
 Incoming label: 800001, Outgoing label: 800000
```

#### *Router P0 Status: Main Router*

**Purpose** Verify connectivity.

**Action** user@P0> show interfaces terse lo0

| Interface | Admin | Link | Proto | Local                                                                                                 | Remote             |
|-----------|-------|------|-------|-------------------------------------------------------------------------------------------------------|--------------------|
| lo0       | up    | up   |       |                                                                                                       |                    |
| lo0.0     | up    | up   | inet  | 10.255.41.175<br>127.0.0.1                                                                            | --> 0/0<br>--> 0/0 |
|           |       |      | iso   |                                                                                                       |                    |
|           |       |      | inet6 | 47.0005.80ff.f800.0000.0108.0003.0102.5501.4175.00<br>fe80::2a0:a5ff:fe12:2b09<br>feee::10:255:14:175 |                    |
| lo0.1     | up    | up   | inet  | 10.10.10.11                                                                                           | --> 0/0            |
|           |       |      | iso   | 47.1111.1111.1111.1112.00                                                                             |                    |
| lo0.2     | up    | up   | inet  | 10.20.20.21                                                                                           | --> 0/0            |
|           |       |      | iso   | 47.2222.2222.2222.2223.00                                                                             |                    |
| lo0.16383 | up    | up   | inet  |                                                                                                       |                    |

user@P0> show ospf neighbor

| Address   | Interface  | State | ID            | Pri | Dead |
|-----------|------------|-------|---------------|-----|------|
| 10.31.2.1 | fe-1/1/3.3 | Full  | 10.255.41.173 | 128 | 34   |
| 10.31.3.2 | so-1/2/0.3 | Full  | 10.255.41.179 | 128 | 37   |

#### *Router P0 Status: Main Router*

**Purpose** Verify routing protocols operation.

**Action** user@P0> show interfaces terse lo0

| Interface | Admin | Link | Proto | Local                                              | Remote  |
|-----------|-------|------|-------|----------------------------------------------------|---------|
| lo0       | up    | up   |       |                                                    |         |
| lo0.0     | up    | up   | inet  | 10.255.41.175                                      | --> 0/0 |
|           |       |      |       | 127.0.0.1                                          | --> 0/0 |
|           |       |      | iso   |                                                    |         |
|           |       |      |       | 47.0005.80ff.f800.0000.0108.0003.0102.5501.4175.00 |         |
|           |       |      | inet6 | fe80::2a0:a5ff:fe12:2b09                           |         |
|           |       |      |       | feee::10:255:14:175                                |         |
| lo0.1     | up    | up   | inet  | 10.10.10.11                                        | --> 0/0 |
|           |       |      | iso   | 47.1111.1111.1111.1112.00                          |         |
| lo0.2     | up    | up   | inet  | 10.20.20.21                                        | --> 0/0 |
|           |       |      | iso   | 47.2222.2222.2222.2223.00                          |         |
| lo0.16383 | up    | up   | inet  |                                                    |         |

user@P0> show ospf neighbor

| Address   | Interface  | State | ID            | Pri | Dead |
|-----------|------------|-------|---------------|-----|------|
| 10.31.2.1 | fe-1/1/3.3 | Full  | 10.255.41.173 | 128 | 34   |
| 10.31.3.2 | so-1/2/0.3 | Full  | 10.255.41.179 | 128 | 37   |

### ***Router P0 Status: Logical System LS1***

**Purpose** Verify routing protocols operation.

**Action** user@P0> show isis adjacency logical-system LS1

| Interface  | System | L State | Hold (secs) | SNPA          |
|------------|--------|---------|-------------|---------------|
| fe-1/1/3.1 | PE1    | 2 Up    | 21          | 0:90:69:9:4:1 |
| fe-1/1/3.1 | PE1    | 1 Up    | 24          | 0:90:69:9:4:1 |
| so-1/2/0.1 | PE2    | 3 Up    | 25          |               |

user@P0> show bgp summary logical-system LS1

BGP is not running

user@P0> show route protocol isis logical-system LS1

inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, \* = Both

10.10.10.10/32 \* [IS-IS/15] 00:09:15, metric 10  
 > to 10.11.2.1 via fe-1/1/3.1  
 10.10.10.12/32 \* [IS-IS/15] 00:09:39, metric 10  
 > to 10.11.3.2 via so-1/2/0.1

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)

### ***Router P0 Status: Logical System LS2***

**Purpose** Verify routing protocols operation.

**Action** user@P0> show bgp summary logical-system LS2  
BGP is not running

user@P0> show isis adjacency logical-system LS2

| Interface  | System | L | State | Hold (secs) | SNPA          |
|------------|--------|---|-------|-------------|---------------|
| fe-1/1/3.0 | PE1    | 2 | Up    | 24          | 0:90:69:9:4:1 |
| fe-1/1/3.0 | PE1    | 1 | Up    | 23          | 0:90:69:9:4:1 |
| so-1/2/0.0 | PE2    | 3 | Up    | 24          |               |

user@P0> show route protocol isis logical-system LS2

inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

```
10.20.20.20/32 *[IS-IS/15] 00:09:44, metric 10
 > to 10.1.2.1 via fe-1/1/3.0
10.20.20.22/32 *[IS-IS/15] 00:09:45, metric 10
 > to 10.1.3.2 via so-1/2/0.0
```

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)

***Router PE2 Status: Main Router***

**Purpose** Verify routing protocols operation.

```

Action user@PE2> show ospf neighbor
 Address Interface State ID Pri Dead
10.31.4.2 fe-0/2/3.0 Full 10.255.41.180 128 38
10.31.3.1 so-1/2/0.3 Full 10.255.41.175 128 36

user@PE2> show interfaces terse lo0
Interface Admin Link Proto Local Remote
lo0 up up
lo0.0 up up inet 10.255.41.179 --> 0/0
 127.0.0.1 --> 0/0
 iso
47.0005.80ff.f800.0000.0108.0003.0102.5501.4179.00
 inet6 fe80::2a0:a5ff:fe12:29ff
 feee::10:255:14:179
lo0.1 up up inet 10.10.10.12 --> 0/0
 iso 47.1111.1111.1111.1113.00
lo0.2 up up inet 10.20.20.22 --> 0/0
 iso 47.2222.2222.2222.2224.00
lo0.16383 up up inet

user@PE2> show bgp summary
Groups: 1 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 1 1 1 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last
Up/DwnState|#Active/Received/Damped...
10.255.41.175 500 24 27 0 0 11:46 0/0/0
0/0/0
10.255.41.173 500 25 25 0 0 11:11 1/1/0
0/0/0

user@PE2> show route protocol ospf

inet.0: 20 destinations, 22 routes (19 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.41.175/32 *[OSPF/10] 00:00:20, metric 1
> via so-1/2/0.3
10.255.41.180/32 [OSPF/10] 00:00:20, metric 1
> to 10.31.4.2 via fe-0/2/3.0
10.255.41.173/32 *[OSPF/10] 00:00:20, metric 2
> via so-1/2/0.3
10.31.2.0/24 *[OSPF/10] 00:00:20, metric 2
> via so-1/2/0.3
10.31.3.0/24 [OSPF/10] 00:00:20, metric 1
> via so-1/2/0.3
224.0.0.5/32 *[OSPF/10] 00:13:46, metric 1
MultiRecv

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

user@PE2> show route protocol bgp

inet.0: 20 destinations, 22 routes (19 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.41.177/32 *[BGP/170] 00:11:23, localpref 100, from 10.255.41.173
AS path: I

```

> via so-1/2/0.3

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

***Router PE2 Status: Logical System LS1***

**Purpose** Verify routing protocols operation.

```

Action user@PE2> show bgp summary logical-system LS1
Groups: 2 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 0 0 0 0 0 0 0
inet.2 0 0 0 0 0 0 0
bgp.l3vpn.0 4 4 0 0 0 0 0
bgp.l3vpn.2 0 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.10.10.10 100 29 31 0 0 11:25 Establ
bgp.l3vpn.0: 4/4/0
bgp.l3vpn.2: 0/0/0
blue.inet.0: 2/2/0
red.inet.0: 2/2/0
10.21.4.2 300 27 28 0 0 11:40 Establ
blue.inet.0: 1/1/0

```

### Red VPN

```

user@PE2> show route logical-system LS1 table red
red.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.11.1.0/24 * [BGP/170] 00:12:02, localpref 100, from 10.10.10.10
AS path: I
> via so-1/2/0.1, label-switched-path to_10.10.10.10
10.11.1.100/32 * [BGP/170] 00:12:02, MED 1, localpref 100, from 10.10.10.10
AS path: I
> via so-1/2/0.1, label-switched-path to_10.10.10.10
10.11.4.0/24 * [Direct/0] 00:13:22
> via fe-0/2/1.0
10.11.4.1/32 * [Local/0] 00:13:29
Local via fe-0/2/1.0
10.11.4.100/32 * [OSPF/10] 00:12:35, metric 1
> to 10.11.4.2 via fe-0/2/1.0
224.0.0.5/32 * [OSPF/10] 00:15:02, metric 1
MultiRecv

```

### Blue VPN

```

user@PE2> show route logical-system LS1 table blue
blue.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.21.1.0/24 * [BGP/170] 00:13:12, localpref 100, from 10.10.10.10
AS path: I
> via so-1/2/0.1, label-switched-path to_10.10.10.10
10.21.1.100/32 * [BGP/170] 00:13:12, localpref 100, from 10.10.10.10
AS path: 200 I
> via so-1/2/0.1, label-switched-path to_10.10.10.10
10.21.4.0/24 * [Direct/0] 00:14:32
> via fe-0/2/2.0
10.21.4.1/32 * [Local/0] 00:14:39
Local via fe-0/2/2.0
10.21.4.100/32 * [BGP/170] 00:13:27, localpref 100
AS path: 300 I
> to 10.21.4.2 via fe-0/2/2.0

user@PE2> show mpls lsp logical-system LS1
Ingress LSP: 1 sessions
To From State Rt ActivePath P LSPname
10.10.10.10 10.10.10.12 Up 0 * to_10.10.10.10
Total 1 displayed, Up 1, Down 0

```

```
Egress LSP: 1 sessions
To From State Rt Style Labelin Labelout LSPname
10.10.10.12 10.10.10.10 Up 0 1 FF 3 - to_10.10.10.12
Total 1 displayed, Up 1, Down 0
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

```
user@PE2> show rsvp session logical-system LS1
Ingress RSVP: 1 sessions
To From State Rt Style Labelin Labelout LSPname
10.10.10.10 10.10.10.12 Up 0 1 FF - 100000 to_10.10.10.10
Total 1 displayed, Up 1, Down 0
Egress RSVP: 1 sessions
To From State Rt Style Labelin Labelout LSPname
10.10.10.12 10.10.10.10 Up 0 1 FF 3 - to_10.10.10.12
Total 1 displayed, Up 1, Down 0
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

***Router PE2 Status: Logical System LS2***

**Purpose** Verify routing protocols operation.



**Action** user@PE2> show vpls connections logical-system LS2

Layer-2 VPN Connections:

Legend for connection status (St)

OR -- out of range WE -- intf encaps != instance encaps  
 EI -- encapsulation invalid Dn -- down  
 EM -- encapsulation mismatch VC-Dn -- Virtual circuit down  
 CM -- control-word mismatch -> -- only outbound conn is up  
 CN -- circuit not provisioned <- -- only inbound conn is up  
 OL -- no outgoing label Up -- operational  
 NC -- intf encaps not CCC/TCC XX -- unknown  
 NP -- intf h/w not present

Legend for interface status

Up -- operational  
 Dn -- down

Instance: new

Local site: newPE (2)

| connection-site | Type | St | Time last up         | # Up trans |
|-----------------|------|----|----------------------|------------|
| 1               | rmt  | Up | Jul 16 14:05:25 2003 | 1          |

Local interface: vt-1/1/0.40960, Status: Up, Encapsulation: VPLS  
 Remote PE: 10.20.20.20, Negotiated control-word: No  
 Incoming label: 800000, Outgoing label: 800001

user@PE2> show bgp summary logical-system LS2

Groups: 1 Peers: 1 Down peers: 0

| Table       | Tot Paths | Act Paths | Suppressed | History | Damp | State | Pending |
|-------------|-----------|-----------|------------|---------|------|-------|---------|
| bgp.12vpn.0 | 1         | 1         | 0          | 0       | 0    | 0     | 0       |

| Peer        | AS  | InPkt | OutPkt | OutQ | Flaps | Last         |
|-------------|-----|-------|--------|------|-------|--------------|
| 10.20.20.20 | 400 | 29    | 31     | 0    | 0     | 13:29 Establ |

Up/DwnState|#Active/Received/Damped...  
 bgp.12vpn.0: 1/1/0  
 new.12vpn.0: 1/1/0

user@PE2> show mpls lsp logical-system LS2

Ingress LSP: 1 sessions

| To          | From        | State | Rt | ActivePath | P | LSPname        |
|-------------|-------------|-------|----|------------|---|----------------|
| 10.20.20.20 | 10.20.20.22 | Up    | 0  |            | * | to_10.20.20.20 |

Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

| To          | From        | State | Rt | Style | Labelin | Labelout | LSPname        |
|-------------|-------------|-------|----|-------|---------|----------|----------------|
| 10.20.20.22 | 10.20.20.20 | Up    | 0  | 1 FF  | 3       | -        | to_10.20.20.22 |

Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

user@PE2> show rsvp session logical-system LS2

Ingress RSVP: 1 sessions

| To          | From        | State | Rt | Style | Labelin | Labelout | LSPname        |
|-------------|-------------|-------|----|-------|---------|----------|----------------|
| 10.20.20.20 | 10.20.20.22 | Up    | 0  | 1 FF  | -       | 100016   | to_10.20.20.20 |

Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions

| To          | From        | State | Rt | Style | Labelin | Labelout | LSPname        |
|-------------|-------------|-------|----|-------|---------|----------|----------------|
| 10.20.20.22 | 10.20.20.20 | Up    | 0  | 1 FF  | 3       | -        | to_10.20.20.22 |

Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions  
Total 0 displayed, Up 0, Down 0

#### ***Router CE5 Status***

**Purpose** Verify connectivity.

**Action** user@CE5> show route table  
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both  
10.11.1.0/24 \* [OSPF/150] 00:19:47, metric 0, tag 3489661028  
> to 10.11.4.1 via fe-0/3/1.0  
10.11.1.100/32 \* [OSPF/10] 00:19:47, metric 2  
> to 10.11.4.1 via fe-0/3/1.0  
10.11.4.0/24 \* [Direct/0] 00:21:12  
> via fe-0/3/1.0  
10.11.4.2/32 \* [Local/0] 00:21:24  
Local via fe-0/3/1.0  
10.11.4.100/32 \* [Direct/0] 00:22:37  
> via lo0.0  
224.0.0.5/32 \* [OSPF/10] 00:22:44, metric 1  
MultiRecv

#### ***Router CE6 Status***

**Purpose** Verify connectivity.

**Action** user@CE6> show route table  
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both  
10.21.1.0/24 \* [BGP/170] 00:19:53, localpref 100  
AS path: 100 I  
> to 10.21.4.1 via fe-0/3/2.0  
10.21.1.100/32 \* [BGP/170] 00:19:53, localpref 100  
AS path: 100 200 I  
> to 10.21.4.1 via fe-0/3/2.0  
10.21.4.0/24 \* [Direct/0] 00:21:16  
> via fe-0/3/2.0  
10.21.4.2/32 \* [Local/0] 00:21:28  
Local via fe-0/3/2.0  
10.21.4.100/32 \* [Direct/0] 00:22:41  
> via lo0.0

#### ***Router CE7 Status***

**Purpose** Verify connectivity.

**Action** user@CE7> show route table

```
inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.0/24 *[Direct/0] 00:21:03
 > via fe-0/3/0.0
10.1.1.2/32 *[Local/0] 00:21:15
 Local via fe-0/3/0.0
```

### *Logical System Administrator Verification Output*

**Purpose** Because logical system administrators only have access to the configuration information of the logical systems to which they are assigned, the verification output is limited to these logical systems as well. The following output shows what the logical system administrator **LS1-admin** in this example configuration would see.

To verify that each pair of CE routers has end-to-end connectivity, issue the **ping** command on Routers CE1, CE2, and CE3:

**Action** From CE1, ping CE5 (the Red VPN).

From CE2, ping CE6 (the Blue VPN).

From CE3, ping CE7 (the VPLS).

```
user@CE1> ping 10.11.4.100
PING 10.11.4.100 (10.11.4.100): 56 data bytes
64 bytes from 10.11.4.100: icmp_seq=0 ttl=252 time=1.216 ms
64 bytes from 10.11.4.100: icmp_seq=1 ttl=252 time=1.052 ms
^C
--- 10.11.4.100 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.052/1.134/1.216/0.082 ms
```

```
user@CE2> ping 10.21.4.100
PING 10.21.4.100 (10.21.4.100): 56 data bytes
64 bytes from 10.21.4.100: icmp_seq=0 ttl=252 time=1.205 ms
64 bytes from 10.21.4.100: icmp_seq=1 ttl=252 time=1.021 ms
^C
--- 10.21.4.100 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.021/1.113/1.205/0.092 ms
```

```
user@CE3> ping 10.1.1.2
PING 10.1.1.2 (10.1.1.2): 56 data bytes
64 bytes from 10.1.1.2: icmp_seq=0 ttl=255 time=1.186 ms
64 bytes from 10.1.1.2: icmp_seq=1 ttl=255 time=1.091 ms
64 bytes from 10.1.1.2: icmp_seq=2 ttl=255 time=1.081 ms
^C
--- 10.1.1.2 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.081/1.119/1.186/0.047 ms
```

**Related Documentation** • [Introduction to Logical Systems on page 3259](#)

### **Example: Configuring a Virtualized Data Center**

---

- [Two-Tiered Virtualized Data Center Solution for Large Enterprise Networks on page 3478](#)
- [Requirements of a Two-Tiered Virtualized Data Center for Large Enterprise Networks on page 3479](#)
- [Example: Configuring a Two-Tiered Virtualized Data Center for Large Enterprise Networks on page 3481](#)

#### ***Two-Tiered Virtualized Data Center Solution for Large Enterprise Networks***

The following describes a Juniper Networks two-tiered, high-speed, multiservice virtualized data center (VDC). A two-tiered architecture meets the low latency requirements of a virtualized server environment and supports the overlying security mandate to maintain controlled segmentation between various business units.

#### ***Network Traffic Segmentation***

Juniper Networks VDC design uses virtualization technologies such as virtual LANs (VLANs), virtual routers, virtual route forwarders, inter-virtual route forwarding, and logical systems to provide flexible traffic isolation.

A fully redundant two-tiered data center design consists of Juniper Networks EX Series Ethernet Switches at the access layer for server connectivity, MX Series 3D Universal Edge Routers as a collapsed LAN aggregation/core layer, and clustered SRX Series Services Gateways to provide firewall security services across the data center trust boundaries.

#### ***Flexibility***

The Juniper Networks VDC design uses 802.1Q VLANs, MPLS, BGP, Virtual Router Redundancy Protocol (VRRP), Traffic Engineering, and Fast Reroute to provide design flexibility while maintaining a standards-based approach. The design can also support a virtual private LAN service (VPLS).

#### ***Security***

The Juniper Networks VDC design uses security zones to implement the policy enforcement points. The SRX cluster is responsible for all stateful packet inspection for traffic that crosses business unit trust boundaries as well as all ingress and egress traffic for the data center.

The Juniper Networks Junos operating system is configured with different administrator accounts for each logical system that supports confined access to network resources and can be customized for individual business units.

#### ***Access and Availability***

In the Juniper Networks VDC design, described in [“Example: Configuring a Two-Tiered Virtualized Data Center for Large Enterprise Networks” on page 3481](#), top-of-rack (TOR) EX Series switches provide access to the servers and provide redundancy.

All uplinks from the TOR switches are 802.1Q trunk links that are terminated directly into each of the MX Series devices that make up the Point of Delivery (POD) at the aggregation/core layer.

A VRRP instance is defined on each VLAN within the MX Series device to act as the default router for all server hosts in a given VLAN. To allow for VRRP to work properly, each bridge domain is extended between each MX Series device through an interconnection link. The MX Series device uses an integrated routing and bridging (IRB) interface as the Layer 3 interface for each bridge domain, with VRRP configured for redundancy.

A pair of 802.3ad aggregated Ethernet bundles are used between the MX Series devices. Each MX Series device is divided into a number of Logical Systems. Logical systems in the MX Series device are used to define logical trust boundaries within the data center itself and between respective business units.

A clustered pair of SRX Series devices acting as firewalls provide security services across the data center trust boundaries. Virtual routers on the SRX Series devices act as customer edge (CE) routers for each business unit.

A single redundancy group for the data plane is defined on the SRX Series Services Gateways with two redundant Ethernet interfaces as member interfaces. This redundancy group handles the data plane failover of the SRX Series firewall and is configured such that any loss of either northbound or southbound SRX Series interfaces forces a full failover to the secondary node. This failover is essentially a Layer 1 failover, which means that it occurs quickly and does not disrupt the routing topology above it.

### ***Cost-Effective Incremental Scaling***

The Juniper Networks VDC design supports incremental scaling of the network. This allows the VDC to be created with minimum cost to meet the current need.

The access layer can be expanded by adding EX Series switches at the top of rack.

The aggregation/core layer can be expanded by adding additional MX Series devices within a given POD.

The security services can be expanded by adding 4-port 10-Gigabit Ethernet I/O cards (IOCs) and services processing cards (SPCs) in the SRX Series devices. The addition of IOCs increases the 10-Gigabit Ethernet port density. The addition of each SPC card to the chassis adds another 10 Gbps (5 Gbps Internet mix (IMIX)), 2 million sessions, and 100,000 connections per second (CPS) up to a maximum rated capacity for the platform of 150 Gbps (47.5 Gbps IMIX), 10 million sessions, and 350,000 CPS (as measured in Junos OS Release 10.2).

### ***Orchestration and Automation***

The Juniper Networks VDC design uses the Juniper Networks Junos Space management platform. Junos Space includes a portfolio of applications for scaling services, simplifying network operations, and automating support for complex network environments.

In addition, the network devices are configured to support background Secure Copy Protocol (SCP) file transfers, commit scripts, and a file archive site.

### ***Requirements of a Two-Tiered Virtualized Data Center for Large Enterprise Networks***

Large enterprises have certain specific needs for the hosting environment that the design of their data center must meet. This section describes the requirements of a company that operates as a service provider to its individual business units (BUs).

One of the primary requirements of a virtualized data center (VDC) for a large enterprise is the ability to segment the network by business unit. This includes traffic segmentation and administrative control segmentation.

Other requirements include security controls between business units, security controls between the company and the outside world, flexibility to grow and adapt the network, and a robust and cost-effective way to manage the entire network.

### ***Network Traffic Segmentation***

The requirement described here is for network resources to be isolated in several ways. Traffic must be segmented by business units. Traffic flows between network segments must be prohibited except where specifically allowed. Traffic isolation must be controlled at designated policy enforcement points. Network resources must be dedicated to a segment, but the network must have the flexibility to change the allocation of resources.

Segmented resources must be logically grouped according to policies. For example, test traffic must be isolated from production traffic. Traffic must also be isolated according to business entities, contractual requirements, legal or regulatory requirements, risk rating, and corporate standards.

The network segmentation design must not be disruptive to the business, must be integrated with the larger data center and cloud network design, must allow business units to access network resources globally, and must support new business capabilities.

### ***Flexibility***

The network design must be flexible enough to react to business and environment changes with minimal design and re-engineering efforts. The VDC design must be flexible in terms of isolating business unit workloads from other business units and general data center services and applications. The network solution must ensure that the business is minimally impacted when network and segmentation changes take place.

The VDC must be flexible enough to be implemented:

- Within a single data center
- Within a data hall
- Across two or more data centers
- Across two or more data halls within or between data centers
- Between a data center and an external cloud service provider

### ***Security***

The network design must allow business units to be isolated within the hosting environment. In the event of a network security incident, business units must be isolated from the hosting environment and other business units.

Traffic flow between business unit segments must be denied by default and must be explicitly permitted only at policy enforcement points owned and controlled by the data center service provider.

The policy enforcement point must include access control capabilities and might include threat protection capabilities.

#### ***Access and Availability***

The VDC must provide access to common data center services such as computation, storage, security, traffic management, operations, and applications. The network must operate across multiple global service providers and must deliver optimal, predictable, and consistent performance across the network. The VDC must be implemented across data center business units.

The network solution must meet business unit availability requirements as defined in service-level agreements.

#### ***Cost-Effective Incremental Scaling***

The VDC design must be cost effective for the business to run and must enable new business capabilities. It must be possible to implement the network solution in an incremental manner with minimal impact to the business.

#### ***Orchestration and Automation***

The VDC design must include a management system that supports automation for provisioning, availability and workload monitoring, and reporting. Workload and availability reports must be available by business unit.

#### ***Example: Configuring a Two-Tiered Virtualized Data Center for Large Enterprise Networks***

This example provides a step-by-step procedure for configuring a two-tiered virtualized data center for large enterprise networks.

- [Requirements on page 3481](#)
- [Configuring a Two-Tiered Virtualized Data Center Overview on page 3481](#)
- [Configuring the Access Layer on page 3484](#)
- [Configuring the Aggregation Layer in the Trusted Logical Systems on page 3488](#)
- [Configuring the Core Layer in the Untrusted Logical Systems on page 3496](#)
- [Configuring the Security Device on page 3501](#)

#### ***Requirements***

This example uses the following hardware and software components:

- Two MX Series 3D Universal Edge Routers running Junos OS Release 10.2 or later
- Six EX Series Ethernet Switches running Junos OS Release 10.2 or later
- Two SRX Series Services Gateways running Junos OS Release 10.4 or later

#### ***Configuring a Two-Tiered Virtualized Data Center Overview***

This example provides a step-by-step procedure for configuring a two-tiered virtualized data center for large enterprises. The steps in the example follow the data path from an interface connected to a server in BU2 using VLAN 17, to Logical System Trust1, through Virtual Router MX-VR2, through Virtual Router SRX-VR2, through VRF2 in the Logical System Untrust, and out to the core network.

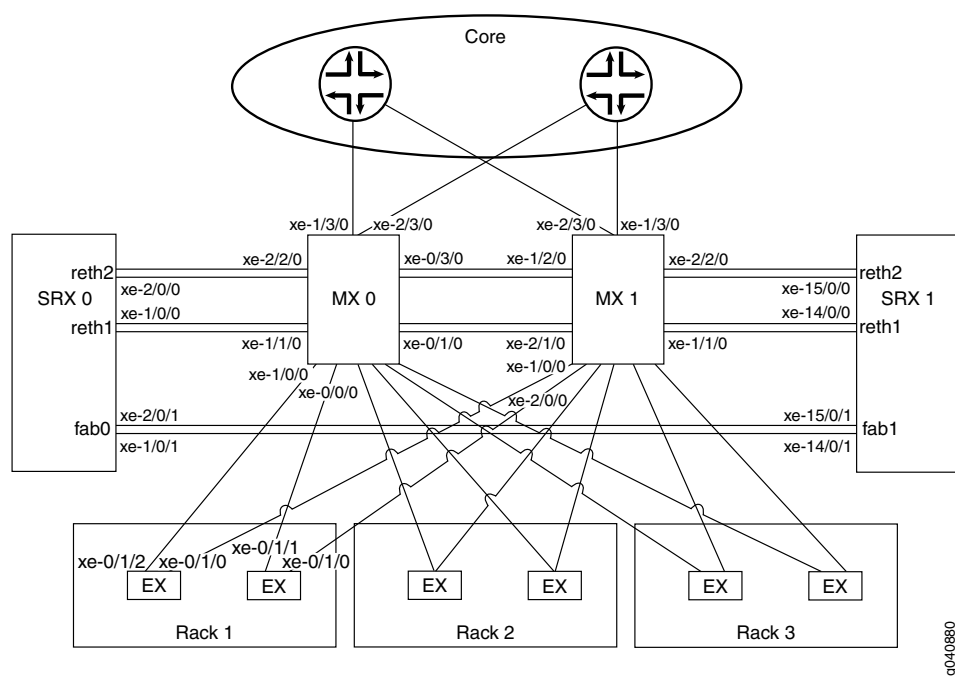
The core network in this example simultaneously supports IP-based routing and MPLS-based label switching. The virtual routers on the SRX Series device perform the functions of customer edge (CE) routers. The VPN routing and forwarding (VRF) routing instances on the MX Series devices perform the functions of service provider edge (PE) routers. The OSPF protocol serves as the interior gateway protocol to carry routes to the PE router loopback addresses that are used as the BGP next-hop address for the IP-based and MPLS-based networks supported by this example.



**NOTE:** The steps in this example are representative of the entire network configuration. The example does not show every step for every virtual device.

The physical connections used in this example are shown in [Figure 76 on page 3482](#).


**Figure 76: Virtualized Data Center Physical Topology**



The logical connections used in this example are shown in [Figure 77 on page 3483](#).



Figure 77: Virtualized Data Center Logical Topology



In the logical topology illustration:

- Users access the data center across the enterprise core network shown at the top.
- Virtual routers configured in Logical System Untrust on the MX Series devices forward the traffic to separate virtual routers configured in the Untrusted security zone on the SRX Series devices. These virtual routers act as edge routers for the various business units.
- Virtual routers configured on the active SRX Series device forward the traffic to the Trusted security zones.
- Virtual routers configured in separate logical systems on the MX Series devices forward the traffic to a bridge domain of VLANs configured on the EX Series devices.
- Business unit 1 requires additional separation. In this case, the virtual router (VR) configured on the SRX Series device forwards the traffic directly to the bridge domain on the EX Series devices.
- The EX Series devices switch the traffic to the data center server.

- The SRX Series devices apply security policy to all traffic traversing the untrust to trust boundary and all traffic forwarded between logical systems.
- The SRX Series devices are configured in an active/passive cluster so that only one node in the cluster is active on the data forwarding plane at a time.
- The SRX Series devices are configured with a single redundancy group for the data plane. The redundancy group uses two Ethernet interfaces (**reth1** and **reth2** in [Figure 76 on page 3482](#)) as member interfaces.

### **Configuring the Access Layer**

Configure the access layer by doing the following:

- [Configuring Interfaces on page 3484](#)
- [Configuring VLANs in the Access Layer on page 3485](#)
- [Configuring a Redundant Trunk Group and Disabling the Spanning Tree Protocol for the Trunk Interfaces on page 3486](#)
- [Configuring Management Automation on page 3487](#)

### **Configuring Interfaces**

#### **Step-by-Step Procedure**

This procedure explains how to configure the physical, logical, and network management interfaces for the access layer devices. This procedure shows a representative sample of the configuration. The example does not show the configuration for every interface.

1. Configure the access layer server-facing 10-Gigabit Ethernet interfaces.

This example configures the **ge-0/0/17** interface with VLAN ID 17.

Include the **member** statement and specify VLAN ID 17 at the **[edit interfaces ge-0/0/17 unit 0 family ethernet-switching vlan]** hierarchy level.

```
[edit interfaces ge-0/0/17 unit 0]
user@ex# set family ethernet-switching vlan members 17
```

Repeat this step for every server-facing interface by using the appropriate interface name and VLAN number.

2. Configure the 10-Gigabit Ethernet trunk interfaces from the EX Series device to the two MX Series devices.

This example configures the **xe-0/1/2** and **xe-0/1/0** interfaces.

Include the **port-mode** statement and specify the **trunk** option at the **[edit interfaces xe-0/1/2 unit 0 family ethernet-switching]** and **[edit interfaces xe-0/1/0 unit 0 family ethernet-switching]** hierarchy levels.

Include the **members** statement and specify the **all** option at the **[edit interfaces xe-0/1/2 unit 0 family ethernet-switching vlan]** and **[edit interfaces xe-0/1/0 unit 0 family ethernet-switching]** hierarchy levels.

```
[edit interfaces xe-0/1/2 unit 0]
user@ex# set family ethernet-switching port-mode trunk
user@ex# set family ethernet-switching vlan members all

[edit interfaces xe-0/1/0 unit 0]
```

```
user@ex# set family ethernet-switching port-mode trunk
user@ex# set family ethernet-switching vlan members all
```

Repeat this step for every 10-Gigabit Ethernet trunk interface by using the appropriate interface name.

3. Enable the IPv4 address family for the loopback logical interface.

Include the **family** statement and specify the **inet** option to enable IPv4 at the **[edit interfaces lo0 unit 0]** hierarchy level.

```
[edit interfaces lo0 unit 0]
user@ex# set family inet
```

Repeat this step for every EX Series device by using the appropriate address for that device.

4. Configure the EX Series device management Ethernet interface.

This example configures the **unit 0** logical interface.

Include the **family** statement and specify the **inet** option at the **[edit me0 unit 0]** hierarchy level.

Include the **address** statement and specify **10.8.108.19/24** as the IPv4 address at the **[edit interfaces me0 unit 0 family inet]** hierarchy level.

```
[edit interfaces me0 unit 0]
user@ex# set family inet address 10.8.108.19/24
```

Repeat this step for every EX Series device by using the appropriate management interface address for that device.

### *Configuring VLANs in the Access Layer*

#### **Step-by-Step Procedure**

This procedure explains how to configure the VLAN names and tag IDs and associate trunk interfaces with one of the access layer devices. This procedure shows a representative sample of the configuration. The example does not show the configuration for every VLAN.

1. Configure the VLAN name and tag ID (number) for each VLAN on the EX Series device.

This example configures a VLAN with the name **vlan17** and tag ID **17**.

Include the **vlan-id** statement and specify **17** as the VLAN tag ID at the **[edit vlans vlan17]** hierarchy level.

```
[edit vlans vlan17]
user@ex# set vlan-id 17
```

Repeat this step for every VLAN on each EX Series device by using the appropriate VLAN names and tag IDs.

2. Associate the logical trunk interfaces with each VLAN on the EX Series device.

This example associates logical interfaces **xe-0/1/0.0** and **xe-0/1/2.0** with **vlan17**.

Include the **interface** statement and specify **xe-0/1/0.0** at the **[edit vlans vlan17]** hierarchy level.

Include the **interface** statement and specify **xe-0/1/2.0** at the **[edit vlans vlan17]** hierarchy level.

```
[edit vlans vlan17]
user@ex# set interface xe-0/1/0.0
user@ex# set interface xe-0/1/2.0
```

Repeat this step for every VLAN on each EX Series device by using the appropriate trunk interface names.

### ***Configuring a Redundant Trunk Group and Disabling the Spanning Tree Protocol for the Trunk Interfaces***

**Step-by-Step Procedure** This procedure explains how to configure a redundant trunk group and disable the Rapid Spanning Tree Protocol (RSTP) on the trunk interfaces.

1. Configure the trunk interfaces as a redundant trunk group.

This example configures the **xe-0/1/0.0** and **xe-0/1/2.0** trunk interfaces in a redundant trunk group named **rtgroup1**.

Include the **interface** statement at the **[edit ethernet-switching-options redundant-trunk-group group rtgroup1]** hierarchy level and specify each trunk interface name.

Include the **primary** statement at the **[edit ethernet-switching-options redundant-trunk-group group rtgroup1 xe-0/1/2.0]** hierarchy level.

```
[edit ethernet-switching-options redundant-trunk-group group rtgroup1]
user@ex# set interface xe-0/1/0.0
user@ex# set interface xe-0/1/2.0 primary
```

Repeat this step for every redundant trunk group by using the appropriate interface names.

2. Disable RSTP on the trunk interfaces.

On an EX Series device, RSTP is enabled by default. RSTP cannot be enabled on the same interface as routing.

This example disables RSTP on the **xe-0/1/0.0** and **xe-0/1/2.0** trunk interfaces.

Include the **disable** statement at the **[edit protocols rstp interface xe-0/1/0.0]** and **[edit protocols rstp interface xe-0/1/2.0]** hierarchy levels.

```
[edit protocols rstp]
user@ex# set interface xe-0/1/0.0 disable
user@ex# set interface xe-0/1/2.0 disable
```

Repeat this step for every core-facing trunk interface by using the appropriate interface name.

### Configuring Management Automation

**Step-by-Step Procedure** This procedure explains how to configure static routes to the management network, a known host to support background Secure Copy Protocol (SCP) file transfers, a commit script, and an event archive site.

1. Configure static routes so the Ethernet management interface can reach the management network.

Include the **route** statement, and specify **10.8.0.0/16** as the IPv4 subnet address of the management network at the **[edit routing-options static]** hierarchy level.

Include the **next-hop** statement, and specify the IPv4 host address of the next-hop router at the **[edit routing-options static route 10.8.0.0/16]** hierarchy level.

```
[edit routing-options static]
user@ex# set route 10.8.0.0/16 next-hop 10.8.108.254
```

Repeat this step for every Ethernet management interface on the EX Series devices.

2. Configure an SSH known host.

Include the **host** statement, and specify the IPv4 address and RSA host key options for trusted servers at the **[edit security ssh-known-hosts]** hierarchy level. In this example, the RSA host key is truncated to make it easier to read.

```
[edit security ssh-known-hosts]
user@ex# set host 127.0.0.1 rsa-key AAAAB3NzaC1yc2
```

Repeat this step for every EX Series device.

3. Configure outbound SSH to support Juniper Message Bundle (JMB) transfers to Juniper Support Systems (JSS).

In this example, the client ID is configured as **00187D0B670D**.

Include the **client** statement, specify **00187D0B670D** as the client ID, and specify **10.8.7.32** as the IPv4 address at the **[edit system services outbound-ssh]** hierarchy level.

Include the **port** statement and specify **7804** as the TCP port at the **[edit system services outbound-ssh client 00187D0B670D 10.8.7.32]** hierarchy level.

Include the **device-id** statement and specify **FA022D** as the device ID at the **[edit system services outbound-ssh client 00187D0B670D]** hierarchy level.

Include the **secret** statement at the **[edit system services outbound-ssh client 00187D0B670D ]** hierarchy level.

Include the **services** statement and specify **netconf** as the available service at the **[edit system services outbound-ssh client 00187D0B670D ]** hierarchy level.

```
[edit system services outbound-ssh client 00187D0B670D]
user@ex# set 10.8.7.32 port 7804
user@ex# set device-id FA022D
user@ex# set secret "9-9w4aik.QznDj9A0BEhrlKMxN"
user@ex# set services netconf
```

Repeat this step for every EX Series device.

4. Configure a commit script.

In this example the script file name is **jais-activate-scripts.slax**.

Include the **allow-transients** statement at the **[edit system scripts commit]** hierarchy level.

Include the **optional** statement at the **[edit system scripts commit file jais-activate-scripts.slax]** hierarchy level.

```
[edit system scripts commit]
user@ex# set allow-transients
user@ex# set file jais-activate-scripts.slax optional
```

Repeat this step for every EX Series devices.

5. Configure an event archive site.

In this example, the archive URL is the local **/var/tmp/** directory, and the name given to the destination is **juniper-aim**.

Include the **archive-sites** statement and specify the archive URL at the **[edit event-options destinations juniper-aim]** hierarchy level.

```
[edit event-options destinations juniper-aim]
user@ex# set archive-sites "scp://admin@127.0.0.1:/var/tmp" password "12345"
```

Repeat this step for every EX Series device.

### ***Configuring the Aggregation Layer in the Trusted Logical Systems***

Configure the aggregation layer by doing the following:

- [Configuring Interfaces in the Trusted Logical Systems on page 3488](#)
- [Configuring VLANs in the Aggregation Layer on page 3491](#)
- [Configuring the Virtual Router Routing Instance on page 3492](#)
- [Configuring Management Interfaces on page 3493](#)
- [Configuring Logical System Administrator Accounts on page 3494](#)
- [Configuring Management Automation on page 3495](#)

### ***Configuring Interfaces in the Trusted Logical Systems***

#### **Step-by-Step Procedure**

This procedure explains how to configure the physical, logical, and Layer 3 routing interfaces for the logical system in the trusted security zone of the aggregation layer. This procedure shows a representative sample of the configuration. The example does not show the configuration for every interface.

1. Enable flexible VLAN tagging on the physical interfaces.

This example configures physical interface **xe-1/0/0**.

Include the **encapsulation** statement and specify the **flexible-ethernet-services** option at the **[edit interfaces xe-1/0/0]** hierarchy level.

Include the **flexible-vlan-tagging** statement at the **[edit interfaces xe-1/0/0]** hierarchy level.

```
[edit interfaces xe-1/0/0]
user@mx# set encapsulation flexible-ethernet-services
user@mx# set flexible-vlan-tagging
```

Repeat this step for every physical interface connected to the EX series, SRX Series, and MX Series devices using the appropriate interface name.

2. Configure the 10-Gigabit Ethernet interfaces connected to the EX Series access layer device.

This example configures logical interface 17 on the **xe-1/0/0** interface under the logical system named **Trust1**.

Include the **encapsulation** statement and specify the **vlan-bridge** option at the **[edit logical-systems Trust1 interfaces xe-1/0/0 unit 17]** hierarchy level.

Include the **vlan-id** statement and specify 17 as the VLAN ID at the **[edit logical-systems Trust1 interfaces xe-1/0/0 unit 17]** hierarchy level.

```
[edit logical-systems Trust1 interfaces xe-1/0/0 unit 17]
user@mx# set encapsulation vlan-bridge
user@mx# set vlan-id 17
```

Repeat this step for every interface connected to the access layer devices by using the appropriate interface name, logical interface number, VLAN ID, and logical system name.

3. Configure the 10-Gigabit Ethernet interfaces connected to the other MX Series device shown in [Figure 76 on page 3482](#).

This example configures logical interface 17 on the **xe-0/1/0** interface.

Include the **encapsulation** statement and specify the **vlan-bridge** option at the **[edit logical-systems Trust1 interfaces xe-0/1/0 unit 17]** hierarchy level.

Include the **vlan-id** statement and specify 17 as the VLAN tag ID at the **[edit logical-systems Trust1 interfaces xe-0/1/0 unit 17]** hierarchy level.

```
[edit logical-systems Trust1 interfaces xe-0/1/0 unit 17]
user@mx# set encapsulation vlan-bridge
user@mx# set vlan-id 17
```

Repeat this step for every interface connected to the other MX Series device shown in [Figure 76 on page 3482](#) by using the appropriate interface name, logical interface number, VLAN ID, and logical system name.

4. Configure the 10-Gigabit Ethernet interface connected to the SRX Series device.

This example configures logical interface 15 on the **xe-1/1/0** interface. Include the **encapsulation** statement and specify the **vlan-bridge** option at the **[edit logical-systems Trust1 interfaces xe-1/1/0 unit 15]** hierarchy level.

Include the **vlan-id** statement and specify 15 as the VLAN tag ID at the **[edit logical-systems Trust1 interfaces xe-1/1/0 unit 15]** hierarchy level.

```
[edit logical-systems Trust1 interfaces xe-1/1/0 unit 15]
user@mx# set encapsulation vlan-bridge
user@mx# set vlan-id 15
```

Repeat this step for every interface connected to the SRX Series device by using the appropriate interface name, logical interface number, VLAN ID, and logical system name.

5. Configure the Layer 3 integrated routing and bridging (IRB) interface address.

This example configures the **unit 17** logical interface with **10.17.2.2/24** as the IPv4 address under the logical system named **Trust1**. Include the **address** statement and specify **10.17.2.2/24** as the IPv4 address at the **[edit logical-systems Trust1 interfaces irb unit 17 family inet]** hierarchy level.

```
[edit logical-systems Trust1 interfaces irb unit 17 family inet]
user@mx# set address 10.17.2.2/24
```

Repeat this step for every Layer 3 IBR by using the appropriate logical interface name and IPv4 address.

6. Configure the IRB interface to participate in Virtual Router Redundancy Protocol (VRRP).

This example configures the **unit 17** logical interface with **17** as the VRRP group name.

Include the **virtual-address** statement and specify **10.17.2.1** as the IPv4 address of the virtual router at the **[edit logical-systems Trust1 interfaces irb unit 17 family inet address 10.17.2.2/24 vrrp-group 17]** hierarchy level.

Include the **accept-data** statement at the **[edit logical-systems Trust1 interfaces irb unit 17 family inet address 10.17.2.2/24 vrrp-group 17]** hierarchy level so the interface will accept packets destined for the virtual IP address.

Include the **priority** statement and specify **200** as the router's priority at the **[edit logical-systems Trust1 interfaces irb unit 17 family inet address 10.17.2.2/24 vrrp-group 17]** hierarchy level.

Include the **fast-interval** statement and specify **200** as the interval between VRRP advertisements at the **[edit logical-systems Trust1 interfaces irb unit 17 family inet address 10.17.2.2/24 vrrp-group 17]** hierarchy level.

Include the **preempt** statement at the **[edit logical-systems Trust1 interfaces irb unit 17 family inet address 10.17.2.2/24 vrrp-group 17]** hierarchy level.

```
[edit logical-systems Trust1 interfaces irb unit 17 family inet address 10.17.2.2/24
vrrp-group 17]
user@mx# set virtual-address 10.17.2.1
user@mx# set accept-data
user@mx# set priority 200
user@mx# set fast-interval 200
user@mx# set preempt
```

Repeat this step for every Layer 3 IBR interface by using the appropriate logical interface name, IPv4 address, VRRP group name, and priority.



### Configuring VLANs in the Aggregation Layer

**Step-by-Step Procedure** This procedure explains how to configure the VLAN names and tag IDs and associate trunk interfaces and Layer 3 routing interfaces with each VLAN. This procedure shows a representative sample of the configuration. The example does not show the configuration for every VLAN.

1. Configure the VLAN name and tag ID (number) for each VLAN on the MX Series device.

This example configures a VLAN with the name **vlan17** and tag ID **17** in the Logical System **Trust1**. Include the **vlan-id** statement and specify **17** as the VLAN ID at the **[edit logical-systems Trust1 bridge-domains vlan17]** hierarchy level.

```
[edit logical-systems Trust1 bridge-domains vlan17]
user@mx# set vlan-id 17
```

Repeat this step for every VLAN on each MX Series device by using the appropriate VLAN names and tag IDs.

2. Associate the logical trunk interfaces with each VLAN on the MX Series device.

This example associates logical interface **xe-1/0/0.17** that is connected to the EX Series device and logical interface **xe-0/1/0.17** that is connected to the other MX Series device with **vlan17**.

Include the **interface** statement and specify **xe-1/0/0.17** at the **[edit logical-systems Trust1 bridge-domains vlan17]** hierarchy level.

Include the **interface** statement and specify **xe-0/1/0.17** at the **[edit logical-systems Trust1 bridge-domains vlan17]** hierarchy level.

```
[edit logical-systems Trust1 bridge-domains vlan17]
user@mx# set interface xe-1/0/0.17
user@mx# set interface xe-0/1/0.17
```

Repeat this step for every server-facing VLAN on each MX Series device by using the appropriate trunk interface names.

3. Associate a Layer 3 interface with each VLAN on the MX Series device.

This example associates the **irb.17** interface with **vlan17**.

Include the **routing-interface** statement and specify **irb.17** at the **[edit logical-systems Trust1 bridge-domains vlan17]** hierarchy level.

```
[edit logical-systems Trust1 bridge-domains vlan17]
user@mx# set routing-interface irb.17
```

Repeat this step for every server-facing VLAN on each MX Series device by using the appropriate Layer 3 interface name.

4. Associate the logical interfaces with each interconnection VLAN on the MX Series device.

This example associates logical interface **xe-1/1/0.15** that is connected to the SRX Series device and logical interface **xe-0/1/0.15** that is connected to the other MX Series device with **vlan15**.

Include the **interface** statement and specify **xe-1/1/0.15** at the **[edit logical-systems Trust1 bridge-domains vlan15]** hierarchy level.

Include the **interface** statement and specify **xe-0/1/0.15** at the **[edit logical-systems Trust1 bridge-domains vlan15]** hierarchy level.

```
[edit logical-systems Trust1 bridge-domains vlan15]
user@mx# set interface xe-1/1/0.15
user@mx# set interface xe-0/1/0.15
```

Repeat this step for every interconnect VLAN on each MX Series device by using the appropriate interconnect interface names.

5. Associate a Layer 3 interface with each interconnection VLAN on the MX Series device to support active participation in the OSPF protocol.

This example associates the **irb.15** interface with **vlan15**.

Include the **routing-interface** statement and specify **irb.15** at the **[edit logical-systems Trust1 bridge-domains vlan15]** hierarchy level.

```
[edit logical-systems Trust1 bridge-domains vlan15]
user@mx# set routing-interface irb.15
```

Repeat this step for every server-facing VLAN on each MX Series device by using the appropriate Layer 3 interface name.

### ***Configuring the Virtual Router Routing Instance***

#### **Step-by-Step Procedure**

This procedure explains how to configure a single virtual router routing instance. This procedure shows a representative sample of the example configuration. The example does not show the configuration for every device.

1. Configure the routing instance type.

This example configures the routing instance with the name **MX-VR2**. Include the **instance-type** statement and specify **virtual-router** as the type at the **[edit logical-systems Trust1 routing-instances MX-VR2]** hierarchy level.

```
[edit logical-systems Trust1 routing-instances MX-VR2]
user@mx# set instance-type virtual-router
```

Repeat this step for every virtual router in each MX Series device by using the appropriate virtual router name.

2. Add the IRB interfaces used by the virtual router routing instance.

Include the **interface** statement and specify the name of each IRB interface at the **[edit routing-instances MX-VR2]** hierarchy level.

```
[edit logical-systems Trust1 routing-instances MX-VR2]
user@mx# set interface irb.15
user@mx# set interface irb.16
user@mx# set interface irb.17
user@mx# set interface irb.18
user@mx# set interface irb.1002
```

Repeat this step for every virtual router in each MX Series device by using the appropriate interface names.

3. Configure the IGP protocol active interface used by the virtual router routing instance so the routing tables can be populated with the routes to the servers.

This example configures one IRB interface to actively participate in the OSPF protocol area **0.0.0.0**.

Include the **interface** statement and specify the name of the IRB interface at the **[edit logical-systems Trust1 routing-instances MX-VR2 protocols ospf area 0.0.0.0]** hierarchy level.

```
[edit routing-instances MX-VR2 protocols ospf area 0.0.0.0]
user@mx# set interface irb.15
```

Repeat this step for every virtual router in each MX Series device by using the appropriate virtual router name.

4. Configure the interior gateway protocol passive interfaces that are associated with each VLAN within the virtual router routing instance.

This example configures the IRB interfaces to passively participate in the OSPF protocol area **0.0.0.0**.

Include the **passive** statement at the **[edit logical-systems Trust1 routing-instances MX-VR2 protocols ospf area 0.0.0.0 interface *irb-name*]** hierarchy level.

```
[edit logical-systems Trust1 routing-instances MX-VR2 protocols ospf area 0.0.0.0]
user@mx# set interface irb.16 passive
user@mx# set interface irb.17 passive
user@mx# set interface irb.18 passive
user@mx# set interface irb.1002 passive
```

Repeat this step for every virtual router in each MX Series device by using the appropriate virtual router name.

5. Configure the logical system router identifier.

Include the **router-id** statement and specify **10.200.11.101** as the router identifier at the **[edit logical-systems Trust1 routing-instances MX-VR2 routing-options]** hierarchy level.

```
[edit logical-systems Trust1 routing-instances MX-VR2 routing-options]
user@mx# set router-id 10.200.11.101
```

Repeat this step for every virtual router in each MX Series device by using the appropriate router identifier.

### *Configuring Management Interfaces*

#### **Step-by-Step Procedure**

This procedure explains how to configure static routes to the management network and the IPv4 address family for the loopback logical interface. This procedure shows a representative sample of the configuration. The example does not show the configuration for every interface.

1. Configure static routes so the Ethernet management interface can reach the management network.

Include the **route** statement and specify **10.0.0.0/8** as the IPv4 subnet address of the management network at the **[edit routing-options static]** hierarchy level.

Include the **next-hop** statement, and specify the IPv4 host address of the next-hop router at the **[edit routing-options static route 10.0.0.0/8]** hierarchy level.

Include the **retain** and **no-readvertise** statements at the **[edit routing-options static route 10.0.0.0/8]** hierarchy level.

```
[edit routing-options static]
user@mx# set route 10.0.0.0/8 next-hop 10.8.3.254
user@mx# set route 10.0.0.0/8 retain
user@mx# set route 10.0.0.0/8 no-readvertise
```

Repeat this step for every MX Series device.

2. Configure the MX Series device management Ethernet interface. This example configures the **unit 0** logical interface.

Include the **family** statement and specify the **inet** option at the **[edit fxp0 unit 0]** hierarchy level.

Include the **address** statement and specify **10.8.3.212/24** as the IPv4 address at the **[edit interfaces fxp0 unit 0]** hierarchy level.

```
[edit interfaces fxp0 unit 0]
user@mx# set family inet address 10.8.3.212/24
```

Repeat this step for every MX Series device by using the appropriate management interface address for that device.

3. Configure the loopback logical interface.

Include the **family** statement and specify the **inet** option at the **[edit interfaces lo0 unit 0]** hierarchy level.

```
[edit interfaces lo0 unit 0]
user@mx# set family inet
```

Repeat this step for every MX Series device.

### ***Configuring Logical System Administrator Accounts***

#### **Step-by-Step Procedure**

This procedure explains how to configure administrator account classes that are confined to the context of the logical system to which they are assigned and administrator accounts for each logical system.

1. Create administrator account classes.

In this example, the **trust1-admin** user class is created with **all** permissions for the **Trust1** logical system.

Include the **class** statement and specify **trust1-admin** as the class name at the **[edit system login]** hierarchy level.

Include the **logical-system** statement and specify **Trust1** as the logical system name at the **[edit system login class trust1-admin]** hierarchy level.

Include the **permissions** statement and specify the **all** option at the **[edit system login class trust1-admin]** hierarchy level.

```
[edit system]
```

```
user@mx# set login class trust1-admin logical-system Trust1
user@mx# set login class trust1-admin permissions all
```

Repeat this step for the trust2-admin and untrust-admin classes on each MX Series device by using the appropriate logical-system name.

2. Create administrator accounts that correspond to each logical system in the MX Series device.

In this example, the **trust1** user account is created and assigned the **trust1-admin** class.

Include the **class** statement and specify **trust1-admin** as the user class at the **[edit system login user trust1]** hierarchy level.

Include the **encrypted-password** statement and enter the encrypted password string at the **[edit system login user trust1 authentication]** hierarchy level.

```
[edit system]
user@mx# set login user trust1 class trust1-admin
user@mx# set login user trust1 authentication encrypted-password 12345
```

Repeat this step for the trust2 and untrust user accounts on each MX Series device.

### *Configuring Management Automation*

**Step-by-Step Procedure** This procedure explains how to configure a known host to support background SCP file transfers, a commit script, and an archive site.

1. Configure a commit script.

In this example, the script file name is **jais-activate-scripts.slax**.

Include the **allow-transients** statement at the **[edit system scripts commit]** hierarchy level.

Include the **optional** statement at the **[edit system scripts commit file jais-activate-scripts.slax]** hierarchy level.

```
[edit system scripts commit]
user@mx# set allow-transients
user@mx# set file jais-activate-scripts.slax optional
```

2. Configure an event archive site.

In this example the archive URL is the local **/var/tmp/** directory, and the name given to the destination is **juniper-aim**.

Include the **archive-sites** statement and specify the archive URL at the **[edit event-options destinations juniper-aim]** hierarchy level.

```
[edit event-options destinations juniper-aim]
user@mx# set archive-sites "scp://admin@127.0.0.1:/var/tmp" password "12345"
```

### ***Configuring the Core Layer in the Untrusted Logical Systems***

Configure the core layer by doing the following:

- [Configuring Interfaces in the Untrusted Logical Systems on page 3496](#)
- [Configuring VLANs in the Core Layer on page 3497](#)
- [Configuring Protocols in the Untrusted Logical System on page 3498](#)

### ***Configuring Interfaces in the Untrusted Logical Systems***

#### **Step-by-Step Procedure**

This procedure explains how to configure the physical, logical, and Layer 3 routing interfaces for the logical system in the untrusted security zone of the core layer. This procedure shows a representative sample of the configuration. The example does not show the configuration for every interface.

1. Configure the 10-Gigabit redundant Ethernet interfaces connected to the other MX Series device shown in [Figure 76 on page 3482](#).

This example configures logical interface **19** on the **xe-0/3/0** interface under the logical system named **Untrust** to participate in VLAN 19. Include the **encapsulation** statement and specify the **vlan-bridge** option at the **[edit logical-systems Untrust interfaces xe-0/3/0 unit 19]** hierarchy level.

Include the **vlan-id** statement and specify **19** as the VLAN tag ID at the **[edit logical-systems Untrust interfaces xe-0/3/0 unit 19]** hierarchy level.

```
[edit logical-systems Untrust interfaces xe-0/3/0 unit 19]
user@mx# set encapsulation vlan-bridge
user@mx# set vlan-id 19
```

Repeat this step for every redundant Ethernet interface connected to the other MX Series device by using the appropriate interface name, logical interface number, VLAN ID, and logical system name.

2. Configure the 10-Gigabit Ethernet interfaces connected to the SRX Series device.

This example configures logical interface **19** on the **xe-2/2/0** interface under the logical system named **Untrust** to participate in VLAN 19.

Include the **encapsulation** statement and specify the **vlan-bridge** option at the **[edit logical-systems Untrust interfaces xe-2/2/0 unit 19]** hierarchy level.

Include the **vlan-id** statement and specify **19** as the VLAN tag ID at the **[edit logical-systems Untrust interfaces xe-2/2/0 unit 19]** hierarchy level.

```
[edit logical-systems Untrust interfaces xe-2/2/0 unit 19]
user@mx# set encapsulation vlan-bridge
user@mx# set vlan-id 19
```

Repeat this step for every redundant Ethernet interface connected to the SRX Series device by using the appropriate interface name, logical interface number, VLAN ID, and logical system name.

3. Configure the 10-Gigabit Ethernet interfaces connected to the IP-based/MPLS-based core network.

This example configures logical interface **0** on the **xe-1/3/0** interface under the logical system named **Untrust**.

Include the **address** statement and specify **10.200.4.1/30** as the IPv4 address at the **[edit logical-systems Untrust interfaces xe-1/3/0 unit 0 family inet]** hierarchy level.

Include the **family** statement and specify the **mpls** option at the **[edit logical-systems Untrust interfaces xe-1/3/0 unit 0]** hierarchy level.

```
[edit logical-systems Untrust interfaces xe-1/3/0 unit 0]
user@mx# set family inet address 10.200.4.1/30
user@mx# set family mpls
```

Repeat this step for every 10-Gigabit Ethernet interface connected to the service provider network by using the appropriate interface name, logical interface number, IPv4 address, and logical system name.

4. Configure the Layer 3 IRB interface address.

This example configures the **unit 19** logical interface that participates in VLAN 19 with **10.19.2.1/24** as the IPv4 address under the logical system named **Untrust**.

Include the **address** statement and specify **10.19.2.1/24** as the IPv4 address at the **[edit logical-systems Untrust interfaces irb unit 19 family inet]** hierarchy level.

```
[edit logical-systems Untrust interfaces irb unit 19 family inet]
user@mx# set address 10.19.2.1/24
```

Repeat this step for every Layer 3 IRB interface by using the appropriate logical interface name and IPv4 address.

5. Configure an IP address for the loopback logical interface of the Logical System **Untrust**.

Include the **address** statement and specify **10.200.11.1/32** as the IPv4 address at the **[edit logical-systems Untrust interfaces lo0 unit 1 family inet]** hierarchy level.

```
[edit logical-systems Untrust interfaces lo0 unit 1 family inet]
user@mx# set address 10.200.11.1/32
```

Repeat this step for every MX Series device by using the appropriate IPv4 address.

### *Configuring VLANs in the Core Layer*

#### **Step-by-Step Procedure**

This procedure explains how to configure the VLAN names and tag IDs and associate interfaces and Layer 3 routing interfaces with each core interconnect VLAN. This procedure shows a representative sample of the configuration. The example does not show the configuration for every VLAN.

1. Configure the VLAN name and tag ID (number) for each core interconnect VLAN on the MX Series device.

This example configures a VLAN with the name **vlan14** and tag ID **14** in the Logical System **Untrust**.

Include the **vlan-id** statement and specify **14** as the VLAN ID at the **[edit logical-systems Untrust bridge-domains vlan14]** hierarchy level.

```
[edit logical-systems Untrust bridge-domains vlan14]
user@mx# set vlan-id 14
```

Repeat this step for every VLAN on each MX Series device by using the appropriate VLAN names and tag IDs.

2. Associate the logical interfaces with each VLAN on the MX Series device.

This example associates logical interface **xe-0/3/0.14** that is connected to the other MX Series device and **xe-2/2/0.14** that is connected to the SRX Series device with **vlan14**.

Include the **interface** statement and specify **xe-0/3/0.14** at the **[edit logical-systems Untrust bridge-domains vlan14]** hierarchy level.

Include the **interface** statement and specify **xe-2/2/0.14** at the **[edit logical-systems Untrust bridge-domains vlan14]** hierarchy level.

```
[edit logical-systems Untrust bridge-domains vlan14]
user@mx# set interface xe-0/3/0.14
user@mx# set interface xe-2/2/0.14
```

Repeat this step for every core interconnect VLAN on each MX Series device by using the appropriate interface names.

3. Associate a Layer 3 interface with each VLAN on the MX Series device.

This example associates the **irb.14** interface with **vlan14**.

Include the **routing-interface** statement and specify **irb.14** at the **[edit logical-systems Untrust bridge-domains vlan14]** hierarchy level.

```
[edit logical-systems Untrust bridge-domains vlan14]
user@mx# set routing-interface irb.14
```

Repeat this step for every core interconnect VLAN on each MX Series device by using the appropriate Layer 3 interface name.

### *Configuring Protocols in the Untrusted Logical System*

#### **Step-by-Step Procedure**

This procedure explains how to configure the BGP, MPLS, RSVP, and OSPF protocols for the Logical System Untrust. This procedure shows a representative sample of the configuration. The example does not show the configuration for every device.

1. Add interfaces to the OSPF protocol on the MX Series device.

This example adds logical interfaces **xe-1/3/0.0** and **lo0.1** to the OSPF protocol used in the core network.

Include the **interface** statement and specify the **xe-1/3/0.0** and **lo0.1** interfaces at the **[edit logical-systems Untrust protocols ospf area 0.0.0.0]** hierarchy level.

```
[edit logical-systems Untrust protocols ospf area 0.0.0.0]
user@mx# set interface xe-1/3/0.0
user@mx# set interface lo0.1
```

Repeat this step for every 10-Gigabit Ethernet interface connected to the core layer devices by using the appropriate interface name.



2. Configure the Generic Router Encapsulation (GRE) tunnel.

This example enables a dynamic GRE tunnel named **GRE1**.

Include the **gre** statement to specify the tunnel type at the **[edit logical-systems Untrust routing-options dynamic-tunnel GRE1]** hierarchy level.

Include the **source-address** statement and specify **10.200.11.1** as the IPv4 source address at the **[edit logical-systems Untrust routing-options dynamic-tunnel GRE1]** hierarchy level.

Include the **destination-networks** statement and specify **0.0.0.0/0** as the destination prefix at the **[edit logical-systems Untrust routing-options dynamic-tunnel GRE1]** hierarchy level.

```
[edit logical-systems Untrust routing-options dynamic-tunnel GRE1]
user@mx# set source-address 10.200.11.1
user@mx# set gre
user@mx# set destination-networks 0.0.0.0/0
```

Repeat this step for each MX Series device by using the appropriate source address.

3. Configure the Logical System local autonomous system number and router identifier.

Include the **autonomous-system** statement and specify **65000** as the autonomous system number at the **[edit logical-systems Untrust routing-options]** hierarchy level.

Include the **router-id** statement and specify **10.200.11.101** as the router identifier at the **[edit logical-systems Untrust routing-options]** hierarchy level.

```
[edit logical-systems Untrust]
user@mx# set routing-options autonomous-system 65000
user@mx# set routing-options router-id 10.200.11.101
```

Repeat this step for each MX Series device by using the appropriate router identifier and autonomous system number 65000.

4. Configure the internal BGP peer group.

Include the **type** statement and specify the **internal** option at the **[edit logical-systems Untrust protocols bgp group int]** hierarchy level.

Include the **local-address** statement and specify the router ID (10.200.11.1) of Logical System Untrust as the local address at the **[edit logical-systems Untrust protocols bgp group int]** hierarchy level.

Include the **unicast** statement at the **[edit logical-systems Untrust protocols bgp group int family inet]** and **[edit logical-systems Untrust protocols bgp group int family inet-vpn]** hierarchy levels.

Include the **local-as** statement and specify **65000** as the local autonomous system number at the **[edit logical-systems Untrust protocols bgp group int]** hierarchy level.

Include the **peer-as** statement and specify **65000** as the peer autonomous system number at the **[edit logical-systems Untrust protocols bgp group int]** hierarchy level.

Include the **neighbor** statement and specify the neighbor IPv4 addresses at the **[edit logical-systems Untrust protocols bgp group int]** hierarchy level.

The neighbor addresses are the router ID addresses of the other MX Series device in the local data center, MX Series devices in a remote data center, and routers located in the IP-based/MPLS-based core network.

```
[edit logical-systems Untrust protocols bgp group int]
user@mx# set type internal
user@mx# set local-address 10.200.11.1
user@mx# set family inet unicast
user@mx# set family inet-vpn unicast
user@mx# set local-as 65000
user@mx# set peer-as 65000
user@mx# set neighbor 10.200.11.2
user@mx# set neighbor 10.200.11.3
user@mx# set neighbor 10.200.11.4
```

Repeat this step for every MX Series device.

5. Add interfaces to the MPLS protocol used in the service provider core network.

This example adds the **xe-1/3/0.0** and **xe-2/3/0.0** interfaces that are connected to the service provider core network.

Include the **interface** statement and specify the **xe-1/3/0.0** and **xe-2/3/0.0** interfaces at the **[edit logical-systems Untrust protocols mpls]** hierarchy level.

```
[edit logical-systems Untrust protocols mpls]
user@mx# set interface xe-1/3/0.0
user@mx# set interface xe-2/3/0.0
```

Repeat this step for every MX Series device.

6. Create an MPLS LSP to the router that is located in the MPLS-based core network.

This example creates an LSP named **to-core-router**.

Include the **to** statement and specify **10.200.11.3** as the IPv4 address of the core router at the **[edit logical-systems Untrust protocols mpls label-switched-path to-core-router]** hierarchy level.

Include the **no-cspf** statement at the **[edit logical-systems Untrust protocols mpls]** hierarchy level.

```
[edit logical-systems Untrust protocols mpls]
user@mx# set label-switched-path to-core-router to 10.200.11.3
user@mx# set no-cspf
```

Repeat this step for every MX Series device.

7. Add interfaces to the RSVP protocol used in the MPLS-based core network.

Include the **interface** statement and specify the **xe-1/3/0.0** and **xe-2/3/0.0** interfaces at the **[edit logical-systems Untrust protocols rsvp]** hierarchy level.

```
[edit logical-systems Untrust protocols rsvp]
user@mx# set interface xe-1/3/0.0
user@mx# set interface xe-2/3/0.0
```

Repeat this step for every MX Series device.

### *Configuring the Security Device*

The following procedures explain how to configure the redundant Ethernet interfaces, node cluster, security zones, security policies, and routing policies for the trusted security zone of the access layer.

- [Configuring the Redundant Ethernet Interface Link Aggregation Group on page 3501](#)
- [Configuring the SRX Series Cluster on page 3502](#)
- [Creating Security Zones and Configuring the In-Bound Traffic Policy Action on page 3503](#)
- [Configuring the Security Zone Policies on page 3504](#)
- [Creating the Routing Policies on page 3506](#)
- [Configuring the Virtual Router Routing Instance on page 3508](#)
- [Results on page 3511](#)

### *Configuring the Redundant Ethernet Interface Link Aggregation Group*

#### **Step-by-Step Procedure**

This procedure explains how to configure the redundant Ethernet interface link aggregation group. This procedure shows a representative sample of the configuration. The example does not show the configuration for every interface.

1. Configure the number of aggregated Ethernet interfaces supported on the node.

This example enables support for two interfaces.

Include the **device-count** statement and specify **2** as the number of interfaces supported at the **[edit chassis aggregated-devices ethernet]** hierarchy level.

```
[edit chassis aggregated-devices ethernet]
user@srx# set device-count 2
```

Repeat this step for every SRX Series device by using the appropriate device count.

2. Assign 10-Gigabit Ethernet child interfaces to the redundant Ethernet (reth) parent interface.

This example assigns the **xe-1/0/0** 10-Gigabit Ethernet child interface to the **reth1** parent interface on Node0.

Include the **redundant-parent** statement and specify **reth1** as the parent interface at the **[edit interfaces xe-1/0/0 gigether-options]** hierarchy level.

```
[edit interfaces xe-1/0/0 gigether-options]
user@srx# set redundant-parent reth1
```

Repeat this step for every redundant Ethernet interface by using the appropriate interface name and redundant parent name.

3. Configure the redundant Ethernet parent interface options.

This example configures the **reth1** redundant parent interface.

Include the **redundancy-group** statement and specify **1** as the group number at the **[edit interfaces reth1 redundant-ether-options]** hierarchy level.

Include the **vlan-tagging** statement at the **[edit interfaces reth1]** hierarchy level.

```
[edit interfaces reth1]
user@srx# set redundant-ether-options redundancy-group 1
user@srx# set vlan-tagging
```

Repeat this step for every redundant parent interface by using the appropriate redundant parent name and redundancy group number.

4. Configure the redundant Ethernet parent logical interfaces.

This example configures the **unit 15** logical interface.

Include the **address** statement and specify **10.15.2.2/24** as the IPv4 address at the **[edit interfaces reth1 unit 15 family inet]** hierarchy level.

Include the **vlan-id** statement and specify **15** as the VLAN identifier at the **[edit interfaces reth1 unit 15]** hierarchy level.

```
[edit interfaces reth1 unit 15]
user@srx# set family inet address 10.15.2.2/24
user@srx# set vlan-id 15
```

Repeat this step for every redundant parent interface by using the appropriate redundant parent name, IPv4 address, and VLAN identifier.

### *Configuring the SRX Series Cluster*

#### **Step-by-Step Procedure**

This procedure explains how to configure fabric connections between the nodes in the cluster. This procedure shows a representative sample of the configuration. The example does not show the configuration for every interface.

1. Configure the 10-Gigabit Ethernet interface to serve as the fabric between the cluster nodes.

This example configures **xe-1/0/1** as the child fabric interface and **fab0** as the parent fabric interface. The connection is from SRX0 to SRX1.

Include the **member-interfaces** statement and specify the **xe-1/0/1** interface at the **[edit interfaces fab0 fabric-options]** hierarchy level.

```
[edit interfaces fab0 fabric-options]
user@srx# set member-interfaces xe-1/0/1
```

Repeat this step for every 10-Gigabit Ethernet interface that is part of the cluster fabric by using the appropriate child interface name and parent interface name.

2. Configure the number of redundant Ethernet interfaces that the cluster supports.

This example configures **4** as the number of interfaces.

Include the **reth-count** statement and specify **4** as the number of interfaces at the **[edit chassis cluster]** hierarchy level.

```
[edit chassis cluster]
user@srx# set reth-count 4
```

Repeat this step for every SRX Series device in the cluster.

3. Configure the node priority for the redundancy group to determine which node is primary and which is secondary.

This example configures **node 0** with a higher priority.

Include the **priority** statement and specify **200** at the **[edit chassis cluster redundancy-group 1 node 0]** hierarchy level.

Include the **priority** statement and specify **100** at the **[edit chassis cluster redundancy-group 1 node 1]** hierarchy level.

```
[edit chassis cluster redundancy-group 1]
user@srx# set node 0 priority 200
user@srx# set node 1 priority 100
```

Repeat this step for every redundancy group on every SRX Series device in the cluster.

4. Allow a node with a higher priority to initiate a failover to become the primary node for the redundancy group.

Include the **preempt** statement at the **[edit chassis cluster redundancy-group 1]** hierarchy level.

```
[edit chassis cluster redundancy-group 1]
user@srx# set preempt
```

Repeat this step for every redundancy group on every SRX Series device in the cluster.

5. Enable control link recovery to be done automatically.

Include the **control-link-recovery** statement at the **[edit chassis cluster]** hierarchy level.

```
[edit chassis cluster]
user@srx# set control-link-recovery
```

Repeat this step for every redundancy group on every SRX Series device in the cluster.

6. Enable interface monitoring to monitor the health of the interfaces and trigger redundancy group failover.

This example configures the **xe-1/0/0** interface with a weight of **255**.

Include the **weight** statement at the **[edit chassis cluster redundancy-group 1 interface-monitor xe-1/0/0]** hierarchy level.

```
[edit chassis cluster redundancy-group 1 interface-monitor xe-1/0/0]
user@srx# set weight 255
```

Repeat this step for every redundancy group interface on every SRX Series device in the cluster.

### ***Creating Security Zones and Configuring the In-Bound Traffic Policy Action***

#### **Step-by-Step Procedure**

This procedure explains how to configure the trusted and untrusted security zones on the SRX Series device. This procedure shows a representative sample of the configuration. The example does not show the configuration for every zone.

1. Assign a redundant Ethernet logical interface to a trusted zones.

This example assigns the **reth1.15** interface to the **Trust2** zone.

Include the **interfaces** statement and specify **reth1.15** as the interface in the zone at the **[edit security zones security-zone Trust2]** hierarchy level.

```
[edit security zones security-zone Trust2]
user@srx# set interfaces reth1.15
```

Repeat this step for every trusted security zone by using the appropriate zone name and redundant Ethernet logical interface name.

2. Assign a redundant Ethernet logical interface to the untrusted zones.

This example assigns the **reth2.19** interface to the **Untrust2** zone.

Include the **interfaces** statement and specify **reth2.19** as the interface in the zone at the **[edit security zones security-zone Untrust2]** hierarchy level.

```
[edit security zones security-zone Untrust2]
user@srx# set interfaces reth2.19
```

Repeat this step for every untrusted security zone by using the appropriate zone name and redundant Ethernet logical interface name.

3. Enable all inbound system services traffic in the trusted security zone.

This example enables all services for the **Trust2** zone.

Include the **system-services** statement and specify the **all** option at the **[edit security zones security-zone Trust2 host-inbound-traffic]** hierarchy level.

```
[edit security zones security-zone Trust2 host-inbound-traffic]
user@srx# set system-services all
```

Repeat this step for every security zone on the SRX Series device where system services are allowed.

4. Enable all protocols for inbound traffic in the trusted security zone.

This example enables all protocols for the **Trust2** zone.

Include the **protocols** statement and specify the **all** option at the **[edit security zones security-zone Trust2 host-inbound-traffic]** hierarchy level.

```
[edit security zones security-zone Trust2 host-inbound-traffic]
user@srx# set protocols all
```

Repeat this step for every security zone on the SRX Series device where all protocols are allowed for inbound traffic.

### ***Configuring the Security Zone Policies***

#### **Step-by-Step Procedure**

This procedure explains how to configure the security zone policies on the SRX Series device. This procedure shows a representative sample of the configuration. The example does not show the configuration for every policy.

1. Define which zone traffic is coming from and which zone traffic is going to for the policy being created.

This example defines the from zone as **Trust2** and the to zone as **Untrust2**.

On a single command line, include the **from-zone** statement and specify **Trust2**, include the **to-zone** statement and specify **Untrust2**, include the **policy** statement and specify **denyftp** as the policy name, and included the **match** statement at the **[edit security policies]** hierarchy level.

```
[edit security policies]
user@srx# set from-zone Trust2 to-zone Untrust2 policy denyftp match
```

Repeat this step for every policy that controls traffic between zones.

2. Configure the policy match criteria for denying traffic.

This example matches the Junos OS FTP application from any source to any destination address in a policy named **denyftp**.

Include the **source-address** statement and specify **any** as the IPv4 address at the **[edit security policies from-zone Trust2 to-zone Untrust2 policy denyftp match]** hierarchy level.

Include the **destination-address** statement and specify **any** as the IPv4 address at the **[edit security policies from-zone Trust2 to-zone Untrust2 policy denyftp match]** hierarchy level.

Include the **application** statement and specify **junos-ftp** as the application at the **[edit security policies from-zone Trust2 to-zone Untrust2 policy denyftp match]** hierarchy level.

```
[edit security policies from-zone Trust2 to-zone Untrust2 policy denyftp match]
user@srx# set source-address any
user@srx# set destination-address any
user@srx# set application junos-ftp
```

Repeat this step for every protocol matching policy by using the correct protocol.

3. Block specific applications from passing from the Trust2 zone to the Untrust2 zone.

This example denies the Junos OS FTP application from the **Trust2** zone to the **Untrust2** zone.

Include the **deny** statement at the **[edit security policies from-zone Trust2 to-zone Untrust2 policy denyftp then]** hierarchy level.

```
[edit security policies from-zone Trust2 to-zone Untrust2 policy denyftp then]
user@srx# set deny
```

Repeat this step for every deny policy.

4. Configure the policy match criteria for allowing traffic.

This example matches any application from any source to any destination address in a policy named **allow\_all**.

Include the **source-address** statement and specify **any** as the IPv4 address at the **[edit security policies from-zone Trust2 to-zone Untrust2 policy allow\_all match]** hierarchy level.

Include the **destination-address** statement and specify **any** as the IPv4 address at the **[edit security policies from-zone Trust2 to-zone Untrust2 policy allow\_all match]** hierarchy level.

Include the **application** statement and specify **any** as the application at the **[edit security policies from-zone Trust2 to-zone Untrust2 policy allow\_all match]** hierarchy level.

```
[edit security policies from-zone Trust2 to-zone Untrust2 policy allow_all match]
user@srx# set source-address any
user@srx# set destination-address any
user@srx# set application any
```

Repeat this step for every application matching policy.

5. Permit any application traffic to pass from the Trust2 zone to the Untrust2 zone.

This example allows any application traffic from the **Trust2** zone to the **Untrust2** zone.

Include the **permit** statement at the **[edit security policies from-zone Trust2 to-zone Untrust2 policy allow\_all then]** hierarchy level.

```
[edit security policies from-zone Trust2 to-zone Untrust2 policy allow_all then]
user@srx# set permit
```

Repeat this step for every permit policy.

### *Creating the Routing Policies*

#### **Step-by-Step Procedure**

This procedure explains how to create the routing policies on the SRX Series device that can be applied to the appropriate routing instances. This procedure shows a representative sample of the configuration. The example does not show the configuration for every policy.

1. Create a policy to set the local preference for BGP routes to 120.

This example creates a policy named **local-pref-120** that sets the BGP local preference value for received routes advertised by BGP to 120.

Include the **protocol** statement and specify **bgp** as the value at the **[edit policy-options policy-statement local-pref-120 term term1 from]** hierarchy level.

Include the **local-preference** statement and specify **120** as the value at the **[edit policy-options policy-statement local-pref-120 term term1 then]** hierarchy level.

```
[edit policy-options policy-statement local-pref-120]
user@srx# set term term1 from protocol bgp
user@srx# set term term1 then local-preference 120
```

Repeat this step for each SRX Series device.

2. Configure the match criteria for a policy named **default-ospf** to accept all aggregate (generated) routes.

Include the **protocol** statement and specify **aggregate** as the protocol to match at the **[edit policy-options policy-statement default-ospf term term1 from]** hierarchy level.

Include the **route-filter** statement and specify **0.0.0.0/0 exact** as the match criteria at the **[edit policy-options policy-statement default-ospf term term1 from]** hierarchy level.



```
[edit policy-options policy-statement default-ospf term term1 from]
user@srx# set protocol aggregate
user@srx# set route-filter 0.0.0.0/0 exact
```

Repeat this step for each SRX Series device.

3. Configure the action for a policy to set the metric to **0**, and set the external route type to **1**.

This example configures a policy named **default-ospf** that sets the metric to **0**, sets the external route to type **1**, and accepts aggregate routes into the routing table.

Include the **metric** statement and specify **0** as the external type at the **[edit policy-options policy-statement default-ospf term term1 then]** hierarchy level.

Include the **type** statement and specify **1** as the external route type at the **[edit policy-options policy-statement default-ospf term term1 then external]** hierarchy level.

Include the **accept** statement at the **[edit policy-options policy-statement default-ospf term term1 then]** hierarchy level.

```
[edit policy-options policy-statement default-ospf term term1 then]
user@srx# set metric 0
user@srx# set external type 1
user@srx# set accept
```

Repeat this step for each SRX Series device.

4. Create a policy that accepts OSPF routes with specified prefixes.

This example creates a policy named **trust2-ebgp-out** that accepts OSPF routes with the route prefixes that correspond to the subnets for each trust VLAN.

Include the **protocol** statement and specify **ospf** as the protocol at the **[edit policy-options policy-statement trust2-ebgp-out term term1 from]** hierarchy level.

Include the **route-filter** statement and specify the VLAN subnet addresses and the **exact** match keyword at the **[edit policy-options policy-statement trust2-ebgp-out term term1 from]** hierarchy level.

Include the **accept** statement at the **[edit policy-options policy-statement trust2-ebgp-out term term1 then]** hierarchy level.

```
[edit policy-options policy-statement trust2-ebgp-out term term1]
user@srx# set from protocol ospf
user@srx# set from route-filter 10.16.2.0/24 exact
user@srx# set from route-filter 10.17.2.0/24 exact
user@srx# set from route-filter 10.18.2.0/24 exact
user@srx# set then accept
```

Repeat this step for each SRX Series device.

5. Create a policy that accepts BGP routes if the route type is external.

This example creates a policy named **check-bgp-routes** that accepts BGP routes only if the route type is external.

Include the **protocol** statement and specify **bgp** as the protocol at the **[edit policy-options policy-statement check-bgp-routes term term1 from]** hierarchy level.

Include the **route-type** statement and specify the **external** option at the **[edit policy-options policy-statement check-bgp-routes term term1 from]** hierarchy level.

Include the **accept** statement at the **[edit policy-options policy-statement check-bgp-routes term term1 then]** hierarchy level.

```
[edit policy-options policy-statement check-bgp-routes term term1]
user@srx# set from protocol bgp
user@srx# set from route-type external
user@srx# set then accept
```

Repeat this step for each SRX Series device.

6. Create a policy that accepts routes from other virtual router routing instances.

This example creates a policy named **from\_srx\_vr1** that accepts routes from routing instance **SRX-VR1**.

Include the **instance** statement and specify **SRX-VR1** as the routing instance name at the **[edit policy-options policy-statement from\_srx\_vr1 term term1 from]** hierarchy level.

Include the **accept** statement at the **[edit policy-options policy-statement from\_srx\_vr1 term term1 then]** hierarchy level.

```
[edit policy-options policy-statement from_srx_vr1 term term1]
user@srx# set from instance SRX-VR1
user@srx# set then accept
```

Repeat this step for each virtual router in each SRX Series device.

### *Configuring the Virtual Router Routing Instance*

#### **Step-by-Step Procedure**

This procedure explains how to configure a single virtual router routing instance. This procedure shows a representative sample of the example configuration. The example does not show the configuration for every virtual router routing instance.

1. Configure the routing instance type.

This example configures the routing instance with the name **SRX-VR2**.

Include the **instance-type** statement and specify **virtual-router** as the type at the **[edit routing-instances SRX-VR2]** hierarchy level.

```
[edit routing-instances SRX-VR2]
user@srx# set instance-type virtual-router
```

Repeat this step for every virtual router in each SRX Series device by using the appropriate virtual router name.

2. Add the redundant Ethernet interfaces used by the virtual router routing instance.

This example adds **reth1.15** and **reth2.19** interfaces to the **SRX-VR2** routing instance.

Include the **interface** statement and specify the name of the redundant Ethernet interface at the **[edit routing-instances SRX-VR2]** hierarchy level.

```
[edit routing-instances SRX-VR2]
user@srx# set interface reth1.15
```

```
user@srx# set interface reth2.19
```

Repeat this step for every virtual router in each SRX Series device by using the appropriate virtual router name and interface names.

3. Configure the routing options used by the virtual router routing instance.

This example configures the autonomous system number and enables the graceful restart feature on the **SRX-VR2** routing instance.

Include the **autonomous-system** statement and specify **65019** as the autonomous system number at the **[edit routing-instances SRX-VR2 routing-options]** hierarchy level.

Include the **graceful-restart** statement at the **[edit routing-instances SRX-VR2 routing-options]** hierarchy level.

```
[edit routing-instances SRX-VR2 routing-options]
user@srx# set autonomous-system 65019
user@srx# set graceful-restart
```

Repeat this step for every virtual router in each SRX Series device by using the appropriate virtual router name and interface names.

4. Apply the routing policy that accepts external BGP routes and uses them as generated routes for the routing instance.

This example applies the policy named **check-bgp-routes** to the **SRX-VR2** routing instance.

Include the **policy** statement and specify **check-bgp-routes** at the **[edit routing-instances SRX-VR2 routing-options generate route 0.0.0.0/0]** hierarchy level.

Include the **graceful-restart** statement at the **[edit routing-instances SRX-VR2 routing-options]** hierarchy level.

```
[edit routing-instances SRX-VR2 routing-options]
user@srx# set generate route 0.0.0.0/0 policy
user@srx# set graceful-restart
```

Repeat this step for every virtual router in each SRX Series device by using the appropriate virtual router name and interface names.

5. Apply the routing policy that accepts routes from other routing instances.

This example applies the policy named **from\_srx\_vr1** to the **SRX-VR2** routing instance.

Include the **instance-import** statement and specify **from\_srx\_vr1** at the **[edit routing-instances SRX-VR2 routing-options]** hierarchy level.

```
[edit routing-instances SRX-VR2 routing-options]
user@srx# set instance-import from_srx_vr1
```

Repeat this step for every virtual router in each SRX Series device except the **SRX-VR1** instance.

6. Configure the IGP protocol export policy used by the virtual router routing instance in the trusted security zone.

This example configures the **default-ospf** policy.

Include the **export** statement and specify **default-ospf** as the policy name at the **[edit routing-instances SRX-VR2 protocols ospf]** hierarchy level.

```
[edit routing-instances SRX-VR2 protocols ospf]
user@srx# set export default-ospf
```

Repeat this step for every virtual router in each SRX Series device by using the appropriate virtual router name and policy name.

7. Configure the IGP protocol active and passive interfaces used by the virtual router routing instance in the trusted security zone.

This example configures the **reth1.15** redundant Ethernet interface to actively participate in the OSPF protocol area 0.0.0.0, and the **reth2.19** redundant Ethernet interface to passively participate.

Include the **interface** statement and specify **reth1.15** at the **[edit routing-instances SRX-VR2 protocols ospf area 0.0.0.0]** hierarchy level.

Include the **interface** statement and specify **reth2.19** at the **[edit routing-instances SRX-VR2 protocols ospf area 0.0.0.0]** hierarchy level.

Include the **passive** statement at the **[edit routing-instances SRX-VR2 protocols ospf area 0.0.0.0 reth2.19]** hierarchy level.

```
[edit routing-instances SRX-VR2 protocols ospf area 0.0.0.0]
user@srx# set interface reth1.15
user@srx# set interface reth2.19 passive
```

Repeat this step for every virtual router in each SRX Series device by using the appropriate virtual router name and interface names.

8. Configure the BGP protocol peer groups used by the virtual router routing instance in the untrusted security zone.

Include the **type** statement and specify the **external** option at the **[edit routing-instances SRX-VR2 protocols bgp group MX0-vrf]** hierarchy level.

Include the **peer-as** statement and specify **65000** as the peer autonomous system number at the **[edit routing-instances SRX-VR2 protocols bgp group MX0-vrf]** hierarchy level.

Include the **neighbor** statement and specify **10.19.2.1** as the IPv4 neighbor address at the **[edit routing-instances SRX-VR2 protocols bgp group MX0-vrf]** hierarchy level. The neighbor address is the IRB Logical interface address of the VRF routing instance on the MX Series device.

```
[edit routing-instances SRX-VR2 protocols bgp group MX0-vrf]
user@srx# set type external
user@srx# set peer-as 65000
user@srx# set neighbor 10.19.2.1
```

Repeat this step for every virtual router in each SRX Series device by using the appropriate virtual router name, instance type, neighbor address, and peer AS number.

9. Configure the BGP protocol peer groups export and import policies used by the virtual router routing instance in the untrusted security zone.

Include the **export** statement and specify **trust2-ebgp-out** as the export policy name at the **[edit routing-instances SRX-VR2 protocols bgp group MX0-vrf]** hierarchy level.

Include the **import** statement and specify **local-pref-120** as the import policy name at the **[edit routing-instances SRX-VR2 protocols bgp group MX0-vrf]** hierarchy level.

```
[edit routing-instances SRX-VR2 protocols bgp group MX0-vrf]
user@srx# set export trust2-ebgp-out
user@srx# set import local-pref-120
```

Repeat this step for every virtual router in each SRX Series device by using the appropriate virtual router name, export policy, and import policy.

### Results

The configuration steps of this example have been completed. The following section is for your reference.

The relevant sample configuration for the EX Series device follows.

```
EX Series Device system {
 scripts {
 commit {
 allow-transients;
 file jais-activate-scripts.slax {
 optional;
 }
 }
 }
 services {
 ftp;
 ssh;
 telnet;
 outbound-ssh {
 client 00187D0B670D {
 device-id FA022D;
 secret "9-9w4aik.QznDj9A0BEhrlKMxN"; ## SECRET-DATA
 services netconf;
 10.8.7.32 port 7804;
 }
 }
 }
 }
 interfaces {
 ge-0/0/17 {
 unit 0 {
 family ethernet-switching {
 vlan {
 members 17;
 }
 }
 }
 }
 xe-0/1/0 {
 unit 0 {
 family ethernet-switching {
```

```
 port-mode trunk;
 vlan {
 members all;
 }
 }
}
xe-0/1/2 {
 unit 0 {
 enable;
 family ethernet-switching {
 port-mode trunk;
 vlan {
 members all;
 }
 }
 }
}
lo0 {
 unit 0 {
 family inet {
 address 127.0.0.1/32;
 }
 }
}
me0 {
 unit 0 {
 family inet {
 address 10.8.108.19/24;
 }
 }
}
}
event-options {
 destinations {
 juniper-aim {
 archive-sites {
 "scp://admin@127.0.0.1:/var/tmp" password "9u3KWOEcrevL7-eKaZ"; ##
 SECRET-DATA
 }
 }
 }
}
routing-options {
 static {
 route 10.8.0.0/16 next-hop 10.8.108.254;
 }
}
protocols {
 rstp {
 interface xe-0/1/0.0 {
 disable;
 }
 interface xe-0/1/2.0 {
 disable;
 }
 }
}
```

```

 }
 }
 security {
 ssh-known-hosts {
 host 127.0.0.1 {
 rsa-key AAAAB3NzaC1yc2;
 }
 }
 }
 ethernet-switching-options {
 redundant-trunk-group {
 group rtgroup1 {
 interface xe-0/1/0.0;
 interface xe-0/1/2.0 {
 primary;
 }
 }
 }
 }
}
vlands {
 vlan17 {
 vlan-id 17;
 interface {
 xe-0/1/0.0;
 xe-0/1/2.0;
 }
 }
}
}

```

The relevant sample configuration for the MX Series device follows.

```

MX Series Device groups {
 re0 {
 system {
 host-name MX0;
 }
 }
 re1 {
 system {
 host-name MX0re1;
 }
 }
 }
 apply-groups [re0 re1];
 system {
 scripts {
 commit {
 allow-transients;
 file jais-activate-scripts.slax {
 optional;
 }
 }
 }
 login {
 class trust1-admin {
 logical-system Trust1;
 }
 }
 }
}

```

```
 permissions all;
 }
 user trust1 {
 uid 2000;
 class trust1-admin;
 authentication {
 encrypted-password 12345; ## SECRET-DATA
 }
 }
}
logical-systems {
 Trust1 {
 interfaces {
 xe-0/1/0 {
 unit 17 {
 encapsulation vlan-bridge;
 vlan-id 17;
 }
 }
 xe-1/0/0 {
 unit 17 {
 encapsulation vlan-bridge;
 vlan-id 17;
 }
 }
 xe-1/1/0 {
 unit 15 {
 encapsulation vlan-bridge;
 vlan-id 15;
 }
 }
 }
 irb {
 unit 15 {
 family inet {
 address 10.15.2.3/24;
 }
 }
 unit 17 {
 family inet {
 address 10.17.2.2/24 {
 vrrp-group 17 {
 virtual-address 10.17.2.1;
 priority 200;
 fast-interval 200;
 preempt;
 accept-data;
 }
 }
 }
 }
 }
 }
}
routing-instances {
 MX-VR2 {
 instance-type virtual-router;
```





```
xe-2/2/0 {
 unit 19 {
 encapsulation vlan-bridge;
 vlan-id 19;
 }
}
irb {
 unit 19 {
 family inet {
 address 10.19.2.1/24;
 }
 }
}
lo0 {
 unit 1 {
 family inet {
 address 10.200.11.1/32;
 }
 }
}
}
protocols {
 rsvp {
 interface xe-1/3/0.0;
 interface xe-2/3/0.0;
 }
 mpls {
 no-cspf;
 label-switched-path to-core-router {
 to 10.200.11.3;
 }
 interface xe-1/3/0.0;
 interface xe-2/3/0.0;
 }
 bgp {
 group int {
 type internal;
 local-address 10.200.11.1;
 family inet {
 unicast;
 }
 family inet-vpn {
 unicast;
 }
 peer-as 65000;
 local-as 65000;
 neighbor 10.200.11.2;
 neighbor 10.200.11.3;
 neighbor 10.200.11.4;
 }
 }
 ospf {
 area 0.0.0.0 {
 interface xe-1/3/0.0;
 interface lo0.1;
 }
 }
}
```

```

 }
 }
 routing-options {
 router-id 10.200.11.101;
 autonomous-system 65000;
 dynamic-tunnels {
 GRE1 {
 source-address 10.200.11.1;
 gre;
 destination-networks {
 0.0.0.0/0;
 }
 }
 }
 }
}
bridge-domains {
 vlan14 {
 vlan-id 14;
 interface xe-0/3/0.14;
 interface xe-2/2/0.14;
 routing-interface irb.14;
 }
}
}
}
interfaces {
 xe-0/1/0 {
 flexible-vlan-tagging;
 encapsulation flexible-ethernet-services;
 }
 xe-0/3/0 {
 flexible-vlan-tagging;
 encapsulation flexible-ethernet-services;
 }
 xe-1/0/0 {
 flexible-vlan-tagging;
 encapsulation flexible-ethernet-services;
 }
 xe-1/1/0 {
 flexible-vlan-tagging;
 encapsulation flexible-ethernet-services;
 }
 xe-2/2/0 {
 flexible-vlan-tagging;
 encapsulation flexible-ethernet-services;
 }
 fxp0 {
 unit 0 {
 family inet {
 address 10.8.3.212/24;
 }
 }
 }
}
lo0 {
 unit 0 {
 family inet;
 }
}

```

```
 }
 }
}
event-options {
 destinations {
 juniper-aim {
 archive-sites {
 "scp://admin@127.0.0.1://var/tmp" password "9DyimfQFnCpOF3re"; ##
 SECRET-DATA
 }
 }
 }
}
routing-options {
 static {
 route 10.0.244.8/30 next-hop 10.0.134.10;
 route 10.0.0.0/8 {
 next-hop 10.8.3.254;
 retain;
 no-readvertise;
 }
 }
}
```

The relevant sample configuration for the SRX Series device follows.

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SRX Series Device</b> | <pre>system {   host-name srx0;   chassis {     cluster {       control-link-recovery;       reth-count 4;       redundancy-group 1 {         node 0 priority 200;         node 1 priority 100;         preempt;         interface-monitor {           xe-1/0/0 weight 255;         }       }     }   } } interfaces {   xe-1/0/0 {     gigether-options {       redundant-parent reth1;     }   }   fab0 {     fabric-options {       member-interfaces {         xe-1/0/1;       }     }   } }</pre> |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

```
lo0 {
 unit 0 {
 family inet {
 address 127.0.0.1/32;
 }
 }
}
reth1 {
 vlan-tagging;
 redundant-ether-options {
 redundancy-group 1;
 }
 unit 15 {
 family inet {
 address 10.15.2.2/24;
 }
 vlan-id 15;
 }
}
}
policy-options {
 policy-statement check-bgp-routes {
 term term1 {
 from {
 protocol bgp;
 route-type external;
 }
 then accept;
 }
 }
 policy-statement default-ospf {
 term term1 {
 from {
 protocol aggregate;
 route-filter 0.0.0.0/0 exact;
 }
 then {
 metric 0;
 external {
 type 1;
 }
 }
 accept;
 }
 }
}
policy-statement from_srx_vr1 {
 term term1 {
 from instance SRX-VR1;
 then accept;
 }
}
policy-statement local-pref-120 {
 term term1 {
 from protocol bgp;
 then {
 local-preference 120;
 }
 }
}
```

```
 }
 }
}
policy-statement trust2-ebgp-out {
 term term1 {
 from {
 protocol ospf;
 route-filter 10.16.2.0/24 exact;
 route-filter 10.17.2.0/24 exact;
 route-filter 10.18.2.0/24 exact;
 }
 then accept;
 }
}
}
security {
 policies {
 from-zone Trust2 to-zone Untrust2 {
 policy denyftp {
 match {
 source-address any;
 destination-address any;
 application junos-ftp;
 }
 then {
 deny;
 }
 }
 policy allow_all {
 match {
 source-address any;
 destination-address any;
 application any;
 }
 then {
 permit;
 }
 }
 }
 }
}
zones {
 security-zone Trust2 {
 host-inbound-traffic {
 system-services {
 all;
 }
 protocols {
 all;
 }
 }
 interfaces {
 reth1.15;
 }
 }
 security-zone Untrust2 {
 interfaces reth2.19;
 }
}
```

```

 }
 }
}
routing-instances {
 SRX-VR2 {
 instance-type virtual-router;
 interface reth1.15;
 interface reth2.19;
 routing-options {
 graceful-restart;
 autonomous-system 65019;
 instance-import from_srx_vr1;
 }
 protocols {
 bgp {
 group MX0-vrf {
 import local-pref-120;
 export trust2-ebgp-out;
 }
 }
 ospf {
 export default-ospf;
 area 0.0.0.0 {
 interface reth1.15;
 interface reth2.19 {
 passive;
 }
 }
 }
 }
 }
}
}
}
}
}

```

Related Documentation • [Introduction to Logical Systems on page 3259](#)

## Configuration Statements

- [\[edit logical-systems\] Hierarchy Level on page 3521](#)

### [\[edit logical-systems\] Hierarchy Level](#)

As indicated in the following hierarchy, you can include at this hierarchy level several of the hierarchies that can be included at the **[edit]** hierarchy level. However, some statements in a subhierarchy are not valid for logical systems. To learn which statements can be included under **[edit logical-systems *logical-system-name*]** on your device, issue the **set ?** command at the hierarchy level of interest.

```

logical-systems {
 logical-system-name {
 access {
 address-assignment {
 ... same statements as in the address-assignment subhierarchy in [edit access]
 Hierarchy Level ...
 }
 }
 }
}

```

```

}
access-profile profile-name;
bridge-domains {
 ... (MX Series only) same statements as in [edit bridge-domains] Hierarchy Level ...
}
bridge-domains {
 ... (MX Series only) same statements as in [edit bridge-domains] Hierarchy Level ...
}
firewall {
 ... same statements as in several subhierarchies in [edit firewall] Hierarchy Level ...
}
forwarding-options {
 ... same statements as in [edit forwarding-options dhcp-relay] Hierarchy Level ...
}
interfaces {
 interface-name {
 unit logical-unit-number {
 ... some of the statements in the unit subhierarchy in [edit interfaces] Hierarchy
 Level ...
 }
 }
}
policy-options {
 ... same statements as in [edit policy-options] Hierarchy Level on page 360 ...
}
protocols {
 ... same statements as in [edit protocols] Hierarchy Level ...
}
routing-instances {
 ... most statements in [edit routing-instances] Hierarchy Level ...
}
routing-options {
 ... most statements in [edit routing-options] Hierarchy Level ...
}
services {
 mobile-ip {
 ... same statements as in [edit services mobile-ip] Hierarchy Level ...
 }
}
switch-options {
 ... (MX Series only) same statements as in [edit switch-options] Hierarchy Level ...
}
system {
 services {
 dhcp-local-server {
 ... same statements as in the services dhcp-local-server subhierarchy in [edit
 system] Hierarchy Level ...
 }
 }
 syslog {
 ... most statements in syslog subhierarchy in [edit system] Hierarchy Level...
 }
}
}
}
}

```



**Related  
Documentation**

- [Notational Conventions Used in Junos OS Configuration Hierarchies](#)

## Administration

---

- [Operational Commands on page 3523](#)

## Operational Commands

- [Operational-Mode Commands on page 3523](#)

### Operational-Mode Commands

---

- [Overview of Junos OS CLI Operational Mode Commands on page 3523](#)
- [Example: Running Operational-Mode Commands on Logical Systems on page 3526](#)
- [Example: Viewing BGP Trace Files on Logical Systems on page 3527](#)
- [Example: Configuring System Logging on Logical Systems on page 3532](#)

#### ***Overview of Junos OS CLI Operational Mode Commands***

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

- [CLI Command Categories on page 3523](#)
- [Commonly Used Operational Mode Commands on page 3524](#)

#### ***CLI Command Categories***

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

- **Commands for controlling the CLI environment**—Some set commands in the **set** hierarchy configure the CLI display screen. For information about these commands, see *Understanding the Junos OS CLI Modes, Commands, and Statement Hierarchies*.
- **Commands for monitoring and troubleshooting**—The following commands display information and statistics about the software and test network connectivity. Detailed command descriptions are provided in the *Junos OS Interfaces Command Reference*.
  - **clear**—Clear statistics and protocol database information.
  - **mtrace**—Trace mtrace packets from source to receiver.
  - **monitor**—Perform real-time debugging of various software components, including the routing protocols and interfaces.
  - **ping**—Determine the reachability of a remote network host.
  - **show**—Display the current configuration and information about interfaces, routing protocols, routing tables, routing policy filters, system alarms, and the chassis.
  - **test**—Test the configuration and application of policy filters and autonomous system (AS) path regular expressions.
  - **traceroute**—Trace the route to a remote network host.

- Commands for connecting to other network systems—The **ssh** command opens Secure Shell connections, and the **telnet** command opens telnet sessions to other hosts on the network. For information about these commands, see the *Junos OS Operational Mode Commands*.
- Commands for copying files—The **copy** command copies files from one location on the router or switch to another, from the router or switch to a remote system, or from a remote system to the router or switch. For information about these commands, see the *Junos OS Operational Mode Commands*.
- Commands for restarting software processes—The commands in the **restart** hierarchy restart the various Junos OS processes, including the routing protocol, interface, and SNMP. For information about these commands, see the *Junos OS Operational Mode Commands*.
- A command—**request**—for performing system-level operations, including stopping and rebooting the router or switch and loading Junos OS images. For information about this command, see the *Junos OS Operational Mode Commands*.
- A command—**start**—to exit the CLI and start a UNIX shell. For information about this command, see the *Junos OS Operational Mode Commands*.
- A command—**configure**—for entering configuration mode, which provides a series of commands that configure Junos OS, including the routing protocols, interfaces, network management, and user access. For information about the CLI configuration commands, see *Understanding Junos OS CLI Configuration Mode*.
- A command—**quit**—to exit the CLI. For information about this command, see the *Junos OS Operational Mode Commands*.
- For more information about the CLI operational mode commands, see the *Junos OS Operational Mode Commands* and the *Junos OS Operational Mode Commands*.

#### Commonly Used Operational Mode Commands

Table 267 on page 3524 lists some operational commands you may find useful for monitoring router or switch operation. For a complete description of operational commands, see the Junos OS command references.



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

Table 267: Commonly Used Operational Mode Commands

| Items to Check   | Description                                          | Command             |
|------------------|------------------------------------------------------|---------------------|
| Software version | Versions of software running on the router or switch | <b>show version</b> |
| Log files        | Contents of the log files                            | <b>monitor</b>      |
|                  | Log files and their contents and recent user logins  | <b>show log</b>     |

Table 267: Commonly Used Operational Mode Commands (*continued*)

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

### **Example: Running Operational-Mode Commands on Logical Systems**

This example shows how to set the CLI to a specified logical system view, run operational-mode commands for the logical system, and then return to the main router view.

- [Requirements on page 3526](#)
- [Overview on page 3526](#)
- [Configuration on page 3526](#)

#### **Requirements**

You must have the **view** privilege for the logical system.

#### **Overview**

For some operational-mode commands, you can include a **logical-system** option to narrow the output of the command or to limit the operation of the command to the specified logical system. For example, the **show route** command has a **logical-system** option. To run this command on a logical system called LS3, you can use **show route logical-system LS3**. However, some commands, such as **show interfaces**, do not have a **logical-system** option. For commands like this, you need another approach.

You can place yourself into the context of a specific logical system. To configure a logical system context, issue the **set cli logical-system logical-system-name** command.

When the CLI is in logical system context mode and you enter an operational-mode command, the output of the command displays information related to the logical system only.

#### **Configuration**

##### **Step-by-Step Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To set the CLI to a specific logical system context:

1. From the main router, configure the logical system.  

```
[edit]
user@host# set logical-systems LS3
```
2. If you are done configuring the device, commit the configuration.  

```
[edit]
user@host# commit
user@host# exit
```
3. Set the CLI to view the logical system.  

```
user@host> set cli logical-system LS3
Logical system: LS3
user@host:LS3>
```
4. Run an operational-mode command.

```

user@host:LS3> show interfaces terse
Interface Admin Link Proto Local Remote
lt-1/2/0
lt-1/2/0.3 up up inet 10.0.2.1/30

```

5. Enter configuration mode to edit the logical system configuration.

```

user@host:LS3> edit
Entering configuration mode

user@host:LS3#

```

6. Exit configuration mode to return to operational mode.

```

user@host:LS3# exit
Exiting configuration mode

```

7. Clear the logical system view to return to the main router view.

```

user@host:LS3> clear cli logical-system
Cleared default logical system

user@host>

```

8. To achieve the same effect when using a Junos XML protocol client application, include the `<set-logical-system>` tag.

```

<rpc>
<set-logical-system>
<logical-system>LS1</logical-system>
</set-logical-system>
</rpc>

```

### **Example: Viewing BGP Trace Files on Logical Systems**

This example shows how to list and view files that are stored on a logical system.

- [Requirements on page 3527](#)
- [Overview on page 3528](#)
- [Configuration on page 3528](#)
- [Verification on page 3532](#)

### **Requirements**

- You must have the **view** privilege for the logical system.
- Configure a network, such as the BGP network shown in “[Example: Configuring Internal BGP Peering Sessions on Logical Systems](#)” on page 3365.

### Overview

Logical systems have their individual directory structure created in the `/var/logical-systems/logical-system-name` directory. It contains the following subdirectories:

- `/config`—Contains the active configuration specific to the logical system.
- `/log`—Contains system log and tracing files specific to the logical system.

To maintain backward compatibility for the log files with previous versions of Junos OS, a symbolic link (symlink) from the `/var/logs/logical-system-name` directory to the `/var/logical-systems/logical-system-name` directory is created when a logical system is configured.

- `/tmp`—Contains temporary files specific to the logical system.

The file system for each logical system enables logical system users to view trace logs and modify logical system files. Logical system administrators have full access to view and modify all files specific to the logical system.

Logical system users and administrators can save and load configuration files at the logical-system level using the **save** and **load** configuration mode commands. In addition, they can also issue the **show log**, **monitor**, and **file** operational mode commands at the logical-system level.

This example shows how to configure and view a BGP trace file on a logical system. The steps can be adapted to apply to trace operations for any Junos OS hierarchy level that supports trace operations.



**TIP:** To view a list of hierarchy levels that support tracing operations, enter the **help apropos traceoptions** command in configuration mode.

### Configuration

- [Configuring Trace Operations on page 3529](#)
- [Viewing the Trace File on page 3529](#)
- [Deactivating and Reactivating Trace Logging on page 3531](#)
- [Results on page 3532](#)

#### CLI Quick Configuration

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

```
set logical-systems A protocols bgp group internal-peers traceoptions file bgp-log
set logical-systems A protocols bgp group internal-peers traceoptions file size 10k
set logical-systems A protocols bgp group internal-peers traceoptions file files 2
set logical-systems A protocols bgp group internal-peers traceoptions flag update detail
```

### Configuring Trace Operations

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure the trace operations:

1. Configure trace operations on the logical system.  
  

```
[edit logical-systems A protocols bgp group internal-peers]
user@host# set traceoptions file bgp-log
user@host# set traceoptions file size 10k
user@host# set traceoptions file files 2
user@host# set traceoptions flag update detail
```
2. If you are done configuring the device, commit the configuration.  
  

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

### Viewing the Trace File

**Step-by-Step Procedure** To view the trace file:

1. In operational mode on the main router, list the directories on the logical system.  
  

```
user@host> file list /var/logical-systems/A
/var/logical-systems/A:
config/
log/
tmp/
```
2. In operational mode on the main router, list the log files on the logical system.  
  

```
user@host> file list /var/logical-systems/A/log/
/var/logical-systems/A/log:
bgp-log
```
3. View the contents of the **bgp-log** file.  
  

```
user@host> file show /var/logical-systems/A/log/bgp-log
Aug 10 17:12:01 trace_on: Tracing to "/var/log/A/bgp-log" started
Aug 10 17:14:22.826182 bgp_peer_mgmt_clear:5829: NOTIFICATION sent to
192.163.6.4 (Internal AS 17): code 6 (Cease) subcode 4 (Administratively
Reset), Reason: Management session cleared BGP neighbor
Aug 10 17:14:22.826445 bgp_send: sending 21 bytes to 192.163.6.4 (Internal
AS 17)
Aug 10 17:14:22.826499
Aug 10 17:14:22.826499 BGP SEND 192.168.6.5+64965 -> 192.163.6.4+179
Aug 10 17:14:22.826559 BGP SEND message type 3 (Notification) length 21
Aug 10 17:14:22.826598 BGP SEND Notification code 6 (Cease) subcode 4
(Administratively Reset)
Aug 10 17:14:22.831756 bgp_peer_mgmt_clear:5829: NOTIFICATION sent to
192.168.40.4 (Internal AS 17): code 6 (Cease) subcode 4 (Administratively
Reset), Reason: Management session cleared BGP neighbor
Aug 10 17:14:22.831851 bgp_send: sending 21 bytes to 192.168.40.4 (Internal
AS 17)
Aug 10 17:14:22.831901
Aug 10 17:14:22.831901 BGP SEND 192.168.6.5+53889 -> 192.168.40.4+179
```

```
Aug 10 17:14:22.831959 BGP SEND message type 3 (Notification) length 21
Aug 10 17:14:22.831999 BGP SEND Notification code 6 (Cease) subcode 4
(Administratively Reset)
...
```

4. Filter the output of the log file.

```
user@host> file show /var/logical-systems/A/log/bgp-log | match "flags 0x40"
Aug 10 17:14:54.867460 BGP SEND flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.867595 BGP SEND flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.867650 BGP SEND flags 0x40 code NextHop(3): 192.168.6.5
Aug 10 17:14:54.867692 BGP SEND flags 0x40 code LocalPref(5): 100
Aug 10 17:14:54.884529 BGP RECV flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.884581 BGP RECV flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.884628 BGP RECV flags 0x40 code NextHop(3): 192.168.6.4
Aug 10 17:14:54.884667 BGP RECV flags 0x40 code LocalPref(5): 100
Aug 10 17:14:54.911377 BGP RECV flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.911422 BGP RECV flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.911466 BGP RECV flags 0x40 code NextHop(3): 192.168.40.4
Aug 10 17:14:54.911507 BGP RECV flags 0x40 code LocalPref(5): 100
Aug 10 17:14:54.916008 BGP SEND flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.916054 BGP SEND flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.916100 BGP SEND flags 0x40 code NextHop(3): 192.168.6.5
Aug 10 17:14:54.916143 BGP SEND flags 0x40 code LocalPref(5): 100
Aug 10 17:14:54.920304 BGP RECV flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.920348 BGP RECV flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.920393 BGP RECV flags 0x40 code NextHop(3): 10.0.0.10
Aug 10 17:14:54.920434 BGP RECV flags 0x40 code LocalPref(5): 100
```

5. View the tracing operations in real time.

```
user@host> clear bgp neighbor logical-system A
Cleared 2 connections
```



**CAUTION:** Clearing the BGP neighbor table is disruptive in a production environment.

6. Run the **monitor start** command with an optional **match** condition.

```
user@host> monitor start A/bgp-log | match 0.0.0.0/0
Aug 10 19:21:40.773467 BGP RECV 0.0.0.0/0
Aug 10 19:21:40.773685 bgp_rcv_nlrri: 0.0.0.0/0
Aug 10 19:21:40.773778 bgp_rcv_nlrri: 0.0.0.0/0 belongs to meshgroup
Aug 10 19:21:40.773832 bgp_rcv_nlrri: 0.0.0.0/0 qualified bnp->ribact 0x0
12afcb 0x0
```

7. Pause the **monitor** command by pressing Esc-Q.

To unpause the output, press Esc-Q again.

8. Halt the **monitor** command by pressing Enter and typing **monitor stop**.

[Enter]

```
user@host> monitor stop
```

9. When you are finished troubleshooting, consider deactivating trace logging to avoid any unnecessary impact to system resources.

```
[edit protocols bgp group internal-peers]
```

```
user@host:A# deactivate traceoptions
```

```
user@host:A# commit
```



When configuration is deactivated, it appears in the configuration with the **inactive** tag. To reactivate trace operations, use the **activate** configuration-mode statement.

```
[edit protocols bgp group internal-peers]
user@host:A# show

type internal;
inactive: traceoptions {
 file bgp-log size 10k files 2;
 flag update detail;
 flag all;
}
local-address 192.168.6.5;
export send-direct;
neighbor 192.163.6.4;
neighbor 192.168.40.4;
```

10. To reactivate trace operations, use the **activate** configuration-mode statement.

```
[edit protocols bgp group internal-peers]
user@host:A# activate traceoptions
user@host:A# commit
```

### *Deactivating and Reactivating Trace Logging*

#### **Step-by-Step Procedure**

To deactivate and reactivate the trace file:

1. When you are finished troubleshooting, consider deactivating trace logging to avoid an unnecessary impact to system resources.

```
[edit protocols bgp group internal-peers]
user@host:A# deactivate traceoptions
user@host:A# commit
```

When configuration is deactivated, the statement appears in the configuration with the **inactive** tag.

```
[edit protocols bgp group internal-peers]
user@host:A# show

type internal;
inactive: traceoptions {
 file bgp-log size 10k files 2;
 flag update detail;
 flag all;
}
local-address 192.168.6.5;
export send-direct;
neighbor 192.163.6.4;
neighbor 192.168.40.4;
```

2. To reactivate logging, use the **activate** configuration-mode statement.

```
[edit protocols bgp group internal-peers]
user@host:A# activate traceoptions
user@host:A# commit
```

### Results

From configuration mode, confirm your configuration by entering the **show logical-systems A protocols bgp group internal-peers** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show logical-systems A protocols bgp group internal-peers
traceoptions {
 file bgp-log size 10k files 2;
 flag update detail;
}
```

### Verification

Confirm that the configuration is working properly.

#### *Verifying That the Trace Log File Is Operating*

**Purpose** Make sure that events are being written to the log file.

**Action** user@host:A> **show log bgp-log**  
Aug 12 11:20:57 trace\_on: Tracing to "/var/log/A/bgp-log" started

#### *Example: Configuring System Logging on Logical Systems*

This example shows how to configure system logging on logical systems and how to view the logs.

- [Requirements on page 3532](#)
- [Overview on page 3533](#)
- [Configuration on page 3533](#)
- [Verification on page 3534](#)

### Requirements

This example has the following requirements:

- You must have the **view** privilege for the logical system.
- Junos OS Release 11.4 or later.

### Overview

Each logical system has its individual directory structure created in the `/var/logical-systems/logical-system-name` directory. This directory contains the following subdirectories:

- `/config`—Contains the active configuration specific to the logical system.
- `/log`—Contains system log and tracing files specific to the logical system.

To maintain backward compatibility for the log files with previous versions of Junos OS, a symbolic link (symlink) from the `/var/log/logical-system-name` directory to the `/var/logical-systems/logical-system-name` directory is created when a logical system is configured.

- `/tmp`—Contains temporary files specific to the logical system.

The file system for each logical system enables logical system users to view trace logs and modify logical system files. Logical system administrators have full access to view and modify all files specific to the logical system.

Logical system users and administrators can save and load configuration files at the logical system level using the **save** and **load** configuration mode commands. In addition, they can issue the **show log**, **monitor**, and **file** operational mode commands at the logical system level.

This example shows how to configure system logging on a logical system.

### Configuration

#### CLI Quick Configuration

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

```
set logical-systems lsys1 system syslog host 10.209.10.69 ftp critical
set logical-systems lsys1 system syslog allow-duplicates
set logical-systems lsys1 system syslog file lsys1-file1 daemon error
set logical-systems lsys1 system syslog file lsys1-file1 firewall critical
```

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure system logging:

1. Configure trace operations on the logical system.

```
[edit logical-systems lsys1 system syslog]
user@host# set host 10.209.10.69 ftp critical
user@host# set allow-duplicates
user@host# set file lsys1-file1 daemon error
user@host# set file lsys1-file1 firewall critical
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
user@host# exit
```

### Results

From configuration mode, confirm your configuration by entering the **show logical-systems** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show logical-systems
lsys1 {
 system {
 syslog {
 host 10.209.10.69 {
 ftp critical;
 }
 allow-duplicates;
 file lsys1-file1 {
 daemon error;
 firewall critical;
 }
 }
 }
}
```

### Verification

Confirm that the configuration is working properly.

#### Verifying That the System Log File Is Operating

**Purpose** Make sure that events are being written to the log file.

#### Action



**TIP:** To make entries in the system log, you can use the **start shell** command and then use the **logger** shell command. For example: **logger -e "firewall\_crit" -p firewall.crit -l lsys1 TEST**

```
user@host> show log lsys1/lsys1-file1
Sep 7 14:15:46 host clear-log[2752]: logfile cleared
Sep 7 14:19:04 host logger: % -: firewall_crit: TEST
...
```

```
user@host> file show /var/logical-systems/lsys1/log/lsys1-file1
Sep 7 14:19:04 host logger: % -: firewall_crit: TEST
...
```

**Related Documentation**

- [Introduction to Logical Systems on page 3259](#)

## CHAPTER 14

# MPLS

- [Overview on page 3535](#)
- [Configuration on page 3540](#)
- [Administration on page 3617](#)

### Overview

---

- [LDP on page 3535](#)

### LDP

- [LDP Introduction on page 3535](#)
- [Junos OS LDP Protocol Implementation on page 3536](#)
- [LDP Operation on page 3536](#)
- [Label Operations on page 3536](#)
- [LDP Message Types on page 3538](#)
- [Discovery Messages on page 3538](#)
- [Session Messages on page 3538](#)
- [Advertisement Messages on page 3538](#)
- [Notification Messages on page 3539](#)
- [LDP Session Protection on page 3539](#)
- [LDP Graceful Restart on page 3539](#)

#### LDP Introduction

---

The Label Distribution Protocol (LDP) is a protocol for distributing labels in non-traffic-engineered applications. LDP allows routers to establish label-switched paths (LSPs) through a network by mapping network-layer routing information directly to data link layer-switched paths.

These LSPs might have an endpoint at a directly attached neighbor (comparable to IP hop-by-hop forwarding), or at a network egress node, enabling switching through all intermediary nodes. LSPs established by LDP can also traverse traffic-engineered LSPs created by RSVP.

LDP associates a forwarding equivalence class (FEC) with each LSP it creates. The FEC associated with an LSP specifies which packets are mapped to that LSP. LSPs are extended through a network as each router chooses the label advertised by the next hop for the FEC and splices it to the label it advertises to all other routers. This process forms a tree of LSPs that converge on the egress router.

### Junos OS LDP Protocol Implementation

---

The Junos OS implementation of LDP supports LDP version 1. The Junos OS supports a simple mechanism for tunneling between routers in an interior gateway protocol (IGP), to eliminate the required distribution of external routes within the core. The Junos OS allows an MPLS tunnel next hop to all egress routers in the network, with only an IGP running in the core to distribute routes to egress routers. Edge routers run BGP but do not distribute external routes to the core. Instead, the recursive route lookup at the edge resolves to an LSP switched to the egress router. No external routes are necessary on the transit LDP routers.

### LDP Operation

---

You must configure LDP for each interface on which you want LDP to run. LDP creates LSP trees rooted at each egress router for the router ID address that is the subsequent BGP next hop. The ingress point is at every router running LDP. This process provides an inet.3 route to every egress router. If BGP is running, it will attempt to resolve next hops by using the inet.3 table first, which binds most, if not all, of the BGP routes to MPLS tunnel next hops.

Two adjacent routing devices running LDP become neighbors. If the routing devices are connected by more than one interface, they become neighbors on each interface. When LDP routing devices become neighbors, they establish an LDP session to exchange label information. If per-router labels are in use on both routing devices, only one LDP session is established between them, even if they are neighbors on multiple interfaces. For this reason, an LDP session is not related to a particular interface.

LDP operates in conjunction with a unicast routing protocol. LDP installs LSPs only when both LDP and the routing protocol are enabled. For this reason, you must enable both LDP and the routing protocol on the same set of interfaces. If this is not done, LSPs might not be established between each egress routing device and all ingress routing devices, which might result in loss of BGP-routed traffic.

You can apply policy filters to labels received from and distributed to other routing devices through LDP. Policy filters provide you with a mechanism to control the establishment of LSPs.

For LDP to run on an interface, MPLS must be enabled on a logical interface on that interface. For more information, see the *Junos® OS Network Interfaces*.

### Label Operations

---


Figure 78 on page 3537 depicts an LDP LSP being tunneled through an RSVP LSP. (For definitions of label operations, see *Label Description*.) The shaded inner oval represents the RSVP domain, whereas the outer oval depicts the LDP domain. RSVP establishes an LSP through routers B, C, D, and E, with the sequence of labels L3, L4. LDP establishes

an LSP through Routers A, B, E, F, and G, with the sequence of labels L1, L2, L5. LDP views the RSVP LSP between Routers B and E as a single hop.

When the packet arrives at Router A, it enters the LSP established by LDP, and a label (L1) is pushed onto the packet. When the packet arrives at Router B, the label (L1) is swapped with another label (L2). Because the packet is entering the traffic-engineered LSP established by RSVP, a second label (L3) is pushed onto the packet.


This outer label (L3) is swapped with a new label (L4) at the intermediate router (C) within the RSVP LSP tunnel, and when the penultimate router (D) is reached, the top label is popped. Router E swaps the label (L2) with a new label (L5), and the penultimate router for the LDP-established LSP (F) pops the last label.

**Figure 78: Swap and Push When LDP LSPs Are Tunneled Through RSVP LSPs**



[Figure 79 on page 3537](#) depicts a double push label operation (L1L2). A double push label operation is used when the ingress router (A) for both the LDP LSP and the RSVP LSP tunneled through it is the same device. Note that Router D is the penultimate hop for the LDP-established LSP, so L2 is popped from the packet by Router D.

**Figure 79: Double Push When LDP LSPs Are Tunneled Through RSVP LSPs**



## LDP Message Types

---

LDP uses the message types described in the following sections to establish and remove mappings and to report errors. All LDP messages have a common structure that uses a type, length, and value (TLV) encoding scheme.

- [Discovery Messages on page 3538](#)
- [Session Messages on page 3538](#)
- [Advertisement Messages on page 3538](#)
- [Notification Messages on page 3539](#)

## Discovery Messages

---

Discovery messages announce and maintain the presence of a router or switch in a network. They indicate their presence in a network by sending hello messages periodically. Hello messages are transmitted as UDP packets to the LDP port at the group multicast address for all routers on the subnet.

LDP uses the following discovery procedures:

- Basic discovery—A router or switch periodically sends LDP link hello messages through an interface. LDP link hello messages are sent as UDP packets addressed to the LDP discovery port. Receipt of an LDP link hello message on an interface identifies an adjacency with the LDP peer router or switch.
- Extended discovery—LDP sessions between routers or switches not directly connected are supported by LDP extended discovery. A router or switch periodically sends LDP targeted hello messages to a specific address. Targeted hello messages are sent as UDP packets addressed to the LDP discovery port at the specific address. The targeted router or switch decides whether to respond to or ignore the targeted hello message. A targeted router or switch that chooses to respond does so by periodically sending targeted hello messages to the initiating router or switch.

## Session Messages

---

Session messages establish, maintain, and terminate sessions between LDP peers. When a router or switch establishes a session with another router or switch learned through the hello message, it uses the LDP initialization procedure over TCP transport. When the initialization procedure is completed successfully, the two routers or switches are LDP peers and can exchange advertisement messages.

## Advertisement Messages

---

Advertisement messages create, change, and delete label mappings for forwarding equivalence classes (FECs). Requesting a label or advertising a label mapping to a peer is a decision made by the local router or switch. In general, the router or switch requests a label mapping from a neighboring router or switch when it needs one and advertises a label mapping to a neighboring router or switch when it wants the neighbor to use a label.



### Notification Messages

---

Notification messages provide advisory information and signal error information. LDP sends notification messages to report errors and other events of interest. There are two kinds of LDP notification messages:

- Error notifications, which signal fatal errors. If a router or switch receives an error notification from a peer for an LDP session, it terminates the LDP session by closing the TCP transport connection for the session and discarding all label mappings learned through the session.
- Advisory notifications, which pass information to a router or switch about the LDP session or the status of some previous message received from the peer.

### LDP Session Protection

---

LDP session protection is based on the LDP targeted hello functionality defined in RFC 5036, *LDP Specification*, and is supported by the Junos OS as well as the LDP implementations of most other vendors. It involves sending unicast User Datagram Protocol (UDP) hello packets to a remote neighbor address and receiving similar packets from the neighbor router.

If you configure LDP session protection on a router or switch, the LDP sessions are maintained as follows:

1. An LDP session is established between a router or switch and a remote neighboring router or switch.
2. If all of the direct links between the routers or switches go down, the LDP session remains up so long as there is IP connectivity between the routers based on another connection over the network.
3. When the direct link between the routers or switches is reestablished, the LDP session is not restarted. They simply exchange LDP hellos with each other over the direct link. They can then begin forwarding LDP-signaled MPLS packets using the original LDP session.

By default, LDP targeted hellos are set to the remote neighbor so long as the LDP session is up, even if there are no more link neighbors to that router or switch. You can also specify the duration you would like to maintain the remote neighbor connection in the absence of link neighbors. When the last link neighbor for a session goes down, the Junos OS starts an LDP session protection timer. If this timer expires before any of the link neighbors come back up, the remote neighbor connection is taken down and the LDP session is terminated. If you configure a different value for the timer while it is currently running, the Junos OS updates the timer to the specified value without disrupting the current state of the LDP session.

### LDP Graceful Restart

---

LDP graceful restart enables a router or switch whose LDP control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers

or switches. It also enables a router or switch on which helper mode is enabled to assist a neighboring router or switch that is attempting to restart LDP.

During session initialization, a router or switch advertises its ability to perform LDP graceful restart or to take advantage of a neighbor performing LDP graceful restart by sending the graceful restart TLV. This TLV contains two fields relevant to LDP graceful restart: the reconnect time and the recovery time. The values of the reconnect and recovery times indicate the graceful restart capabilities supported by the router or switch.

When a router or switch discovers that a neighboring router or switch is restarting, it waits until the end of the recovery time before attempting to reconnect. The recovery time is the length of time a router or switch waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This time period is also typically the length of time that a neighboring router or switch maintains its information about the restarting router or switch, allowing it to continue to forward traffic.

You can configure LDP graceful restart in both the master instance for the LDP protocol and for a specific routing instance. You can disable graceful restart at the global level for all protocols, at the protocol level for LDP only, and on a specific routing instance. LDP graceful restart is disabled by default, because at the global level, graceful restart is disabled by default. However, helper mode (the ability to assist a neighboring router or switch attempting a graceful restart) is enabled by default.

The following are some of the behaviors associated with LDP graceful restart:

- Outgoing labels are not maintained in restarts. New outgoing labels are allocated.
- When a router or switch is restarting, no label-map messages are sent to neighbors that support graceful restart until the restarting router or switch has stabilized (label-map messages are immediately sent to neighbors that do not support graceful restart). However, all other messages (keepalive, address-message, notification, and release) are sent as usual. Distributing these other messages prevents the router from distributing incomplete information.
- Helper mode and graceful restart are independent. You can disable graceful restart in the configuration, but still allow the router or switch to cooperate with a neighbor attempting to restart gracefully.

## Configuration

---

- [Configuration Tasks on page 3540](#)
- [Configuration Statements on page 3567](#)

### Configuration Tasks

- [Minimum LDP Configuration on page 3541](#)
- [Enabling and Disabling LDP on page 3541](#)
- [Configuring the LDP Timer for Hello Messages on page 3542](#)
- [Configuring the Delay Before LDP Neighbors Are Considered Down on page 3543](#)
- [Enabling Strict Targeted Hello Messages for LDP on page 3544](#)

- [Configuring the Interval for LDP Keepalive Messages on page 3544](#)
- [Configuring the LDP Keepalive Timeout on page 3545](#)
- [Configuring LDP Route Preferences on page 3545](#)
- [Configuring LDP Graceful Restart on page 3545](#)
- [Filtering Inbound LDP Label Bindings on page 3548](#)
- [Filtering Outbound LDP Label Bindings on page 3550](#)
- [Specifying the Transport Address Used by LDP on page 3552](#)
- [Configuring the Prefixes Advertised into LDP from the Routing Table on page 3552](#)
- [Configuring FEC Deaggregation on page 3553](#)
- [Configuring Policers for LDP FECs on page 3554](#)
- [Configuring LDP IPv4 FEC Filtering on page 3554](#)
- [Configuring BFD for LDP LSPs on page 3555](#)
- [Configuring ECMP-Aware BFD for LDP LSPs on page 3558](#)
- [Configuring a Failure Action for the BFD Session on an LDP LSP on page 3558](#)
- [Configuring the Holddown Interval for the BFD Session on page 3559](#)
- [Configuring OAM Ingress Policies for LDP on page 3559](#)
- [Configuring LDP LSP Traceroute on page 3560](#)
- [Collecting LDP Statistics on page 3561](#)
- [Tracing LDP Protocol Traffic on page 3563](#)
- [Standard Firewall Filter Match Conditions for MPLS Traffic on page 3565](#)

---

### Minimum LDP Configuration

To enable LDP on a single interface, include the **ldp** statement and specify the interface using the **interface** statement. This is the minimum LDP configuration. All other LDP configuration statements are optional.

```
ldp {
 interface interface-name;
}
```

To enable LDP on all interfaces, specify **all** for *interface-name*.

For a list of hierarchy levels at which you can include these statements, see the statement summary sections.

---

### Enabling and Disabling LDP

LDP is routing-instance-aware. To enable LDP on a specific interface, include the following statements:

```
ldp {
 interface interface-name;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections.

To enable LDP on all interfaces, specify **all** for *interface-name*.

If you have configured interface properties on a group of interfaces and want to disable LDP on one of the interfaces, include the **interface** statement with the **disable** option:

```
interface interface-name {
 disable;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section.

### Configuring the LDP Timer for Hello Messages

---

LDP hello messages enable LDP nodes to discover one another and to detect the failure of a neighbor or the link to the neighbor. Hello messages are sent periodically on all interfaces where LDP is enabled.

There are two types of LDP hello messages:

- Link hello messages—Sent through the LDP interface as UDP packets addressed to the LDP discovery port. Receipt of an LDP link hello message on an interface identifies an adjacency with the LDP peer router.
- Targeted hello messages—Sent as UDP packets addressed to the LDP discovery port at a specific address. Targeted hello messages are used to support LDP sessions between routers that are not directly connected. A targeted router determines whether to respond or ignore a targeted hello message. A targeted router that chooses to respond does so by periodically sending targeted hello messages back to the initiating router.

By default, LDP sends hello messages every 5 seconds for link hello messages and every 15 seconds for targeted hello messages. You can configure the LDP timer to alter how often both types of hello messages are sent. However, you cannot configure a time for the LDP timer that is greater than the LDP hold time. For more information, see [“Configuring the Delay Before LDP Neighbors Are Considered Down” on page 3543](#).

#### Configuring the LDP Timer for Link Hello Messages

To modify how often LDP sends link hello messages, specify a new link hello message interval for the LDP timer using the **hello-interval** statement:

```
hello-interval seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

#### Configuring the LDP Timer for Targeted Hello Messages

To modify how often LDP sends targeted hello messages, specify a new targeted hello message interval for the LDP timer by configuring the **hello-interval** statement as an option for the **targeted-hello** statement:

```
targeted-hello {
 hello-interval seconds;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

### Configuring the Delay Before LDP Neighbors Are Considered Down

The hold time determines how long an LDP node should wait for a hello message before declaring a neighbor to be down. This value is sent as part of a hello message so that each LDP node tells its neighbors how long to wait. The values sent by each neighbor do not have to match.

The hold time should normally be at least three times the hello interval. The default is 15 seconds for link hello messages and 45 seconds for targeted hello messages. However, it is possible to configure an LDP hold time that is close to the value for the hello interval.



**NOTE:** By configuring an LDP hold time close to the hello interval (less than three times the hello interval), LDP neighbor failures might be detected more quickly. However, this also increases the possibility that the router or switch might declare an LDP neighbor down that is still functioning normally. For more information, see [“Configuring the LDP Timer for Hello Messages” on page 3542](#).

The LDP hold time is also negotiated automatically between LDP peers. When two LDP peers advertise different LDP hold times to one another, the smaller value is used. If an LDP peer router or switch advertises a shorter hold time than the value you have configured, the peer router's or switch's advertised hold time is used. This negotiation can affect the LDP keepalive interval as well.

If the local LDP hold time is not shortened during LDP peer negotiation, the user-configured keepalive interval is left unchanged. However, if the local hold time is reduced during peer negotiation, the keepalive interval is recalculated. If the LDP hold time has been reduced during peer negotiation, the keepalive interval is reduced to one-third of the new hold time value. For example, if the new hold-time value is 45 seconds, the keepalive interval is set to 15 seconds.

This automated keepalive interval calculation can cause different keepalive intervals to be configured on each peer router or switch. This enables the routers or switches to be flexible in how often they send keepalive messages, because the LDP peer negotiation ensures they are sent more frequently than the LDP hold time.

When you reconfigure the hold-time interval, changes do not take effect until after the session is reset. The hold time is negotiated when the LDP peering session is initiated and cannot be renegotiated as long as the session is up (required by RFC 5036, *LDP Specification*). To manually force the LDP session to reset, issue the **clear ldp session** command.

### Configuring the LDP Hold Time for Link Hello Messages

To modify how long an LDP node should wait for a link hello message before declaring the neighbor down, specify a new time in seconds using the **hold-time** statement:

```
hold-time seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

### ***Configuring the LDP Hold Time for Targeted Hello Messages***

To modify how long an LDP node should wait for a targeted hello message before declaring the neighbor down, specify a new time in seconds using the **hold-time** statement as an option for the **targeted-hello** statement:

```
targeted-hello {
 hold-time seconds;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

---

### **Enabling Strict Targeted Hello Messages for LDP**

Use strict targeted hello messages to prevent LDP sessions from being established with remote neighbors that have not been specifically configured. If you configure the **strict-targeted-hellos** statement, an LDP peer does not respond to targeted hello messages coming from a source that is not one of its configured remote neighbors. Configured remote neighbors can include:

- Endpoints of RSVP tunnels for which LDP tunneling is configured
- Layer 2 circuit neighbors

If an unconfigured neighbor sends a hello message, the LDP peer ignores the message and logs an error (with the **error** trace flag) indicating the source. For example, if the LDP peer received a targeted hello from the Internet address 10.0.0.1 and no neighbor with this address is specifically configured, the following message is printed to the LDP log file:

```
LDP: Ignoring targeted hello from 10.0.0.1
```

To enable strict targeted hello messages, include the **strict-targeted-hellos** statement:

```
strict-targeted-hellos;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

---

### **Configuring the Interval for LDP Keepalive Messages**

The keepalive interval determines how often a message is sent over the session to ensure that the keepalive timeout is not exceeded. If no other LDP traffic is sent over the session in this much time, a keepalive message is sent. The default is 10 seconds. The minimum value is 1 second.

The value configured for the keepalive interval can be altered during LDP session negotiation if the value configured for the LDP hold time on the peer router or switch is lower than the value configured locally. For more information, see [“Configuring the Delay Before LDP Neighbors Are Considered Down” on page 3543](#).

To modify the keepalive interval, include the **keepalive-interval** statement:

**keepalive-interval** *seconds*;

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

### Configuring the LDP Keepalive Timeout

After an LDP session is established, messages must be exchanged periodically to ensure that the session is still working. The keepalive timeout defines the amount of time that the neighbor LDP node waits before deciding that the session has failed. This value is usually set to at least three times the keepalive interval. The default is 30 seconds.

To modify the keepalive interval, include the **keepalive-timeout** statement:

**keepalive-timeout** *seconds*;

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The value configured for the **keepalive-timeout** statement is displayed as the hold time when you issue the **show ldp session detail** command.

### Configuring LDP Route Preferences

When several protocols calculate routes to the same destination, route preferences are used to select which route is installed in the forwarding table. The route with the lowest preference value is selected. The preference value can be a number in the range 0 through 255. By default, LDP routes have a preference value of 9.

To modify the route preferences, include the **preference** statement:

**preference** *preference*;

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

### Configuring LDP Graceful Restart

When you alter the graceful restart configuration at either the **[edit routing-options graceful-restart]** or **[edit protocols ldp graceful-restart]** hierarchy levels, any running LDP session is automatically restarted to apply the graceful restart configuration. This behavior mirrors the behavior of BGP when you alter its graceful restart configuration.

By default, graceful restart helper mode is enabled, but graceful restart is disabled. Thus, the default behavior of a router or switch is to assist neighboring routers or switches attempting a graceful restart, but not to attempt a graceful restart itself.

To configure LDP graceful restart, see the following sections:

- [Enabling Graceful Restart on page 3546](#)
- [Disabling LDP Graceful Restart or Helper Mode on page 3546](#)
- [Configuring Reconnect Time on page 3547](#)
- [Configuring Recovery Time and Maximum Recovery Time on page 3547](#)

### **Enabling Graceful Restart**

To enable LDP graceful restart, you also need to enable graceful restart on the router or switch. To enable graceful restart, include the **graceful-restart** statement:

```
graceful-restart;
```

You can include this statement at the following hierarchy levels:

- [edit routing-options]
- [edit logical-systems *logical-system-name* routing-options]

The **graceful-restart** statement enables graceful restart for all protocols supporting this feature on the route or switch. For more information about graceful restart, see the *Junos OS Routing Protocols Configuration Guide*.

By default, LDP graceful restart is enabled when you enable graceful restart at both the LDP protocol level and on all the routing instances. However, you can disable both LDP graceful restart and LDP graceful restart helper mode.

### **Disabling LDP Graceful Restart or Helper Mode**

To disable LDP graceful restart and recovery, include the **disable** statement:

```
ldp {
 graceful-restart {
 disable;
 }
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

You can disable helper mode at the LDP protocols level only. You cannot disable helper mode for a specific routing instance. To disable LDP helper mode, include the **helper-disable** statement:

```
ldp {
 graceful-restart {
 helper-disable;
 }
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The following LDP graceful restart configurations are possible:

- LDP graceful restart and helper mode are both enabled.
- LDP graceful restart is disabled but helper mode is enabled. A router or switch configured in this way cannot restart gracefully but can help a restarting neighbor.
- LDP graceful restart and helper mode are both disabled. The router or switch does not use LDP graceful restart or the graceful restart type, length, and value (TLV) sent in



the initialization message. The router or switch behaves as a router or switch that cannot support LDP graceful restart.

A configuration error is issued if you attempt to enable graceful restart and disable helper mode.

### ***Configuring Reconnect Time***

After the LDP connection between neighbors fails, neighbors wait a certain amount of time for the gracefully restarting router or switch to resume sending LDP messages. After the wait period, the LDP session can be reestablished. You can configure the wait period in seconds. This value is included in the fault tolerant session TLV sent in LDP initialization messages when LDP graceful restart is enabled.

Suppose that Router A and Router B are LDP neighbors. Router A is the restarting Router. The reconnect time is the time that Router A tells Router B to wait after Router B detects that Router A restarted.

To configure the reconnect time, include the **reconnect-time** statement:

```
graceful-restart {
 reconnect-time seconds;
}
```

You can set the reconnect time to a value in the range from 30 through 300 seconds. By default, it is 60 seconds.

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

### ***Configuring Recovery Time and Maximum Recovery Time***

The recovery time is the amount of time a router or switch waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This period is also typically the amount of time that a neighboring router maintains its information about the restarting router, allowing it to continue to forward traffic.

To prevent a neighboring router or switch from being adversely affected if it receives a false value for the recovery time from the restarting router or switch, you can configure the maximum recovery time on the neighboring router or switch. A neighboring router or switch maintains its state for the shorter of the two times. For example, Router A is performing an LDP graceful restart. It has sent a recovery time of 900 seconds to neighboring Router B. However, Router B has its maximum recovery time configured at 400 seconds. Router B will only wait for 400 seconds before it purges its LDP information from Router A.

To configure recovery time, include the **recovery-time** statement and the **maximum-neighbor-recovery-time** statement:

```
graceful-restart {
 maximum-neighbor-recovery-time seconds;
 recovery-time seconds;
}
```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

### Filtering Inbound LDP Label Bindings

You can filter received LDP label bindings, applying policies to accept or deny bindings advertised by neighboring routers. To configure received-label filtering, include the **import** statement:

```
import [policy-names];
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The named policy (configured at the **[edit policy-options]** hierarchy level) is applied to all label bindings received from all LDP neighbors. All filtering is done with **from** statements. [Table 268 on page 3548](#) lists the only **from** operators that apply to LDP received-label filtering.

**Table 268: from Operators That Apply to LDP Received-Label Filtering**

| from Operator       | Description                                                                                |
|---------------------|--------------------------------------------------------------------------------------------|
| <b>interface</b>    | Matches on bindings received from a neighbor that is adjacent over the specified interface |
| <b>neighbor</b>     | Matches on bindings received from the specified LDP router ID                              |
| <b>next-hop</b>     | Matches on bindings received from a neighbor advertising the specified interface address   |
| <b>route-filter</b> | Matches on bindings with the specified prefix                                              |

If a binding is filtered, it still appears in the LDP database, but is not considered for installation as part of a label-switched path (LSP).

Generally, applying policies in LDP can be used only to block the establishment of LSPs, not to control their routing. This is because the path that an LSP follows is determined by unicast routing, and not by LDP. However, when there are multiple equal-cost paths to the destination through different neighbors, you can use LDP filtering to exclude some of the possible next hops from consideration. (Otherwise, LDP chooses one of the possible next hops at random.)

LDP sessions are not bound to interfaces or interface addresses. LDP advertises only per-router or per-switch (not per-interface) labels; so if multiple parallel links exist between two routers or switches, only one LDP session is established, and it is not bound to a single interface. When a router or switch has multiple adjacencies to the same neighbor, take care to ensure that the filter does what is expected. (Generally, using **next-hop** and **interface** is not appropriate in this case.)

If a label has been filtered (meaning that it has been rejected by the policy and is not used to construct an LSP), it is marked as filtered in the database:

```
user@host> show ldp database
Input label database, 10.10.255.1:0-10.10.255.6:0
Label Prefix
3 10.10.255.6/32 (Filtered)
Output label database, 10.10.255.1:0-10.10.255.6:0
Label Prefix
3 10.10.255.1/32 (Filtered)
```

For more information about how to configure policies for LDP, see the *Routing Policy Configuration Guide*.

### **Examples: Filtering Inbound LDP Label Bindings**

Accept only /32 prefixes from all neighbors:

```
[edit]
protocols {
 ldp {
 import only-32;
 ...
 }
}
policy-options {
 policy-statement only-32 {
 term first {
 from {
 route-filter 0.0.0.0/0 upto /31;
 }
 then reject;
 }
 then accept;
 }
}
```

Accept 131.108/16 or longer from router ID 10.10.255.2 and accept all prefixes from all other neighbors:

```
[edit]
protocols {
 ldp {
 import nosy-neighbor;
 ...
 }
}
policy-options {
 policy-statement nosy-neighbor {
 term first {
 from {
 neighbor 10.10.255.2;
 route-filter 131.108.0.0/16 orlonger accept;
 route-filter 0.0.0.0/0 orlonger reject;
 }
 }
 then accept;
 }
}
```

```
}
}
```

### Filtering Outbound LDP Label Bindings

You can configure export policies to filter LDP outbound labels. You can filter outbound label bindings by applying routing policies to block bindings from being advertised to neighboring routers or switches. To configure outbound label filtering, include the **export** statement:

```
export [policy-name];
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The named export policy (configured at the **[edit policy-options]** hierarchy level) is applied to all label bindings transmitted to all LDP neighbors. The only **from** operator that applies to LDP outbound label filtering is **route-filter**, which matches bindings with the specified prefix. The only **to** operators that apply to outbound label filtering are the operators in [Table 269 on page 3550](#).

**Table 269: to Operators for LDP Outbound-Label Filtering**

| to Operator      | Description                                                                          |
|------------------|--------------------------------------------------------------------------------------|
| <b>interface</b> | Matches on bindings sent to a neighbor that is adjacent over the specified interface |
| <b>neighbor</b>  | Matches on bindings sent to the specified LDP router ID                              |
| <b>next-hop</b>  | Matches on bindings sent to a neighbor advertising the specified interface address   |

If a binding is filtered, the binding is not advertised to the neighboring router or switch, but it can be installed as part of an LSP on the local router or switch. You can apply policies in LDP to block the establishment of LSPs, but not to control their routing. The path an LSP follows is determined by unicast routing, not by LDP.

LDP sessions are not bound to interfaces or interface addresses. LDP advertises only per-router or per-switch (not per-interface) labels. If multiple parallel links exist between two routers or switches, only one LDP session is established, and it is not bound to a single interface.

Do not use the **next-hop** and **interface** operators when a router or switch has multiple adjacencies to the same neighbor.

Filtered labels are marked in the database:

```
user@host> show ldp database
Input label database, 10.10.255.1:0-10.10.255.3:0
Label Prefix
100007 10.10.255.2/32
3 10.10.255.3/32
Output label database, 10.10.255.1:0-10.10.255.3:0
Label Prefix
```

```
3 10.10.255.1/32
100001 10.10.255.6/32 (Filtered)
```

For more information about how to configure policies for LDP, see the *Routing Policy Configuration Guide*.

### **Examples: Filtering Outbound LDP Label Bindings**

Block transmission of the route for 10.10.255.6/32 to any neighbors:

```
[edit protocols]
ldp {
 export block-one;
}
policy-options {
 policy-statement block-one {
 term first {
 from {
 route-filter 10.10.255.6/32 exact;
 }
 then reject;
 }
 then accept;
 }
}
```

Send only 131.108/16 or longer to router ID 10.10.255.2, and send all prefixes to all other routers or switches:

```
[edit protocols]
ldp {
 export limit-lsps;
}
policy-options {
 policy-statement limit-lsps {
 term allow-one {
 from {
 route-filter 131.108.0.0/16 orlonger;
 }
 to {
 neighbor 10.10.255.2;
 }
 then accept;
 }
 term block-the-rest {
 to {
 neighbor 10.10.255.2;
 }
 then reject;
 }
 then accept;
 }
}
```

### Specifying the Transport Address Used by LDP

---

You can control the transport address used by LDP. The transport address is the address used for the TCP session over which LDP is running. To configure transport address control, include the **transport-address** statement:

**transport-address** (router-id | interface);

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

If you specify the **router-id** option, the address of the router or switch identifier is used as the transport address (unless otherwise configured, the router or switch identifier is typically the same as the loopback address). If you specify the **interface** option, the interface address is used as the transport address for any LDP sessions to neighbors that can be reached over that interface. Note that the identifier is used as the transport address by default.

You cannot specify the **interface** option when there are multiple parallel links to the same LDP neighbor, because the LDP specification requires that the same transport address be advertised on all interfaces to the same neighbor. If LDP detects multiple parallel links to the same neighbor, it disables interfaces to that neighbor one by one until the condition is cleared, either by disconnecting the neighbor on an interface or by specifying the **router-id** option.

### Configuring the Prefixes Advertised into LDP from the Routing Table

---

You can control the set of prefixes that are advertised into LDP and cause the router or switch to be the egress router or switch for those prefixes. By default, only the loopback address is advertised into LDP. To configure the set of prefixes from the routing table to be advertised into LDP, include the **egress-policy** statement:

**egress-policy** policy-name;

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.



**NOTE:** If you configure an egress policy for LDP that does not include the loopback address, it is no longer advertised in LDP. To continue to advertise the loopback address, you need to explicitly configure it as a part of the LDP egress policy.

---

The named policy (configured at the **[edit policy-options]** or **[edit logical-systems logical-system-name policy-options]** hierarchy level) is applied to all routes in the routing table. Those routes that match the policy are advertised into LDP. You can control the set of neighbors to which those prefixes are advertised by using the **export** statement. Only **from** operators are considered; you can use any valid **from** operator. For more information, see the *Junos OS Routing Protocols Configuration Guide*.

**Example: Configuring the Prefixes Advertised into LDP**

Advertise all connected routes into LDP:

```
[edit protocols]
ldp {
 egress-policy connected-only;
}
policy-options {
 policy-statement connected-only {
 from {
 protocol direct;
 }
 then accept;
 }
}
```

**Configuring FEC Deaggregation**

When an LDP egress router or switch advertises multiple prefixes, the prefixes are bound to a single label and aggregated into a single forwarding equivalence class (FEC). By default, LDP maintains this aggregation as the advertisement traverses the network.

Normally, because an LSP is not split across multiple next hops and the prefixes are bound into a single LSP, load-balancing across equal-cost paths does not occur. You can, however, load-balance across equal-cost paths if you configure a load-balancing policy and deaggregate the FECs.

Deaggregating the FECs causes each prefix to be bound to a separate label and become a separate LSP.

To configure deaggregated FECs, include the **deaggregate** statement:

```
deaggregate;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

For all LDP sessions, you can configure deaggregated FECs only globally.

Deaggregating a FEC allows the resulting multiple LSPs to be distributed across multiple equal-cost paths and distributes LSPs across the multiple next hops on the egress segments but installs only one next hop per LSP.

To aggregate FECs, include the **no-deaggregate** statement:

```
no-deaggregate;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

For all LDP sessions, you can configure aggregated FECs only globally.

**Related  
Documentation**

- *Configuring Load Balancing Across RSVP LSPs*
- *Configuring Protocol-Independent Load Balancing in Layer 3 VPNs*

- [Configuring VPLS Load Balancing on page 5348](#)
- *Example: Load Balancing BGP Traffic*

### Configuring Policers for LDP FECs

---

You can configure the Junos OS to track and police traffic for LDP FECs. LDP FEC policers can be used to do any of the following:

- Track or police the ingress traffic for an LDP FEC.
- Track or police the transit traffic for an LDP FEC.
- Track or police LDP FEC traffic originating from a specific forwarding class.
- Track or police LDP FEC traffic originating from a specific virtual routing and forwarding (VRF) site.
- Discard false traffic bound for a specific LDP FEC.

To police traffic for an LDP FEC, you must first configure a filter. Specifically, you need to configure either the **interface** statement or the **interface-set** statement at the **[edit firewall family protocol-family filter filter-name term term-name from]** hierarchy level. The **interface** statement allows you to match the filter to a single interface. The **interface-set** statement allows you to match the filter to multiple interfaces.

For more information on how to configure the **interface** statement, the **interface-set** statement, and policers for LDP FECs, see the *Routing Policy Configuration Guide*.

Once you have configured the filters, you need to include them in the **policing** statement configuration for LDP. To configure policers for LDP FECs, include the **policing** statement:

```
policing {
 fec fec-address {
 ingress-traffic filter-name;
 transit-traffic filter-name;
 }
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The **policing** statement includes the following options:

- **fec**—Specify the FEC address for the LDP FEC you want to police.
- **ingress-filter**—Specify the name of the ingress traffic filter.
- **transit-traffic**—Specify the name of the transit traffic filter.

### Configuring LDP IPv4 FEC Filtering

---

By default, when a targeted LDP session is established, the Junos OS always exchanges both the IPv4 forwarding equivalence classes (FECs) and the Layer 2 circuit FECs over the targeted LDP session. For an LDP session to an indirectly connected neighbor, you



might only want to export Layer 2 circuit FECs to the neighbor if the session was specifically configured to support Layer 2 circuits or VPLS.

In a mixed vendor network where all non-BGP prefixes are advertised into LDP, the LDP database can become large. For this type of environment, it can be useful to prevent the advertisement of IPv4 FECs over LDP sessions formed because of Layer 2 circuit or LDP VPLS configuration. Similarly, it can be useful to filter any IPv4 FECs received in this sort of environment.

If all the LDP neighbors associated with an LDP session are Layer 2 only, you can configure the Junos OS to advertise only Layer 2 circuit FECs by configuring the **l2-smart-policy** statement. This feature also automatically filters out the IPv4 FECs received on this session. If you have configured an explicit export or import policy, this feature is disabled.

If one of the LDP session's neighbors is formed because of a discovered adjacency or if the adjacency is formed because of an LDP tunneling configuration on one or more RSVP LSPs, the IPv4 FECs are advertised and received using the default behavior.

To prevent LDP from exporting IPv4 FECs over LDP sessions with Layer 2 neighbors only and to filter out IPv4 FECs received over such sessions, include the **l2-smart-policy** statement:

```
l2-smart-policy;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary for this statement.

### Configuring BFD for LDP LSPs

You can configure Bidirectional Forwarding Detection (BFD) for LDP LSPs. The BFD protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the router or switch stops receiving a reply after a specified interval. BFD works with a wide variety of network environments and topologies. The failure detection timers for BFD have shorter time limits than the failure detection mechanisms of static routes, providing faster detection.

An error is logged whenever a BFD session for a path fails. The following shows how BFD for LDP LSP log messages might appear:

```
RPD_LDP_BFD_UP: LDP BFD session for FEC 10.255.16.14/32 is up
RPD_LDP_BFD_DOWN: LDP BFD session for FEC 10.255.16.14/32 is down
```

You can also configure BFD for RSVP LSPs, as described in *Configuring BFD for MPLS IPv4 LSPs*.

The BFD failure detection timers are adaptive and can be adjusted to be more or less aggressive. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear**

**bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

To enable BFD for LDP LSPs, include the **oam** and **bfd-liveness-detection** statements:

```
oam {
 bfd-liveness-detection {
 detection-time threshold milliseconds;
 ecmp;
 failure-action {
 remove-nexthop;
 remove-route;
 }
 holddown-interval seconds;
 ingress-policy ingress-policy-name;
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 minimum-transmit-interval milliseconds;
 multiplier detection-time-multiplier;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 }
}
fec fec-address {
 bfd-liveness-detection {
 detection-time threshold milliseconds;
 ecmp;
 failure-action {
 remove-nexthop;
 remove-route;
 }
 holddown-interval milliseconds;
 ingress-policy ingress-policy-name;
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 minimum-transmit-interval milliseconds;
 multiplier detection-time-multiplier;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 }
 version (0 | 1 | automatic);
}
no-bfd-liveness-detection;
periodic-traceroute {
 disable;
 exp exp-value;
 fanout fanout-value;
 frequency minutes;
 paths number-of-paths;
 retries retry-attempts;
 source address;
```

```

 ttl ttl-value;
 wait seconds;
 }
}
lsp-ping-interval seconds;
periodic-traceroute {
 disable;
 exp exp-value;
 fanout fanout-value;
 frequency minutes;
 paths number-of-paths;
 retries retry-attempts;
 source address;
 ttl ttl-value;
 wait seconds;
}
}

```

You can enable BFD for the LDP LSPs associated with a specific forwarding equivalence class (FEC) by configuring the FEC address using the **fec** option at the **[edit protocols ldp]** hierarchy level. Alternatively, you can configure an Operation Administration and Management (OAM) ingress policy to enable BFD on a range of FEC addresses. For more information, see [“Configuring OAM Ingress Policies for LDP” on page 3559](#).

You cannot enable BFD LDP LSPs unless their equivalent FEC addresses are explicitly configured or OAM is enabled on the FECs using an OAM ingress policy. If BFD is not enabled for any FEC addresses, the BFD session will not come up.

You can configure the **oam** statement at the following hierarchy levels:

- **[edit protocols ldp]**
- **[edit logical-systems *logical-system-name* protocols ldp]**

The **oam** statement includes the following options:

- **fec**—Specify the FEC address. You must either specify a FEC address or configure an OAM ingress policy to ensure that the BFD session comes up.
- **lsp-ping-interval**—Specify the duration of the LSP ping interval in seconds. To issue a ping on an LDP-signaled LSP, use the **ping mpls ldp** command. For more information, see the *Junos OS Operational Mode Commands*.

The **bfd-liveness-detection** statement includes the following options:

- **ecmp**—Cause LDP to establish BFD sessions for all ECMP paths configured for the specified FEC. If you configure the **ecmp** option, you must also configure the **periodic-traceroute** statement for the specified FEC. If you do not do so, the commit operation fails. You can configure the **periodic-traceroute** statement at the global hierarchy level (**[edit protocols ldp oam]**) while only configuring the **ecmp** option for a specific FEC (**[edit protocols ldp oam fec address bfd-liveness-detection]**).
- **holddown-interval**—Specify the duration the BFD session should remain up before adding the route or next hop. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.

- **minimum-interval**—Specify the minimum transmit and receive interval. If you configure the **minimum-interval** option, you do not need to configure the **minimum-receive-interval** option or the **minimum-transmit-interval** option.
- **minimum-receive-interval**—Specify the minimum receive interval. The range is from 1 through 255,000 milliseconds.
- **minimum-transmit-interval**—Specify the minimum transmit interval. The range is from 1 through 255,000 milliseconds.
- **multiplier**—Specify the detection time multiplier. The range is from 1 through 255.

---

### Configuring ECMP-Aware BFD for LDP LSPs

---

When you configure BFD for a FEC, a BFD session is established for only one active local next-hop for the router or switch. However, you can configure multiple BFD sessions, one for each FEC associated with a specific equal-cost multipath (ECMP) path. For this to function properly, you also need to configure LDP LSP periodic traceroute. (See [“Configuring LDP LSP Traceroute” on page 3560](#).) LDP LSP traceroute is used to discover ECMP paths. A BFD session is initiated for each ECMP path discovered. Whenever a BFD session for one of the ECMP paths fails, an error is logged.

LDP LSP traceroute is run periodically to check the integrity of the ECMP paths. The following might occur when a problem is discovered:

- If the latest LDP LSP traceroute for a FEC differs from the previous traceroute, the BFD sessions associated with that FEC (the BFD sessions for address ranges that have changed from previous run) are brought down and new BFD sessions are initiated for the destination addresses in the altered ranges.
- If the LDP LSP traceroute returns an error (for example, a timeout), all the BFD sessions associated with that FEC are torn down.

To configure LDP to establish BFD sessions for all ECMP paths configured for the specified FEC, include the **ecmp** statement.

**ecmp;**

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Along with the **ecmp** statement, you must also include the **periodic-traceroute** statement, either in the global LDP OAM configuration (at the **[edit protocols ldp oam]** or **[edit logical-systems logical-system-name protocols ldp oam]** hierarchy level) or in the configuration for the specified FEC (at the **[edit protocols ldp oam fec address]** or **[edit logical-systems logical-system-name protocols ldp oam fec address]** hierarchy level). Otherwise, the commit operation fails.

---

### Configuring a Failure Action for the BFD Session on an LDP LSP

---

You can configure route and next-hop properties in the event of a BFD session failure event on an LDP LSP. The failure event could be an existing BFD session that has gone down or could be a BFD session that never came up. LDP adds back the route or next hop when the relevant BFD session comes back up.

You can configure one of the following failure action options for the **failure-action** statement in the event of a BFD session failure on the LDP LSP:

- **remove-nexthop**—Removes the route corresponding to the next hop of the LSP's route at the ingress node when a BFD session failure event is detected.
- **remove-route**—Removes the route corresponding to the LSP from the appropriate routing tables when a BFD session failure event is detected. If the LSP is configured with ECMP and a BFD session corresponding to any path goes down, the route is removed.

To configure a failure action in the event of a BFD session failure on an LDP LSP, include either the **remove-nexthop** option or the **remove-route** option for the **failure-action** statement:

```
failure-action {
 remove-nexthop;
 remove-route;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

### Configuring the Holddown Interval for the BFD Session

You can specify the duration the BFD session should be up before adding a route or next hop by configuring the **holddown-interval** statement at either the **[edit protocols ldp oam bfd-liveness-detection]** hierarchy level or at the **[edit protocols ldp oam fec address bfd-liveness-detection]** hierarchy level. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.

```
holddown-interval seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

### Configuring OAM Ingress Policies for LDP

Using the **ingress-policy** statement, you can configure an Operation, Administration, and Management (OAM) policy to choose which forwarding equivalence classes (FECs) need to have OAM enabled. If the FEC passes through the policy or if the FEC is explicitly configured, OAM is enabled for a FEC. For FECs chosen using a policy, the BFD parameters configured under **[edit protocols ldp oam bfd-liveness-detection]** are applied.

You configure the OAM ingress policy at the **[edit policy-options]** hierarchy level. To configure an OAM ingress policy, include the **ingress-policy** statement:

```
ingress-policy ingress-policy-name;
```

You can configure this statement at the following hierarchy levels:

- **[edit protocols ldp oam]**
- **[edit logical-systems logical-system-name protocols ldp oam]**

## Configuring LDP LSP Traceroute

---

You can trace the route followed by an LDP-sigaled LSP. LDP LSP traceroute is based on RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*. This feature allows you to periodically trace all paths in a FEC. The FEC topology information is stored in a database accessible from the CLI.

A topology change does not automatically trigger a trace of an LDP LSP. However, you can manually initiate a traceroute. If the traceroute request is for an FEC that is currently in the database, the contents of the database are updated with the results.

The periodic traceroute feature applies to all FECs specified by the **oam** statement configured at the **[edit protocols ldp]** hierarchy level. To configure periodic LDP LSP traceroute, include the **periodic-traceroute** statement:

```
periodic-traceroute {
 disable;
 exp exp-value;
 fanout fanout-value;
 frequency minutes;
 paths number-of-paths;
 retries retry-attempts;
 source address;
 ttl ttl-value;
 wait seconds;
}
```

You can configure this statement at the following hierarchy levels:

- **[edit protocols ldp oam]**
- **[edit protocols ldp oam fec address]**

You can configure the **periodic-traceroute** statement by itself or with any of the following options:

- **exp**—Specify the class of service to use when sending probes.
- **fanout**—Specify the maximum number of next hops to search per node.
- **frequency**—Specify the interval between traceroute attempts.
- **paths**—Specify the maximum number of paths to search.
- **retries**—Specify the number of attempts to send a probe to a specific node before giving up.
- **source**—Specify the IPv4 source address to use when sending probes.
- **ttl**—Specify the maximum time-to-live value. Nodes that are beyond this value are not traced.
- **wait**—Specify the wait interval before resending a probe packet.

## Collecting LDP Statistics

LDP traffic statistics show the volume of traffic that has passed through a particular FEC on a router.

When you configure the **traffic-statistics** statement at the **[edit protocols ldp]** hierarchy level, the LDP traffic statistics are gathered periodically and written to a file. You can configure how often statistics are collected (in seconds) by using the **interval** option. The default collection interval is 5 minutes. You must configure an LDP statistics file; otherwise, LDP traffic statistics are not gathered. If the LSP goes down, the LDP statistics are reset.

To collect LDP traffic statistics, include the **traffic-statistics** statement:

```
traffic-statistics {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 interval interval;
 no-penultimate-hop;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

This section includes the following topics:

- [LDP Statistics Output on page 3561](#)
- [Disabling LDP Statistics on the Penultimate-Hop Router on page 3562](#)
- [LDP Statistics Limitations on page 3562](#)

### LDP Statistics Output

The following sample output is from an LDP statistics file:

| FEC               | Type    | Packets | Bytes | Shared |
|-------------------|---------|---------|-------|--------|
| 10.255.350.448/32 | Transit | 0       | 0     | No     |
|                   | Ingress | 0       | 0     | No     |
| 10.255.350.450/32 | Transit | 0       | 0     | Yes    |
|                   | Ingress | 0       | 0     | No     |
| 10.255.350.451/32 | Transit | 0       | 0     | No     |
|                   | Ingress | 0       | 0     | No     |
| 220.220.220.1/32  | Transit | 0       | 0     | Yes    |
|                   | Ingress | 0       | 0     | No     |
| 220.220.220.2/32  | Transit | 0       | 0     | Yes    |
|                   | Ingress | 0       | 0     | No     |
| 220.220.220.3/32  | Transit | 0       | 0     | Yes    |
|                   | Ingress | 0       | 0     | No     |

May 28 15:02:05, read 12 statistics in 00:00:00 seconds

The LDP statistics file includes the following columns of data:

- **Bytes**—Number of bytes of data passed by the FEC since its LSP came up.
- **FEC**—FEC for which LDP traffic statistics are collected.
- **Packets**—Number of packets passed by the FEC since its LSP came up.

- **read**—This number (which appears next to the date and time) might differ from the actual number of the statistics displayed. Some of the statistics are summarized before being displayed.
- **Shared**—A **Yes** value indicates that several prefixes are bound to the same label (for example, when several prefixes are advertised with an egress policy). The LDP traffic statistics for this case apply to all the prefixes and should be treated as such.
- **Type**—Type of traffic originating from a router or switch, either **Ingress** (originating from this router or switch) or **Transit** (forwarded through this router or switch).

### *Disabling LDP Statistics on the Penultimate-Hop Router*

Gathering LDP traffic statistics at the penultimate-hop router or switch can consume excessive system resources, on next-hop routes in particular. This problem is exacerbated if you have configured the **deaggregate** statement in addition to the **traffic-statistics** statement. For routers or switches reaching their limit of next-hop route usage, we recommend configuring the **no-penultimate-hop** option for the **traffic-statistics** statement:

```
traffic-statistics {
 no-penultimate-hop;
}
```

For a list of hierarchy levels at which you can configure the **traffic-statistics** statement, see the statement summary section for this statement.



**NOTE:** When you configure the **no-penultimate-hop** option, no statistics are available for the FECs that are the penultimate hop for this router or switch.

Whenever you include or remove this option from the configuration, the LDP sessions are taken down and then restarted.

The following sample output is from an LDP statistics file showing routers or switches on which the **no-penultimate-hop** option is configured:

| FEC               | Type    | Packets             | Bytes | Shared |
|-------------------|---------|---------------------|-------|--------|
| 10.255.245.218/32 | Transit | 0                   | 0     | No     |
|                   | Ingress | 4                   | 246   | No     |
| 10.255.245.221/32 | Transit | statistics disabled |       |        |
|                   | Ingress | statistics disabled |       |        |
| 13.1.1.0/24       | Transit | statistics disabled |       |        |
|                   | Ingress | statistics disabled |       |        |
| 13.1.3.0/24       | Transit | statistics disabled |       |        |
|                   | Ingress | statistics disabled |       |        |

### *LDP Statistics Limitations*

The following are issues related to collecting LDP statistics by configuring the **traffic-statistics** statement:

- You cannot clear the LDP statistics.
- If you shorten the specified interval, a new LDP statistics request is issued only if the statistics timer expires later than the new interval.



- A new LDP statistics collection operation cannot start until the previous one has finished. If the interval is short or if the number of LDP statistics is large, the time gap between the two statistics collections might be longer than the interval.

When an LSP goes down, the LDP statistics are reset.

### Tracing LDP Protocol Traffic

The following sections describe how to configure the trace options to examine LDP protocol traffic:

- [Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels on page 3563](#)
- [Tracing LDP Protocol Traffic Within FECs on page 3564](#)
- [Examples: Tracing LDP Protocol Traffic on page 3564](#)

#### *Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels*

To trace LDP protocol traffic, you can specify options in the global **traceoptions** statement at the **[edit routing-options]** hierarchy level, and you can specify LDP-specific options by including the **traceoptions** statement:

```
traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Use the **file** statement to specify the name of the file that receives the output of the tracing operation. All files are placed in the directory `/var/log`. We recommend that you place LDP-tracing output in the file **ldp-log**.

The following trace flags display the operations associated with the sending and receiving of various LDP messages. Each can carry one or more of the following modifiers:

- **address**—Trace the operation of address and address withdrawal messages.
- **binding**—Trace label-binding operations.
- **error**—Trace error conditions.
- **event**—Trace protocol events.
- **initialization**—Trace the operation of initialization messages.
- **label**—Trace the operation of label request, label map, label withdrawal, and label release messages.
- **notification**—Trace the operation of notification messages.
- **packets**—Trace the operation of address, address withdrawal, initialization, label request, label map, label withdrawal, label release, notification, and periodic messages. This modifier is equivalent to setting the **address**, **initialization**, **label**, **notification**, and **periodic** modifiers.

You can also configure the **filter** flag modifier with the **match-on address** sub-option for the **packets** flag. This allows you to trace based on the source and destination addresses of the packets.

- **path**—Trace label-switched path operations.
- **periodic**—Trace the operation of hello and keepalive messages.
- **route**—Trace the operation of route messages.
- **state**—Trace protocol state transitions.

### *Tracing LDP Protocol Traffic Within FECs*

LDP associates a forwarding equivalence class (FEC) with each LSP it creates. The FEC associated with an LSP specifies which packets are mapped to that LSP. LSPs are extended through a network as each router or switch chooses the label advertised by the next hop for the FEC and splices it to the label it advertises to all other routers or switches.

You can trace LDP protocol traffic within a specific FEC and filter LDP trace statements based on an FEC. This is useful when you want to trace or troubleshoot LDP protocol traffic associated with an FEC. The following trace flags are available for this purpose: **route**, **path**, and **binding**.

The following example illustrates how you might configure the LDP **traceoptions** statement to filter LDP trace statements based on an FEC:

```
[edit protocols ldp traceoptions]
set flag route filter match-on fec policy "filter-policy-for-ldp-fec";
```

This feature has the following limitations:

- The filtering capability is only available for FECs composed of IP version 4 (IPv4) prefixes.
- Layer 2 circuit FECs cannot be filtered.
- When you configure both route tracing and filtering, MPLS routes are not displayed (they are blocked by the filter).
- Filtering is determined by the policy and the configured value for the **match-on** option. When configuring the policy, be sure that the default behavior is always **reject**.
- The only **match-on** option is **fec**. Consequently, the only type of policy you should include is a route-filter policy.

### *Examples: Tracing LDP Protocol Traffic*

Trace LDP path messages in detail:

```
[edit]
protocols {
 ldp {
 traceoptions {
 file ldp size 10m files 5;
 flag path;
```

```

 }
 }
}

```

Trace all LDP outgoing messages:

```

[edit]
protocols {
 ldp {
 traceoptions {
 file ldp size 10m files 5;
 flag packets;
 }
 }
}

```

Trace all LDP error conditions:

```

[edit]
protocols {
 ldp {
 traceoptions {
 file ldp size 10m files 5;
 flag error;
 }
 }
}

```

Trace all LDP incoming messages and all label-binding operations:

```

[edit]
protocols {
 ldp {
 traceoptions {
 file ldp size 10m files 5 world-readable;
 flag packets receive;
 flag binding;
 }
 interface all {
 }
 }
}

```

Trace LDP protocol traffic for an FEC associated with the LSP:

```

[edit]
protocols {
 ldp {
 traceoptions {
 flag route filter match-on fec policy filter-policy-for-ldp-fec;
 }
 }
}

```

### Standard Firewall Filter Match Conditions for MPLS Traffic

You can configure a standard stateless firewall filter with match conditions for MPLS traffic (**family mpls**).



**NOTE:** The input-list *filter-names* and output-list *filter-names* statements for firewall filters for the mpls protocol family are supported on all interfaces with the exception of management interfaces and internal Ethernet interfaces (fxp or em0), loopback interfaces (lo0), and USB modem interfaces (umd).

Table 270 on page 3566 describes the *match-conditions* you can configure at the [edit firewall family mpls filter *filter-name* term *term-name* from] hierarchy level.

**Table 270: Standard Firewall Filter Match Conditions for MPLS Traffic**

| Match Condition                                | Description                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>apply-groups</b>                            | Specify which groups to inherit configuration data from. You can specify more than one group name. You must list them in order of inheritance priority. The configuration data in the first group takes priority over the data in subsequent groups.                                                                                                                                                                   |
| <b>apply-groups-except</b>                     | Specify which groups not to inherit configuration data from. You can specify more than one group name.                                                                                                                                                                                                                                                                                                                 |
| <b>exp <i>number</i></b>                       | Experimental (EXP) bit number or range of bit numbers in the MPLS header. For <i>number</i> , you can specify one or more values from 0 through 7 in decimal, binary, or hexadecimal format.<br><br><b>NOTE:</b> This match condition is not supported on PTX series packet transport switches.                                                                                                                        |
| <b>exp-except <i>number</i></b>                | Do not match on the EXP bit number or range of bit numbers in the MPLS header. For <i>number</i> , you can specify one or more values from 0 through 7.<br><br><b>NOTE:</b> This match condition is not supported on PTX series packet transport switches.                                                                                                                                                             |
| <b>forwarding-class <i>class</i></b>           | Forwarding class. Specify <b>assured-forwarding</b> , <b>best-effort</b> , <b>expedited-forwarding</b> , or <b>network-control</b> .                                                                                                                                                                                                                                                                                   |
| <b>forwarding-class-except <i>class</i></b>    | Do not match on the forwarding class. Specify <b>assured-forwarding</b> , <b>best-effort</b> , <b>expedited-forwarding</b> , or <b>network-control</b> .                                                                                                                                                                                                                                                               |
| <b>interface <i>interface-name</i></b>         | Interface on which the packet was received. You can configure a match condition that matches packets based on the interface on which they were received.<br><br><b>NOTE:</b> If you configure this match condition with an interface that does not exist, the term does not match any packet.                                                                                                                          |
| <b>interface-set <i>interface-set-name</i></b> | Match the interface on which the packet was received to the specified interface set.<br><br>To define an interface set, include the <b>interface-set</b> statement at the [edit firewall] hierarchy level.<br><br><b>NOTE:</b> This match condition is not supported on PTX series packet transport switches.<br><br>For more information, see “Filtering Packets Received on an Interface Set Overview” on page 4515. |
| <b>ip-version <i>number</i></b>                | (Interfaces on Enhanced Scaling flexible PIC concentrators [FPCs] on supported T Series routers only) Inner IP version. To match MPLS-tagged IPv4 packets, match on the text synonym <b>ipv4</b> .<br><br><b>NOTE:</b> This match condition is not supported on PTX series packet transport switches.                                                                                                                  |

Table 270: Standard Firewall Filter Match Conditions for MPLS Traffic (*continued*)

| Match Condition                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>loss-priority level</b>        | <p>Match the packet loss priority (PLP) level.</p> <p>Specify a single level or multiple levels: <b>low</b>, <b>medium-low</b>, <b>medium-high</b>, or <b>high</b>.</p> <p>Supported on M120 and M320 routers; M7i and M10i routers with the Enhanced CFEB (CFEB-E); and MX Series routers and EX Series switches.</p> <p>For IP traffic on M320, MX Series, T Series routers with Enhanced II Flexible PIC Concentrators (FPCs) and EX Series switches, you must include the <b>tri-color</b> statement at the <b>[edit class-of-service]</b> hierarchy level to commit a PLP configuration with any of the four levels specified. If the <b>tri-color</b> statement is not enabled, you can only configure the <b>high</b> and <b>low</b> levels. This applies to all protocol families.</p> <p>For information about the <b>tri-color</b> statement and for information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>Junos OS Class of Service Configuration Guide</i>.</p> |
| <b>loss-priority-except level</b> | <p>Do not match the PLP level. For details, see the <b>loss-priority</b> match condition.</p> <p><b>NOTE:</b> This match condition is not supported on PTX series packet transport switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

#### Related Documentation

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Standard Firewall Filter Terminating Actions on page 4742](#)
- [Standard Firewall Filter Nonterminating Actions on page 4744](#)

## Configuration Statements

- [\[edit protocols bgp\] Hierarchy Level on page 3567](#)
- [\[edit protocols ldp\] Hierarchy Level on page 3573](#)
- [\[edit protocols mpls\] Hierarchy Level on page 3575](#)

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

Several statements in the **[edit protocols mpls]** hierarchy are valid at numerous locations within it. To make the complete hierarchy easier to read, the repeated statements are listed in “[Common BGP Family Options](#)” on page 369 and that section is referenced at the appropriate locations in “[Complete \[edit protocols bgp\] Hierarchy](#)” on page 369.

- [Common BGP Family Options on page 3567](#)
- [Complete \[edit protocols bgp\] Hierarchy on page 3568](#)

### **Common BGP Family Options**

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit protocols bgp\] Hierarchy](#)” on page 369 instead of the statements being repeated.

- **[edit protocols bgp family inet (any | flow | labeled-unicast | multicast | unicast)]**
- **[edit protocols bgp family inet6 (any | labeled-unicast | multicast | unicast)]**

- [edit protocols bgp family (inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) signaling]
- [edit protocols bgp family inet-vpn (any | flow | multicast | unicast)]
- [edit protocols bgp family inet6-vpn (any | multicast | unicast)]
- [edit protocols bgp family iso-vpn unicast]

The common BGP family options are as follows:

```
accepted-prefix-limit {
 maximum number;
 teardown <percentage> <idle-timeout (forever | minutes)>;
}
damping;
loops number;
prefix-limit {
 maximum number;
 teardown <percentage> <idle-timeout (forever | minutes)>;
}
rib-group group-name;
topology name {
 community {
 target identifier;
 }
}
```

#### **Complete [edit protocols bgp] Hierarchy**

The statement hierarchy listed in this section can also be included at the [edit logical-systems *logical-system-name*] hierarchy level.

```
protocols {
 bgp {
 disable;
 accept-remote-nexthop;
 advertise-external <conditional>;
 advertise-from-main-vpn-tables;
 advertise-inactive;
 (advertise-peer-as | no-advertise-peer-as);
 authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
 authentication-key key;
 authentication-key-chain key-chain;
 bfd-liveness-detection {
 authentication {
 algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
 meticulous-keyed-sha-1 | simple-password);
 key-chain key-chain-name;
 loose-check;
 }
 detection-time {
 threshold milliseconds;
 }
 holddown-interval milliseconds;
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 }
 }
}
```

```

no-adaptation;
session-mode (automatic | multihop | single-hop);
transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
}
version (1 | automatic);
}
cluster cluster-identifier;
damping;
description text-description;
export [policy-names];
family family-name {
 ... the family subhierarchies appear after the main [edit protocols bgp] hierarchy ...
}
graceful-restart {
 disable;
 restart-time seconds;
 stale-routes-time seconds;
}
group group-name {
 ... the group subhierarchy appears after the main [edit protocols bgp] hierarchy ...
}
hold-time seconds;
idle-after-switch-over (seconds | forever);
import [policy-names];
include-mp-next-hop;
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <loops number> <alias> <private>;
local-interface interface-name;
local-preference local-preference;
log-updown;
metric-out (metric | igp (delay-med-update | offset) | minimum-igp offset);
mtu-discovery;
multihop {
 no-nexthop-change;
 ttl ttl-value;
}
no-agggregator-id;
no-client-reflect;
out-delay seconds;
outbound-route-filter {
 bgp-orf-cisco-mode;
 prefix-based {
 accept {
 inet;
 inet6;
 }
 }
}
}
passive;
path-selection {
 always-compare-med;
 as-path-ignore;

```

```
cisco-non-deterministic;
external-router-id;
med-plus-igp {
 igp-multiplier number;
 med-multiplier number;
}
}
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
 file filename <files number> <size maximum-file-size> <world-readable |
 no-world-readable>;
 flag flag <flag-modifier> <disable>;
}
vpn-apply-export;
}

bgp {
 family inet {
 (any | multicast) {
 ... statements in Common BGP Family Options on page 369 ...
 }
 flow {
 ... statements in Common BGP Family Options on page 369 PLUS ...
 no-validate [validation-procedure-names];
 }
 labeled-unicast {
 ... statements in Common BGP Family Options on page 369 PLUS ...
 add-path {
 receive;
 send {
 path-count number;
 prefix-policy [policy-names];
 }
 }
 aggregate-label {
 community community-name;
 }
 aigp [disable];
 explicit-null connected-only;
 per-group-label;
 per-prefix-label;
 resolve-vpn;
 rib (inet.3 | inet6.3);
 traffic-statistics {
 file filename <files number> <size maximum-file-size> <world-readable |
 no-world-readable>;
 interval seconds;
 }
 }
 }
 unicast {
 ... statements in Common BGP Family Options on page 369 PLUS ...
 add-path {
 receive;
```



```

 send {
 path-count number;
 prefix-policy [policy-names];
 }
 }
 topology name {
 community target identifier;
 }
}
}

bgp {
 family inet6 {
 (any | multicast) {
 ... statements in Common BGP Family Options on page 369 ...
 }
 labeled-unicast {
 ... statements in Common BGP Family Options on page 369 PLUS ...
 add-path {
 receive;
 send {
 path-count number;
 prefix-policy [policy-names];
 }
 }
 aggregate-label {
 community community-name;
 }
 aigp [disable];
 explicit-null;
 per-group-label;
 traffic-statistics {
 file filename <files number> <size maximum-file-size> <world-readable |
 no-world-readable>;
 interval seconds;
 }
 }
 }
 unicast {
 ... statements in Common BGP Family Options on page 369 PLUS ...
 topology name {
 community target identifier;
 }
 }
}

bgp {
 family (inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) {
 signaling {
 ... statements in Common BGP Family Options on page 369 ...
 }
 }
}

bgp {

```

```
family inet-vpn {
 (any | multicast | unicast) {
 ... statements in Common BGP Family Options on page 369 PLUS ...
 aggregate-label <community community-name>;
 }
 flow {
 ... statements in Common BGP Family Options on page 369 ...
 }
}

bgp {
 family inet6-vpn {
 (any | multicast | unicast) {
 ... statements in Common BGP Family Options on page 369 PLUS ...
 aggregate-label <community community-name>;
 }
 }
}

bgp {
 family iso-vpn {
 unicast {
 ... statements in Common BGP Family Options on page 369 PLUS ...
 aggregate-label <community community-name>;
 }
 }
}

bgp {
 family route-target {
 accepted-prefix-limit {
 maximum number;
 teardown <percentage> <idle-timeout (forever | minutes)>;
 }
 advertise-default;
 external-paths number;
 prefix-limit {
 maximum number;
 teardown <percentage> <idle-timeout (forever | minutes)>;
 }
 proxy-generate <route-target-policy route-target-policy-name>;
 }
}

bgp {
 group group-name {
 ... same statements as at the [edit protocols bgp] hierarchy level PLUS ...
 allow [all ip-prefix</prefix-length>];
 as-override;
 multipath <multiple-as>;
 neighbor address {
 ... the neighbor subhierarchy appears after the main [edit protocols bgp group
 group-name] hierarchy ...
 }
 type (external | internal);
 }
}
```

```

... BUT NOT ...
 disable; # NOT valid at this level
 group group-name { ... } # NOT valid at this level
 path-selection { ... } # NOT valid at this level
}

group group-name {
 neighbor address {
 ... same statements as at the [edit protocols bgp] hierarchy level PLUS ...
 as-override;
 multipath <multiple-as>;
 ... BUT NOT ...
 disable; # NOT valid at this level
 group group-name { ... } # NOT valid at this level
 neighbor address { ... } # NOT valid at this level
 path-selection { ... } # NOT valid at this level
 }
}
}
}
}

```

#### Related Documentation

- [Notational Conventions Used in Junos OS Configuration Hierarchies](#)
- [\[edit protocols\] Hierarchy Level](#)

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

The following statement hierarchy can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```

protocols {
 ldp {
 (deaggregate | no-deaggregate);
 dod-request-policy [policy-names];
 egress-policy [policy-names];
 explicit-null;
 export [policy-names];
 graceful-restart {
 disable;
 helper-disable;
 maximum-neighbor-reconnect-time seconds;
 maximum-neighbor-recovery-time seconds;
 reconnect-time seconds;
 recovery-time seconds;
 }
 igp-synchronization holddown-interval seconds;
 import [policy-names];
 interface interface-name {
 (allow-subnet-mismatch | no-allow-subnet-mismatch);
 disable;
 hello-interval seconds;
 hold-time seconds;
 transport-address (interface | router-id);
 }
 keepalive-interval seconds;
 }
}

```

```
keepalive-timeout seconds;
l2-smart-policy;
log-updown {
 trap disable;
}
next-hop {
 merged {
 policy [policy-names];
 }
}
no-forwarding;
oam {
 ... the oam subhierarchy appears after the main [edit protocols ldp] hierarchy ...
}
policing {
 fec class-address {
 ingress-traffic filter-name;
 transit-traffic filter-name;
 }
}
preference preference;
session destination-address {
 authentication-algorithm algorithm;
 authentication-key key;
 authentication-key-chain key-chain;
 downstream-on-demand;
}
session-protection <timeout seconds>;
strict-targeted-hellos;
targeted-hello {
 hello-interval seconds;
 hold-time seconds;
}
traceoptions {
 file filename <files number> <size maximum-file-size> <world-readable |
 no-world-readable>;
 flag flag <flag-modifier> <disable>;
}
track-igp-metric;
traffic-statistics {
 file filename <files number> <size maximum-file-size> <world-readable |
 no-world-readable>;
 interval seconds;
 no-penultimate-hop;
}
transport-address (address | interface | router-id);
}

ldp {
 oam {
 bfd-liveness-detection {
 detection-time {
 threshold milliseconds;
 }
 }
 ecmp;
 failure-action (remove-nexthop | remove-route);
 }
}
```

```

 holddown-interval milliseconds;
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 version (1 | automatic);
 }
}
fec class-address {
 bfd-liveness-detection {
 ... same statements as at the [edit protocols ldp oam bfd-liveness-detection]
 hierarchy level ...
 }
 no-bfd-liveness-detection;
 periodic-traceroute {
 ... same statements as at the [edit protocols ldp oam periodic-traceroute]
 hierarchy level PLUS ...
 }
 disable;
}
}
ingress-policy [policy-names];
periodic-traceroute {
 exp cos-value;
 fanout next-hops;
 frequency minutes;
 paths number;
 retries number;
 source address;
 ttl number;
 wait seconds;
}
}
}
}

```

#### Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

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

- [Complete \[edit protocols mpls\] Hierarchy on page 3575](#)

#### **Complete [edit protocols mpls] Hierarchy**

The statement hierarchy listed in this section can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

```

protocols {
 mpls {
 disable;
 interface (interface-name | all) {
 always-mark-connection-protection-tlv

```

```
 disable;
 admin-group [group-names];
 srlg srlg-name
 static {
 protection-revert-time seconds
 }
 switch-away-lsp;
 }
 ipv6-tunneling;
 priority setup-priority hold-priority;
 traceoptions {
 file filename <files number> <size maximum-file-size> <world-readable |
 no-world-readable>;
 flag flag;
 }
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
  - *[edit protocols] Hierarchy Level*

---

## allow-subnet-mismatch

|                                 |                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | allow-subnet-mismatch;                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i> ],<br>[edit protocols ldp interface <i>interface-name</i> ]                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3.                                                                                                                                                                                                              |
| <b>Description</b>              | Ignore the LDP subnet check. For Junos OS Release 8.4 and later releases, an LDP source address subnet check was added for the neighbor establishment procedure. The source address in the LDP link hello packet is matched against the interface address. |
| <b>Default</b>                  | The source address in the LDP link hello packet is matched against the interface address.                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Miscellaneous LDP Properties</i></li></ul>                                                                                                                                                          |

## authentication-algorithm

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authentication-algorithm <i>algorithm</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ldp session <i>session-address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp session <i>session-address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit protocols ldp session <i>session-address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ldp session <i>session-address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced for BGP in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure an authentication algorithm type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b><i>algorithm</i></b>—Specify one of the following types of authentication algorithms:</p> <ul style="list-style-type: none"> <li>• <b>aes-128-cmac-96</b>—Cipher-based message authentication code (AES128, 96 bits).</li> <li>• <b>hmac-sha-1-96</b>—Hash-based message authentication code (SHA1, 96 bits).</li> <li>• <b>md5</b>—Message digest 5.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding Route Authentication</i></li> <li>• <i>Example: Configuring Route Authentication for BGP</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## authentication-key (Protocols LDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | authentication-key <i>md5-authentication-key</i> ;                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp session <i>address</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp session <i>address</i> ],<br>[edit protocols ldp session <i>address</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp session <i>address</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure the MD5 authentication signature. The maximum length of the authentication signature is 69 characters.                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Miscellaneous LDP Properties</i></li></ul>                                                                                                                                                                                                                                                                                   |



## bfd-liveness-detection (Protocols LDP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> bfd-liveness-detection {   detection-time threshold <i>milliseconds</i>;   ecmp;   failure-action {     remove-nexthop;     remove-route;   }   holddown-interval <i>seconds</i>;   minimum-interval <i>milliseconds</i>;   minimum-receive-interval <i>milliseconds</i>;   minimum-transmit-interval <i>milliseconds</i>;   multiplier <i>detection-time-multiplier</i>;   no-adaptation;   transmit-interval {     minimum-interval <i>milliseconds</i>;     threshold <i>milliseconds</i>;   } }</pre>                                         |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols ldp oam],<br/> [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec address],<br/> [edit protocols ldp oam],<br/> [edit protocols ldp oam fec address]</p>                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Support for the <b>bfd-liveness-detection</b> statement at the [edit protocols ldp oam fec address] hierarchy level and the <b>ecmp</b> option added in Junos OS Release 9.0.</p> <p>Support for the <b>failure-action</b> statement with the <b>remove-nexthop</b> and <b>remove-route</b> options and the <b>holddown-interval</b> statement added in Junos OS Release 9.4.</p>                                                                                                               |
| <b>Description</b>         | <p>Enable Bidirectional Forwarding Detection (BFD) for all MPLS LSPs or for just a specific LSP.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>             | <p><b>minimum-interval</b>—Minimum transmit and receive interval.<br/> <b>Range:</b> 50 through 255,000 milliseconds<br/> <b>Default:</b> 50</p> <p><b>minimum-receive-interval</b>—Minimum receive interval.<br/> <b>Range:</b> 50 through 255,000 milliseconds<br/> <b>Default:</b> 50</p> <p><b>minimum-transmit-interval</b>—Minimum transmit interval.<br/> <b>Range:</b> 50 through 255,000 milliseconds<br/> <b>Default:</b> 50</p> <p><b>multiplier</b>—Detection time multiplier.<br/> <b>Range:</b> 50 through 255<br/> <b>Default:</b> 3</p> |

The other options are explained separately.

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

**Related Documentation**

- [Configuring BFD for LDP LSPs on page 3555](#)

---

## deaggregate

---

**Syntax** deaggregate | no-deaggregate;

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols ldp],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],  
[edit protocols ldp],  
[edit routing-instances *routing-instance-name* protocols ldp]

**Release Information** Statement introduced before Junos OS Release 7.4.

**Description** Control forwarding equivalence class (FEC) deaggregation on the router. The use of the **deaggregate** statement in LDP is a standard practice that we recommend for LDP deployments.

**Default** Deaggregation is disabled on the router or switch.

**Options** **deaggregate**—Deaggregate FECs.  
**no-deaggregate**—Aggregate FECs.

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

**Related Documentation**

- [Configuring FEC Deaggregation on page 3553](#)

## disable (Protocols LDP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart],<br>[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options graceful-restart],<br>[edit protocols ldp graceful-restart],<br>[edit protocols ldp interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options graceful-restart] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Explicitly disable LDP on an interface, or explicitly disable LDP graceful restart.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Default</b>                  | LDP is enabled on interfaces configured with the LDP <b>interface</b> statement. LDP graceful restart is automatically enabled when graceful restart is enabled under the <b>[edit routing-options]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling and Disabling LDP on page 3541</a></li> <li>• <a href="#">Configuring LDP Graceful Restart on page 3545</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## dod-request-policy

---

|                                 |                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>dod-request-policy <i>dod-request-policy-name</i>;</code>                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit protocols ldp]                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2.                                                                                                                             |
| <b>Description</b>              | Specify the name of the LDP downstream on demand request policy. LDP sends label request messages only for those FECs matching in the downstream on demand request policy. |
| <b>Options</b>                  | <i>dod-request-policy-name</i> —Specify the name of the downstream on demand request policy.                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring LDP Downstream on Demand</i></li></ul>                                                                     |

## downstream-on-demand

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>downstream-on-demand;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical systems <i>logical-system-name</i> protocols ldp session <i>session-address</i> ],<br>[edit protocols ldp session <i>session-address</i> ]                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2.                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Enable LDP downstream on demand on the LDP session. LDP is widely deployed in downstream unsolicited advertisement mode. As service providers integrate the access and aggregation networks into a single MPLS domain, LDP downstream on demand is needed to distribute the bindings between access and aggregation networks to minimize the workload for the access node (AN) control plane and to avoid the storage of tens of thousands of label bindings from upstream aggregation nodes. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring LDP Downstream on Demand</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                        |

## ecmp

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ecmp;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp oam bfd-liveness-detection],<br>[edit logical-systems <i>logical-system-name</i> protocols ldp oam fec address<br>bfd-liveness-detection],<br>[edit protocols ldp oam bfd-liveness-detection],<br>[edit protocols ldp oam fec address bfd-liveness-detection]                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Allows LDP to establish BFD sessions for all ECMP paths configured for the specified FEC. If you configure the <b>ecmp</b> statement, you must also configure the <b>periodic-traceroute</b> statement for the specified FEC. If you do not do so, the commit operation fails. You can configure the <b>periodic-traceroute</b> statement at the global hierarchy level ([edit protocols ldp oam]) while only configuring the <b>ecmp</b> statement for a specific FEC ([edit protocols ldp oam fec address bfd-liveness-detection]). |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring ECMP-Aware BFD for LDP LSPs on page 3558</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                              |

## egress-policy

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>egress-policy [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                   |
| <b>Description</b>              | Control the prefixes advertised into LDP.                                                                                                                                                                                                                                           |
| <b>Default</b>                  | Only the loopback address is advertised.                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more routing policies.                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Prefixes Advertised into LDP from the Routing Table on page 3552</a></li> </ul>                                                                                                                                |

## explicit-null (Protocols LDP)

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | explicit-null;                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                   |
| <b>Description</b>              | Advertise label 0 to the egress router of a label-switched path (LSP).                                                                                                                                                                                                              |
| <b>Default</b>                  | If you do not include the <b>explicit-null</b> statement in the MPLS configuration, label 3 (implicit null) is advertised.                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Miscellaneous LDP Properties</a></li></ul>                                                                                                                                                                          |

## export (Protocols LDP)

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | export [ <i>policy-names</i> ];                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                   |
| <b>Description</b>              | Apply policy filters to outbound LDP label bindings. Filters are applied to all label bindings from all neighbors.                                                                                                                                                                  |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more routing policies.                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Filtering Outbound LDP Label Bindings on page 3550</a></li></ul>                                                                                                                                                                |

## failure-action (Protocols LDP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>failure-action {     remove-nexthop;     remove-route; }</pre>                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols ldp oam bfd-livenessss-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>address</i> bfd-livenessss-detection],</p> <p>[edit protocols ldp oam bfd-livenessss-detection],</p> <p>[edit protocols ldp oam fec <i>address</i> bfd-livenessss-detection]</p>                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure route and next-hop properties in the event of a BFD session failure event on an LDP LSP. The failure event could be an existing BFD session that has gone down or could be a BFD session that never came up. LDP adds back the route or next hop when the relevant BFD session comes back up.                                                                                                                            |
| <b>Options</b>                  | <p><b>remove-nexthop</b>—Remove a route corresponding to a next hop of the LSP's route at the ingress node when a BFD session failure event is detected.</p> <p><b>remove-route</b>—Remove the route corresponding to an LSP from the appropriate routing tables when a BFD session failure event is detected. If the LSP is configured with ECMP and a BFD session corresponding to any path goes down, the route is removed.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring a Failure Action for the BFD Session on an LDP LSP on page 3558</a></li> </ul>                                                                                                                                                                                                                                                                                    |

## fec

---

**Syntax**    `fec fec-address {  
              bfd-liveness-detection {  
                  detection-time threshold milliseconds;  
                  ecmp;  
                  failure-action {  
                      remove-nexthop;  
                      remove-route;  
                  }  
                  holddown-interval milliseconds;  
                  ingress-policy ingress-policy-name;  
                  minimum-interval milliseconds;  
                  minimum-receive-interval milliseconds;  
                  minimum-transmit-interval milliseconds;  
                  multiplier detection-time-multiplier;  
                  no-adaptation;  
                  transmit-interval {  
                      minimum-interval milliseconds;  
                      threshold milliseconds;  
                  }  
                  version (0 | 1 | automatic);  
              }  
              no-bfd-liveness-detection;  
              periodic-traceroute {  
                  disable;  
                  exp exp-value;  
                  fanout fanout-value;  
                  frequency minutes;  
                  paths number-of-paths;  
                  retries retry-attempts;  
                  source address;  
                  ttl ttl-value;  
                  wait seconds;  
              }  
          }`

**Hierarchy Level**    [edit logical-systems *logical-systems-name* protocols ldp oam],  
                          [edit protocols ldp oam]

**Release Information**    Statement introduced in Junos OS Release 8.5.  
                              Statement introduced in Junos OS Release 12.2 for EX Series switches.

**Description**    Allows you to configure BFD for a specific LDP forwarding equivalence class (FEC).


**Options**    *fec-address*—Specify the FEC address.  
  
              The other statements are explained separately.

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



**Related Documentation** • [Configuring BFD for LDP LSPs on page 3555](#)

## graceful-restart (Protocols LDP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>graceful-restart {   disable;   helper-disable;   maximum-neighbor-recovery-time <i>value</i>;   reconnect-time <i>seconds</i>;   recovery-time <i>value</i>; }</pre>                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols ldp],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br/> [edit protocols ldp],<br/> [edit routing-instances <i>routing-instance-name</i> protocols ldp]</p>                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Enable LDP graceful restart on the LDP master protocol instance or for a specific routing instance.                                                                                                                                                                                                                                                                                                                                                                                  |
|                                 | <div>  <p><b>NOTE:</b> When you alter the graceful restart configuration at either the [edit routing-options graceful-restart] or [edit protocols ldp graceful-restart] hierarchy levels, any running LDP session is automatically restarted to apply the graceful restart configuration. This behavior mirrors the behavior of BGP when you alter its graceful restart configuration.</p> </div> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.<br/> routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | • <a href="#">Configuring LDP Graceful Restart on page 3545</a>                                                                                                                                                                                                                                                                                                                                                                                                                      |

## hello-interval (Protocols LDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hello-interval <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols ldp targeted-hello],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  ldp interface <i>interface-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  ldp targeted-hello],</code><br><code>[edit protocols ldp interface <i>interface-name</i>],</code><br><code>[edit protocols ldp targeted-hello],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols ldp targeted-hello]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Support for LDP targeted hellos added in Junos OS Release 9.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Control the LDP timer that regulates how often hello messages are sent. You can control the rate both link hello messages and targeted hello messages are sent depending on the hierarchy level at which you configure the <b>hello-interval</b> statement.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <b><i>seconds</i></b> —Length of time between transmission of hello packets.<br><b>Range:</b> 1 through 65,535 seconds<br><b>Default:</b> 5 seconds for link hello messages, 15 seconds for targeted hello messages                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the LDP Timer for Hello Messages on page 3542</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## helper-disable (LDP)

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | helper-disable;                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart],<br>[edit protocols ldp graceful-restart],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Disable helper mode for LDP graceful restart. When helper mode is disabled, a router cannot help a neighboring router that is attempting to restart LDP.                                                                                                                                                                                                |
| <b>Default</b>                  | Helper mode is enabled by default on all routing protocols (including LDP) that support graceful restart.                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring LDP Graceful Restart on page 3545</a></li> </ul>                                                                                                                                                                                                                                       |

## holddown-interval

|                                 |                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | holddown-interval <i>holddown-interval</i> ;                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp oam bfd-liveness-detection],<br>[edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>address</i> bfd-liveness-detection],<br>[edit protocols ldp oam bfd-liveness-detection],<br>[edit protocols ldp oam fec <i>address</i> bfd-liveness-detection] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Specify how long the BFD session should be up before adding the route or next hop. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.                                                                                                                                    |
| <b>Options</b>                  | <p><b><i>holddown-interval</i></b>—Number of seconds the BFD session should remain up before adding the route or next hop.</p> <p><b>Default:</b> 0 seconds</p> <p><b>Range:</b> 0 through 65,535 seconds</p>                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Holddown Interval for the BFD Session on page 3559</a></li> </ul>                                                                                                                                                                                                  |

## hold-time (Protocols LDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hold-time seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols ldp targeted-hello],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  ldp interface <i>interface-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  ldp targeted-hello],</code><br><code>[edit protocols ldp interface <i>interface-name</i>],</code><br><code>[edit protocols ldp targeted-hello],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols ldp targeted-hello]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Support for LDP targeted hellos added in Junos OS Release 9.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Specify how long an LDP node should wait for a hello message before declaring a neighbor to be down. This value is sent as part of a hello message so that each LDP node tells its neighbors how long to wait. You can specify times for both link hello messages and targeted hello messages depending on the hierarchy level at which you configure the <b>hold-time</b> statement.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <b>seconds</b> —Hold-time value.<br><b>Range:</b> 1 through 65,535 seconds<br><b>Default:</b> 15 seconds for link hello messages, 45 seconds for targeted hello messages                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Delay Before LDP Neighbors Are Considered Down on page 3543</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## igp-synchronization

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>igp-synchronization holddown-interval <i>seconds</i>;</code>                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.5.                                                                                                                                                                                                                                       |
| <b>Description</b>              | Configure the time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational. For large networks with numerous FECs, you might need to configure a longer value to allow enough time for the LDP label databases to be exchanged.   |
| <b>Options</b>                  | <b>holddown-interval <i>seconds</i></b> —Time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational.<br><b>Default:</b> 10 seconds<br><b>Range:</b> 10 through 60 seconds                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Miscellaneous LDP Properties</a></li> </ul>                                                                                                                                                                        |

## import (Protocols LDP)

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                   |
| <b>Description</b>              | Apply policy filters to received LDP label bindings. Filters are applied to all label bindings from all neighbors.                                                                                                                                                                  |
| <b>Options</b>                  | <b><i>policy-names</i></b> —Name of one or more routing policies.                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Filtering Inbound LDP Label Bindings on page 3548</a></li> </ul>                                                                                                                                                               |

## ingress-policy

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ingress-policy [ <i>ingress-policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <code>[edit logical-system <i>logical-system-name</i> protocols ldp oam],</code><br><code>[edit protocols ldp oam]</code>                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Configure an Operation, Administration, and Management (OAM) policy to choose which forwarding equivalence classes (FECs) need to have OAM enabled. If the FEC passes through the policy or if the FEC is explicitly configured, OAM is enabled for a FEC. For FECs chosen using a policy, the BFD parameters configured under <code>[edit protocols ldp oam bfd-liveness-detection]</code> are applied. |
| <b>Options</b>                  | <i>ingress-policy-names</i> —Specify the names of the ingress policies.                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring OAM Ingress Policies for LDP on page 3559</a></li></ul>                                                                                                                                                                                                                                                                                  |

## interface (Protocols LDP)

|                                 |                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>interface <i>interface-name</i> {     disable;     hello-interval <i>seconds</i>;     hold-time <i>seconds</i>;     transport-address (interface   loopback); }</pre>                                                                                                              |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]</pre> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                       |
| <b>Description</b>              | Enable LDP on one or more router or switch interfaces.                                                                                                                                                                                                                                  |
| <b>Default</b>                  | LDP is disabled on all interfaces.                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><i>interface-name</i>—Name of an interface. To configure all interfaces, specify <b>all</b>.</p> <p>The remaining statements are explained separately.</p>                                                                                                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling and Disabling LDP on page 3541</a></li> </ul>                                                                                                                                                                             |

## keepalive-interval

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>keepalive-interval seconds;</code>                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                   |
| <b>Description</b>              | Set the keepalive interval value.                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>seconds</b> —Keepalive value.<br><b>Range:</b> 1 through 65,535<br><b>Default:</b> 10 seconds                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Interval for LDP Keepalive Messages on page 3544</a></li></ul>                                                                                                                                                  |

## keepalive-timeout

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>keepalive-timeout seconds;</code>                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                   |
| <b>Description</b>              | Set the keepalive timeout value. The keepalive timeout defines the amount of time that the neighbor LDP node waits before determining that the session has failed.                                                                                                                  |
| <b>Options</b>                  | <b>seconds</b> —Keepalive timeout value.<br><b>Range:</b> 1 through 65,535<br><b>Default:</b> 30 seconds                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the LDP Keepalive Timeout on page 3545</a></li></ul>                                                                                                                                                                |



## l2-smart-policy

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | l2-smart-policy;                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.                                                                                                                                                                                                                                       |
| <b>Description</b>              | Prevent LDP from exporting IPv4 FECs over sessions with Layer 2 neighbors only. IPv4 FECs received over such sessions are filtered out.                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring LDP IPv4 FEC Filtering on page 3554</a></li> </ul>                                                                                                                                                                 |

## label-withdrawal-delay

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | label-withdrawal-delay <i>seconds</i> ;                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1.                                                                                                                                                                                                                                       |
| <b>Description</b>              | Delay the withdrawal of labels to reduce router workload during IGP convergence.                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>seconds</b>—Configure the number of seconds to wait before withdrawing labels for the LDP LSPs.</p> <p><b>Default:</b> 60 seconds</p> <p><b>Range:</b> 0 through 300 seconds</p>                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Miscellaneous LDP Properties</a></li> </ul>                                                                                                                                                                        |

## ldp

---

```
Syntax ldp {
 (deaggregate | no-deaggregate);
 egress-policy [policy-names];
 explicit-null;
 export [policy-names];
 graceful-restart {
 disable;
 helper-disable;
 maximum-neighbor-recovery-time seconds;
 reconnect-time seconds;
 recovery-time seconds;
 }
 import [policy-names];
 interface (interface-name | all) {
 disable;
 hello-interval seconds;
 hold-time seconds;
 transport-address (interface | router-id);
 }
 keepalive-interval seconds;
 keepalive-timeout seconds;
 log-updown {
 trap disable;
 }
 no-forwarding;
 oam {
 bfd-liveness-detection {
 detection-time threshold milliseconds;
 ecmp;
 failure-action {
 remove-nexthop;
 remove-route;
 }
 holddown-interval milliseconds;
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 minimum-transmit-interval milliseconds;
 multiplier detection-time-multiplier;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 }
 }
 fec fec-address {
 bfd-liveness-detection {
 detection-time threshold milliseconds;
 ecmp;
 failure-action {
 remove-nexthop;
 remove-route;
 }
 }
 }
}
```

```

 holddown-interval milliseconds;
 ingress-policy ingress-policy-name;
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 minimum-transmit-interval milliseconds;
 multiplier detection-time-multiplier;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 version (0 | 1 | automatic);
}
no-bfd-liveness-detection;
periodic-traceroute {
 disable;
 exp exp-value;
 fanout fanout-value;
 frequency minutes;
 paths number-of-paths;
 retries retry-attempts;
 source address;
 ttl ttl-value;
 wait seconds;
}
}
ingress-policy ingress-policy-name;
periodic-traceroute {
 disable;
 exp exp-value;
 fanout fanout-value;
 frequency minutes;
 paths number-of-paths;
 retries retry-attempts;
 source address;
 ttl ttl-value;
 wait seconds;
}
}
p2mp;
policing {
 fec fec-address {
 ingress-traffic filter-name;
 transit-traffic filter-name;
 }
}
}
preference preference;
session address {
 authentication-algorithm algorithm;
 authentication-key authentication-key;
 authentication-key-chain key-chain-name;
}
strict-targeted-hellos;
traceoptions {
 file filename <files number <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;

```

```
}
track-igp-metric;
traffic-statistics {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 interval interval;
 no-penultimate-hop;
}
transport-address (address | interface | router-id);
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*  
protocols],  
[edit protocols],  
[edit routing-instances *routing-instance-name* protocols]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 11.1 for EX Series switches.

**Description** Enable LDP routing on the router or switch.

You must include the **ldp** statement in the configuration to enable LDP on the router or switch.

**Default** LDP is disabled on the router.

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

**Related Documentation**

- [Minimum LDP Configuration on page 3541](#)
- [Enabling and Disabling LDP on page 3541](#)

## ldp-p2mp

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ldp-p2mp;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> provider-tunnel],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> provider-tunnel selective wildcard-group-inet wildcard-source],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> provider-tunnel selective wildcard-group-inet6 wildcard-source],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> provider-tunnel selective group <i>group-prefix</i> wildcard-source],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> provider-tunnel selective group <i>group-prefix</i> source <i>source-prefix</i>],</p> <p>[edit routing-instances <i>instance-name</i> provider-tunnel],</p> <p>[edit routing-instances <i>instance-name</i> provider-tunnel selective wildcard-group-inet wildcard-source],</p> <p>[edit routing-instances <i>instance-name</i> provider-tunnel selective wildcard-group-inet6 wildcard-source],</p> <p>[edit routing-instances <i>instance-name</i> provider-tunnel selective group <i>group-prefix</i> wildcard-source],</p> <p>[edit routing-instances <i>instance-name</i> provider-tunnel selective group <i>group-prefix</i> source <i>source-prefix</i>]</p> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Specify a point-to-multipoint provider tunnel with LDP signalling for an MBGP MVPN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Point-to-Multipoint LDP LSPs as the Data Plane for Intra-AS MBGP MVPNs</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## log-updown (Protocols LDP)

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | log-updown {<br>trap disable;<br>}                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                   |
| <b>Description</b>              | Disable LDP traps on the router, logical system, or routing instance.                                                                                                                                                                                                               |
| <b>Options</b>                  | <b>trap disable</b> —Disable LDP traps.<br><b>Default:</b> LDP traps are enabled on the router or switch.                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Miscellaneous LDP Properties</i></li></ul>                                                                                                                                                                                   |

## make-before-break (LDP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>make-before-break {     timeout <i>seconds</i>;     switchover-delay <i>seconds</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit protocols ldp]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Configures make before break (MBB) for multicast LDP (MLDP) link protection to ensure minimum packet loss when attempting to signal a new label-switched path (LSP) before tearing down the old LSP path.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>timeout <i>seconds</i></b>—Specify a value to change a make -before-break timeout for point-to-multipoint LSPs. Even if an MBB acknowledgment is not received for a point-to-multipoint LSP before the specified timeout period expires, the label-switching router (LSR) performs an MBB switchover from the old LSR to the new upstream LSR.</p> <p><b>Range:</b> 1 through 300 seconds</p> <p><b>Default:</b> 30 seconds</p> <p><b>switchover-delay <i>seconds</i></b>—Specify a value to change switchover delay for a point-to-multipoint LSP from the old LSR to the new upstream LSR. If an MBB acknowledgment is received on a point of local repair (PLR) router, the PLR waits for the specified seconds to switch its upstream LSR from the old LSR to the new LSR.</p> <p><b>Range:</b> 1 through 300 seconds</p> <p><b>Default:</b> 30 seconds</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Multicast LDP Link Protection</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## maximum-neighbor-recovery-time

---

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>maximum-neighbor-recovery-time seconds;</code>                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart],<br>[edit protocols ldp graceful-restart],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4. Statement changed from <b>maximum-recovery-time</b> to <b>maximum-neighbor-recovery-time</b> in Junos OS Release 9.1.                                                                                                                                                                                 |
| <b>Description</b>              | Specify the maximum amount of time to wait before giving up an attempt to gracefully restart.                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <b>seconds</b> —Configure the maximum recovery time, in seconds.<br><b>Range:</b> 120 through 1800 seconds<br><b>Default:</b> 140 seconds                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Recovery Time and Maximum Recovery Time on page 3547</a></li><li>• <a href="#">Configuring Graceful Restart Options for LDP on page 2135</a></li><li>• <a href="#">no-strict-lsa-checking on page 2172</a></li><li>• <a href="#">recovery-time on page 2175</a></li></ul>               |



## no-forwarding

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-forwarding;                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                   |
| <b>Description</b>              | Do not add ingress routes to the inet.0 routing table even if <b>traffic-engineering bgp-igp</b> (configured at the [edit protocols mpls] hierarchy level) is enabled.                                                                                                              |
| <b>Default</b>                  | The <b>no-forwarding</b> statement is disabled. Ingress routes are added to the inet.0 routing table instead of the inet.3 routing table when <b>traffic-engineering bgp-igp</b> is enabled.                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Miscellaneous LDP Properties</i></li> <li>• <i>Configuring Virtual-Router Routing Instances in VPNs</i></li> </ul>                                                                                                          |

## oam (Protocols LDP)

---

**Syntax**    oam {  
              bfd-liveness-detection {  
                  detection-time threshold *milliseconds*;  
                  ecmp;  
                  failure-action {  
                      remove-nexthop;  
                      remove-route;  
                  }  
                  holddown-interval *milliseconds*;  
                  ingress-policy *ingress-policy-name*;  
                  minimum-interval *milliseconds*;  
                  minimum-receive-interval *milliseconds*;  
                  minimum-transmit-interval *milliseconds*;  
                  multiplier *detection-time-multiplier*;  
                  no-adaptation;  
                  transmit-interval {  
                      minimum-interval *milliseconds*;  
                      threshold *milliseconds*;  
                  }  
                  version (0 | 1 | automatic);  
              }  
              fec *fec-address*;  
              ingress-policy *ingress-policy-name*;  
              lsp-ping-interval *seconds*;  
              periodic-traceroute {  
                  disable;  
                  exp *exp-value*;  
                  fanout *fanout-value*;  
                  frequency *minutes*;  
                  paths *number-of-paths*;  
                  retries *retry-attempts*;  
                  source *address*;  
                  ttl *ttl-value*;  
                  wait *seconds*;  
              }  
          }

**Hierarchy Level**    [edit logical-systems *logical-system-name* protocols ldp]  
                          [edit protocols ldp]

**Release Information**    Statement introduced in Junos OS Release 7.6.  
                          lsp-ping-interval option introduced in Junos OS Release 9.4.

**Description**    Configure Operation, Administration, and Maintenance (OAM) and Bidirectional Forwarding Detection (BFD) protocol for LDP.

**Options**    **fec *fec-address***—Specify the forwarding equivalence class (FEC) address. You must either specify a FEC address or configure an OAM ingress policy to ensure that the BFD session comes up.

**lsp-ping-interval *seconds***—Specify the duration of the LSP ping interval in seconds. To issue a ping on an LDP-signaled LSP, use the **ping mpls ldp** command.

**Default:** 60 seconds

**Range:** 30 through 3,600 seconds

The remaining statements are explained separately.

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

**Related Documentation**

- [Configuring BFD for LDP LSPs on page 3555](#)

## p2mp (Protocols LDP)

**Syntax** p2mp;

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols ldp],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],  
 [edit protocols ldp],  
 [edit routing-instances *routing-instance-name* protocols ldp]

**Release Information** Statement introduced in Junos OS Release 11.2.

**Description** Enable point-to-multipoint MPLS LSPs in an LDP-signaled LSP.

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

**Related Documentation**

- [Example: Configuring Point-to-Multipoint LDPLSPs as the Data Plane for Intra-AS MBGP MVPNs](#)
- [Point-to-Multipoint LSPs Overview on page 3407](#)

## periodic-traceroute

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>periodic-traceroute {<br/>  disable;<br/>  exp <i>exp-value</i>;<br/>  fanout <i>fanout-value</i>;<br/>  frequency <i>minutes</i>;<br/>  paths <i>number-of-paths</i>;<br/>  retries <i>retry-attempts</i>;<br/>  source <i>address</i>;<br/>  ttl <i>ttl-value</i>;<br/>  wait <i>seconds</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols ldp oam],<br>[edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>fec-address</i> ],<br>[edit protocols ldp oam],<br>[edit protocols ldp oam fec <i>fec-address</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.4.<br>Support added at the [edit protocols ldp oam] and [edit logical-systems <i>logical-system-name</i> protocols ldp oam] hierarchy levels in Junos OS Release 9.0.<br>Statement introduced in Junos OS Release 12.2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>         | Enable tracing of forwarding equivalence classes (FECs) for LDP LSPs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable tracing for a specific FEC. This option is available at the [edit protocols ldp oam fec <i>fec-address</i> periodic-traceroute] and [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>fec-address</i> periodic-traceroute] hierarchy levels only.</p> <p><b>exp <i>exp-value</i></b>—(Optional) Specify the class of service to use when sending probes.<br/><b>Default:</b> 7<br/><b>Range:</b> 0 through 7</p> <p><b>fanout <i>fanout-value</i></b>—(Optional) Specify the maximum number of next hops to search per node.<br/><b>Default:</b> 16<br/><b>Range:</b> 1 through 16</p> <p><b>frequency <i>minutes</i></b>—(Optional) Specify the interval between traceroute attempts.<br/><b>Default:</b> 60 minutes<br/><b>Range:</b> 15 through 120 minutes</p> <p><b>paths <i>number-of-paths</i></b>—(Optional) Specify the maximum number of paths to search.<br/><b>Default:</b> 3<br/><b>Range:</b> 1 through 255</p> |

**retries** *retry-attempts*—(Optional) Specify the number of attempts to send a probe to a specific node before giving up.

**Default:** 3

**Range:** 1 through 9

**source address**—(Optional) Specify the IPv4 source address to use when sending probes.

**ttl value**—(Optional) Specify the maximum time-to-live value. Nodes that are beyond this value are not traced.

**Default:** 64

**Range:** 1 through 255

**wait seconds**—(Optional) Specify the wait interval before resending a probe packet.

**Default:** 10 seconds

**Range:** 5 though 15 seconds

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

|                              |                                                                                                               |
|------------------------------|---------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Configuring LDP LSP Traceroute on page 3560</a></li></ul> |
|------------------------------|---------------------------------------------------------------------------------------------------------------|

## policing (Protocols LDP)

---

|                                 |                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>policing {<br/>    fec <i>fec-address</i> {<br/>        ingress-traffic <i>filter-name</i>;<br/>        transit-traffic <i>filter-name</i>;<br/>    }<br/>}</pre>                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp]                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Enable policing of forwarding equivalence classes (FECs) for LDP.                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>fec <i>fec-address</i></b>—Specify the address for the FEC.</p> <p><b>ingress-traffic <i>filter-name</i></b>—Specify the name of the filter for policing ingress FEC traffic.</p> <p><b>transit-traffic <i>filter-name</i></b>—Specify the name of the filter for policing transit FEC traffic.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Policers for LDP FECs on page 3554</a></li></ul>                                                                                                                                                                                          |

## preference (Protocols LDP)

|                                 |                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>preference <i>preference</i>;</code>                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Set the route preference level for LDP routes.                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <i>preference</i> —Preferred value.<br><b>Range:</b> 0 through 255<br><b>Default:</b> 9                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring LDP Route Preferences on page 3545</a></li> </ul>                                                                                                                                                                                                                                    |

## reconnect-time

|                                 |                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>reconnect-time <i>seconds</i>;</code>                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp <a href="#">graceful-restart</a> ],<br>[edit protocols ldp <a href="#">graceful-restart</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp <a href="#">graceful-restart</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1.                                                                                                                                                                                                                          |
| <b>Description</b>              | Specify the length of time required to reestablish a Label Distribution Protocol (LDP) session after graceful restart.                                                                                                                                                 |
| <b>Options</b>                  | <i>seconds</i> —Time required for reconnection.<br><b>Range:</b> 30 through 300<br><b>Default:</b> 60 seconds                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring LDP Graceful Restart on page 3545 on LDP Configuration Guide</a></li> <li>• <a href="#">Configuring Graceful Restart Options for LDP on page 2135</a></li> </ul>                                      |

## recovery-time

---

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>recovery-time seconds;</code>                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart],<br>[edit protocols ldp graceful-restart],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Specify the amount of time a router waits for LDP to restart gracefully.                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <b>seconds</b> —Configure the recovery time, in seconds.<br><b>Range:</b> 120 through 1800 seconds<br><b>Default:</b> 140 seconds                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Recovery Time and Maximum Recovery Time on page 3547</a></li></ul>                                                                                                                                                                                                                      |

## session (ldp)

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>session address {<br/>    authentication-algorithm <i>algorithm</i>;<br/>    authentication-key <i>authentication-key</i>;<br/>    authentication-key-chain <i>key-chain-name</i>;<br/>}</code>                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br><b>authentication-algorithm</b> statement introduced in Junos OS Release 7.6.                                                                                                                                                  |
| <b>Description</b>              | Specify the address for the remote end of the LDP session.<br><br>The remaining statements are explained separately.                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Miscellaneous LDP Properties</a></li></ul>                                                                                                                                                                          |



## session-protection

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | session-protection {<br>timeout <i>seconds</i> ;<br>}                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp]                                                                                                                        |
| <b>Description</b>              | Configure when an LDP session is torn down and resigaled after the router or switch stops receiving hello messages from a neighboring router or switch. You might want to modify this behavior to prevent an LDP session from being unnecessarily terminated and reestablished. The LDP session remains up for the duration specified as long as the routers or switches maintain IP network connectivity. |
| <b>Options</b>                  | <b>timeout <i>seconds</i></b> —Time in seconds before the LDP session is torn down and resigaled.<br><b>Range:</b> 1 through 65,535 seconds                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Miscellaneous LDP Properties</i></li> </ul>                                                                                                                                                                                                                                                                                                        |

## strict-targeted-hellos

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | strict-targeted-hellos;                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                   |
| <b>Description</b>              | Prevent LDP sessions from being established with remote neighbors that have not been specifically configured. LDP peers will not respond to targeted hellos coming from a source that is not one of the configured remote neighbors.                                                |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling Strict Targeted Hello Messages for LDP on page 3544</a></li> </ul>                                                                                                                                                    |

## targeted-hello

---

|                                 |                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | targeted-hello {<br>hello-interval <i>seconds</i> ;<br>hold-time <i>seconds</i> ;<br>}                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <i>ldp</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>ldp</i> ],<br>[edit protocols <i>ldp</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <i>ldp</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.5.                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Specify the LDP timer and LDP hold time for targeted hellos.                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the LDP Timer for Hello Messages on page 3542</a></li><li>• <a href="#">Configuring the Delay Before LDP Neighbors Are Considered Down on page 3543</a></li></ul>                                                                                   |

## traceoptions (Protocols LDP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols ldp],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br/> [edit protocols ldp],<br/> [edit routing-instances <i>routing-instance-name</i> protocols ldp]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.<br/> <b>match-on address</b> option for the <b>filter</b> flag modifier added in Junos OS Release 10.4.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>         | LDP protocol-level trace options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Default</b>             | The default LDP protocol-level trace options are inherited from the routing protocols <b>traceoptions</b> statement included at the [edit routing-options] hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>ldp-log</b>. We recommend that you place LDP tracing output in the file <b>ldp-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p><b>Range:</b> 2 through 1000<br/> <b>Default:</b> 2 files</p> <p>If you specify a maximum number of files, you must also include the <b>size</b> statement to specify the maximum file size.</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <ul style="list-style-type: none"> <li>• <b>address</b>—Operation of address and address withdrawal messages</li> <li>• <b>binding</b>—Label-binding operations</li> <li>• <b>error</b>—Error conditions</li> <li>• <b>event</b>—Protocol events</li> <li>• <b>initialization</b>—Operation of initialization messages</li> </ul> |

- **label**—Operation of label request, label map, label withdrawal, and label release messages
- **notification**—Operation of notification messages
- **packets**—Equivalent to setting **address**, **initialization**, **label**, **notification**, and **periodic** flags (see also the **filter** flag modifier)
- **path**—Label-switched path operations
- **periodic**—Operation of hello and keepalive messages
- **route**—Operation of route messages
- **state**—Protocol state transitions

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **disable**—Disable this trace flag.
- **filter**—Filter to apply to this flag. The **filter** flag modifier can be applied only to the **route**, **path**, and **binding** flags. This flag modifier has the following options:
  - **match-on**—Match on argument specified. The **match-on** option has the following suboptions:
    - **address**—Filter based on the source and destination addresses of packets. Available for the **packets** flag option only.
    - **fec**—Filter based on the FEC associated with the traced object.
  - **policy *policy-name***—Specify the filter policy.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

**no-world-readable**—(Optional) Prevent all users from reading the log file.

**size *size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

**world-readable**—(Optional) Enable any user to read the log file.

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

**Related Documentation**

- [Tracing LDP Protocol Traffic on page 3563](#)
- *Network Management Configuration Guide*

## track-igp-metric

**Syntax** track-igp-metric;

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols ldp],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],  
[edit protocols ldp],  
[edit routing-instances *routing-instance-name* protocols ldp]

**Release Information** Statement introduced before Junos OS Release 7.4.

**Description** Cause the IGP route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1).

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

**Related Documentation**

- *Configuring Miscellaneous LDP Properties*

## traffic-statistics (Protocols LDP)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traffic-statistics {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;<br/>    interval <i>seconds</i>;<br/>    no-penultimate-hop;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>         | LDP traffic statistics display the amount of traffic passed through a router or switch for a particular FEC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>             | <p><b>file <i>filename</i></b>—Name of the file to receive the output of the LDP statistics operation. Enclose the name within quotation marks. All files are placed in the directory <code>/var/log</code>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of LDP statistics files. When a statistics file named <b><i>ldp-stat</i></b> reaches its maximum size, it is renamed <b><i>ldp-stat.0</i></b>, then <b><i>ldp-stat.1</i></b>, and so on, until the maximum number of LDP statistics files is reached. Then the oldest file is overwritten.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 2 files</p> <p>If you specify a maximum number of files, you also must include the <b>size</b> statement to specify the maximum file size.</p> <p><b>interval <i>seconds</i></b>—(Optional) Specify the interval at which the statistics are polled and written to the file.</p> <p><b>Default:</b> 300 seconds (5 minutes)</p> <p><b>no-penultimate-hop</b>—(Optional) Do not collect traffic statistics on the penultimate hop router.</p> <p><b>no-world-readable</b>—(Optional) Prevent all users from reading the log file.</p> <p><b>size <i>size</i></b>—(Optional) Maximum size of each statistics file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a statistics file named <b><i>ldp-stat</i></b> reaches this size, it is renamed <b><i>ldp-stat.0</i></b>. When <b><i>ldp-stat</i></b> again reaches this size, <b><i>ldp-stat.0</i></b> is renamed <b><i>ldp-stat.1</i></b> and <b><i>ldp-stat</i></b> is renamed <b><i>ldp-stat.0</i></b>. This renaming scheme continues until the maximum number of statistics files is reached. Then the oldest statistics file is overwritten.</p> <p><b>Syntax:</b> <i>xk</i> to specify KB, <i>xm</i> to specify MB, or <i>xg</i> to specify GB</p> <p><b>Range:</b> 10 KB through the maximum file size supported on your system</p> <p><b>Default:</b> 1 MB</p> |

If you specify a maximum file size, you also must also include the **files** statement to specify the maximum number of files.

**world-readable**—(Optional) Enable log file access for all users.

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

**Related Documentation**

- [Collecting LDP Statistics on page 3561](#)

## transport-address

|                                 |                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | transport-address (interface   router-id);                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols ldp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],<br>[edit protocols ldp interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Enable control of the transport address used by LDP.                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>                  | router-id                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b>interface</b> —The first IP address on the interface is used as the transport address.<br><b>router-id</b> —The router identifier is used as the transport address.                                                                                                                                                                                |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Specifying the Transport Address Used by LDP on page 3552</a></li> </ul>                                                                                                                                                                                                                         |

## Administration

- [Operational Commands on page 3617](#)

## Operational Commands

## ping mpls ldp

---

**Syntax**    ping mpls ldp *fec*  
              <count *count*>  
              <destination *address*>  
              <detail>  
              <exp *forwarding-class*>  
              <instance *routing-instance-name*>  
              <logical-system (all | *logical-system-name*)>  
              <p2mp root-addr *ip-address* lsp-id *identifier*>  
              <size *bytes*>  
              <source *source-address*>  
              <sweep>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                          Command introduced in Junos OS Release 9.0 for EX Series switches.  
                          **size** and **sweep** options introduced in Junos OS Release 9.6.  
                          **instance** option introduced in Junos OS Release 10.0.  
                          **p2mp**, **root-address**, and **lsp-id** options introduced in Junos OS Release 11.2.

**Description**    Check the operability of MPLS LDP-signaled label-switched path (LSP) connections.  
                  Type Ctrl+c to interrupt a **ping mpls** command.

**Options**    **count** *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through **1,000,000**. The default value is **5**.

**destination** *address*—(Optional) Specify an address other than the default (**127.0.0.1/32**) for the ping echo requests. The address can be anything within the **127/8** subnet.

**detail**—(Optional) Display detailed information about the echo requests sent and received.

**exp** *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

**fec**—Ping an LDP-signaled LSP using the forwarding equivalence class (FEC) prefix and length.

**instance** *routing-instance-name*—(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.

**logical-system** (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

**p2mp root-addr** *ip-address* **lsp-id** *identifier*—(Optional) Ping the end points of a point-to-multipoint LSP. Enter the IP address of the point-to-multipoint LSP root and the ID number of the point-to-multipoint LSP.

**size** *bytes*—(Optional) Size of the LSP ping request packet (**88** through **65468** bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 88-byte minimum.



**source *source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

**Additional Information** If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

You can configure the ping interval for the **ping mpls ldp** command by specifying a new time in seconds using the **lsp-ping-interval** statement at the **[edit protocols ldp oam]** hierarchy level. For more information, see the *Junos OS MPLS Applications Configuration Guide*.

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

**Required Privilege Level** network

**List of Sample Output** [ping mpls ldp fec count on page 3619](#)  
[ping mpls ldp p2mp root-addr lsp-id on page 3619](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with error codes are not counted in the received packets count. They are accounted for separately.

## Sample Output

### ping mpls ldp fec count

```
user@host> ping mpls ldp 10.255.245.222 count 10
!!!xxx...x--- 1sping statistics ---10 packets transmitted, 3 packets received,
70% packet loss 4 packets received with error status, not counted as received.
```

### ping mpls ldp p2mp root-addr lsp-id

```
user@host> ping mpls ldp p2mp root-addr 10.1.1.1/32 lsp-id 1 count 1
Request for seq 1, to interface 71, no label stack.
Request for seq 1, to interface 70, label 299786
Reply for seq 1, egress 10.1.1.3, return code: Egress-ok, time: 18.936 ms
 Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
 Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
Reply for seq 1, egress 10.1.1.4, return code: Egress-ok, time: 18.936 ms
 Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
 Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
```

```
Reply for seq 1, egress 10.1.1.5, return code: Egress-ok, time: 18.936 ms
Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
```

## show ldp database

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show ldp database &lt;brief   detail   extensive&gt; &lt;inet   l2circuit&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;session <i>session</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Display entries in the Label Distribution Protocol (LDP) database.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b>none</b>—Display standard information about all entries in the LDP database for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>inet   l2circuit</b>—(Optional) Display only IPv4 or Layer 2 circuit bindings.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display routing instance information for the specified instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>session <i>session</i></b>—(Optional) Display database for the specified session only. <b><i>session</i></b> is the destination address of the LDP session.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>    | <a href="#">show ldp database on page 3623</a><br><a href="#">show ldp database l2circuit detail on page 3623</a><br><a href="#">show ldp database session on page 3624</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>            | Table 271 on page 3621 describes the output fields for the <b>show ldp database</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

**Table 271: show ldp database Output Fields**

| Field Name                       | Field Description                                                                | Level of Output |
|----------------------------------|----------------------------------------------------------------------------------|-----------------|
| <b>Input label database</b>      | Label received from the other router.                                            | All levels      |
| <b>Output label database</b>     | Label advertised to the other router.                                            | All levels      |
| <b><i>session-identifier</i></b> | Session identifier, which includes the local and remote label space identifiers. | All levels      |
| <b>Label</b>                     | Label binding to a route prefix.                                                 | All levels      |

Table 271: show ldp database Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Prefix                          | <p>Route prefix. It can be either the IP prefix or the Layer 2 encapsulation type in the format <b>L2CKT control word status encapsulation-type vc-number</b>, for example, <b>L2CKT CtlfWord FRAME RELAY VC 2</b></p> <ul style="list-style-type: none"> <li>• <b>control-word-status</b>—Displays whether the use of the control word has been negotiated for this virtual circuit: <ul style="list-style-type: none"> <li>• NoCtrlWord</li> <li>• CtrlWord</li> </ul> </li> <li>• <b>encapsulation-type</b>—Encapsulation type: <ul style="list-style-type: none"> <li>• FRAME RELAY</li> <li>• ATM AAL5</li> <li>• ATM CELL</li> <li>• VLAN</li> <li>• ETHERNET</li> <li>• CISCO_HDLC</li> <li>• PPP</li> </ul> </li> <li>• <b>VC number</b>—Virtual circuit number. It can have any numeric value.</li> <li>• <b>(Stale)</b>—When you display the LDP database for the neighbor of a restarting router, the bindings learned from the restarting neighbor are displayed as (Stale). Stale bindings are deleted if they are not refreshed within the recovery time.</li> </ul> | All levels      |
| MTU                             | MTU of the Layer 2 circuit. MTU is displayed for all encapsulation types except ATM cell encapsulations.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | detail          |
| VCCV Control Channel types      | <p>Virtual Circuit Connection Verification (VCCV) control channel types</p> <ul style="list-style-type: none"> <li>• MPLS router alert label</li> <li>• MPLS PW label with TTL=1</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | extensive       |
| VCCV Control Verification types | The only valid VCCV control verification type is <b>LSP ping</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | extensive       |
| TDM payload size                | Size of the Time Division Multiplex (TDM) payload.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels      |
| TDM bitrate                     | Bit rate for the TDM traffic.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels      |
| Requested VLAN ID               | (VLANs) VLAN identifier of the Layer 2 circuit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | detail          |
| Cell bundle size                | (ATM cell encapsulations) Maximum number of cells that the Layer 2 circuit can receive in a packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | detail          |

Table 271: show ldp database Output Fields (*continued*)

| Field Name   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>State</b> | State of the label binding: <ul style="list-style-type: none"> <li>• <b>Active</b>—Label binding has been installed and distributed appropriately. A label binding is almost always in this state.</li> <li>• <b>New</b>—New label that has not yet been distributed. <ul style="list-style-type: none"> <li>• <b>MapRcv</b>—Waiting to receive a label mapping message.</li> <li>• <b>MapSend</b>—Waiting to send a label mapping message.</li> <li>• <b>RelRcv</b>—Waiting to receive a label release message.</li> <li>• <b>RelRsnd</b>—Waiting to receive a label release message before resending label mapping message.</li> <li>• <b>RelSend</b>—Waiting to send a label release message.</li> <li>• <b>ReqSend</b>—Waiting to send a label request message.</li> <li>• <b>W/dSend</b>—Waiting to send a label withdrawal message.</li> </ul> </li> </ul> | <b>detail</b>   |
| <b>Age</b>   | Time elapsed since the binding was created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b>   |

## Sample Output

### show ldp database

```

user@host> show ldp database
Input label database, 10.255.245.222:0--10.255.245.221:0
 Label Prefix
 3 10.255.245.221/32 (Stale)
100018 10.255.245.222/32
100011 L2CKT FRAME RELAY VC 11
Output label database, 10.255.245.222:0--10.255.245.221:0
 Label Prefix
 3 10.255.245.221/32
100018 10.255.245.222/32
100011 L2CKT FRAME RELAY VC 1

```

### show ldp database l2circuit detail

```

user@host> show ldp database l2circuit detail
Input label database, 10.255.245.44:0--10.255.245.45:0
 Label Prefix
 100176 L2CKT CtrlWord ATM CELL (VC Mode) VC 100
 Cell bundle size: 80
 State: Active
 Age: 9:48
 100256 L2CKT CtrlWord FRAME RELAY VC 101
 MTU: 4470
 State: Active
 Age: 9:48

Output label database, 10.255.245.44:0--10.255.245.45:0
 Label Prefix
 100048 L2CKT CtrlWord ATM CELL (VC Mode) VC 100
 Cell bundle size: 80
 State: Active
 Age: 9:48
 100112 L2CKT CtrlWord FRAME RELAY VC 101

```

MTU: 4470  
State: Active  
Age: 9:48

#### show ldp database session

```
user@host> show ldp database session 10.1.1.195
Input label database, 10.0.0.194:0--10.1.1.195:0
 Label Prefix
100002 10.255.245.197/32
100003 10.255.245.196/32
100004 10.0.0.194/32
 3 10.1.1.195/32
100000 L2CKT NoCtrlWord FRAME RELAY VC 1
100001 L2CKT CtrlWord FRAME RELAY VC 2
Output label database, 10.0.0.194:0--10.1.1.195:0
 Label Prefix
100003 10.255.245.197/32
100004 10.1.1.195/32
100002 10.255.245.196/32
 3 10.0.0.194/32
100000 L2CKT CtrlWord FRAME RELAY VC 2
100001 L2CKT NoCtrlWord FRAME RELAY VC 1
```

## show ldp session

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show ldp session &lt;brief   detail   extensive&gt; &lt;destination&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Display information about Label Distribution Protocol (LDP) sessions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b>none</b>—Display standard information about all LDP sessions for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>destination</b>—(Optional) Restrict LDP session display to the specified address.</p> <p><b>instance instance-name</b>—(Optional) Display routing instance information for the specified instance. If <b>instance-name</b> is omitted, information is displayed for the master instance.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">clear ldp session</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <a href="#">show ldp session brief on page 3628</a><br><a href="#">show ldp session detail on page 3628</a><br><a href="#">show ldp session extensive on page 3629</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>            | Table 272 on page 3625 describes the output fields for the <b>show ldp session</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 272: show ldp session Output Fields

| Field Name | Field Description                                                                                                                                                                                                                                                                           | Level of Output  |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Address    | Transport address of the session.                                                                                                                                                                                                                                                           | any              |
| State      | State of the session: <b>Nonexistent</b> , <b>Connecting</b> , <b>Initialized</b> , <b>OpenRec</b> , <b>OpenSent</b> , <b>Operational</b> , or <b>Closing</b> . The states correspond to the state diagram specified in Internet Draft LDP Specification draft-ietf-mpls-rfc3036bis-01.txt. | any              |
| Connection | TCP connection state: <b>Closed</b> , <b>Opening</b> , or <b>Open</b> .                                                                                                                                                                                                                     | any              |
| Hold time  | Time remaining until the session will be closed, in seconds.                                                                                                                                                                                                                                | any              |
| Session ID | LDP identifiers of the peers of this session.                                                                                                                                                                                                                                               | detail extensive |

Table 272: show ldp session Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output         |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Next keepalive</b>         | Time until next keepalive is sent, in seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b> |
| <b>Active</b>                 | Whether the local router or switch is playing the active role in the session and during session establishment.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b> |
| <b>Passive</b>                | Whether the local router or switch is playing the passive role in the session and during session establishment.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b> |
| <b>Maximum PDU</b>            | Maximum protocol data unit (PDU) size (packet size) for the session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b> |
| <b>Hold time</b>              | Time remaining until the session will be closed, in seconds. This value corresponds to the one configured using the <b>keepalive-timeout</b> statement configured at the <b>[edit protocols ldp]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b> |
| <b>Neighbor count</b>         | Number of neighbors that are contributing to the session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b> |
| <b>Keepalive interval</b>     | Keepalive interval, in seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b> |
| <b>Connect retry interval</b> | TCP connection retry interval, in seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Local address</b>          | Local transport address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b> |
| <b>Remote address</b>         | Remote transport address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b> |
| <b>Up for</b>                 | Time that this session has been up.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b> |
| <b>Last down</b>              | Time since the session last went down.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b> |
| <b>Reason</b>                 | Reason the session went down: <ul style="list-style-type: none"> <li>• Aborted graceful restart</li> <li>• Authentication key was changed</li> <li>• Bad type length value (TLV)</li> <li>• Bad protocol data unit (PDU) packets</li> <li>• Command-line interface (CLI) command</li> <li>• Connect time expired</li> <li>• Connection error</li> <li>• Connection reset</li> <li>• Error during initialization</li> <li>• Hold time expired</li> <li>• No adjacency or all adjacencies down</li> <li>• Notification received</li> <li>• Received notification from peer</li> <li>• Unexpected End of File (EOF)</li> <li>• Unknown reason</li> </ul> | <b>detail extensive</b> |



Table 272: show ldp session Output Fields (*continued*)

| Field Name                          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output  |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Number of session flaps             | Number of times the session changes from up to down.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | detail extensive |
| Restarting                          | LDP is in the process of gracefully restarting.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | detail extensive |
| Capabilities advertised             | LDP capabilities advertised to a peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | detail extensive |
| Capabilities received               | LDP capabilities received from a peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | detail extensive |
| Protection                          | Information about the status of MPLS LDP session protection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | detail extensive |
| restart complete in <i>nnn msec</i> | Amount of time (in milliseconds) remaining until graceful restart is declared complete.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | detail extensive |
| Local                               | <p>Information about graceful restart for the local end of an LDP session. Graceful restart and helper mode are independent.</p> <ul style="list-style-type: none"> <li>• <b>Restart</b>—Status of the graceful restart feature at the local end of the LDP session: <b>enabled</b> or <b>disabled</b>.</li> <li>• <b>Helper mode</b>—Status of the helper mode feature at the local end of the LDP session: <b>enabled</b> or <b>disabled</b>. When this feature is enabled, the local end of the LDP session can help the restarting router with its LDP restart procedures.</li> <li>• <b>Reconnect time</b>—Amount of time to wait from when a restart is initiated until the router can exchange LDP messages with its neighbors. The default is <b>60000 msec</b> and is not configurable. (<b>Reconnect timeout</b> refers to "FT Reconnect timeout" in draft-ietf-mpls-ldp-restart-06, <i>Internet Draft Graceful Restart Mechanism for LDP</i>.)</li> </ul> | detail extensive |
| Remote                              | <p>Information about graceful restart at the remote end of an LDP session. Graceful restart and helper mode are independent.</p> <ul style="list-style-type: none"> <li>• <b>Restart</b>—Status of the graceful restart feature at the remote end of the LDP session: <b>enabled</b> or <b>disabled</b>.</li> <li>• <b>Helper mode</b>—Status of the helper mode feature at the remote end of the LDP session: <b>enabled</b> or <b>disabled</b>. When this feature is enabled, the remote end of the LDP session can help the restarting router with its LDP restart procedures.</li> <li>• <b>Reconnect time</b>—Amount of time in milliseconds from when a restart is initiated until the remote router can exchange LDP messages with its neighbors.</li> </ul>                                                                                                                                                                                                  | detail extensive |
| Local maximum recovery time         | Amount of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | detail extensive |
| Next-hop addresses received         | Next-hop addresses received on the session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | detail extensive |
| Queue depth                         | Number of messages that are queued for sending to the peers in the group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | extensive        |

Table 272: show ldp session Output Fields (*continued*)

| Field Name   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Message type | <p>Type of message being sent:</p> <ul style="list-style-type: none"> <li>• <b>Initialization</b>—Session initialization negotiation messages sent by an LSR to an LDP peer when the transport connection is established.</li> <li>• <b>Keepalive</b>—Keepalive timer messages sent by an LSR to an LDP peer to keep the session active when there is no information or PDU exchanged between them.</li> <li>• <b>Notification</b>—Notification messages (such as state of the LDP session) or error information (such as bad PDU length) sent by an LSR to an LDP peer.</li> <li>• <b>Address</b>—Message sent by an LSR to an LDP peer to advertise interface addresses.</li> <li>• <b>Address withdraw</b>—Message sent by an LSR to an LDP peer to withdraw a previously advertised interface address.</li> <li>• <b>Label mapping</b>—Message sent by an LSR to an LDP peer to advertise label mapping for a forwarding equivalence class (FEC).</li> <li>• <b>Label request</b>—Message sent by an LSR to an LDP peer to request a label mapping for an FEC.</li> <li>• <b>Label withdraw</b>—Message sent by an LSR to an LDP peer to withdraw a previously advertised FEC-label mapping.</li> <li>• <b>Label release</b>—Message sent by an LSR to an LDP peer to notify the peer that a specific FEC-label mapping has been released.</li> <li>• <b>Label abort</b>—Message sent by an LSR to an LDP peer to abort a label request message.</li> <li>• <b>Total</b>—Messages sent and received during the lifetime of the session.</li> <li>• <b>Last 5 seconds</b>—Messages sent and received during the current session.</li> </ul> | extensive       |

## Sample Output

### show ldp session brief

```

user@host> show ldp session brief
 Address State Connection Hold time
10.255.72.160 Operational Open 21
10.255.72.164 Operational Open 20
10.255.72.172 Operational Open 21

```

### show ldp session detail

```

user@host> show ldp session detail
Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 27
Session ID: 192.168.0.2:0--192.168.0.3:0
Next keepalive in 7 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: discovered
Keepalive interval: 10, Connect retry interval: 1
Local address: 192.168.0.2, Remote address: 192.168.0.3
Up for 00:00:02
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000

```

```

Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream unsolicited
Negotiated Label Advertisement mode: Downstream unsolicited
Nonstop routing state: Not in sync
Next-hop addresses received:
 10.0.0.5
 10.0.0.33

```

### show ldp session extensive

```

user@host> show ldp session extensive
Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 22
 Session ID: 192.168.0.2:0--192.168.0.3:0
 Next keepalive in 2 seconds
 Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
 Neighbor types: discovered
 Keepalive interval: 10, Connect retry interval: 1
 Local address: 192.168.0.2, Remote address: 192.168.0.3
 Up for 00:05:37
 Capabilities advertised: none
 Capabilities received: none
 Protection: disabled
 Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
 Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
 Local maximum neighbor reconnect time: 120000 msec
 Local maximum neighbor recovery time: 240000 msec
 Local Label Advertisement mode: Downstream unsolicited
 Remote Label Advertisement mode: Downstream unsolicited
 Negotiated Label Advertisement mode: Downstream unsolicited
 Nonstop routing state: Not in sync
 Next-hop addresses received:
 10.0.0.5
 10.0.0.33
 Queue depth: 0

```

| Message type     | Total |          | Last 5 seconds |          |
|------------------|-------|----------|----------------|----------|
|                  | Sent  | Received | Sent           | Received |
| Initialization   | 1     | 1        | 0              | 0        |
| Keepalive        | 33    | 33       | 1              | 1        |
| Notification     | 0     | 0        | 0              | 0        |
| Address          | 1     | 1        | 0              | 0        |
| Address withdraw | 0     | 0        | 0              | 0        |
| Label mapping    | 7     | 5        | 0              | 0        |
| Label request    | 0     | 0        | 0              | 0        |
| Label withdraw   | 3     | 1        | 0              | 0        |
| Label release    | 1     | 3        | 0              | 0        |
| Label abort      | 0     | 0        | 0              | 0        |

## show ldp traffic-statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show ldp traffic-statistics &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;p2mp&gt;</pre>                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p><b>p2mp</b> option added in Junos OS Release 11.2.</p>                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Display Label Distribution Protocol (LDP) traffic statistics.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>none</b>—Display LDP traffic statistics for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display LDP traffic statistics for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>p2mp</b>—(Optional) Display only the data traffic statistics for a point-to-multipoint LSP.</p> |
| <b>Additional Information</b>   | To obtain output from this command, you must configure the <b>traffic-statistics</b> statement for the LDP protocol. For more information, see the <i>Junos MPLS Applications Configuration Guide</i> .                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">clear ldp statistics</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>List of Sample Output</b>    | <a href="#">show ldp traffic-statistics on page 3631</a><br><a href="#">show ldp traffic-statistics p2mp on page 3631</a>                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>            | <p><a href="#">Table 273 on page 3630</a> lists the output fields for the <b>show ldp traffic-statistics</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                             |

**Table 273: show ldp traffic-statistics Output Fields**

| Field Name   | Field Description                                                                                                                                                                                      |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Message type | LDP message types.                                                                                                                                                                                     |
| FEC          | <p>Forwarding equivalence class (FEC) for which LDP traffic statistics are collected.</p> <p>For P2MP LSPs, FEC appears as a combination of root address and the LSP ID (<b>root_addr:lsp_id</b>).</p> |
| Type         | Type of traffic originating from a router, either <b>Ingress</b> (originating from this router) or <b>Transit</b> (forwarded through this router).                                                     |
| Packets      | Number of packets passed by the FEC since its LSP came up.                                                                                                                                             |

Table 273: show ldp traffic-statistics Output Fields (*continued*)

| Field Name     | Field Description                                                                                                                                                                                                                                                                                                            |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Bytes</b>   | Number of bytes of data passed by the FEC since its LSP came up.                                                                                                                                                                                                                                                             |
| <b>Shared</b>  | Whether a label is shared by prefixes: <b>Yes</b> or <b>No</b> . A <b>Yes</b> value indicates that several prefixes are bound to the same label (for example, when several prefixes are advertised with an egress policy). The LDP traffic statistics for this case apply to all the prefixes and should be treated as such. |
| <b>Nexthop</b> | The next hop address for P2MP LSPs.                                                                                                                                                                                                                                                                                          |

## Sample Output

### show ldp traffic-statistics

```
user@host> show ldp traffic-statistics
```

| FEC               | Type    | Packets | Bytes | Shared |
|-------------------|---------|---------|-------|--------|
| 10.35.3.0/30      | Transit | 0       | 0     | Yes    |
|                   | Ingress | 0       | 0     | No     |
| 10.35.10.1/32     | Transit | 0       | 0     | Yes    |
|                   | Ingress | 0       | 0     | No     |
| 10.255.245.214/32 | Transit | 0       | 0     | No     |
|                   | Ingress | 11      | 752   | No     |
| 192.168.37.36/30  | Transit | 0       | 0     | Yes    |
|                   | Ingress | 0       | 0     | No     |

| FEC(root_addr:lsp_id)  | Nexthop      | Packets | Bytes    | Shared |
|------------------------|--------------|---------|----------|--------|
| 10.255.72.160:16777217 | 192.168.8.81 | 152056  | 14597376 | No     |
|                        | 192.168.8.1  | 152056  | 14597376 | No     |
|                        | 192.168.8.65 | 152056  | 14597376 | No     |

### show ldp traffic-statistics p2mp

```
user@host> show ldp traffic-statistics p2mp
```

| FEC(root_addr:lsp_id)  | Nexthop      | Packets | Bytes    | Shared |
|------------------------|--------------|---------|----------|--------|
| 10.255.72.160:16777217 | 192.168.8.81 | 152056  | 14597376 | No     |
|                        | 192.168.8.1  | 152056  | 14597376 | No     |
|                        | 192.168.8.65 | 152056  | 14597376 | No     |

## show ldp session

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show ldp session &lt;brief   detail   extensive&gt; &lt;destination&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Display information about Label Distribution Protocol (LDP) sessions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b>none</b>—Display standard information about all LDP sessions for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>destination</b>—(Optional) Restrict LDP session display to the specified address.</p> <p><b>instance instance-name</b>—(Optional) Display routing instance information for the specified instance. If <b>instance-name</b> is omitted, information is displayed for the master instance.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">clear ldp session</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <a href="#">show ldp session brief on page 3635</a><br><a href="#">show ldp session detail on page 3635</a><br><a href="#">show ldp session extensive on page 3636</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>            | Table 272 on page 3625 describes the output fields for the <b>show ldp session</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 274: show ldp session Output Fields

| Field Name | Field Description                                                                                                                                                                                                                                                                           | Level of Output  |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Address    | Transport address of the session.                                                                                                                                                                                                                                                           | any              |
| State      | State of the session: <b>Nonexistent</b> , <b>Connecting</b> , <b>Initialized</b> , <b>OpenRec</b> , <b>OpenSent</b> , <b>Operational</b> , or <b>Closing</b> . The states correspond to the state diagram specified in Internet Draft LDP Specification draft-ietf-mpls-rfc3036bis-01.txt. | any              |
| Connection | TCP connection state: <b>Closed</b> , <b>Opening</b> , or <b>Open</b> .                                                                                                                                                                                                                     | any              |
| Hold time  | Time remaining until the session will be closed, in seconds.                                                                                                                                                                                                                                | any              |
| Session ID | LDP identifiers of the peers of this session.                                                                                                                                                                                                                                               | detail extensive |

Table 274: show ldp session Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output         |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Next keepalive</b>         | Time until next keepalive is sent, in seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b> |
| <b>Active</b>                 | Whether the local router or switch is playing the active role in the session and during session establishment.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b> |
| <b>Passive</b>                | Whether the local router or switch is playing the passive role in the session and during session establishment.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b> |
| <b>Maximum PDU</b>            | Maximum protocol data unit (PDU) size (packet size) for the session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b> |
| <b>Hold time</b>              | Time remaining until the session will be closed, in seconds. This value corresponds to the one configured using the <b>keepalive-timeout</b> statement configured at the <b>[edit protocols ldp]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b> |
| <b>Neighbor count</b>         | Number of neighbors that are contributing to the session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b> |
| <b>Keepalive interval</b>     | Keepalive interval, in seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b> |
| <b>Connect retry interval</b> | TCP connection retry interval, in seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Local address</b>          | Local transport address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b> |
| <b>Remote address</b>         | Remote transport address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b> |
| <b>Up for</b>                 | Time that this session has been up.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b> |
| <b>Last down</b>              | Time since the session last went down.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b> |
| <b>Reason</b>                 | Reason the session went down: <ul style="list-style-type: none"> <li>• Aborted graceful restart</li> <li>• Authentication key was changed</li> <li>• Bad type length value (TLV)</li> <li>• Bad protocol data unit (PDU) packets</li> <li>• Command-line interface (CLI) command</li> <li>• Connect time expired</li> <li>• Connection error</li> <li>• Connection reset</li> <li>• Error during initialization</li> <li>• Hold time expired</li> <li>• No adjacency or all adjacencies down</li> <li>• Notification received</li> <li>• Received notification from peer</li> <li>• Unexpected End of File (EOF)</li> <li>• Unknown reason</li> </ul> | <b>detail extensive</b> |

Table 274: show ldp session Output Fields (*continued*)

| Field Name                          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output  |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Number of session flaps             | Number of times the session changes from up to down.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | detail extensive |
| Restarting                          | LDP is in the process of gracefully restarting.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | detail extensive |
| Capabilities advertised             | LDP capabilities advertised to a peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | detail extensive |
| Capabilities received               | LDP capabilities received from a peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | detail extensive |
| Protection                          | Information about the status of MPLS LDP session protection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | detail extensive |
| restart complete in <i>nnn msec</i> | Amount of time (in milliseconds) remaining until graceful restart is declared complete.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | detail extensive |
| Local                               | <p>Information about graceful restart for the local end of an LDP session. Graceful restart and helper mode are independent.</p> <ul style="list-style-type: none"> <li>• <b>Restart</b>—Status of the graceful restart feature at the local end of the LDP session: <b>enabled</b> or <b>disabled</b>.</li> <li>• <b>Helper mode</b>—Status of the helper mode feature at the local end of the LDP session: <b>enabled</b> or <b>disabled</b>. When this feature is enabled, the local end of the LDP session can help the restarting router with its LDP restart procedures.</li> <li>• <b>Reconnect time</b>—Amount of time to wait from when a restart is initiated until the router can exchange LDP messages with its neighbors. The default is <b>60000 msec</b> and is not configurable. (<b>Reconnect timeout</b> refers to "FT Reconnect timeout" in draft-ietf-mpls-ldp-restart-06, <i>Internet Draft Graceful Restart Mechanism for LDP</i>.)</li> </ul> | detail extensive |
| Remote                              | <p>Information about graceful restart at the remote end of an LDP session. Graceful restart and helper mode are independent.</p> <ul style="list-style-type: none"> <li>• <b>Restart</b>—Status of the graceful restart feature at the remote end of the LDP session: <b>enabled</b> or <b>disabled</b>.</li> <li>• <b>Helper mode</b>—Status of the helper mode feature at the remote end of the LDP session: <b>enabled</b> or <b>disabled</b>. When this feature is enabled, the remote end of the LDP session can help the restarting router with its LDP restart procedures.</li> <li>• <b>Reconnect time</b>—Amount of time in milliseconds from when a restart is initiated until the remote router can exchange LDP messages with its neighbors.</li> </ul>                                                                                                                                                                                                  | detail extensive |
| Local maximum recovery time         | Amount of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | detail extensive |
| Next-hop addresses received         | Next-hop addresses received on the session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | detail extensive |
| Queue depth                         | Number of messages that are queued for sending to the peers in the group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | extensive        |



Table 274: show ldp session Output Fields (*continued*)

| Field Name   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Message type | <p>Type of message being sent:</p> <ul style="list-style-type: none"> <li>• <b>Initialization</b>—Session initialization negotiation messages sent by an LSR to an LDP peer when the transport connection is established.</li> <li>• <b>Keepalive</b>—Keepalive timer messages sent by an LSR to an LDP peer to keep the session active when there is no information or PDU exchanged between them.</li> <li>• <b>Notification</b>—Notification messages (such as state of the LDP session) or error information (such as bad PDU length) sent by an LSR to an LDP peer.</li> <li>• <b>Address</b>—Message sent by an LSR to an LDP peer to advertise interface addresses.</li> <li>• <b>Address withdraw</b>—Message sent by an LSR to an LDP peer to withdraw a previously advertised interface address.</li> <li>• <b>Label mapping</b>—Message sent by an LSR to an LDP peer to advertise label mapping for a forwarding equivalence class (FEC).</li> <li>• <b>Label request</b>—Message sent by an LSR to an LDP peer to request a label mapping for an FEC.</li> <li>• <b>Label withdraw</b>—Message sent by an LSR to an LDP peer to withdraw a previously advertised FEC-label mapping.</li> <li>• <b>Label release</b>—Message sent by an LSR to an LDP peer to notify the peer that a specific FEC-label mapping has been released.</li> <li>• <b>Label abort</b>—Message sent by an LSR to an LDP peer to abort a label request message.</li> <li>• <b>Total</b>—Messages sent and received during the lifetime of the session.</li> <li>• <b>Last 5 seconds</b>—Messages sent and received during the current session.</li> </ul> | extensive       |

## Sample Output

### show ldp session brief

```

user@host> show ldp session brief
 Address State Connection Hold time
10.255.72.160 Operational Open 21
10.255.72.164 Operational Open 20
10.255.72.172 Operational Open 21

```

### show ldp session detail

```

user@host> show ldp session detail
Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 27
Session ID: 192.168.0.2:0--192.168.0.3:0
Next keepalive in 7 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: discovered
Keepalive interval: 10, Connect retry interval: 1
Local address: 192.168.0.2, Remote address: 192.168.0.3
Up for 00:00:02
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000

```

```

Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream unsolicited
Negotiated Label Advertisement mode: Downstream unsolicited
Nonstop routing state: Not in sync
Next-hop addresses received:
 10.0.0.5
 10.0.0.33

```

### show ldp session extensive

```

user@host> show ldp session extensive
Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 22
 Session ID: 192.168.0.2:0--192.168.0.3:0
 Next keepalive in 2 seconds
 Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
 Neighbor types: discovered
 Keepalive interval: 10, Connect retry interval: 1
 Local address: 192.168.0.2, Remote address: 192.168.0.3
 Up for 00:05:37
 Capabilities advertised: none
 Capabilities received: none
 Protection: disabled
 Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
 Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
 Local maximum neighbor reconnect time: 120000 msec
 Local maximum neighbor recovery time: 240000 msec
 Local Label Advertisement mode: Downstream unsolicited
 Remote Label Advertisement mode: Downstream unsolicited
 Negotiated Label Advertisement mode: Downstream unsolicited
 Nonstop routing state: Not in sync
 Next-hop addresses received:
 10.0.0.5
 10.0.0.33
 Queue depth: 0

```

| Message type     | Total |          | Last 5 seconds |          |
|------------------|-------|----------|----------------|----------|
|                  | Sent  | Received | Sent           | Received |
| Initialization   | 1     | 1        | 0              | 0        |
| Keepalive        | 33    | 33       | 1              | 1        |
| Notification     | 0     | 0        | 0              | 0        |
| Address          | 1     | 1        | 0              | 0        |
| Address withdraw | 0     | 0        | 0              | 0        |
| Label mapping    | 7     | 5        | 0              | 0        |
| Label request    | 0     | 0        | 0              | 0        |
| Label withdraw   | 3     | 1        | 0              | 0        |
| Label release    | 1     | 3        | 0              | 0        |
| Label abort      | 0     | 0        | 0              | 0        |

## CHAPTER 15

# Multicast

- [Overview on page 3637](#)
- [Configuration on page 3656](#)
- [Administration on page 3952](#)

## Overview

---

- [Multicast Overview on page 3637](#)
- [Multicast Protocols Overview on page 3649](#)

## Multicast Overview

- [Multicast Overview on page 3637](#)

### Multicast Overview

---

IP has three fundamental types of addresses: unicast, broadcast, and multicast. A *unicast address* is used to send a packet to a single destination. A *broadcast address* is used to send a datagram to an entire subnetwork. A *multicast address* is used to send a datagram to a set of hosts that can be on different subnetworks and that are configured as members of a multicast group.

A multicast datagram is delivered to destination group members with the same best-effort reliability as a standard unicast IP datagram. This means that multicast datagrams are not guaranteed to reach all members of a group or to arrive in the same order in which they were transmitted. The only difference between a multicast IP packet and a unicast IP packet is the presence of a group address in the IP header destination address field. Multicast addresses use the Class D address format.

Individual hosts can join or leave a multicast group at any time. There are no restrictions on the physical location or the number of members in a multicast group. A host can be a member of more than one multicast group at any time. A host does not have to belong to a group to send packets to members of a group.

routing devices use a group membership protocol to learn about the presence of group members on directly attached subnetworks. When a host joins a multicast group, it transmits a group membership protocol message for the group or groups that it wants to receive and sets its IP process and network interface card to receive frames addressed to the multicast group.

### Comparing Multicast to Unicast

The Junos® operating system (Junos OS) routing protocol process supports a wide variety of routing protocols. These routing protocols carry network information among routing devices not only for *unicast* traffic streams sent between one pair of clients and servers, but also for *multicast* traffic streams containing video, audio, or both, between a single server source and many client receivers. The routing protocols used for multicast differ in many key ways from unicast routing protocols.

Information is delivered over a network by three basic methods: unicast, broadcast, and multicast.

The differences among unicast, broadcast, and multicast can be summarized as follows:

- Unicast: One-to-one, from one source to one destination.
- Broadcast: One-to-all, from one source to all possible destinations.
- Multicast: One-to-many, from one source to multiple destinations expressing an interest in receiving the traffic.



**NOTE:** This list does not include a special category for many-to-many applications, such as online gaming or videoconferencing, where there are many sources for the same receiver and where receivers often double as sources. Many-to-many is a service model that repeatedly employs one-to-many multicast and therefore requires no unique protocol. The original multicast specification, RFC 1112, supports both the any-source multicast (ASM) many-to-many model and the source-specific multicast (SSM) one-to-many model.

With unicast traffic, many streams of IP packets that travel across networks flow from a single source, such as a website server, to a single destination such as a client PC. Unicast traffic is still the most common form of information transfer on networks.

Broadcast traffic flows from a single source to all possible destinations reachable on the network, which is usually a LAN. Broadcasting is the easiest way to make sure traffic reaches its destinations.

Television networks use broadcasting to distribute video and audio. Even if the television network is a cable television (CATV) system, the source signal reaches all possible destinations, which is the main reason that some channels' content is scrambled. Broadcasting is not feasible on the Internet because of the enormous amount of unnecessary information that would constantly arrive at each end user's device, the complexities and impact of scrambling, and related privacy issues.

Multicast traffic lies between the extremes of unicast (one source, one destination) and broadcast (one source, all destinations). Multicast is a "one source, many destinations" method of traffic distribution, meaning only the destinations that explicitly indicate their need to receive the information from a particular source receive the traffic stream.

On an IP network, because destinations (clients) do not often communicate directly with sources (servers), the routing devices between source and destination must be able to determine the topology of the network from the unicast or multicast perspective to avoid routing traffic haphazardly. Multicast routing devices replicate packets received on one input interface and send the copies out on multiple output interfaces.

In IP multicast, the source and destination are almost always hosts and not routing devices. Multicast routing devices distribute the multicast traffic across the network from source to destinations. The multicast routing device must find multicast sources on the network, send out copies of packets on several interfaces, prevent routing loops, connect interested destinations with the proper source, and keep the flow of unwanted packets to a minimum. Standard multicast routing protocols provide most of these capabilities, but some router architectures cannot send multiple copies of packets and so do not support multicasting directly.

### ***IP Multicast Uses***

Multicast allows an IP network to support more than just the unicast model of data delivery that prevailed in the early stages of the Internet. Multicast, originally defined as a host extension in RFC 1112 in 1989, provides an efficient method for delivering traffic flows that can be characterized as one-to-many or many-to-many.

Unicast traffic is not strictly limited to data applications. Telephone conversations, wireless or not, contain digital audio samples and might contain digital photographs or even video and still flow from a single source to a single destination. In the same way, multicast traffic is not strictly limited to multimedia applications. In some data applications, the flow of traffic is from a single source to many destinations that require the packets, as in a news or stock ticker service delivered to many PCs. For this reason, the term *receiver* is preferred to *listener* for multicast destinations, although both terms are common.

Network applications that can function with unicast but are better suited for multicast include collaborative groupware, teleconferencing, periodic or “push” data delivery (stock quotes, sports scores, magazines, newspapers, and advertisements), server or website replication, and distributed interactive simulation (DIS) such as war simulations or virtual reality. Any IP network concerned with reducing network resource overhead for one-to-many or many-to-many data or multimedia applications with multiple receivers benefits from multicast.

If unicast were employed by radio or news ticker services, each radio or PC would have to have a separate traffic session for each listener or viewer at a PC (this is actually the method for some Web-based services). The processing load and bandwidth consumed by the server would increase linearly as more people “tune in” to the server. This is extremely inefficient when dealing with the global scale of the Internet. Unicast places the burden of packet duplication on the server and consumes more and more backbone bandwidth as the number of users grows.

If broadcast were employed instead, the source could generate a single IP packet stream using a broadcast destination address. Although broadcast eliminates the server packet duplication issue, this is not a good solution for IP because IP broadcasts can be sent only to a single subnetwork, and IP routing devices normally isolate IP subnetworks on

separate interfaces. Even if an IP packet stream could be addressed to literally go everywhere, and there were no need to “tune” to any source at all, broadcast would be extremely inefficient because of the bandwidth strain and need for uninterested hosts to discard large numbers of packets. Broadcast places the burden of packet rejection on each host and consumes the maximum amount of backbone bandwidth.


For radio station or news ticker traffic, multicast provides the most efficient and effective outcome, with none of the drawbacks and all of the advantages of the other methods. A single source of multicast packets finds its way to every *interested* receiver. As with broadcast, the transmitting host generates only a single stream of IP packets, so the load remains constant whether there is one receiver or one million. The network routing devices replicate the packets and deliver the packets to the proper receivers, but only the replication role is a new one for routing devices. The links leading to subnets consisting of entirely uninterested receivers carry no multicast traffic. Multicast minimizes the burden placed on sender, network, and receiver.

### ***IP Multicast Terminology***

Multicast has its own particular set of terms and acronyms that apply to IP multicast routing devices and networks. [Figure 80 on page 3641](#) depicts some of the terms commonly used in an IP multicast network.

In a multicast network, the key component is the *routing device*, which is able to replicate packets and is therefore multicast-capable. The routing devices in the IP multicast network, which has exactly the same topology as the unicast network it is based on, use a *multicast routing protocol* to build a *distribution tree* that connects receivers (preferred to the multimedia implications of listeners, but listeners is also used) to *sources*. In multicast terminology, the distribution tree is *rooted at the source* (the root of the distribution tree is the source). The interface on the routing device leading toward the source is the *upstream* interface, although the less precise terms *incoming* or *inbound* interface are used as well. To keep bandwidth use to a minimum, it is best for only one upstream interface on the routing device to receive multicast packets. The interface on the routing device leading toward the receivers is the *downstream* interface, although the less precise terms *outgoing* or *outbound* interface are used as well. There can be 0 to  $N-1$  downstream interfaces on a routing device, where  $N$  is the number of logical interfaces on the routing device. To prevent looping, the upstream interface must never receive copies of downstream multicast packets.

Figure 80: Multicast Terminology in an IP Network



Routing loops are disastrous in multicast networks because of the risk of repeatedly replicated packets. One of the complexities of modern multicast routing protocols is the need to avoid routing loops, packet by packet, much more rigorously than in unicast routing protocols. Three multicast strategies—reverse-path forwarding (RPF), shortest-path tree (SPT), and administrative scoping—help prevent routing loops by defining routing paths in different ways.

#### ***Reverse-Path Forwarding for Loop Prevention***

The routing device's multicast forwarding state runs more logically based on the reverse path, from the receiver back to the root of the distribution tree. In RPF, every multicast packet received must pass an RPF check before it can be replicated or forwarded on any interface. When it receives a multicast packet on an interface, the routing device verifies that the *source* address in the multicast IP packet is the *destination* address for a unicast IP packet back to the source.

If the outgoing interface found in the unicast routing table is the same interface that the multicast packet was received on, the packet passes the RPF check. Multicast packets that fail the RPF check are dropped, because the incoming interface is not on the shortest path back to the source. routing devices can build and maintain separate tables for RPF purposes.

#### ***Shortest-Path Tree for Loop Prevention***

The distribution tree used for multicast is rooted at the source and is the shortest-path tree (SPT), but this path can be long if the source is at the periphery of the network. Providing a *shared tree* on the backbone as the distribution tree locates the multicast source more centrally in the network. Shared distribution trees with roots in the core network are created and maintained by a multicast routing device operating as a rendezvous point (RP), a feature of sparse mode multicast protocols.

### ***Administrative Scoping for Loop Prevention***

Scoping limits the routing devices and interfaces that can forward a multicast packet. Multicast scoping is *administrative* in the sense that a range of multicast addresses is reserved for scoping purposes, as described in RFC 2365, *Administratively Scoped IP Multicast*. routing devices at the boundary must filter multicast packets and ensure that packets do not stray beyond the established limit.

### ***Multicast Leaf and Branch Terminology***

Each subnetwork with hosts on the routing device that has at least one interested receiver is a *leaf* on the distribution tree. routing devices can have multiple leaves on different interfaces and must send a copy of the IP multicast packet out on each interface with a leaf. When a new leaf subnetwork is added to the tree (that is, the interface to the host subnetwork previously received no copies of the multicast packets), a new *branch* is built, the leaf is joined to the tree, and replicated packets are sent out on the interface. The number of leaves on a particular interface does not affect the routing device. The action is the same for one leaf or a hundred.



**NOTE:** On Juniper Networks security devices, if the maximum number of leaves on a multicast distribution tree is exceeded, multicast sessions are created up to the maximum number of leaves, and any multicast sessions that exceed the maximum number of leaves are ignored. The maximum number of leaves on a multicast distribution tree is device specific.

---

When a branch contains no leaves because there are no interested hosts on the routing device interface leading to that IP subnetwork, the branch is *pruned* from the distribution tree, and no multicast packets are sent out that interface. Packets are replicated and sent out multiple interfaces only where the distribution tree branches at a routing device, and no link ever carries a duplicate flow of packets.

Collections of hosts all receiving the same stream of IP packets, usually from the same multicast source, are called *groups*. In IP multicast networks, traffic is delivered to multicast groups based on an IP multicast address, or *group address*. The groups determine the location of the leaves, and the leaves determine the branches on the multicast network.

### ***IP Multicast Addressing***

Multicast uses the Class D IP address range (224.0.0.0 through 239.255.255.255). Class D addresses are commonly referred to as *multicast addresses* because the entire classful address concept is obsolete. Multicast addresses can never appear as the source address in an IP packet and can only be the destination of a packet.

Multicast addresses usually have a prefix length of /32, although other prefix lengths are allowed. Multicast addresses represent logical groupings of receivers and not physical collections of devices. Blocks of multicast addresses can still be described in terms of prefix length in traditional notation, but only for convenience. For example, the multicast address range from 232.0.0.0 through 232.255.255.255 can be written as 232.0.0.0/8 or 232/8.



Internet service providers (ISPs) do not typically allocate multicast addresses to their customers because multicast addresses relate to content, not to physical devices. Receivers are not assigned their own multicast addresses, but need to know the multicast address of the content. Sources need to be assigned multicast addresses only to produce the content, not to identify their place in the network. Every source and receiver still needs an ordinary, unicast IP address.

Multicast addressing most often references the receivers, and the source of multicast content is usually not even a member of the multicast group for which it produces content. If the source needs to monitor the packets it produces, monitoring can be done locally, and there is no need to make the packets traverse the network.

Many applications have been assigned a range of multicast addresses for their own use. These applications assign multicast addresses to sessions created by that application. You do not usually need to statically assign a multicast address, but you can do so.

### ***Multicast Addresses***

Multicast host group addresses are defined to be the IP addresses whose high-order four bits are 1110, giving an address range from 224.0.0.0 through 239.255.255.255, or simply 224.0.0.0/4. (These addresses also are referred to as Class D addresses.)

The Internet Assigned Numbers Authority (IANA) maintains a list of registered IP multicast groups. The base address 224.0.0.0 is reserved and cannot be assigned to any group. The block of multicast addresses from 224.0.0.1 through 224.0.0.255 is reserved for local wire use. Groups in this range are assigned for various uses, including routing protocols and local discovery mechanisms.

The range from 239.0.0.0 through 239.255.255.255 is reserved for administratively scoped addresses. Because packets addressed to administratively scoped multicast addresses do not cross configured administrative boundaries, and because administratively scoped multicast addresses are locally assigned, these addresses do not need to be unique across administrative boundaries.

### ***Layer 2 Frames and IPv4 Multicast Addresses***

Multicasting on a LAN is a good place to start an investigation of multicasting at Layer 2. At Layer 2, multicast deals with media access control (MAC) frames and addresses instead of IPv4 or IPv6 packets and addresses. Consider a single LAN, without routing devices, with a multicast source sending to a certain group. The rest of the hosts are receivers interested in the multicast group's content. So the multicast source host generates packets with its unicast IP address as the source, and the multicast group address as the destination.

Which MAC addresses are used on the frame containing this packet? The packet source address—the unicast IP address of the host originating the multicast content—translates easily and directly to the MAC address of the source. But what about the packet's destination address? This is the IP multicast group address. Which destination MAC address for the frame corresponds to the packet's multicast group address?

One option is for LANs simply to use the LAN broadcast MAC address, which guarantees that the frame is processed by every station on the LAN. However, this procedure defeats

the whole purpose of multicast, which is to limit the circulation of packets and frames to interested hosts. Also, hosts might have access to many multicast groups, which multiplies the amount of traffic to noninterested destinations. Broadcasting frames at the LAN level to support multicast groups makes no sense.

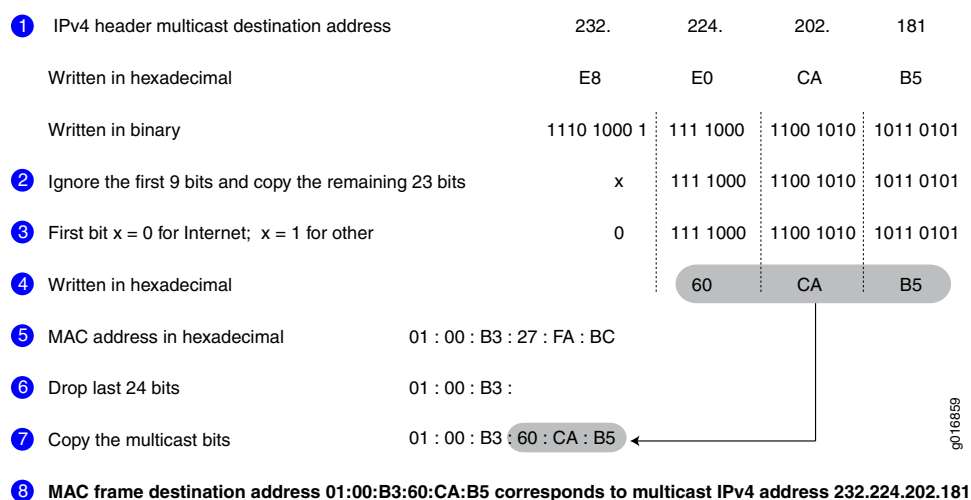
However, there is an easy way to effectively use Layer 2 frames for multicast purposes. The MAC address has a bit that is set to 0 for unicast (the LAN term is *individual address*) and set to 1 to indicate that this is a multicast address. Some of these addresses are reserved for multicast groups of specific vendors or MAC-level protocols. Internet multicast applications use the range 0x01-00-5E-00-00-00 to 0x01-00-5E-FF-FF-FF. Multicast receivers (hosts running TCP/IP) listen for frames with one of these addresses when the application joins a multicast group. The host stops listening when the application terminates or the host leaves the group at the packet layer (Layer 3).

This means that 3 bytes, or 24 bits, are available to map IPv4 multicast addresses at Layer 3 to MAC multicast addresses at Layer 2. However, all IPv4 addresses, including multicast addresses, are 32 bits long, leaving 8 IP address bits left over. Which method of mapping IPv4 multicast addresses to MAC multicast addresses minimizes the chance of “collisions” (that is, two different IP multicast groups at the packet layer mapping to the same MAC multicast address at the frame layer)?

First, it is important to realize that all IPv4 multicast addresses begin with the same 4 bits (1110), so there are really only 4 bits of concern, not 8. A LAN must not drop the last bits of the IPv4 address because these are almost guaranteed to be host bits, depending on the subnet mask. But the high-order bits, the leftmost address bits, are almost always network bits, and there is only one LAN (for now).

One other bit of the remaining 24 MAC address bits is reserved (an initial 0 indicates an Internet multicast address), so the 5 bits following the initial 1110 in the IPv4 address are dropped. The 23 remaining bits are mapped, one for one, into the last 23 bits of the MAC address. An example of this process is shown in [Figure 81 on page 3644](#).

**Figure 81: Converting MAC Addresses to Multicast Addresses**



Note that this process means that there are  $32 (2^5)$  IPv4 multicast addresses that could map to the same MAC multicast addresses. For example, multicast IPv4 addresses 224.8.7.6 and 229.136.7.6 translate to the same MAC address (0x01-00-5E-08-07-06). This is a real concern, and because the host could be interested in frames sent to both of the those multicast groups, the IP software must reject one or the other.



**NOTE:** This “collision” problem does not exist in IPv6 because of the way IPv6 handles multicast groups, but it is always a concern in IPv4. The procedure for placing IPv6 multicast packets inside multicast frames is nearly identical to that for IPv4, except for the MAC destination address 0x3333 prefix (and the lack of “collisions”).

Once the MAC address for the multicast group is determined, the host's operating system essentially orders the LAN interface card to join or leave the multicast group. Once joined to a multicast group, the host accepts frames sent to the multicast address as well as the host's unicast address and ignores other multicast group's frames. It is possible for a host to join and receive multicast content from more than one group at the same time, of course.

### ***Multicast Interface Lists***

To avoid multicast routing loops, every multicast routing device must always be aware of the interface that leads to the source of that multicast group content by the shortest path. This is the upstream (incoming) interface, and packets are never to be forwarded back toward a multicast source. All other interfaces are potential downstream (outgoing) interfaces, depending on the number of branches on the distribution tree.

routing devices closely monitor the status of the incoming and outgoing interfaces, a process that determines the *multicast forwarding state*. A routing device with a multicast forwarding state for a particular multicast group is essentially “turned on” for that group's content. Interfaces on the routing device's outgoing interface list send copies of the group's packets received on the incoming interface list for that group. The incoming and outgoing interface lists might be different for different multicast groups.

The multicast forwarding state in a routing device is usually written in either (S,G) or (\*,G) notation. These are pronounced “ess comma gee” and “star comma gee,” respectively. In (S,G), the S refers to the unicast IP address of the source for the multicast traffic, and the G refers to the particular multicast group IP address for which S is the source. All multicast packets sent from this source have S as the source address and G as the destination address.

The asterisk (\*) in the (\*,G) notation is a wildcard indicating that the state applies to any multicast application source sending to group G. So, if two sources are originating exactly the same content for multicast group 224.1.1.2, a routing device could use (\*,224.1.1.2) to represent the state of a routing device forwarding traffic from both sources to the group.

### ***Multicast Routing Protocols***

Multicast routing protocols enable a collection of multicast routing devices to build (join) distribution trees when a host on a directly attached subnet, typically a LAN, wants to

receive traffic from a certain multicast group, prune branches, locate sources and groups, and prevent routing loops.

There are several multicast routing protocols:

- **Distance Vector Multicast Routing Protocol (DVMRP)**—The first of the multicast routing protocols and hampered by a number of limitations that make this method unattractive for large-scale Internet use. DVMRP is a dense-mode-only protocol, and uses the flood-and-prune or implicit join method to deliver traffic everywhere and then determine where the uninterested receivers are. DVMRP uses source-based distribution trees in the form (S,G), and builds its own multicast routing tables for RPF checks.
- **Multicast OSPF (MOSPF)**—Extends OSPF for multicast use, but only for dense mode. However, MOSPF has an explicit join message, so routing devices do not have to flood their entire domain with multicast traffic from every source. MOSPF uses source-based distribution trees in the form (S,G).
- **Bidirectional PIM mode**—A variation of PIM. Bidirectional PIM builds bidirectional shared trees that are rooted at a rendezvous point (RP) address. Bidirectional traffic does not switch to shortest path trees as in PIM-SM and is therefore optimized for routing state size instead of path length. This means that the end-to-end latency might be longer compared to PIM sparse mode. Bidirectional PIM routes are always wildcard-source (\*G) routes. The protocol eliminates the need for (S,G) routes and data-triggered events. The bidirectional (\*G) group trees carry traffic both upstream from senders toward the RP, and downstream from the RP to receivers. As a consequence, the strict reverse path forwarding (RPF)-based rules found in other PIM modes do not apply to bidirectional PIM. Instead, bidirectional PIM (\*G) routes forward traffic from all sources and the RP. Bidirectional PIM routing devices must have the ability to accept traffic on many potential incoming interfaces. Bidirectional PIM scales well because it needs no source-specific (S,G) state. Bidirectional PIM is recommended in deployments with many dispersed sources and many dispersed receivers.
- **PIM dense mode**—In this mode of PIM, the assumption is that almost all possible subnets have at least one receiver wanting to receive the multicast traffic from a source, so the network is *flooded* with traffic on all possible branches, then pruned back when branches do not express an interest in receiving the packets, explicitly (by message) or implicitly (time-out silence). This is the *dense mode* of multicast operation. LANs are appropriate networks for dense-mode operation. Some multicast routing protocols, especially older ones, support only dense-mode operation, which makes them inappropriate for use on the Internet. In contrast to DVMRP and MOSPF, PIM dense mode allows a routing device to use any unicast routing protocol and performs RPF checks using the unicast routing table. PIM dense mode has an implicit join message, so routing devices use the flood-and-prune method to deliver traffic everywhere and then determine where the uninterested receivers are. PIM dense mode uses source-based distribution trees in the form (S,G), as do all dense-mode protocols. PIM also supports sparse-dense mode, with mixed sparse and dense groups, but there is no special notation for that operational mode. If *sparse-dense mode* is supported, the multicast routing protocol allows some multicast groups to be sparse and other groups to be dense.
- **PIM sparse mode**—In this mode of PIM, the assumption is that very few of the possible receivers want packets from each source, so the network establishes and sends packets

only on branches that have at least one leaf indicating (by message) an interest in the traffic. This multicast protocol allows a routing device to use any unicast routing protocol and performs reverse-path forwarding (RPF) checks using the unicast routing table. PIM sparse mode has an *explicit* join message, so routing devices determine where the interested receivers are and send join messages upstream to their neighbors, building trees from receivers to the rendezvous point (RP). PIM sparse mode uses an RP routing device as the initial source of multicast group traffic and therefore builds distribution trees in the form (\*G), as do all sparse-mode protocols. PIM sparse mode migrates to an (S,G) source-based tree if that path is shorter than through the RP for a particular multicast group's traffic. WANs are appropriate networks for sparse-mode operation, and indeed a common multicast guideline is not to run dense mode on a WAN under any circumstances.

- Core Based Trees (CBT)—Shares all of the characteristics of PIM sparse mode (sparse mode, explicit join, and shared (\*G) trees), but is said to be more efficient at finding sources than PIM sparse mode. CBT is rarely encountered outside academic discussions. There are no large-scale deployments of CBT, commercial or otherwise.
- PIM source-specific multicast (SSM)—Enhancement to PIM sparse mode that allows a client to receive multicast traffic directly from the source, without the help of an RP. Used with IGMPv3 to create a shortest-path tree between receiver and source.
- IGMPv1—The original protocol defined in RFC 1112, *Host Extensions for IP Multicasting*. IGMPv1 sends an explicit join message to the routing device, but uses a timeout to determine when hosts leave a group. Three versions of the Internet Group Management Protocol (IGMP) run between receiver hosts and routing devices.
- IGMPv2—Defined in RFC 2236, *Internet Group Management Protocol, Version 2*. Among other features, IGMPv2 adds an explicit leave message to the join message.
- IGMPv3—Defined in RFC 3376, *Internet Group Management Protocol, Version 3*. Among other features, IGMPv3 optimizes support for a single source of content for a multicast group, or source-specific multicast (SSM). Used with PIM SSM to create a shortest-path tree between receiver and source.
- Bootstrap Router (BSR) and Auto-Rendezvous Point (RP)—Allow sparse-mode routing protocols to find RPs within the routing domain (autonomous system, or AS). RP addresses can also be statically configured.
- Multicast Source Discovery Protocol (MSDP)—Allows groups located in one multicast routing domain to find RPs in other routing domains. MSDP is not used on an RP if all receivers and sources are located in the same routing domain. Typically runs on the same routing device as PIM sparse mode RP. Not appropriate if all receivers and sources are located in the same routing domain.
- Session Announcement Protocol (SAP) and Session Description Protocol (SDP)—Display multicast session names and correlate the names with multicast traffic. SDP is a session directory protocol that advertises multimedia conference sessions and communicates setup information to participants who want to join the session. A client commonly uses SDP to announce a conference session by periodically

multicasting an announcement packet to a well-known multicast address and port using SAP.

- **Pragmatic General Multicast (PGM)**—Special protocol layer for multicast traffic that can be used between the IP layer and the multicast application to add reliability to multicast traffic. PGM allows a receiver to detect missing information in all cases and request replacement information if the receiver application requires it.

The differences among the multicast routing protocols are summarized in [Table 275 on page 3648](#).

**Table 275: Multicast Routing Protocols Compared**

| Multicast Routing Protocol | Dense Mode | Sparse Mode | Implicit Join | Explicit Join | (S,G) SBT  | (*G) Shared Tree |
|----------------------------|------------|-------------|---------------|---------------|------------|------------------|
| DVMRP                      | Yes        | No          | Yes           | No            | Yes        | No               |
| MOSPF                      | Yes        | No          | No            | Yes           | Yes        | No               |
| PIM dense mode             | Yes        | No          | Yes           | No            | Yes        | No               |
| PIM sparse mode            | No         | Yes         | No            | Yes           | Yes, maybe | Yes, initially   |
| Bidirectional PIM          | No         | No          | No            | Yes           | No         | Yes              |
| CBT                        | No         | Yes         | No            | Yes           | No         | Yes              |
| SSM                        | No         | Yes         | No            | Yes           | Yes, maybe | Yes, initially   |
| IGMPv1                     | No         | Yes         | No            | Yes           | Yes, maybe | Yes, initially   |
| IGMPv2                     | No         | Yes         | No            | Yes           | Yes, maybe | Yes, initially   |
| IGMPv3                     | No         | Yes         | No            | Yes           | Yes, maybe | Yes, initially   |
| BSR and Auto-RP            | No         | Yes         | No            | Yes           | Yes, maybe | Yes, initially   |
| MSDP                       | No         | Yes         | No            | Yes           | Yes, maybe | Yes, initially   |

It is important to realize that retransmissions due to a high bit-error rate on a link or overloaded routing device can make multicast as inefficient as repeated unicast. Therefore, there is a trade-off in many multicast applications regarding the session support provided by the Transmission Control Protocol (TCP) (but TCP always resends missing segments), or the simple drop-and-continue strategy of the User Datagram Protocol (UDP) datagram service (but reordering can become an issue). Modern multicast uses UDP almost exclusively.

#### ***T Series Router Multicast Performance***

The Juniper Networks T Series Core Routers handle extreme multicast packet replication requirements with a minimum of router load. Each memory component replicates a

multicast packet twice at most. Even in the worst-case scenario involving maximum fan-out, when 1 input port and 63 output ports need a copy of the packet, the T Series routing platform copies a multicast packet only six times. Most multicast distribution trees are much sparser, so in many cases only two or three replications are necessary. In no case does the T Series architecture have an impact on multicast performance, even with the largest multicast fan-out requirements.

## Multicast Protocols Overview

- [Supported IP Multicast Protocol Standards on page 3649](#)
- [Understanding MLD on page 3650](#)
- [PIM Overview on page 3653](#)

### Supported IP Multicast Protocol Standards

Junos OS substantially supports the following RFCs and Internet drafts, which define standards for IP multicast protocols, including the Distance Vector Multicast Routing Protocol (DVMRP), Internet Group Management Protocol (IGMP), Multicast Listener Discovery (MLD), Multicast Source Discovery Protocol (MSDP), Pragmatic General Multicast (PGM), Protocol Independent Multicast (PIM), Session Announcement Protocol (SAP), and Session Description Protocol (SDP).

- RFC 1112, *Host Extensions for IP Multicasting* (defines IGMP Version 1)
- RFC 2236, *Internet Group Management Protocol, Version 2*
- RFC 2327, *SDP: Session Description Protocol*
- RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*
- RFC 2858, *Multiprotocol Extensions for BGP-4*
- RFC 3031, *Multiprotocol Label Switching Architecture*
- RFC 3376, *Internet Group Management Protocol, Version 3*
- RFC 3590, *Source Address Selection for the Multicast Listener Discovery (MLD) Protocol*
- RFC 4601, *Protocol Independent Multicast – Sparse Mode (PIM-SM): Protocol Specification (Revised)*
- RFC 4607, *Source-Specific Multicast for IP*
- RFC 5015, *Bidirectional Protocol Independent Multicast (BIDIR-PIM)*
- *Using IGMPv3 and MLDv2 for Source-Specific Multicast*
- Internet draft draft-ietf-l3vpn-2547bis-mcast-10.txt, *Multicast in MPLS/BGP IP VPNs*
- Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp-08.txt, *BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs*
- Internet draft draft-ietf-pim-sm-bsr-05.txt, *Bootstrap Router (BSR) Mechanism for PIM*

The scoping mechanism is not supported.

- Internet draft draft-raggarwa-l3vpn-2547-mvpn-00.txt, *Base Specification for Multicast in BGP/MPLS VPNs* (expires December 2004)

The following RFCs and Internet drafts do not define standards, but provide information about multicast protocols and related technologies. The IETF classifies them variously as “Best Current Practice,” “Experimental,” or “Informational.”

- RFC 1075, *Distance Vector Multicast Routing Protocol*
- RFC 2362, *Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification*
- RFC 2365, *Administratively Scoped IP Multicast*
- RFC 2547, *BGP/MPLS VPNs*
- RFC 2974, *Session Announcement Protocol*
- RFC 3208, *PGM Reliable Transport Protocol Specification*
- RFC 3446, *Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*
- RFC 3569, *An Overview of Source-Specific Multicast (SSM)*
- RFC 3618, *Multicast Source Discovery Protocol (MSDP)*
- RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*
- RFC 3973, *Protocol Independent Multicast – Dense Mode (PIM-DM): Protocol Specification (Revised)*
- RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*
- Internet draft draft-ietf-idmr-dvmrp-v3-11.txt, *Distance Vector Multicast Routing Protocol*
- Internet draft draft-ietf-mboned-ssm232-08.txt, *Source-Specific Protocol Independent Multicast in 232/8*
- Internet draft draft-ietf-mmusic-sap-00.txt, *SAP: Session Announcement Protocol*
- Internet draft draft-rosen-vpn-mcast-07.txt, *Multicast in MPLS/BGP VPNs*

Only section 7, “Data MDT: Optimizing flooding,” is supported.

#### **Related Documentation**

- *Accessing Standards Documents on the Internet*

---

### **Understanding MLD**

The Multicast Listener Discovery (MLD) Protocol manages the membership of hosts and routers in multicast groups. IP version 6 (IPv6) multicast routers use MLD to learn, for each of their attached physical networks, which groups have interested listeners. Each routing device maintains a list of host multicast addresses that have listeners for each subnetwork, as well as a timer for each address. However, the routing device does not need to know the address of each listener—just the address of each host. The routing device provides addresses to the multicast routing protocol it uses, which ensures that multicast packets are delivered to all subnetworks where there are interested listeners.



In this way, MLD is used as the transport for the Protocol Independent Multicast (PIM) Protocol.


MLD is an integral part of IPv6 and must be enabled on all IPv6 routing devices and hosts that need to receive IP multicast traffic. The Junos OS supports MLD versions 1 and 2. Version 2 is supported for source-specific multicast (SSM) include and exclude modes.

In include mode, the receiver specifies the source or sources it is interested in receiving the multicast group traffic from. Exclude mode works the opposite of include mode. It allows the receiver to specify the source or sources it is not interested in receiving the multicast group traffic from.

For each attached network, a multicast routing device can be either a querier or a nonquerier. A querier routing device, usually one per subnet, solicits group membership information by transmitting MLD queries. When a host reports to the querier routing device that it has interested listeners, the querier routing device forwards the membership information to the rendezvous point (RP) routing device by means of the receiver's (host's) designated router (DR). This builds the rendezvous-point tree (RPT) connecting the host with interested listeners to the RP routing device. The RPT is the initial path used by the sender to transmit information to the interested listeners. Nonquerier routing devices do not transmit MLD queries on a subnet but can do so if the querier routing device fails.

All MLD-configured routing devices start as querier routing devices on each attached subnet (see [Figure 82 on page 3651](#)). The querier routing device on the right is the receiver's DR.

**Figure 82: Routing Devices Start Up on a Subnet**




To elect the querier routing device, the routing devices exchange query messages containing their IPv6 source addresses. If a routing device hears a query message whose IPv6 source address is numerically lower than its own selected address, it becomes a nonquerier. In [Figure 83 on page 3652](#), the routing device on the left has a source address numerically lower than the one on the right and therefore becomes the querier routing device.




**NOTE:** In the practical application of MLD, several routing devices on a subnet are nonqueriers. If the elected querier routing device fails, query messages are exchanged among the remaining routing devices. The routing device with the lowest IPv6 source address becomes the new querier routing device. The IPv6 Neighbor Discovery Protocol (NDP) implementation drops incoming Neighbor Announcement (NA) messages that have a broadcast or multicast address in the target link-layer address option. This behavior is recommended by RFC 2461.

**Figure 83: Querier Routing Device Is Determined**




The querier routing device sends general MLD queries on the **link-scope all-nodes** multicast address FF02::1 at short intervals to all attached subnets to solicit group membership information (see [Figure 84 on page 3652](#)). Within the query message is the *maximum response delay* value, specifying the maximum allowed delay for the host to respond with a report message.

**Figure 84: General Query Message Is Issued**




If interested listeners are attached to the host receiving the query, the host sends a report containing the host's IPv6 address to the routing device (see [Figure 85 on page 3652](#)). If the reported address is not yet in the routing device's list of multicast addresses with interested listeners, the address is added to the list and a timer is set for the address. If the address is already on the list, the timer is reset. The host's address is transmitted to the RP in the PIM domain.

**Figure 85: Reports Are Received by the Querier Routing Device**




If the host has no interested multicast listeners, it sends a done message to the querier routing device. On receipt, the querier routing device issues a multicast address-specific query containing the last **listener query interval** value to the multicast address of the host. If the routing device does not receive a report from the multicast address, it removes the multicast address from the list and notifies the RP in the PIM domain of its removal (see [Figure 86 on page 3653](#)).

**Figure 86: Host Has No Interested Receivers and Sends a Done Message to Routing Device**



If a done message is not received by the querier routing device, the querier routing device continues to send multicast address-specific queries. If the timer set for the address on receipt of the last report expires, the querier routing device assumes there are no longer interested listeners on that subnet, removes the multicast address from the list, and notifies the RP in the PIM domain of its removal (see [Figure 87 on page 3653](#)).

**Figure 87: Host Address Timer Expires and Address Is Removed from Multicast Address List**



## PIM Overview

The predominant multicast routing protocol in use on the Internet today is Protocol Independent Multicast, or PIM. The type of PIM used on the Internet is PIM sparse mode. PIM sparse mode is so accepted that when the simple term “PIM” is used in an Internet context, some form of sparse mode operation is assumed.

PIM emerged as an algorithm to overcome the limitations of dense-mode protocols such as the Distance Vector Multicast Routing Protocol (DVMRP), which was efficient for dense clusters of multicast receivers, but did not scale well for the larger, sparser, groups encountered on the Internet. The Core Based Trees (CBT) Protocol was intended to support sparse mode as well, but CBT, with its all-powerful core approach, made placement of the core critical, and large conference-type applications (many-to-many) resulted in bottlenecks in the core. PIM was designed to avoid the dense-mode scaling issues of DVMRP and the potential performance issues of CBT at the same time.

PIM is one of the most rapidly evolving specifications on the Internet today. Since its introduction in 1995, PIM has already seen two major revisions to its packet structure (PIM version 1 [PIMv1] and PIM version 2 [PIMv2]), two major RFCs (RFC 2362 obsoleted

RFC 2117), and numerous drafts describing major components of PIM, such as many-to-many trees and source-specific multicast (SSM). Long-lasting RFCs are not a feature of PIM, and virtually all of PIM must be researched, understood, and implemented directly from Internet drafts. In fact, no current RFC describes PIMv1 at all. The drafts have all expired, and PIMv1 was never issued as an official RFC.

PIM itself is not nonstandard or unstable, however. PIM has been a promising multicast routing protocol since its inception, especially PIM sparse mode, the first real sparse-mode multicast routing protocol. Work continues on PIM in a number of areas, from bidirectional trees to network management, and the rapid pace of development makes drafts essential for PIM.

PIMv1 and PIMv2 can coexist on the same routing device and even on the same interface. The main difference between PIMv1 and PIMv2 is the packet format. PIMv1 messages use Internet Group Management Protocol (IGMP) packets, whereas PIMv2 has its own IP protocol number (103) and packet structure. All routing devices connecting to an IP subnet such as a LAN must use the same PIM version. Some PIM implementations can recognize PIMv1 packets and automatically switch the routing device interface to PIMv1. Because the difference between PIMv1 and PIMv2 involves the message format, but not the meaning of the message or how the routing device processes the PIM message, a routing device can easily mix PIMv1 and PIMv2 interfaces.

PIM is used for efficient routing to multicast groups that might span wide-area and interdomain internetworks. It is called “protocol independent” because it does not depend on a particular unicast routing protocol. Junos OS supports bidirectional mode, sparse mode, dense mode, and sparse-dense mode.

PIM operates in several modes: bidirectional mode, sparse mode, dense mode, and sparse-dense mode. In sparse-dense mode, some multicast groups are configured as dense mode (flood-and-prune, [S,G] state) and others are configured as sparse mode (explicit join to rendezvous point [RP], [\*G] state).

PIM drafts also establish a mode known as PIM source-specific mode, or PIM SSM. In PIM SSM there is only one specific source for the content of a multicast group within a given domain.

Because the PIM mode you choose determines the PIM configuration properties, you first must decide whether PIM operates in bidirectional, sparse, dense, or sparse-dense mode in your network. Each mode has distinct operating advantages in different network environments.

- In sparse mode, routing devices must join and leave multicast groups explicitly. Upstream routing devices do not forward multicast traffic to a downstream routing device unless the downstream routing device has sent an explicit request (by means of a join message) to the rendezvous point (RP) routing device to receive this traffic. The RP serves as the root of the shared multicast delivery tree and is responsible for forwarding multicast data from different sources to the receivers.

Sparse mode is well suited to the Internet, where frequent interdomain join messages and prune messages are common.

- Bidirectional PIM is similar to sparse mode, and is especially suited to applications that must scale to support a large number of dispersed sources and receivers. In bidirectional PIM, routing devices build shared bidirectional trees and do not switch to a source-based tree. Bidirectional PIM scales well because it needs no source-specific (S,G) state. Instead, it builds only group-specific (\*,G) state.
- Unlike sparse mode and bidirectional mode, in which data is forwarded only to routing devices sending an explicit PIM join request, dense mode implements a *flood-and-prune* mechanism, similar to the Distance Vector Multicast Routing Protocol (DVMRP). In dense mode, a routing device receives the multicast data on the incoming interface, then forwards the traffic to the outgoing interface list. Flooding occurs periodically and is used to refresh state information, such as the source IP address and multicast group pair. If the routing device has no interested receivers for the data, and the outgoing interface list becomes empty, the routing device sends a PIM prune message upstream.

Dense mode works best in networks where few or no prunes occur. In such instances, dense mode is actually more efficient than sparse mode.

- Sparse-dense mode, as the name implies, allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as “dense” is not mapped to an RP. Instead, data packets destined for that group are forwarded by means of PIM dense mode rules. A group specified as “sparse” is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules. Sparse-dense mode is useful in networks implementing auto-RP for PIM sparse mode.

### **Basic PIM Network Components**

PIM dense mode requires only a multicast source and series of multicast-enabled routing devices running PIM dense mode to allow receivers to obtain multicast content. Dense mode makes sure that all multicast traffic gets everywhere by periodically flooding the network with multicast traffic, and relies on prune messages to make sure that subnets where all receivers are uninterested in that particular multicast group stop receiving packets.

PIM sparse mode is more complicated and requires the establishment of special routing devices called *rendezvous points (RPs)* in the network core. These routing devices are where upstream join messages from interested receivers meet downstream traffic from

the source of the multicast group content. A network can have many RPs, but PIM sparse mode allows only one RP to be active for any multicast group.

If there is only one RP in a routing domain, the RP and adjacent links might become congested and form a single point of failure for all multicast traffic. Thus, multiple RPs are the rule, but the issue then becomes how other multicast routing devices find the RP that is the source of the multicast group the receiver is trying to join. This RP-to-group mapping is controlled by a special *bootstrap router (BSR)* running the PIM BSR mechanism. There can be more than one bootstrap router as well, also for single-point-of-failure reasons.

The bootstrap router does not have to be an RP itself, although this is a common implementation. The bootstrap router's main function is to manage the collection of RPs and allow interested receivers to find the source of their group's multicast traffic. PIM bootstrap messages are sourced from the loopback address, which is always up. The loopback address must be routable. If it is not routable, then the bootstrap router is unable to send bootstrap messages to update the RP domain members. The **show pim bootstrap** command displays only those bootstrap routers that have routable loopback addresses.

PIM SSM can be seen as a subset of a special case of PIM sparse mode and requires no specialized equipment other than that used for PIM sparse mode (and IGMP version 3).

Bidirectional PIM RPs, unlike RPs for PIM sparse mode, do not need to perform PIM Register tunneling or other specific protocol action. Bidirectional PIM RPs implement no specific functionality. RP addresses are simply a location in the network to rendezvous toward. In fact, for bidirectional PIM, RP addresses need not be loopback interface addresses or even be addresses configured on any routing device, as long as they are covered by a subnet that is connected to a bidirectional PIM-capable routing device and advertised to the network.

**Related Documentation** • [Supported IP Multicast Protocol Standards on page 3649](#) in the *Multicast Protocols Configuration Guide*

---

## Configuration

- [Configuration on page 3656](#)
- [Configuration Statements: IGMP on page 3757](#)
- [Configuration Statements: IGMP Snooping on page 3782](#)
- [Configuration Statements: MLD on page 3799](#)
- [Configuration Statements: MSDP on page 3821](#)
- [Configuration Statements: PIM on page 3846](#)

## Configuration

- [Configuring IGMP on page 3657](#)
- [Configuring Multiple Instances of MSDP on page 3679](#)
- [Configuring Basic PIM Settings on page 3680](#)

- [Configuring Multiple Instances of PIM on page 3692](#)
- [Configuring a Designated Router for PIM on page 3692](#)
- [Configuring Static RP on page 3694](#)
- [Configuring PIM Bootstrap Router on page 3701](#)
- [Configuring PIM Auto-RP on page 3705](#)
- [Configuring Embedded RP on page 3709](#)
- [Configuring PIM Filtering on page 3712](#)
- [Configuring PIM and the Bidirectional Forwarding Detection \(BFD\) Protocol on page 3725](#)
- [Configuring PIM Dense Mode on page 3737](#)
- [Configuring PIM Sparse-Dense Mode on page 3740](#)
- [PIM Join Load Balancing on Multipath MVPN Routes Overview on page 3741](#)
- [PIM Snooping for VPLS on page 3745](#)

---

### Configuring IGMP

- [Understanding Group Membership Protocols on page 3657](#)
- [Understanding IGMP on page 3658](#)
- [Configuring IGMP on page 3660](#)
- [Enabling IGMP on page 3661](#)
- [Modifying the IGMP Host-Query Message Interval on page 3662](#)
- [Modifying the IGMP Query Response Interval on page 3663](#)
- [Specifying Immediate-Leave Host Removal for IGMP on page 3664](#)
- [Filtering Unwanted IGMP Reports at the IGMP Interface Level on page 3664](#)
- [Accepting IGMP Messages from Remote Subnetworks on page 3665](#)
- [Modifying the IGMP Last-Member Query Interval on page 3666](#)
- [Modifying the IGMP Robustness Variable on page 3667](#)
- [Limiting the Maximum IGMP Message Rate on page 3668](#)
- [Changing the IGMP Version on page 3668](#)
- [Enabling IGMP Static Group Membership on page 3668](#)
- [Recording IGMP Join and Leave Events on page 3674](#)
- [Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces on page 3676](#)
- [Tracing IGMP Protocol Traffic on page 3677](#)
- [Disabling IGMP on page 3679](#)
- [IGMP and Nonstop Active Routing on page 3679](#)

### ***Understanding Group Membership Protocols***

There is a big difference between the multicast protocols used between host and routing device and between the multicast routing devices themselves. Hosts on a given subnetwork need to inform their routing device only whether or not they are interested in receiving packets from a certain multicast group. The source host needs to inform its

routing devices only that it is the source of traffic for a particular multicast group. In other words, no detailed knowledge of the distribution tree is needed by any hosts; only a group membership protocol is needed to inform routing devices of their participation in a multicast group. Between adjacent routing devices, on the other hand, the multicast routing protocols must avoid loops as they build a detailed sense of the network topology and distribution tree from source to leaf. So, different multicast protocols are used for the host-router portion and the router-router portion of the multicast network.

Multicast group membership protocols enable a routing device to detect when a host on a directly attached subnet, typically a LAN, wants to receive traffic from a certain multicast group. Even if more than one host on the LAN wants to receive traffic for that multicast group, the routing device sends only one copy of each packet for that multicast group out on that interface, because of the inherent broadcast nature of LANs. When the multicast group membership protocol informs the routing device that there are no interested hosts on the subnet, the packets are withheld and that leaf is pruned from the distribution tree.

The Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) Protocol are the standard IP multicast group membership protocols: IGMP and MLD have several versions that are supported by hosts and routing devices:

- IGMPv1—The original protocol defined in RFC 1112. An explicit join message is sent to the routing device, but a timeout is used to determine when hosts leave a group. This process wastes processing cycles on the routing device, especially on older or smaller routing devices.
- IGMPv2—Defined in RFC 2236. Among other features, IGMPv2 adds an explicit leave message to the join message so that routing devices can more easily determine when a group has no interested listeners on a LAN.
- IGMPv3—Defined in RFC 3376. Among other features, IGMPv3 optimizes support for a single source of content for a multicast group, or *source-specific multicast (SSM)*.
- MLDv1—Defined in RFC 2710. MLDv1 is similar to IGMPv2.
- MLDv2—Defined in RFC 3810. MLDv2 similar to IGMPv3.

The various versions of IGMP and MLD are backward compatible. It is common for a routing device to run multiple versions of IGMP and MLD on LAN interfaces. Backward compatibility is achieved by dropping back to the most basic of all versions run on a LAN. For example, if one host is running IGMPv1, any routing device attached to the LAN running IGMPv2 can drop back to IGMPv1 operation, effectively eliminating the IGMPv2 advantages. Running multiple IGMP versions ensures that both IGMPv1 and IGMPv2 hosts find peers for their versions on the routing device.

### ***Understanding IGMP***

The Internet Group Management Protocol (IGMP) manages the membership of hosts and routing devices in multicast groups. IP hosts use IGMP to report their multicast group memberships to any immediately neighboring multicast routing devices. Multicast routing devices use IGMP to learn, for each of their attached physical networks, which groups have members.



IGMP is also used as the transport for several related multicast protocols (for example, Distance Vector Multicast Routing Protocol [DVMRP] and Protocol Independent Multicast version 1 [PIMv1]).

A routing device receives explicit join and prune messages from those neighboring routing devices that have downstream group members. When PIM is the multicast protocol in use, IGMP begins the process as follows:

1. To join a multicast group, G, a host conveys its membership information through IGMP.
2. The routing device then forwards data packets addressed to a multicast group G to only those interfaces on which explicit join messages have been received.
3. A designated router (DR) sends periodic join and prune messages toward a group-specific rendezvous point (RP) for each group for which it has active members. One or more routing devices are automatically or statically designated as the RP, and all routing devices must explicitly join through the RP.
4. Each routing device along the path toward the RP builds a wild card (any-source) state for the group and sends join and prune messages toward the RP.

The term *route entry* is used to refer to the state maintained in a routing device to represent the distribution tree.

A route entry can include such fields as:

- source address
- group address
- incoming interface from which packets are accepted
- list of outgoing interfaces to which packets are sent
- timers
- flag bits

The wild card route entry's incoming interface points toward the RP.

The outgoing interfaces point to the neighboring downstream routing devices that have sent join and prune messages toward the RP as well as the directly connected hosts that have requested membership to group G.

5. This state creates a shared, RP-centered, distribution tree that reaches all group members.

IGMP is an integral part of IP and must be enabled on all routing devices and hosts that need to receive IP multicast traffic.

For each attached network, a multicast routing device can be either a querier or a nonquerier. The querier routing device periodically sends general query messages to solicit group membership information. Hosts on the network that are members of a multicast group send report messages. When a host leaves a group, it sends a leave group message.

IGMP version 3 (IGMPv3) supports inclusion and exclusion lists. Inclusion lists enable you to specify which sources can send to a multicast group. This type of multicast group is called a source-specific multicast (SSM) group, and its multicast address is 232/8.

IGMPv3 provides support for source filtering. For example, a routing device can specify particular routing devices from which it accepts or rejects traffic. With IGMPv3, a multicast routing device can learn which sources are of interest to neighboring routing devices.

Exclusion mode works the opposite of an inclusion list. It allows any source but the ones listed to send to the SSM group.

IGMPv3 interoperates with versions 1 and 2 of the protocol. However, to remain compatible with older IGMP hosts and routing devices, IGMPv3 routing devices must also implement versions 1 and 2 of the protocol. IGMPv3 supports the following membership-report record types: mode is allowed, allow new sources, and block old sources.

### **Configuring IGMP**

Before you begin:

1. Determine whether the routing device is directly attached to any multicast sources. Receivers must be able to locate these sources.
2. Determine whether the routing device is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
3. Determine whether to configure multicast to use sparse, dense, or sparse-dense mode. Each mode has different configuration considerations.
4. Determine the address of the RP if sparse or sparse-dense mode is used.
5. Determine whether to locate the RP with the static configuration, BSR, or auto-RP method.
6. Determine whether to configure multicast to use its own RPF routing table when configuring PIM in sparse, dense, or sparse-dense mode.
7. Configure the SAP and SDP protocols to listen for multicast session announcements. See *Configuring the Session Announcement Protocol*.

To configure the Internet Group Management Protocol (IGMP), include the **igmp** statement:

```
igmp {
 accounting;
 interface interface-name {
 disable;
 (accounting | no-accounting);
 group-policy [policy-names];
 immediate-leave;
 oif-map map-name;
 promiscuous-mode;
 ssm-map ssm-map-name;
 static {
 group multicast-group-address {
 exclude;
```

```

 group-count number;
 group-increment increment;
 source ip-address {
 source-count number;
 source-increment increment;
 }
}
}
version version;
}
query-interval seconds;
query-last-member-interval seconds;
query-response-interval seconds;
robust-count number;
traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
}
}

```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

By default, IGMP is enabled on all interfaces on which you configure Protocol Independent Multicast (PIM), and on all broadcast interfaces on which you configure the Distance Vector Multicast Routing Protocol (DVMRP).



**NOTE:** You can configure IGMP on an interface without configuring PIM. PIM is generally not needed on IGMP downstream interfaces. Therefore, only one “pseudo PIM interface” is created to represent all IGMP downstream (IGMP-only) interfaces on the routing device. This reduces the amount of routing device resources, such as memory, that are consumed. You must configure PIM on upstream IGMP interfaces to enable multicast routing, perform reverse-path forwarding for multicast data packets, populate the multicast forwarding table for upstream interfaces, and in the case of bidirectional PIM and PIM sparse mode, to distribute IGMP group memberships into the multicast routing domain.

### Enabling IGMP

The Internet Group Management Protocol (IGMP) manages multicast groups by establishing, maintaining, and removing groups on a subnet. Multicast routing devices use IGMP to learn which groups have members on each of their attached physical networks. IGMP must be enabled for the routing device to receive IPv4 multicast packets. IGMP is only needed for IPv4 networks, because multicast is handled differently in IPv6 networks. IGMP is automatically enabled on all IPv4 interfaces on which you configure PIM and on all IPv4 broadcast interfaces when you configure DVMRP.

If IGMP is not running on an interface—either because PIM and DVMRP are not configured on the interface or because IGMP is explicitly disabled on the interface—you can explicitly enable IGMP.

To explicitly enable IGMP:

1. If PIM and DVMRP are not running on the interface, explicitly enable IGMP by including the interface name.

```
[edit protocols igmp]
user@host# set interface fe-0/0/0.0
```

2. See if IGMP is disabled on any interfaces. In the following example, IGMP is disabled on a Gigabit Ethernet interface.

```
[edit protocols igmp]
user@host# show
interface fe-0/0/0.0;
interface ge-1/0/0.0 {
 disable;
}
```

3. Enable IGMP on the interface by deleting the **disable** statement.

```
[edit protocols igmp]
delete interface ge-1/0/0.0 disable
```

4. Verify the configuration.

```
[edit protocols igmp]
user@host# show
interface fe-0/0/0.0;
interface ge-1/0/0.0;
```

5. Verify the operation of IGMP on the interfaces by checking the output of the **show igmp interface** command.

### ***Modifying the IGMP Host-Query Message Interval***

The objective of IGMP is to keep routing devices up to date with group membership of the entire subnet. Routing devices need not know who all the members are, only that members exist. Each host keeps track of which multicast groups are subscribed to. On each link, one routing device is elected the querier. The IGMP querier routing device periodically sends general host-query messages on each attached network to solicit membership information. The messages are sent to the all-systems multicast group address, 224.0.0.1.

The query interval, the response interval, and the robustness variable are related in that they are all variables that are used to calculate the group membership timeout. The group membership timeout is the number of seconds that must pass before a multicast routing device determines that no more members of a host group exist on a subnet. The group membership timeout is calculated as the (robustness variable x query-interval) + (query-response-interval). If no reports are received for a particular group before the group membership timeout has expired, the routing device stops forwarding remotely-originated multicast packets for that group onto the attached network.

By default, host-query messages are sent every 125 seconds. You can change this interval to change the number of IGMP messages sent on the subnet.

To modify the query interval:

1. Configure the interval.

```
[edit protocols igmp]
user@host# set query-interval 200
```

The value can be from 1 through 1024 seconds.

2. Verify the configuration by checking the IGMP Query Interval field in the output of the **show igmp interface** command.
3. Verify the operation of the query interval by checking the Membership Query field in the output of the **show igmp statistics** command.

### *Modifying the IGMP Query Response Interval*

The query response interval is the maximum amount of time that can elapse between when the querier routing device sends a host-query message and when it receives a response from a host. Configuring this interval allows you to adjust the burst peaks of IGMP messages on the subnet. Set a larger interval to make the traffic less bursty. Bursty traffic refers to an uneven pattern of data transmission: sometimes a very high data transmission rate, whereas at other times a very low data transmission rate.

The query response interval, the host-query interval, and the robustness variable are related in that they are all variables that are used to calculate the group membership timeout. The group membership timeout is the number of seconds that must pass before a multicast routing device determines that no more members of a host group exist on a subnet. The group membership timeout is calculated as the (robustness variable x query-interval) + (query-response-interval). If no reports are received for a particular group before the group membership timeout has expired, the routing device stops forwarding remotely originated multicast packets for that group onto the attached network.

The default query response interval is 10 seconds. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify the query response interval:

1. Configure the interval.

```
[edit protocols igmp]
user@host# set query-response-interval 0.4
```

2. Verify the configuration by checking the IGMP Query Response Interval field in the output of the **show igmp interface** command.
3. Verify the operation of the query interval by checking the Membership Query field in the output of the **show igmp statistics** command.

### ***Specifying Immediate-Leave Host Removal for IGMP***

The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.

The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.

When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.

When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.



**NOTE:** Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the routing device only knows about the one interested host and does not know about the others.

To enable immediate leave on an interface:

1. Configure immediate leave on the IGMP interface.

```
[edit protocols IGMP]
user@host# set interface ge-0/0/0.1 immediate-leave
```

2. Verify the configuration by checking the Immediate Leave field in the output of the `show igmp interface` command.

### ***Filtering Unwanted IGMP Reports at the IGMP Interface Level***

Suppose you need to limit the subnets that can join a certain multicast group. The **group-policy** statement enables you to filter unwanted IGMP reports at the interface level. When this statement is enabled on a routing device running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), after the routing device receives an IGMP report, the routing device compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report if the policy matches the defined address or network).

You define the policy to match only IGMP group addresses (for IGMPv2) by using the policy's **route-filter** statement to match the group address. You define the policy to match IGMP (source, group) addresses (for IGMPv3) by using the policy's **route-filter** statement to match the group address and the policy's **source-address-filter** statement to match the source address.

To filter unwanted IGMP reports:

1. Configure an IGMPv2 policy.

```
[edit policy-statement reject_policy_v2]
user@host# set from route-filter 224.1.1.1/32 exact
user@host# set from route-filter 239.0.0.0/8 orlonger
user@host# set then reject
```

2. Configure an IGMPv3 policy.

```
[edit policy-statement reject_policy_v3]
user@host# set from route-filter 224.1.1.1/32 exact
user@host# set from route-filter 239.0.0.0/8 orlonger
user@host# set from source-address-filter 10.0.0.0/8 orlonger
user@host# set from source-address-filter 127.0.0.0/8 orlonger
user@host# set then reject
```

3. Apply the policies to the IGMP interfaces on which you prefer not to receive specific group or (source, group) reports. In this example, **ge-0/0/0.1** is running IGMPv2, and **ge-0/1/1.0** is running IGMPv3.

```
[edit protocols igmp]
user@host# set interface ge-0/0/0.1 group-policy reject_policy_v2
user@host# set interface ge-0/1/1.0 group-policy reject_policy_v3
```

4. Verify the operation of the filter by checking the Rejected Report field in the output of the **show igmp statistics** command.

### *Accepting IGMP Messages from Remote Subnetworks*

By default, IGMP interfaces accept IGMP messages only from the same subnet. Including the **promiscuous-mode** statement enables the routing device to accept IGMP messages from indirectly connected subnets.



**NOTE:** When you enable IGMP on an unnumbered Ethernet interface that uses a /32 loopback address as a donor address, you must configure IGMP promiscuous mode to accept the IGMP packets received on this interface.



**NOTE:** When enabling promiscuous-mode, all routing devices on the ethernet segment must be configured with the promiscuous mode statement. Otherwise, only the interface configured with lowest IPv4 address acts as the querier for IGMP for this Ethernet segment.

To enable IGMP promiscuous mode on an interface:

1. Configure the IGMP interface.

```
[edit protocols igmp]
user@host# set interface ge-0/1/1.0 promiscuous-mode
```

2. Verify the configuration by checking the Promiscuous Mode field in the output of the **show igmp interface** command.
3. Verify the operation of the filter by checking the Rx non-local field in the output of the **show igmp statistics** command.

### ***Modifying the IGMP Last-Member Query Interval***

The last-member query interval is the maximum amount of time between group-specific query messages, including those sent in response to leave-group messages. You can configure this interval to change the amount of time it takes a routing device to detect the loss of the last member of a group.

When the routing device that is serving as the querier receives a leave-group message from a host, the routing device sends multiple group-specific queries to the group being left. The querier sends a specific number of these queries at a specific interval. The number of queries sent is called the last-member query count. The interval at which the queries are sent is called the last-member query interval. Because both settings are configurable, you can adjust the leave latency. The IGMP leave latency is the time between a request to leave a multicast group and the receipt of the last byte of data for the multicast group.

The last-member query count x (times) the last-member query interval = (equals) the amount of time it takes a routing device to determine that the last member of a group has left the group and to stop forwarding group traffic.

The default last-member query interval is 1 second. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify this interval:

1. Configure the time (in seconds) that the routing device waits for a report in response to a group-specific query.

```
[edit protocols igmp]
user@host# set query-last-member-interval 0.1
```

2. Verify the configuration by checking the IGMP Last Member Query Interval field in the output of the **show igmp interfaces** command.



**NOTE:** You can configure the last-member query count by configuring the robustness variable. The two are always equal.

---



### *Modifying the IGMP Robustness Variable*

Fine-tune the IGMP robustness variable to allow for expected packet loss on a subnet. The robust count automatically changes certain IGMP message intervals for IGMPv2 and IGMPv3. Increasing the robust count allows for more packet loss but increases the leave latency of the subnetwork.

When the query routing device receives an IGMP leave message on a shared network running IGMPv2, the query routing device must send an IGMP group query message a specified number of times. The number of IGMP group query messages sent is determined by the robust count.

The value of the robustness variable is also used in calculating the following IGMP message intervals:

- Group member interval—Amount of time that must pass before a multicast routing device determines that there are no more members of a group on a network. This interval is calculated as follows: (robustness variable x query-interval) + (1 x query-response-interval).
- Other querier present interval—The robust count is used to calculate the amount of time that must pass before a multicast routing device determines that there is no longer another multicast routing device that is the querier. This interval is calculated as follows: (robustness variable x query-interval) + (0.5 x query-response-interval).
- Last-member query count—Number of group-specific queries sent before the routing device assumes there are no local members of a group. The number of queries is equal to the value of the robustness variable.

In IGMPv3, a change of interface state causes the system to immediately transmit a state-change report from that interface. In case the state-change report is missed by one or more multicast routing devices, it is retransmitted. The number of times it is retransmitted is the robust count minus one. In IGMPv3, the robust count is also a factor in determining the group membership interval, the older version querier interval, and the other querier present interval.

By default, the robustness variable is set to 2. You might want to increase this value if you expect a subnet to lose packets.

The number can be from 2 through 10.

To change the value of the robustness variable:

1. Configure the robust count.

When you set the robust count, you are in effect configuring the number of times the querier retries queries on the connected subnets.

```
[edit protocols igmp]
user@host# set robust-count 5
```

2. Verify the configuration by checking the IGMP Robustness Count field in the output of the **show igmp interfaces** command.

### ***Limiting the Maximum IGMP Message Rate***

This section describes how to change the limit for the maximum number of IGMP packets transmitted in 1 second by the routing device.

Increasing the maximum number of IGMP packets transmitted per second might be useful on a routing device with a large number of interfaces participating in IGMP.

To change the limit for the maximum number of IGMP packets the routing device can transmit in 1 second, include the **maximum-transmit-rate** statement and specify the maximum number of packets per second to be transmitted.

### ***Changing the IGMP Version***

By default, the routing device runs IGMPv2. Routing devices running different versions of IGMP determine the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version.

To enable source-specific multicast (SSM) functionality, you must configure version 3 on the host and the host's directly connected routing device. If a source address is specified in a multicast group that is statically configured, the version must be set to IGMPv3.

If a static multicast group is configured with the source address defined, and the IGMP version is configured to be version 2, the source is ignored and only the group is added. In this case, the join is treated as an IGMPv2 group join.

If you configure the IGMP version setting at the individual interface hierarchy level, it overrides the **interface all** statement.

If you have already configured the routing device to use IGMP version 1 (IGMPv1) and then configure it to use IGMPv2, the routing device continues to use IGMPv1 for up to 6 minutes and then uses IGMPv2.

To change to IGMPv3 for SSM functionality:

1. Configure the IGMP interface.

```
[edit protocols igmp]
user@host# set interface ge-0/0/0 version 3
```

2. Verify the configuration by checking the version field in the output of the **show igmp interfaces** command. The **show igmp statistics** command has version-specific output fields, such as V1 Membership Report, V2 Membership Report, and V3 Membership Report.

### ***Enabling IGMP Static Group Membership***

You can create IGMP static group membership to test multicast forwarding without a receiver host. When you enable IGMP static group membership, data is forwarded to an interface without that interface receiving membership reports from downstream hosts. The routing device on which you enable static IGMP group membership must be the designated router (DR) for the subnet. Otherwise, traffic does not flow downstream.

When enabling IGMP static group membership, you cannot configure multiple groups using the **group-count**, **group-increment**, **source-count**, and **source-increment** statements if the **all** option is specified as the IGMP interface.

Class-of-service (CoS) adjustment is not supported with IGMP static group membership.

In this example, you create static group 225.1.1.1.

1. On the DR, configure the static groups to be created by including the **static** statement and **group** statement and specifying which IP multicast address of the group to be created. When creating groups individually, you must specify a unique address for each group.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
 static {
 group 225.1.1.1 ;
 }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created.

```
user@host> show igmp group
Interface: fe-0/1/2
Group: 225.1.1.1
Source: 10.0.0.2
Last reported by: Local
Timeout: 0 Type: Static
```



**NOTE:** When you configure static IGMP group entries on point-to-point links that connect routing devices to a rendezvous point (RP), the static IGMP group entries do not generate join messages toward the RP.

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of static groups be automatically created. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately.

In this example, you create three groups.

1. On the DR, configure the number of static groups to be created by including the **group-count** statement and specifying the number of groups to be created.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 group-count 3
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
 static {
 group 225.1.1.1 {
 group-count 3;
 }
 }
}
```

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static groups 225.1.1.1, 225.1.1.2, and 225.1.1.3 have been created.

```
user@host> show igmp group
Interface: fe-0/1/2
 Group: 225.1.1.1
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.2
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.3
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can also configure the group address to be automatically incremented for each group created. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately and when you do not want the group addresses to be sequential.

In this example, you create three groups and increase the group address by an increment of two for each group.

1. On the DR, configure the group address increment by including the **group-increment** statement and specifying the number by which the address should be incremented for each group. The increment is specified in dotted decimal notation similar to an IPv4 address.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 group-count 3 group-increment
0.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
 version 3;
 static {
```

```

 group 225.1.1.1 {
 group-increment 0.0.0.2;
 group-count 3;
 }
 }
}

```

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static groups 225.1.1.1, 225.1.1.3, and 225.1.1.5 have been created.

```

user@host> show igmp group
Interface: fe-0/1/2
 Group: 225.1.1.1
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.3
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.5
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static

```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, and your network is operating in source-specific multicast (SSM) mode, you can also specify that the multicast source address be accepted. This is useful when you want to test forwarding to multicast receivers from a specific multicast source.

If you specify a group address in the SSM range, you must also specify a source.

If a source address is specified in a multicast group that is statically configured, the IGMP version on the interface must be set to IGMPv3. IGMPv2 is the default value.

In this example, you create group 225.1.1.1 and accept IP address 10.0.0.2 as the only source.

1. On the DR, configure the source address by including the **source** statement and specifying the IPv4 address of the source host.

```

[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2

```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```

user@host> show configuration protocol igmp
interface fe-0/1/2.0 {
 version 3;
 static {
 group 225.1.1.1 {
 source 10.0.0.2;
 }
 }
}

```

```
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that source 10.0.0.2 has been accepted.

```
user@host> show igmp group
Interface: fe-0/1/2
 Group: 225.1.1.1
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of multicast sources be automatically accepted. This is useful when you want to test forwarding to multicast receivers from more than one specified multicast source.

In this example, you create group 255.1.1.1 and accept addresses 10.0.0.2, 10.0.0.3, and 10.0.0.4 as the sources.

1. On the DR, configure the number of multicast source addresses to be accepted by including the **source-count** statement and specifying the number of sources to be accepted.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2 source-count
3
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
 version 3;
 static {
 group 225.1.1.1 {
 source 10.0.0.2 {
 source-count 3;
 }
 }
 }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that sources 10.0.0.2, 10.0.0.3, and 10.0.0.4 have been accepted.

```
user@host> show igmp group
Interface: fe-0/1/2
 Group: 225.1.1.1
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.1
 Source: 10.0.0.3
 Last reported by: Local
 Timeout: 0 Type: Static
```

```

Group: 225.1.1.1
Source: 10.0.0.4
Last reported by: Local
Timeout: 0 Type: Static

```

When you configure static groups on an interface on which you want to receive multicast traffic, and specify that a number of multicast sources be automatically accepted, you can also specify the number by which the address should be incremented for each source accepted. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately and you do not want the source addresses to be sequential.

In this example, you create group 225.1.1.1 and accept addresses 10.0.0.2, 10.0.0.4, and 10.0.0.6 as the sources.

1. Configure the multicast source address increment by including the **source-increment** statement and specifying the number by which the address should be incremented for each source. The increment is specified in dotted decimal notation similar to an IPv4 address.

```
[edit protocols igmp]
```

```
user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2 source-count
3 source-increment 0.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp
```

```

interface fe-0/1/2.0 {
 version 3;
 static {
 group 225.1.1.1 {
 source 10.0.0.2 {
 source-count 3;
 source-increment 0.0.0.2;
 }
 }
 }
}

```

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that sources 10.0.0.2, 10.0.0.4, and 10.0.0.6 have been accepted.

```
user@host> show igmp group
```

```

Interface: fe-0/1/2
 Group: 225.1.1.1
 Source: 10.0.0.2
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.1
 Source: 10.0.0.4
 Last reported by: Local
 Timeout: 0 Type: Static
 Group: 225.1.1.1
 Source: 10.0.0.6

```

Last reported by: Local  
Timeout: 0 Type: Static

When you configure static groups on an interface on which you want to receive multicast traffic and your network is operating in source-specific multicast (SSM) mode, you can specify that certain multicast source addresses be excluded.

By default the multicast source address configured in a static group operates in include mode. In include mode the multicast traffic for the group is accepted from the source address configured. You can also configure the static group to operate in exclude mode. In exclude mode the multicast traffic for the group is accepted from any address other than the source address configured.

If a source address is specified in a multicast group that is statically configured, the IGMP version on the interface must be set to IGMPv3. IGMPv2 is the default value.

In this example, you exclude address 10.0.0.2 as a source for group 225.1.1.1.

1. On the DR, configure a multicast static group to operate in exclude mode by including the **exclude** statement and specifying which IPv4 source address to exclude.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 exclude source 10.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
 version 3;
 static {
 group 225.1.1.1 {
 exclude;
 source 10.0.0.2;
 }
 }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group detail** command to verify that static group 225.1.1.1 has been created and that the static group is operating in exclude mode.

```
user@host> show igmp group detail
Interface: fe-0/1/2
Group: 225.1.1.1
Group mode: Exclude
Source: 10.0.0.2
Last reported by: Local
Timeout: 0 Type: Static
```

### ***Recording IGMP Join and Leave Events***

To determine whether IGMP tuning is needed in a network, you can configure the routing device to record IGMP join and leave events. You can record events globally for the routing device or for individual interfaces.



Table 276 on page 3675 describes the recordable IGMP events.

**Table 276: IGMP Event Messages**

| ERRMSG Tag                  | Definition                                                     |
|-----------------------------|----------------------------------------------------------------|
| RPD_IGMP_JOIN               | Records IGMP join events.                                      |
| RPD_IGMP_LEAVE              | Records IGMP leave events.                                     |
| RPD_IGMP_ACCOUNTING_ON      | Records when IGMP accounting is enabled on an IGMP interface.  |
| RPD_IGMP_ACCOUNTING_OFF     | Records when IGMP accounting is disabled on an IGMP interface. |
| RPD_IGMP_MEMBERSHIP_TIMEOUT | Records IGMP membership timeout events.                        |

To enable IGMP accounting:

1. Enable accounting globally or on an IGMP interface. This example shows both options.

```
[edit protocols igmp]
user@host# set accounting
user@host# set interface fe-0/1/0.2 accounting
```

2. Configure the events to be recorded and filter the events to a system log file with a descriptive filename, such as **igmp-events**.

```
[edit system syslog file igmp-events]
user@host# set any info
user@host# set match ".*RPD_IGMP_JOIN.* | .*RPD_IGMP_LEAVE.* |
.*RPD_IGMP_ACCOUNTING.* | .*RPD_IGMP_MEMBERSHIP_TIMEOUT.*"
```

3. Periodically archive the log file.

This example rotates the file size when it reaches 100 KB and keeps three files.

```
[edit system syslog file igmp-events]
user@host# set archive size 100000
user@host# set archive files 3
user@host# set archive archive-sites "ftp://user@host1//var/tmp" password
"anonymous"
user@host# set archive archive-sites "ftp://user@host2//var/tmp" password "test"
user@host# set archive transfer-interval 24
user@host# set archive start-time 2011-01-07:12:30
```

4. You can monitor the system log file as entries are added to the file by running the **monitor start** and **monitor stop** commands.

```
user@host> monitor start igmp-events

*** igmp-events ***
Apr 16 13:08:23 host mgd[16416]: UI_CMDLINE_READ_LINE: User 'user', command
'run monitor start igmp-events '
monitor
```

### ***Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces***

The **group-limit** statement enables you to limit the number of IGMP multicast group joins for logical interfaces. When this statement is enabled on a routing device running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), the limit is applied upon receipt of the group report. Once the group limit is reached, subsequent join requests are rejected.

When configuring limits for IGMP multicast groups, keep the following in mind:

- Each any-source group (\*G) counts as one group toward the limit.
- Each source-specific group (S,G) counts as one group toward the limit.
- Groups in IGMPv3 exclude mode are counted toward the limit.
- Multiple source-specific groups count individually toward the group limit, even if they are for the same group. For example, (S1, G1) and (S2, G1) would count as two groups toward the configured limit.
- Combinations of any-source groups and source-specific groups count individually toward the group limit, even if they are for the same group. For example, (\*, G1) and (S, G1) would count as two groups toward the configured limit.
- Configuring and committing a group limit on a network that is lower than what already exists on the network results in the removal of all groups from the configuration. The groups must then request to rejoin the network (up to the newly configured group limit).
- You can dynamically limit multicast groups on IGMP logical interfaces using dynamic profiles.

Beginning with Junos OS 12.2, you can optionally configure a system log warning threshold for IGMP multicast group joins received on the logical interface. It is helpful to review the system log messages for troubleshooting purposes and to detect if an excessive amount of IGMP multicast group joins have been received on the interface. These log messages convey when the configured group limit has been exceeded, when the configured threshold has been exceeded, and when the number of groups drop below the configured threshold.

The **group-threshold** statement enables you to configure the threshold at which a warning message is logged. The range is 1 through 100 percent. The warning threshold is a percentage of the group limit, so you must configure the **group-limit** statement to configure a warning threshold. For instance, when the number of groups exceed the configured warning threshold, but remain below the configured group limit, multicast groups continue to be accepted, and the device logs the warning message. In addition, the device logs a warning message after the number of groups drop below the configured warning threshold. You can further specify the amount of time (in seconds) between the log messages by configuring the **log-interval** statement. The range is 6 through 32,767 seconds.

You might consider throttling log messages because every entry added after the configured threshold and every entry rejected after the configured limit causes a warning message to be logged. By configuring a log interval, you can throttle the amount of system log warning messages generated for IGMP multicast group joins.

To limit multicast group joins on an IGMP logical interface:

1. Access the logical interface at the IGMP protocol hierarchy level.

```
[edit]
user@host# edit protocols igmp interface interface-name
```

2. Specify the group limit for the interface.

```
[edit protocols igmp interface interface-name]
user@host# set group-limit limit
```

3. (Optional) Configure the threshold at which a warning message is logged.

```
[edit protocols igmp interface interface-name]
user@host# set group-threshold value
```

4. (Optional) Configure the amount of time between log messages.

```
[edit protocols igmp interface interface-name]
user@host# set log-interval seconds
```

To confirm your configuration, use the **show protocols igmp** command. To verify the operation of IGMP on the interface, including the configured group limit and the optional warning threshold and interval between log messages, use the **show igmp interface** command.

### Tracing IGMP Protocol Traffic

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

| Flag                       | Description                                                                       |
|----------------------------|-----------------------------------------------------------------------------------|
| <b>all</b>                 | Trace all operations.                                                             |
| <b>client-notification</b> | Trace notifications.                                                              |
| <b>general</b>             | Trace general flow.                                                               |
| <b>group</b>               | Trace group operations.                                                           |
| <b>host-notification</b>   | Trace host notifications.                                                         |
| <b>leave</b>               | Trace leave group messages (IGMPv2 only).                                         |
| <b>mtrace</b>              | Trace mtrace packets. Use the <b>mtrace</b> command to troubleshoot the software. |
| <b>normal</b>              | Trace normal events.                                                              |
| <b>packets</b>             | Trace all IGMP packets.                                                           |
| <b>policy</b>              | Trace policy processing.                                                          |

| Flag          | Description                                                                         |
|---------------|-------------------------------------------------------------------------------------|
| <b>query</b>  | Trace IGMP membership query messages, including general and group-specific queries. |
| <b>report</b> | Trace membership report messages.                                                   |
| <b>route</b>  | Trace routing information.                                                          |
| <b>state</b>  | Trace state transitions.                                                            |
| <b>task</b>   | Trace task processing.                                                              |
| <b>timer</b>  | Trace timer processing.                                                             |

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on IGMP packets of a particular type. To configure tracing operations for IGMP:

1. (Optional) Configure tracing at the routing options level to trace all protocol packets.

```
[edit routing-options traceoptions]
user@host# set file all-packets-trace
user@host# set flag all
```

2. Configure the filename for the IGMP trace file.

```
[edit protocols igmp traceoptions]
user@host# set file igmp-trace
```

3. (Optional) Configure the maximum number of trace files.

```
[edit protocols igmp traceoptions]
user@host# set file files 5
```

4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols igmp traceoptions]
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols igmp traceoptions]
user@host# set file world-readable
```

6. Configure tracing flags. Suppose you are troubleshooting issues with a particular multicast group. The following example shows how to flag all events for packets associated with the group IP address.

```
[edit protocols igmp traceoptions]
user@host# set flag group | match 232.1.1.2
```

7. View the trace file.

```
user@host> file list /var/log
user@host> file show /var/log/igmp-trace
```

### Disabling IGMP

To disable IGMP on an interface, include the **disable** statement:

```
disable;
```

You can include this statement at the following hierarchy levels:

- [edit protocols igmp interface *interface-name*]
- [edit logical-systems *logical-system-name* protocols igmp interface *interface-name*]

### IGMP and Nonstop Active Routing

Nonstop active routing (NSR) configurations include two Routing Engines that share information so that routing is not interrupted during Routing Engine failover. These NSR configurations include passive support with IGMP in connection with PIM. The master Routing Engine uses IGMP to determine its PIM multicast state, and this IGMP-derived information is replicated on the backup Routing Engine. IGMP on the new master Routing Engine (after failover) relearns the state information quickly through IGMP operation. In the interim, the new master Routing Engine retains the IGMP-derived PIM state as received by the replication process from the old master Routing Engine. This state information times out unless refreshed by IGMP on the new master Routing Engine. No additional IGMP configuration is required.

#### Related Documentation

- *Examples: Configuring MLD*

### Configuring Multiple Instances of MSDP

MSDP instances are supported only for VRF instance types. You can configure multiple instances of MSDP to support multicast over VPNs.

To configure multiple instances of MSDP, include the following statements:

```
routing-instances {
 routing-instance-name {
 interface interface-name;
 instance-type vrf;
 route-distinguisher (as-number:number | ip-address:number);
 vrf-import [policy-names];
 vrf-export [policy-names];
 protocols {
 msdp {
 ... msdp-configuration ...
 }
 }
 }
}
```

You can include the statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]

- Related Documentation**
- *Example: Configuring MSDP in a Routing Instance*
  - *Junos OS MPLS Applications Configuration Guide*
  - *Junos OS VPNs Configuration Guide*

---

## Configuring Basic PIM Settings

- [PIM Configuration Statements on page 3680](#)
- [Changing the PIM Version on page 3682](#)
- [Modifying the PIM Hello Interval on page 3683](#)
- [Preserving Multicast Performance by Disabling Response to the ping Utility on page 3684](#)
- [PIM on Aggregated Interfaces on page 3684](#)
- [Configuring PIM Trace Options on page 3685](#)
- [Disabling PIM on page 3687](#)
- [Verifying a Multicast Configuration on page 3689](#)

### *PIM Configuration Statements*

To configure Protocol Independent Multicast (PIM), include the **pim** statement:

```
pim {
 disable;
 default-vpn-source {
 interface-name interface-name;
 }
 assert-timeout seconds;
 dense-groups {
 addresses;
 }
 dr-election-on-p2p;
 export;
 graceful-restart {
 disable;
 no-bidirectional-mode;
 restart-duration seconds;
 }
 idle-standby-path-switchover-delay seconds;
 import [policy-names];
 interface interface-name {
 bidirectional {
 df-election {
 backoff-period milliseconds;
 offer-period milliseconds;
 robustness-count number;
 }
 }
 }
 import;
 hello-interval seconds;
 mode bidirectional-sparse | bidirectional-sparse-dense | (dense | sparse |
 sparse-dense);
 neighbor-policy [policy-names];
```

```

 override-interval milliseconds;
 priority number;
 propagation-delay milliseconds;
 reset-tracking-bit;
 version version;
}
join-load-balance {
 automatic;
}
join-prune-timeout;
nonstop-routing {
 disable;
}
override-interval milliseconds;
propagation-delay milliseconds;
reset-tracking-bit;
rib-group {
 inet group-name;
 inet6 group-name;
}
rp {
 auto-rp {
 (announce | discovery | mapping);
 (mapping-agent-election | no-mapping-agent-election);
 }
 bidirectional {
 address address {
 group-ranges {
 destination-ip-prefix</prefix-length>;
 }
 hold-time seconds;
 priority number;
 }
 }
 bootstrap {
 family (inet | inet6) {
 export [policy-names];
 import [policy-names];
 priority number;
 }
 }
 bootstrap-export [policy-names];
 bootstrap-import [policy-names];
 bootstrap-priority number;
 dr-register-policy [policy-names];
 embedded-rp {
 group-ranges {
 destination-ip-prefix</prefix-length>;
 }
 maximum-rps limit;
 }
}
local {
 family (inet | inet6) {
 address address;
 anycast-pim {
 rp-set {

```

```

 address address <forward-msdp-sa>;
 }
 local-address address;
}
disable;
group-ranges {
 destination-ip-prefix </prefix-length>;
}
hold-time seconds;
override;
priority number;
}
}
rp-register-policy [policy-names];
standby-path-creation-delay seconds;
static {
 address address {
 override;
 version version;
 group-ranges {
 destination-ip-prefix </prefix-length>;
 }
 spt-threshold {
 infinity [policy-names];
 }
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
 }
}
}
}
}

```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]

By default, PIM is disabled.



**NOTE:** You cannot configure PIM within a nonforwarding instance. If you try to do so, the routing device displays a commit check error and does not complete the configuration commit process.

### Changing the PIM Version

All systems on a subnet must run the same version of PIM.



The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default for rendezvous point (RP) mode (at the **[edit protocols pim rp static address address]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface interface-name]** hierarchy level). Explicitly configured versions override the defaults.

To configure the PIM version, include the **version** statement:

```
version (1 | 2);
```

### ***Modifying the PIM Hello Interval***

Routing devices send hello messages at a fixed interval on all PIM-enabled interfaces. By using hello messages, routing devices advertise their existence as PIM routing devices on the subnet. With all PIM-enabled routing devices advertised, a single designated router for the subnet is established.

When a routing device is configured for PIM, it sends a hello message at a 30-second default interval. The interval range is from 0 through 255. When the interval counts down to 0, the routing device sends another hello message, and the timer is reset. A routing device that receives no response from a neighbor in 3.5 times the interval value drops the neighbor. In the case of a 30-second interval, the amount of time a routing device waits for a response is 105 seconds.

If a PIM hello message contains the hold-time option, the neighbor timeout is set to the hold-time sent in the message. If a PIM hello message does not contain the hold-time option, the neighbor timeout is set to the default hello hold time.

To modify how often the routing device sends hello messages out of an interface:

1. This example shows the configuration for the routing instance. Configure the interface globally or in the routing instance.

```
[edit routing-instances PIM.master protocols pim interface fe-3/0/2.0]
user@host# set hello-interval 255
```

2. Verify the configuration by checking the **Hello Option Holdtime** field in the output of the **show pim neighbors detail** command.

```
user@host> show pim neighbors detail
Instance: PIM.master
Interface: fe-3/0/2.0
Address: 192.168.195.37, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 255 seconds
Hello Option DR Priority: 1
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
Rx Join: Group Source Timeout
225.1.1.1 192.168.195.78 0
225.1.1.1 0

Interface: lo0.0
Address: 10.255.245.91, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 255 seconds
Hello Option DR Priority: 1
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

```
Interface: pd-6/0/0.32768
Address: 0.0.0.0, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 255 seconds
Hello Option DR Priority: 0
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

### ***Preserving Multicast Performance by Disabling Response to the ping Utility***

The ping utility uses ICMP Echo messages to verify connectivity to any device with an IP address. However, in the case of multicast applications, a single ping sent to a multicast address can degrade the performance of routing devices because the stream of packets is replicated multiple times.

You can disable the routing device's response to ping (ICMP Echo) packets sent to multicast addresses. The system responds normally to unicast ping packets.

To disable the routing device's response to ping packets sent to multicast addresses:

1. Include the **no-multicast-echo** statement:

```
[edit system]
user@host# set no-multicast-echo
```

2. Verify the configuration by checking the **echo drops with broadcast or multicast destination address** field in the output of the **show system statistics icmp** command.

```
user@host> show system statistics icmp

icmp:
0 drops due to rate limit
0 calls to icmp_error
0 errors not generated because old message was icmp
Output histogram:
echo reply: 21
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
100 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
echo: 21
21 message responses generated
```

### ***PIM on Aggregated Interfaces***

You can configure several Protocol Independent Multicast (PIM) features on an interface regardless of its PIM mode (bidirectional, sparse, dense, or sparse-dense mode).

If you configure PIM on an aggregated (**ae-** or **as-**) interface, each of the interfaces in the aggregate is included in the multicast output interface list and carries the single stream of replicated packets in a load-sharing fashion. The multicast aggregate interface is “expanded” into its constituent interfaces in the next-hop database.

### Configuring PIM Trace Options

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

| Flag                             | Description                                                                                                                                                              |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>all</b>                       | Trace all operations.                                                                                                                                                    |
| <b>assert</b>                    | Trace assert messages, which are used to resolve which of the parallel routing devices connected to a multiaccess LAN is responsible for forwarding packets to the LAN.  |
| <b>autorp</b>                    | Trace bootstrap, RP, and auto-RP messages.                                                                                                                               |
| <b>bidirectional-df-election</b> | Trace bidirectional PIM designated-forwarder (DF) election events.                                                                                                       |
| <b>bootstrap</b>                 | Trace bootstrap messages, which are sent periodically by the PIM domain's bootstrap routing device and are forwarded, hop by hop, to all routing devices in that domain. |
| <b>general</b>                   | Trace general events.                                                                                                                                                    |
| <b>graft</b>                     | Trace graft and graft acknowledgment messages.                                                                                                                           |
| <b>hello</b>                     | Trace hello packets, which are sent so that neighboring routing devices can discover one another.                                                                        |
| <b>join</b>                      | Trace join messages, which are sent to join a branch onto the multicast distribution tree.                                                                               |
| <b>mdt</b>                       | Trace messages related to multicast data tunnels.                                                                                                                        |
| <b>normal</b>                    | Trace normal events.                                                                                                                                                     |
| <b>nsr-synchronization</b>       | Trace nonstop routing synchronization events                                                                                                                             |
| <b>packets</b>                   | Trace all PIM packets.                                                                                                                                                   |
| <b>policy</b>                    | Trace poison-route-reverse packets.                                                                                                                                      |
| <b>prune</b>                     | Trace prune messages, which are sent to prune a branch off the multicast distribution tree.                                                                              |
| <b>register</b>                  | Trace register and register-stop messages. Register messages are sent to the RP when a multicast source first starts sending to a group.                                 |

| Flag         | Description                        |
|--------------|------------------------------------|
| <b>route</b> | Trace routing information.         |
| <b>rp</b>    | Trace candidate RP advertisements. |
| <b>state</b> | Trace state transitions.           |
| <b>task</b>  | Trace task processing.             |
| <b>timer</b> | Trace timer processing.            |

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on PIM packets of a particular type.

To configure tracing operations for PIM:

1. (Optional) Configure tracing at the [**routing-options** hierarchy level to trace all protocol packets.

```
[edit routing-options traceoptions]
user@host# set file all-packets-trace
user@host# set flag all
```

2. Configure the filename for the PIM trace file.

```
[edit protocols pim traceoptions]
user@host# set file pim-trace
```

3. (Optional) Configure the maximum number of trace files.

```
[edit protocols pim traceoptions]
user@host# set file files 5
```

4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols pim traceoptions]
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols pim traceoptions]
user@host# set file world-readable
```

6. Configure tracing flags.

Suppose you are troubleshooting issues with PIM version 1 control packets that are received on an interface configured for PIM version 2. The following example shows how to trace messages associated with this problem.

```
[edit protocols pim traceoptions]
user@host# set flag packets | match "Rx V1 Require V2"
```

7. View the trace file.

```
user@host> file list /var/log
user@host> file show /var/log/pim-trace
```

### *Disabling PIM*

By default, when configured, the PIM protocol is enabled on all interfaces for all families. If desired, you can disable PIM at the protocol, interface, or family hierarchy levels.

The hierarchy in which you configure PIM is critical. In general, the most specific configuration takes precedence. However, if PIM is disabled at the protocol level, then any disable statements with respect to an interface or family are ignored.

For example, the order of precedence for disabling PIM on a particular interface family is:

1. If PIM is disabled at the **[edit protocols pim interface *interface-name* family]** hierarchy level, then PIM is disabled for that interface family.
2. If PIM is not configured at the **[edit protocols pim interface *interface-name* family]** hierarchy level, but is disabled at the **[edit protocols pim interface *interface-name*]** hierarchy level, then PIM is disabled for all families on the specified interface.
3. If PIM is not configured at either the **[edit protocols pim interface *interface-name* family]** hierarchy level or the **[edit protocols pim interface *interface-name*]** hierarchy level, but is disabled at the **[edit protocols pim]** hierarchy level, then the PIM protocol is disabled globally for all interfaces and all families.

The following sections describe how to disable PIM at the various hierarchy levels.

- [Disabling the PIM Protocol on page 3687](#)
- [Disabling PIM On an Interface on page 3688](#)
- [Disabling PIM for a Family on page 3688](#)
- [Disabling PIM for a Rendezvous Point on page 3689](#)

### *Disabling the PIM Protocol*

You can explicitly disable the PIM protocol. Disabling the PIM protocol disables the protocol for all interfaces and all families. This is accomplished at the **[edit protocols pim]** hierarchy level:

```
[edit protocols]
pim {
 disable;
}
```

To disable the PIM protocol:

1. Include the **disable** statement.  

```
user@host# set protocols pim disable
```
2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.  

```
user@host# run show protocols pim
```

### *Disabling PIM On an Interface*

You can disable the PIM protocol on a per-interface basis. This is accomplished at the **[edit protocols pim interface *interface-name*]** hierarchy level:

```
[edit protocols]
pim {
 interface interface-name {
 disable;
 }
}
```

To disable PIM on an interface:

1. Include the **disable** statement.

```
user@host# set protocols pim interface fe-0/1/0 disable
```

2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

```
user@host# run show protocols pim
```

### *Disabling PIM for a Family*

You can disable the PIM protocol on a per-family basis. This is accomplished at the **[edit protocols pim family]** hierarchy level:

```
[edit protocols]
pim {
 family inet {
 disable;
 }
 family inet6 {
 disable;
 }
}
```

To disable PIM for a family:

1. Include the **disable** statement.

```
user@host# set protocols pim family inet disable
```

```
user@host# set protocols pim family inet6 disable
```

2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

```
user@host# run show protocols pim
```

**Disabling PIM for a Rendezvous Point**

You can disable the PIM protocol for a rendezvous point (RP) on a per-family basis. This is accomplished at the **[edit protocols pim rp local family]** hierarchy level:

```
[edit protocols]
pim {
 rp {
 local {
 family inet {
 disable;
 }
 family inet6 {
 disable;
 }
 }
 }
}
```

To disable PIM for an RP family:

1. Use the **disable** statement.
 

```
user@host# set protocols pim rp local family inet disable
user@host# set protocols pim rp local family inet6 disable
```
2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.
 

```
user@host# run show protocols pim
```

**Verifying a Multicast Configuration**

To verify a multicast configuration, perform these tasks:

- [Verifying SAP and SDP Addresses and Ports on page 3689](#)
- [Verifying the IGMP Version on page 3690](#)
- [Verifying the PIM Mode and Interface Configuration on page 3690](#)
- [Verifying the PIM RP Configuration on page 3691](#)
- [Verifying the RPF Routing Table Configuration on page 3691](#)

**Verifying SAP and SDP Addresses and Ports**

- Purpose** Verify that SAP and SDP are configured to listen on the correct group addresses and ports.
- Action** From the CLI, enter the **show sap listen** command.

**Sample Output**

```
user@host> show sap listen
Group Address Port
224.2.127.254 9875
```

- Meaning** The output shows a list of the group addresses and ports that SAP and SDP listen on. Verify the following information:

- Each group address configured, especially the default **224.2.127.254**, is listed.
- Each port configured, especially the default **9875**, is listed.

#### *Verifying the IGMP Version*

**Purpose** Verify that IGMP version 2 is configured on all applicable interfaces.

**Action** From the CLI, enter the **show igmp interface** command.

#### Sample Output

```
user@host> show igmp interface
Interface: ge-0/0/0.0
 Querier: 192.168.4.36
 State: Up Timeout: 197 Version: 2 Groups: 0

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0
```

**Meaning** The output shows a list of the interfaces that are configured for IGMP. Verify the following information:

- Each interface on which IGMP is enabled is listed.
- Next to **Version**, the number 2 appears.

#### *Verifying the PIM Mode and Interface Configuration*

**Purpose** Verify that PIM sparse mode is configured on all applicable interfaces.

**Action** From the CLI, enter the **show pim interfaces** command.

#### Sample Output

```
user@host> show pim interfaces
Instance: PIM.master
Name Stat Mode IP V State Count DR address
1o0.0 Up Sparse 4 2 DR 0 127.0.0.1
pim.32769 Up Sparse 4 2 P2P 0
```

**Meaning** The output shows a list of the interfaces that are configured for PIM. Verify the following information:

- Each interface on which PIM is enabled is listed.
- The network management interface, either **ge-0/0/0** or **fe-0/0/0**, is *not* listed.
- Under **Mode**, the word **Sparse** appears.



**Verifying the PIM RP Configuration**

**Purpose** Verify that the PIM RP is statically configured with the correct IP address.

**Action** From the CLI, enter the **show pim rps** command.

**Sample Output**

```
user@host> show pim rps
Instance: PIM.master
Address family INET
RP address Type Holdtime Timeout Active groups Group prefixes
192.168.14.27 static 0 None 2 224.0.0.0/4
```

**Meaning** The output shows a list of the RP addresses that are configured for PIM. At least one RP must be configured. Verify the following information:

- The configured RP is listed with the proper IP address.
- Under **Type**, the word **static** appears.

**Verifying the RPF Routing Table Configuration**

**Purpose** Verify that the PIM RPF routing table is configured correctly.

**Action** From the CLI, enter the **show multicast rpf** command.

**Sample Output**

```
user@host> show multicast rpf
Multicast RPF table: inet.0 , 2 entries...
```

**Meaning** The output shows the multicast RPF table that is configured for PIM. If no multicast RPF routing table is configured, RPF checks use **inet.0**. Verify the following information:

- The configured multicast RPF routing table is **inet.0**.
- The **inet.0** table contains entries.

**Related Documentation**

- [Configuring PIM Auto-RP on page 3705](#)
- [Configuring PIM Bootstrap Router on page 3701](#)
- [Configuring PIM Dense Mode on page 3737](#)
- [Configuring a Designated Router for PIM on page 3692](#)
- [Configuring PIM Filtering on page 3712](#)
- [Configuring PIM Sparse-Dense Mode on page 3740](#)
- [Configuring PIM and the Bidirectional Forwarding Detection \(BFD\) Protocol on page 3725](#)
- *Example: Configuring Nonstop Active Routing for PIM*
- *Examples: Configuring PIM RPT and SPT Cutover*

- [Examples: Configuring PIM Sparse Mode](#)
- [Configuring PIM Sparse-Dense Mode on page 3740](#)
- [Configuring PIM and the Bidirectional Forwarding Detection \(BFD\) Protocol on page 3725](#)

---

### Configuring Multiple Instances of PIM

PIM instances are supported only for VRF instance types. You can configure multiple instances of PIM to support multicast over VPNs.

To configure multiple instances of PIM, include the following statements:

```
routing-instances {
 routing-instance-name {
 interface interface-name;
 instance-type vrf;
 protocols {
 pim {
 ... pim-configuration ...
 }
 }
 }
}
```

You can include the statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]

#### Related Documentation

- [Multicast Protocols Configuration Guide](#)
- [Junos OS VPNs Configuration Guide](#)

---

### Configuring a Designated Router for PIM

- [Configuring Interface Priority for PIM Designated Router Selection on page 3692](#)
- [Configuring PIM Designated Router Election on Point-to-Point Links on page 3693](#)

#### **Configuring Interface Priority for PIM Designated Router Selection**

A designated router (DR) sends periodic join messages and prune messages toward a group-specific rendezvous point (RP) for each group for which it has active members. When a Protocol Independent Multicast (PIM) routing device learns about a source, it originates a Multicast Source Discovery Protocol (MSDP) source-address message if it is the DR on the upstream interface.

By default, every PIM interface has the lowest probability (priority 0) of being selected as the DR. Configuring the interface DR priority helps ensure that changing an IP address does not alter your forwarding model.

To configure the interface designated router priority:

1. This example shows the configuration for the routing instance. Configure the interface globally or in the routing instance.

```
[edit routing-instances PIM.master protocols pim interface ge-0/0/0.0 family inet]
user@host# set priority 5
```

2. Verify the configuration by checking the **Hello Option DR Priority** field in the output of the **show pim neighbors detail** command.

```
user@host> show pim neighbors detail

Instance: PIM.master
Interface: ge-0/0/0.0
Address: 192.168.195.37, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 5
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
Rx Join: Group Source Timeout
225.1.1.1 192.168.195.78 0
225.1.1.1 0

Interface: lo0.0
Address: 10.255.245.91, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported

Interface: pd-6/0/0.32768
Address: 0.0.0.0, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 0
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

### ***Configuring PIM Designated Router Election on Point-to-Point Links***

To comply with the latest PIM drafts, enable designated router (DR) election on all PIM interfaces, including point-to-point (P2P) interfaces. (DR election is enabled by default on all other interfaces.) One of the two routing devices might join a multicast group on its P2P link interface. The DR on that link is responsible for initiating the relevant join messages.

To enable DR election on point-to-point interfaces:

1. On both point-to-point link routing devices, configure the routing device globally or in the routing instance. This example shows the configuration for the routing instance.

```
[edit routing-instances PIM.master protocols pim]
user@host# set dr-election-on-p2p
```

2. Verify the configuration by checking the **State** field in the output of the **show pim interfaces** command. The possible values for the **State** field are DR, NotDR, and P2P. When a point-to-point link interface is elected to be the DR, the interface state becomes DR instead of P2P.

3. If the **show pim interfaces** command continues to report the P2P state, consider running the **restart routing** command on both routing devices on the point-to-point link. Then recheck the state.



**CAUTION:** Do not restart a software process unless specifically asked to do so by your Juniper Networks customer support representative. Restarting a software process during normal operation of a routing platform could cause interruption of packet forwarding and loss of data.

[edit]  
user@host# run restart routing

#### Related Documentation

- [Configuring PIM Auto-RP on page 3705](#)
- [Configuring PIM Bootstrap Router on page 3701](#)
- [Configuring PIM Dense Mode on page 3737](#)
- [Configuring PIM Filtering on page 3712](#)
- [Example: Configuring Nonstop Active Routing for PIM](#)
- [Examples: Configuring PIM RPT and SPT Cutover](#)
- [Examples: Configuring PIM Sparse Mode](#)
- [Configuring PIM Sparse-Dense Mode on page 3740](#)
- [Configuring PIM and the Bidirectional Forwarding Detection \(BFD\) Protocol on page 3725](#)
- [Configuring Basic PIM Settings on page 3680](#)

---

### Configuring Static RP

- [Understanding Static RP on page 3694](#)
- [Configuring Local PIM RPs on page 3695](#)
- [Example: Configuring PIM Sparse Mode and RP Static IP Addresses on page 3697](#)
- [Configuring the Static PIM RP Address on the Non-RP Routing Device on page 3699](#)

#### Understanding Static RP

Protocol Independent Multicast (PIM) sparse mode is the most common multicast protocol used on the Internet. PIM sparse mode is the default mode whenever PIM is configured on any interface of the device. However, because PIM must not be configured on the network management interface, you must disable it on that interface.

Each any-source multicast (ASM) group has a shared tree through which receivers learn about new multicast sources and new receivers learn about all multicast sources. The rendezvous point (RP) router is the root of this shared tree and receives the multicast traffic from the source. To receive multicast traffic from the groups served by the RP, the device must determine the IP address of the RP for the source.

You can configure a static rendezvous point (RP) configuration that is similar to static routes. A static configuration has the benefit of operating in PIM version 1 or version 2. When you configure the static RP, the RP address that you select for a particular group must be consistent across all routers in a multicast domain.

One common way for the device to locate RPs is by static configuration of the IP address of the RP. A static configuration is simple and convenient. However, if the statically defined RP router becomes unreachable, there is no automatic failover to another RP router. To remedy this problem, you can use anycast RP.

### **Configuring Local PIM RPs**

Local RP configuration makes the routing device a statically defined RP. Consider statically defining an RP if the network does not have many different RPs defined or if the RP assignment does not change very often. The Junos IPv6 PIM implementation supports only static RP configuration. Automatic RP announcement and bootstrap routers are not available with IPv6.

You can configure a local RP globally or for a routing instance. This example shows how to configure a local RP in a routing instance for IPv4 or IPv6.

To configure the routing device's RP properties:

1. Configure the routing instance as the local RP.

```
[routing-instances VPN-A protocols pim]
user@host# set rp local
```

2. Configure the IP protocol family and IP address.

IPv6 PIM hello messages are sent to every interface on which you configure **family inet6**, whether at the PIM level of the hierarchy or not. As a result, if you configure an interface with both **family inet** at the **[edit interface interface-name]** hierarchy level and **family inet6** at the **[edit protocols pim interface interface-name]** hierarchy level, PIM sends both IPv4 and IPv6 hellos to that interface.

By default, PIM operates in sparse mode on an interface. If you explicitly configure sparse mode, PIM uses this setting for all IPv6 multicast groups. However, if you configure sparse-dense mode, PIM does not accept IPv6 multicast groups as dense groups and operates in sparse mode over them.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set family inet6 address 2001:db8:85a3::8a2e:370:7334
user@host# set family inet address 10.1.2.254
```

3. (IPv4 only) Configure the routing device's RP priority.



**NOTE:** The priority statement is not supported for IPv6, but is included here for informational purposes. The routing device's priority value for becoming the RP is included in the bootstrap messages that the routing device sends. Use a smaller number to increase the likelihood that the routing device becomes the RP for local multicast groups. Each PIM routing device uses the priority value and other factors to determine the candidate RPs for a particular group range. After the set of candidate RPs is distributed, each routing device determines algorithmically the RP from the candidate RP set using a hash function. By default, the priority value is set to 1. If this value is set to 0, the bootstrap router can override the group range being advertised by the candidate RP.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set priority 5
```

4. Configure the groups for which the routing device is the RP.

By default, a routing device running PIM is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12). The following example limits the groups for which this routing device can be the RP.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set group-ranges fec0::/10
user@host# set group-ranges 10.1.2.0/24
```

5. (IPv4 only) Modify the local RP hold time.

If the local routing device is configured as an RP, it is considered a candidate RP for its local multicast groups. For candidate RPs, the hold time is used by the bootstrap router to time out RPs, and applies to the bootstrap RP-set mechanism. The RP hold time is part of the candidate RP advertisement message sent by the local routing device to the bootstrap router. If the bootstrap router does not receive a candidate RP advertisement from an RP within the hold time, it removes that routing device from its list of candidate RPs. The default hold time is 150 seconds.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set hold-time 200
```

6. (Optional) Override dynamic RP for the specified group address range.

If you configure both static RP mapping and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for the given static RP group range, and allow dynamic RP mapping for all other groups.

If you exclude this statement from the configuration and you use both static and dynamic RP mechanisms for different group ranges within the same routing instance, the dynamic RP mapping takes precedence over the static RP mapping, even if static RP is defined for a specific group range.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set override
```

7. Monitor the operation of PIM by running the **show pim** commands. Run **show pim ?** to display the supported commands.

**Example: Configuring PIM Sparse Mode and RP Static IP Addresses**

This example shows how to configure PIM sparse mode and RP static IP addresses.

- [Requirements on page 3697](#)
- [Overview on page 3697](#)
- [Configuration on page 3697](#)
- [Verification on page 3699](#)

**Requirements**

Before you begin:

1. Determine whether the router is directly attached to any multicast sources. Receivers must be able to locate these sources.
2. Determine whether the router is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
3. Determine whether to configure multicast to use sparse, dense, or sparse-dense mode. Each mode has different configuration considerations.
4. Determine the address of the RP if sparse or sparse-dense mode is used.
5. Determine whether to locate the RP with the static configuration, BSR, or auto-RP method.
6. Determine whether to configure multicast to use its own RPF routing table when configuring PIM in sparse, dense, or sparse-dense mode.
7. Configure the SAP and SDP protocols to listen for multicast session announcements.
8. Configure IGMP.

**Overview**

In this example, you set the interface value to **all** and disable the **ge-0/0/0** interface. Then you configure the IP address of the RP as **192.168.14.27**.

**Configuration****CLI Quick Configuration**

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

```
set protocols pim interface all
set protocols pim interface ge-0/0/0 disable
set protocols pim rp static address 192.168.14.27
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure PIM sparse mode and the RP static IP address:

1. Configure PIM.  

```
[edit]
user@host# edit protocols pim
```
2. Set the interface value.  

```
[edit protocols pim]
user@host# set pim interface all
```
3. Disable PIM on the network management interface.  

```
[edit protocols pim interface]
user@host# set pim interface ge-0/0/0 unit 0 disable
```
4. Configure RP.  

```
[edit]
user@host# edit protocols pim rp
```
5. Configure the IP address of the RP.  

```
[edit]
user@host# set static address 192.168.14.27
```

**Results** From configuration mode, confirm your configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show protocols
pim {
 rp {
 static {
 address 192.168.14.27;
 }
 }
}
interface all;
 interface ge-0/0/0.0 {
 disable;
 }
}
```

If you are done configuring the device, enter **commit** from configuration mode.



**Verification**

To confirm that the configuration is working properly, perform these tasks:

- [Verifying SAP and SDP Addresses and Ports on page 3699](#)
- [Verifying the IGMP Version on page 3699](#)
- [Verifying the PIM Mode and Interface Configuration on page 3699](#)

**Verifying SAP and SDP Addresses and Ports**

**Purpose** Verify that SAP and SDP are configured to listen on the correct group addresses and ports.

**Action** From operational mode, enter the **show sap listen** command.

**Verifying the IGMP Version**

**Purpose** Verify that IGMP version 2 is configured on all applicable interfaces.

**Action** From operational mode, enter the **show igmp interface** command.

**Verifying the PIM Mode and Interface Configuration**

**Purpose** Verify that PIM sparse mode is configured on all applicable interfaces.

**Action** From operational mode, enter the **show pim interfaces** command.

**Configuring the Static PIM RP Address on the Non-RP Routing Device**

Consider statically defining an RP if the network does not have many different RPs defined or if the RP assignment does not change very often. The Junos IPv6 PIM implementation supports only static RP configuration. Automatic RP announcement and bootstrap routers are not available with IPv6.

You configure a static RP address on the non-RP routing device. This enables the non-RP routing device to recognize the local statically defined RP. For example, if R0 is a non-RP router and R1 is the local RP router, you configure R0 with the static RP address of R1. The static IP address is the routable address assigned to the loopback interface on R1. In the following example, the loopback address of the RP is 2001:db8:85a3::8a2e:370:7334.

You can configure a static RP address globally or for a routing instance. This example shows how to configure a static RP address in a routing instance for IPv6.

To configure the static RP address:

1. On a non-RP routing device, configure the routing instance to point to the routable address assigned to the loopback interface of the RP.

```
[routing-instances VPN-A protocols pim rp]
user@host# set static address 2001:db8:85a3::8a2e:370:7334
```



**NOTE:** Logical systems are also supported. You can configure a static RP address in a logical system only if the logical system is not directly connected to a source.

2. (Optional) Set the PIM sparse mode version.

For each static RP address, you can optionally specify the PIM version. The default PIM version is version 1.

```
[edit routing-instances VPN-A protocols pim rp]
user@host# set static address 2001:db8:85a3::8a2e:370:7334 version 2
```



**NOTE:** The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIM version 1 is the default for RP mode ([edit **pim rp static address address**]). PIM version 2 is the default for interface mode ([edit **pim interface interface-name**]). Explicitly configured versions override the defaults.

3. (Optional) Set the group address range.

By default, a routing device running PIM is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12). The following example limits the groups for which the 2001:db8:85a3::8a2e:370:7334 address can be the RP.

```
[edit routing-instances VPN-A protocols pim rp]
user@host# set static address 2001:db8:85a3::8a2e:370:7334 group-ranges fec0::/10
```

The RP that you select for a particular group must be consistent across all routers in a multicast domain.

4. (Optional) Override dynamic RP for the specified group address range.

If you configure both static RP mapping and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for the given static RP group range, and allow dynamic RP mapping for all other groups.

If you exclude this statement from the configuration and you use both static and dynamic RP mechanisms for different group ranges within the same routing instance, the dynamic RP mapping takes precedence over the static RP mapping, even if static RP is defined for a specific group range.

```
[edit routing-instances VPN-A protocols pim rp static address
 2001:db8:85a3::8a2e:370:7334]
user@host# set override
```

5. Monitor the operation of PIM by running the **show pim** commands. Run **show pim ?** to display the supported commands.

**Related Documentation**

- [Configuring PIM Auto-RP on page 3705](#)
- [Configuring PIM Bootstrap Router on page 3701](#)
- [Configuring a Designated Router for PIM on page 3692](#)

- *Examples: Configuring PIM Sparse Mode*
- [Configuring Basic PIM Settings on page 3680](#)

### Configuring PIM Bootstrap Router

- [Understanding the PIM Bootstrap Router on page 3701](#)
- [Configuring PIM Bootstrap Properties for IPv4 on page 3701](#)
- [Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 3702](#)
- [Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain on page 3704](#)
- [Example: Configuring PIM BSR Filters on page 3705](#)

#### ***Understanding the PIM Bootstrap Router***

To determine which routing device is the rendezvous point (RP), all routing devices within a PIM sparse-mode domain collect bootstrap messages. A PIM sparse-mode domain is a group of routing devices that all share the same RP router. The domain bootstrap routing device initiates bootstrap messages, which are sent hop by hop within the domain. The routing devices use bootstrap messages to distribute RP information dynamically and to elect a bootstrap routing device when necessary.

#### ***Configuring PIM Bootstrap Properties for IPv4***

For correct operation, every multicast routing device within a PIM domain must be able to map a particular multicast group address to the same Rendezvous Point (RP). The bootstrap routing device mechanism is one way that a multicast routing device can learn the set of group-to-RP mappings. Bootstrap routing devices are supported in IPv4 and IPv6.



**NOTE:** For legacy configuration purposes, there are two sections that describe the configuration of bootstrap routing devices: one section for both IPv4 and IPv6, and this section, which is for IPv4 only. The method described in [“Configuring PIM Bootstrap Properties for IPv4 or IPv6” on page 3702](#) is recommended. A commit error occurs if the same IPv4 bootstrap statements are included in both the IPv4-only and the IPv4-and-IPv6 sections of the hierarchy. The error message is “duplicate IPv4 bootstrap configuration.”

To determine which routing device is the RP, all routing devices within a PIM domain collect bootstrap messages. A PIM domain is a contiguous set of routing devices that implement PIM. All are configured to operate within a common boundary. The domain's bootstrap routing device initiates bootstrap messages, which are sent hop by hop within the domain. The routing devices use bootstrap messages to distribute RP information dynamically and to elect a bootstrap routing device when necessary.

PIM bootstrap messages are sourced from the loopback address, which is always up. The loopback address must be routable; if it is not routable, then the bootstrap routing device is unable to send bootstrap messages to update the RP domain members. See [Configuring the Loopback Interface](#) for information about configuring a loopback interface.

You can configure bootstrap properties globally or for a routing instance. This example shows the global configuration.

To configure the bootstrap routing device properties:

1. Configure the bootstrap priority.

By default, each routing device has a bootstrap priority of 0, which means the routing device can never be the bootstrap routing device. A priority of 0 disables the function for IPv4 and does not cause the routing device to send bootstrap routing device packets with a 0 in the priority field. The routing device with the highest priority value is elected to be the bootstrap routing device. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap routing device. A simple bootstrap configuration assigns a bootstrap priority value to a routing device.

```
[edit protocols pim rp]
user@host# set bootstrap-priority 3
```

2. (Optional) Create import and export policies to control the flow of IPv4 bootstrap messages to and from the RP, and apply the policies to PIM. Import and export policies are useful when some of the routing devices in your PIM domain have interfaces that connect to other PIM domains. Configuring a policy prevents bootstrap messages from crossing domain boundaries. The **bootstrap-import** statement prevents messages from being imported into the RP. The **bootstrap-export** statement prevents messages from being exported from the RP.

```
[edit protocols pim rp]
user@host# set bootstrap-import pim-bootstrap-import
user@host# set bootstrap-export pim-bootstrap-export
```

3. Configure the policies.

```
[edit policy-options policy-statement pim-bootstrap-import]
user@host# set from interface se-0/0/0
user@host# set then reject
```

```
[edit policy-options policy-statement pim-bootstrap-export]
user@host# set from interface se-0/0/0
user@host# set then reject
```

4. Monitor the operation of PIM bootstrap routing devices by running the **show pim bootstrap** command.

### ***Configuring PIM Bootstrap Properties for IPv4 or IPv6***

For correct operation, every multicast routing device within a PIM domain must be able to map a particular multicast group address to the same Rendezvous Point (RP). The bootstrap routing device mechanism is one way that a multicast routing device can learn the set of group-to-RP mappings. Bootstrap routing devices are supported in IPv4 and IPv6.



**NOTE:** For legacy configuration purposes, there are two sections that describe the configuration of bootstrap routing devices: one section for IPv4 only, and this section, which is for both IPv4 and IPv6. The method described in this section is recommended. A commit error occurs if the same IPv4 bootstrap statements are included in both the IPv4-only and the IPv4-and-IPv6 sections of the hierarchy. The error message is “duplicate IPv4 bootstrap configuration.”

To determine which routing device is the RP, all routing devices within a PIM domain collect bootstrap messages. A PIM domain is a contiguous set of routing devices that implement PIM. All devices are configured to operate within a common boundary. The domain's bootstrap routing device initiates bootstrap messages, which are sent hop by hop within the domain. The routing devices use bootstrap messages to distribute RP information dynamically and to elect a bootstrap routing device when necessary.

PIM bootstrap messages are sourced from the loopback address, which is always up. The loopback address must be routable; if it is not routable, then the bootstrap routing device is unable to send bootstrap messages to update the RP domain members. See *Configuring the Loopback Interface* for information about configuring a loopback interface.

You can configure bootstrap properties globally or for a routing instance. This example shows the global configuration.

To configure the bootstrap routing device properties:

1. Configure the bootstrap priority.

By default, each routing device has a bootstrap priority of 0, which means the routing device can never be the bootstrap routing device. The routing device with the highest priority value is elected to be the bootstrap routing device. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap routing device. A simple bootstrap configuration assigns a bootstrap priority value to a routing device.



**NOTE:** In the IPv4-only configuration, specifying a bootstrap priority of 0 disables the bootstrap function and does not cause the routing device to send BSR packets with a 0 in the priority field. In the configuration shown here, specifying a bootstrap priority of 0 does not disable the function, but causes the routing device to send BSR packets with a 0 in the priority field. To disable the bootstrap function in the IPv4 and IPv6 configuration, delete the `bootstrap` statement.

```
user@host# edit protocols pim rp
user@host# set bootstrap family inet priority 3
```

2. (Optional) Create import and export policies to control the flow of bootstrap messages to and from the RP, and apply the policies to PIM. Import and export policies are useful when some of the routing devices in your PIM domain have interfaces that connect to other PIM domains. Configuring a policy prevents bootstrap messages from crossing domain boundaries. The **import** statement prevents messages from being imported

into the RP. The **export** statement prevents messages from being exported from the RP.

```
[edit protocols pim rp]
user@host# set bootstrap family inet import pim-bootstrap-import
user@host# set bootstrap family inet export pim-bootstrap-export
```

3. Configure the policies.

```
[edit policy-options policy-statement pim-bootstrap-import]
user@host# set from interface se-0/0/0
user@host# set then reject
user@host# exit
user@host# edit policy-options policy-statement pim-bootstrap-export
user@host# set from interface se-0/0/0
user@host# set then reject
```

4. Monitor the operation of PIM bootstrap routing devices by running the **show pim bootstrap** command.

#### *Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain*

In this example, the **from interface so-0-1/0 then reject** policy statement rejects bootstrap messages from the specified interface (the example is configured for both IPv4 and IPv6 operation):

```
protocols {
 pim {
 rp {
 bootstrap {
 family inet {
 priority 1;
 import pim-import;
 export pim-export;
 }
 family inet6 {
 priority 1;
 import pim-import;
 export pim-export;
 }
 }
 }
 }
}
policy-options {
 policy-statement pim-import {
 from interface so-0/1/0;
 then reject;
 }
 policy-statement pim-export {
 to interface so-0/1/0;
 then reject;
 }
}
```

**Example: Configuring PIM BSR Filters**

Configure a filter to prevent BSR messages from entering or leaving your network. Add this configuration to all routers:

```
protocols {
 pim {
 rp {
 bootstrap-import no-bsr;
 bootstrap-export no-bsr;
 }
 }
}
policy-options {
 policy-statement no-bsr {
 then reject;
 }
}
```

**Related Documentation**

- [Configuring PIM Auto-RP on page 3705](#)
- [Configuring a Designated Router for PIM on page 3692](#)
- [Examples: Configuring PIM Sparse Mode](#)
- [Configuring Basic PIM Settings on page 3680](#)

**Configuring PIM Auto-RP**

- [Understanding PIM Auto-RP on page 3705](#)
- [Configuring PIM Auto-RP on page 3705](#)

**Understanding PIM Auto-RP**

You can configure a more dynamic way of assigning rendezvous points (RPs) in a multicast network by means of auto-RP. When you configure auto-RP for a , the learns the address of the RP in the network automatically and has the added advantage of operating in PIM version 1 and version 2.

Although auto-RP is a nonstandard (non-RFC-based) function that typically uses dense mode PIM to advertise control traffic, it provides an important failover advantage that simple static RP assignment does not. You can configure multiple s as RP candidates. If the elected RP fails, one of the other preconfigured s takes over the RP functions. This capability is controlled by the auto-RP mapping agent.

**Configuring PIM Auto-RP**

For correct operation, every multicast within a PIM domain must be able to map a particular multicast group address to the same rendezvous point (RP). The auto-RP mechanism is one way that a multicast can learn the set of group-to-RP mappings. Auto-RP automatically distributes mapping information to routing devices. It simplifies use of multiple RPs for different multicast group ranges, thus allowing multiple RPs to act as backups for each other. Auto-RP relies on a to act as the RP mapping agent.

Potential RPs announce themselves to the mapping agent, and the mapping agent resolves any conflicts.

The mapping agent sends the multicast group-RP mapping information to the other s using PIM dense mode. The specific groups used are 224.0.1.39 and .40. The first (.39) is used to advertise, the second (.40) is used for discovery. Because PIM dense mode is necessary to enable auto-RP to work, which in turns enables PIM sparse mode to work, you must configure PIM sparse-dense mode in the PIM domains that use auto-RP.

Although auto-RP is a nonstandard (non-RFC-based) function requiring dense mode PIM to advertise control traffic, it provides an important failover advantage that static RP assignment does not. That is, you can configure multiple routing devices as RP candidates. If the elected RP fails, one of the other preconfigured routing devices takes over the RP functions. This capability is controlled by the auto-RP mapping agent.

Auto-RP operates in PIM version 1 and version 2.

In most cases, how the routing device handles auto-RP discovery, announce, or mapping messages depends on whether the routing device is an RP (configured as local RP) or not. [Table 277 on page 3706](#) shows how the routing device behaves depending on the local RP configuration.

**Table 277: Local RP and Auto-RP Message Types**

| Auto-RP Message Type | Local RP? | Routing Device Behavior                                                                                                                                             |
|----------------------|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| discovery            | No        | Listen for auto-RP mapping messages.                                                                                                                                |
| discovery            | Yes       | Listen for auto-RP mapping messages.                                                                                                                                |
| announce             | No        | Listen for auto-RP mapping messages.                                                                                                                                |
| announce             | Yes       | Listen for auto-RP mapping messages. Send auto-RP announce messages.                                                                                                |
| mapping              | No        | Listen for auto-RP mapping messages. Listen for auto-RP announce messages. If elected mapping agent, send auto-RP mapping messages.                                 |
| mapping              | Yes       | Listen for auto-RP mapping messages. Send auto-RP announce messages. Listen for auto-RP announce messages. If elected mapping agent, send auto-RP mapping messages. |



**NOTE:** If the routing device receives auto-RP announcements split across multiple messages, the routing device loses the information in the previous part of the message as soon as the next part of the message is received.



You can configure auto-RP properties globally or for a routing instance. This example shows the global configuration.

To configure auto-RP properties:

1. Configure PIM in sparse-dense mode on all routing devices in the PIM domain.

```
[edit protocols pim]
user@host# edit
user@host# set interface all mode sparse-dense
```

This configuration allows the routing device to operate in sparse mode for most groups and dense mode for others. The default is to operate in sparse mode unless the routing device is specifically informed of a dense mode group.

2. Configure a routable loopback interface address on all routing devices in the PIM domain.

The routing device joins the auto-RP groups on the configured interfaces and on the loopback interface **lo0.0**. For auto-RP to work correctly, configure a routable IP address on the loopback interface. The router ID is used as the address for auto-RP updates. You cannot use the loopback address 127.0.0.1. Also, you must enable PIM sparse-dense mode on the **lo0.0** interface if you do not specify **interface all**.

```
[edit interfaces lo0.0 unit 0 family inet]
user@host# set address 192.168.0.3 preferred
```

3. Configure the two multicast dense groups on all the routing devices.

Auto-RP requires multicast flooding to announce potential RP candidates and to discover the elected RPs in the network. Multicast flooding occurs through a PIM dense mode model, where group 224.0.1.39 is used for **announce** messages and group 224.0.1.40 is used for **discovery** messages.

```
[edit protocols pim]
user@host# set dense-groups 224.0.1.39/32
user@host# set dense-groups 224.0.1.40/32
```



**TIP:** Step 3 is required. When auto-RP is enabled, the auto-RP announce group (224.0.1.39) and auto-RP-discovery group (224.0.1.40) must be configured explicitly as dense groups. When the auto-RP discovery group is not configured as a dense group, auto-RP is not enabled. When the auto-RP announce group is not configured as a dense group, auto-RP is enabled in the discovery mode only, and mapping and announce modes are disabled.

4. Configure the auto-RP **announce** option.

At least one routing device in the PIM domain must announce auto-RP messages and at least one must map them, or you can configure a routing device to perform both functions.

When a routing device sends announce messages in the network, it is advertising itself as a candidate RP. A routing device configured with this option must also be configured as an RP, or announce messages are not sent.

```
[edit protocols pim rp]
user@host# set local address 192.168.0.1
user@host# set auto-rp announce
```



**NOTE:** You cannot include the `auto-rp announce` option at the `[edit logical-systems logical-system-name routing-instances routing-instance-name protocols pim]` hierarchy level.

#### 5. Configure the auto-RP mapping agent.

The mapping agent sends discovery messages to the network, informing all routing devices in a multicast group of which RP to use. If the mapping agent is also an RP, the **mapping** option also allows the routing device to send auto-RP announcements (mapping on an RP allows the routing device to perform both the announcement and mapping functions).

```
[edit protocols pim rp]
user@host# set auto-rp mapping
```

If the mapping agent is also an RP, configure the mapping agent as a local RP.

```
[edit protocols pim rp]
user@host# set local address 192.168.0.2
```

#### 6. Configure mapping agent election.

If you configure the **mapping** option on more than one routing device in the PIM domain, configure mapping agent election on each potential mapping agent.

Auto-RP specifications state that mapping agents do not send mapping messages if they receive messages from a mapping agent with a higher IP address. However, some vendors' mapping agents continue to announce mappings, even in the presence of higher-addressed mapping agents. In other words, some mapping agents will always send mapping messages.

The default auto-RP operation is to perform mapping agent election. To explicitly configure mapping agent election, you can include the **mapping-agent-election** statement. When this option is configured, the mapping agent will stop sending mapping messages if it receives messages from a mapping agent with a higher IP address.

```
[edit protocols pim rp]
user@host# set auto-rp mapping mapping-agent-election
```

Mapping message suppression is disabled with the **no-mapping-agent-election** statement. When this option is configured, the mapping agent will always send mapping messages even in the presence of higher-addressed mapping agents.

To disable mapping agent election for compatibility with other vendors' equipment, include the **no-mapping-agent-election** statement.

```
[edit protocols pim rp]
user@host# set auto-rp mapping no-mapping-agent-election
```

7. Configure the remaining routing devices in the PIM domain to discover the RP.

Discovery enables the routing devices to receive and process discovery messages from the mapping agent. This is the most basic auto-RP option.

```
[edit protocols pim rp]
user@host# set auto-rp discovery
```

8. Monitor the operation of PIM auto-RP routers by running the following commands:

- `show pim interfaces`
- `show pim rps`
- `show pim rps`

9. Issue the `show pim rps extensive` command to see information about how an RP is learned, what groups it handles, and the number of groups actively using the RP.

```
user@host> show pim rps extensive
RP: 192.168.5.1
Learned from 192.168.5.1 via: auto-rp
Time Active: 00:34:29
Holdtime: 150 with 108 remaining
Device Index: 6
Subunit: 32769
Interface: pd-0/0/0.32769
Group Ranges:
 224.0.0.0/4
Active groups using RP:
 224.2.2.100
 total 1 groups active
Register State for RP:
Group Source FirstHop RP Address StateRP address Type Holdtime
Timeout
```

In the example, the RP at 192.168.5.1 was learned through auto-RP. The RP is able to support all groups in the 224.0.0.0/4 range (all possible groups). The local router has sent PIM control traffic for the 224.2.2.100 group to the RP.

Additionally, the presence of a Tunnel Physical Interface Card (PIC) in an RP router creates a de-encapsulation interface, which allows the RP to receive multicast traffic from the source. This interface is indicated by `pd-0/0/0.32769`.

#### Related Documentation

- [Configuring PIM Bootstrap Router on page 3701](#)
- [Configuring a Designated Router for PIM on page 3692](#)
- [Examples: Configuring PIM Sparse Mode](#)
- [Configuring Basic PIM Settings on page 3680](#)

#### Configuring Embedded RP

- [Understanding Embedded RP for IPv6 Multicast on page 3710](#)
- [Configuring PIM Embedded RP for IPv6 on page 3711](#)

### Understanding Embedded RP for IPv6 Multicast

Global IPv6 multicast between routing domains has been possible only with source-specific multicast (SSM) because there is no way to convey information about IPv6 multicast RPs between PIM sparse mode RPs. In IPv4 multicast networks, this information is conveyed between PIM RPs using MSDP, but there is no IPv6 support in current MSDP standards. IPv6 uses the concept of an embedded RP to resolve this issue without requiring SSM. This feature embeds the RP address in an IPv6 multicast address.

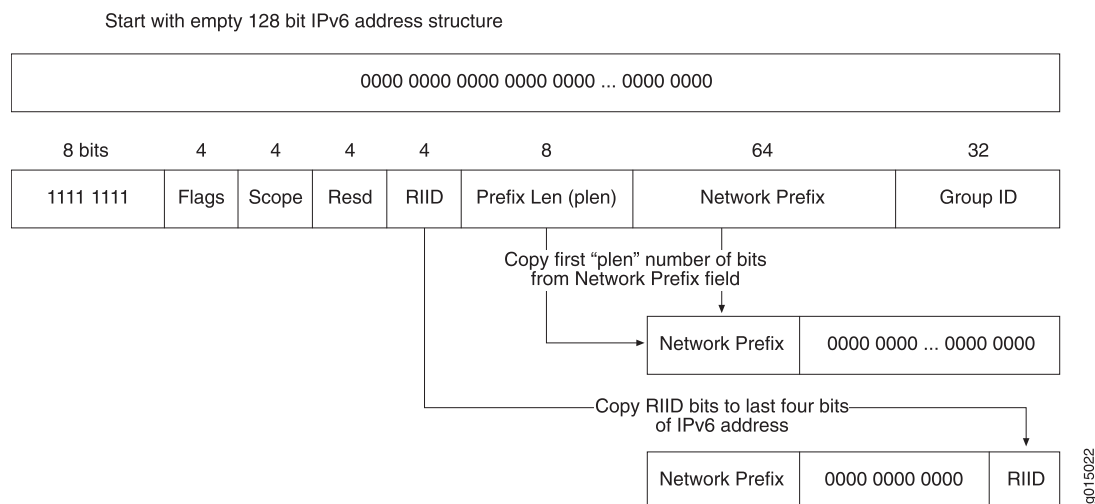
All IPv6 multicast addresses begin with 8 1-bits (1111 1111) followed by a 4-bit flag field normally set to 0011. The flag field is set to 0111 when embedded RP is used. Then the low-order bits of the normally reserved field in the IPv6 multicast address carry the 4-bit RP interface identifier (RIID).

When the IPv6 address of the RP is embedded in a unicast-prefix-based any-source multicast (ASM) address, all of the following conditions must be true:

- The address must be an IPv6 multicast address and have 0111 in the flags field (that is, the address is part of the prefix FF70::/12).
- The 8-bit prefix length (plen) field must not be all 0. An all 0 plen field implies that SSM is in use.
- The 8-bit prefix length field value must not be greater than 64, which is the length of the network prefix field in unicast-prefix-based ASM addresses.

The routing platform derives the value of the interdomain RP by copying the prefix length field number of bits from the 64-bit network prefix field in the received IPv6 multicast address to an empty 128-bit IPv6 address structure and copying the last bits from the 4-bit RIID. For example, if the prefix length field bits have the value 32, then the routing platform copies the first 32 bits of the IPv6 multicast address network prefix field to an all-0 IPv6 address and appends the last four bits determined by the RIID. See [Figure 88 on page 3710](#) for an illustration of this process.

**Figure 88: Extracting the Embedded RP IPv6 Address**



For example, the administrator of IPv6 network 2001:DB8::/32 sets up an RP for the 2001:DB8:BEEF:FEED::/96 subnet. In that case, the received embedded RP IPv6 ASM address has the form:

FF70:y40:2001:DB8:BEEF:FEED::/96

and the derived RP IPv6 address has the form:

2001:DB8:BEEF:FEED::y

where y is the RIID (y cannot be 0).

When configured, the routing platform checks for embedded RP information in every PIM join request received for IPv6. The use of embedded RP does not change the processing of IPv6 multicast and RPs in any way, except that the embedded RP address is used if available and selected for use. There is no need to specify the IPv6 address family for embedded RP configuration because the information can be used only if IPv6 multicast is properly configured on the routing platform.

The following receive events trigger extraction of an IPv6 embedded RP address on the routing platform:

- Multicast Listener Discovery (MLD) report for an embedded RP multicast group address
- PIM join message with an embedded RP multicast group address
- Static embedded RP multicast group address associated with an interface
- Packets sent to an embedded RP multicast group address received on the DR

The embedded RP node discovered through these events is added if it does not already exist on the routing platform. The routing platform chooses the embedded RP as the RP for a multicast group before choosing an RP learned through BSRs or a statically configured RP. The embedded RP is removed whenever all PIM join states using this RP are removed or the configuration changes to remove the embedded RP feature.

### ***Configuring PIM Embedded RP for IPv6***

You configure embedded RP to allow multidomain IPv6 multicast networks to find RPs in other routing domains. Embedded RP embeds an RP address inside PIM join messages and other types of messages sent between routing domains. Global IPv6 multicast between routing domains has been possible only with source-specific multicast (SSM) because there is no way to convey information about IPv6 multicast RPs between PIM sparse mode RPs. In IPv4 multicast networks, this information is conveyed between PIM RPs using MSDP, but there is no IPv6 support in current MSDP standards. IPv6 uses the concept of an embedded RP to resolve this issue without requiring SSM. Thus, embedded RP enables you can deploy IPv6 with any-source multicast (ASM).

Embedded RP is disabled by default.

When you configure embedded RP for IPv6, embedded RPs are preferred to RPs discovered by IPv6 any other way. You configure embedded RP independent of any other IPv6 multicast properties. This feature is applied only when IPv6 multicast is properly configured.

You can configure embedded RP globally or for a routing instance. This example shows the routing instance configuration.

To configure embedded RP for IPv6 PIM sparse mode:

1. Define which multicast addresses or prefixes can embed RP address information. If messages within a group range contain embedded RP information and the group range is not configured, the embedded RP in that group range is ignored. Any valid unicast-prefix-based ASM address can be used as a group range. The default group range is FF70::/12 to FFF0::/12. Messages with embedded RP information that do not match any configured group ranges are treated as normal multicast addresses.

```
[edit routing-instances vpn-A protocols pim rp embedded-rp]
user@host# set group-ranges fec0::/10
```

If the derived RP address is not a valid IPv6 unicast address, it is treated as any other multicast group address and is not used for RP information. Verification fails if the extracted RP address is a local interface, unless the routing device is configured as an RP and the extracted RP address matches the configured RP address. Then the local RP determines whether it is configured to act as an RP for the embedded RP multicast address.

2. Limit the number of embedded RPs created in a specific routing instance. The range is from 1 through 500. The default is 100.

```
[edit routing-instances vpn-A protocols pim rp]
user@host# set maximum-rps 50
```

3. Monitor the operation by running the **show pim rps** and **show pim statistics** commands.

#### Related Documentation

- [Configuring PIM Auto-RP on page 3705](#)
- [Configuring PIM Bootstrap Router on page 3701](#)
- [Configuring a Designated Router for PIM on page 3692](#)
- [Examples: Configuring PIM Sparse Mode](#)
- [Configuring Basic PIM Settings on page 3680](#)

---

#### Configuring PIM Filtering

- [Understanding Multicast Message Filters on page 3713](#)
- [Filtering MAC Addresses on page 3713](#)
- [Filtering RP and DR Register Messages on page 3714](#)
- [Filtering MSDP SA Messages on page 3715](#)
- [Configuring Interface-Level PIM Neighbor Policies on page 3715](#)
- [Filtering Outgoing PIM Join Messages on page 3716](#)
- [Example: Stopping Outgoing PIM Register Messages on a Designated Router on page 3717](#)
- [Filtering Incoming PIM Join Messages on page 3719](#)

- [Example: Rejecting Incoming PIM Register Messages on RP Routers on page 3721](#)
- [Configuring Register Message Filters on a PIM RP and DR on page 3723](#)

### ***Understanding Multicast Message Filters***

Multicast sources and routers generate a considerable number of control messages, especially when using PIM sparse mode. These messages form distribution trees, locate rendezvous points (RPs) and designated routers (DRs), and transition from one type of tree to another. In most cases, this multicast messaging system operates transparently and efficiently. However, in some configurations, more control over the sending and receiving of multicast control messages is necessary.

You can configure multicast filtering to control the sending and receiving of multicast control messages.

To prevent unauthorized groups and sources from registering with an RP router, you can define a routing policy to reject PIM register messages from specific groups and sources and configure the policy on the designated router or the RP router.

- If you configure the reject policy on an RP router, it rejects incoming PIM register messages from the specified groups and sources. The RP router also sends a register stop message by means of unicast to the designated router. On receiving the register stop message, the designated router sends periodic null register messages for the specified groups and sources to the RP router.
- If you configure the reject policy on a designated router, it stops sending PIM register messages for the specified groups and sources to the RP router.



**NOTE:** If you have configured the reject policy on an RP router, we recommend that you configure the same policy on all the RP routers in your multicast network.



**NOTE:** If you delete a group and source address from the reject policy configured on an RP router and commit the configuration, the RP router will register the group and source only when the designated router sends a null register message.

### ***Filtering MAC Addresses***

When a router is exclusively configured with multicast protocols on an interface, multicast sets the interface media access control (MAC) filter to multicast promiscuous mode, and the number of multicast groups is unlimited. However, when the router is not exclusively used for multicasting and other protocols such as OSPF, Routing Information Protocol version 2 (RIPv2), or Network Time Protocol (NTP) are configured on an interface, each of these protocols individually requests that the interface program the MAC filter to pick up its respective multicast group only. In this case, without multicast configured on the interface, the maximum number of multicast MAC filters is limited to 20. For example, the maximum number of interface MAC filters for protocols such as OSPF

(multicast group 224.0.0.5) is 20, unless a multicast protocol is also configured on the interface.

No configuration is necessary for MAC filters.

### ***Filtering RP and DR Register Messages***

You can filter Protocol Independent Multicast (PIM) register messages sent from the designated router (DR) or to the rendezvous point (RP). The PIM RP keeps track of all active sources in a single PIM sparse mode domain. In some cases, more control over which sources an RP discovers, or which sources a DR notifies other RPs about, is desired. A high degree of control over PIM register messages is provided by RP and DR register message filtering. Message filtering also prevents unauthorized groups and sources from registering with an RP router.

Register messages that are filtered at a DR are not sent to the RP, but the sources are available to local users. Register messages that are filtered at an RP arrive from source DRs, but are ignored by the router. Sources on multicast group traffic can be limited or directed by using RP or DR register message filtering alone or together.

If the action of the register filter policy is to discard the register message, the router needs to send a register-stop message to the DR. Register-stop messages are throttled to prevent malicious users from triggering them on purpose to disrupt the routing process.

Multicast group and source information is encapsulated inside unicast IP packets. This feature allows the router to inspect the multicast group and source information before sending or accepting the PIM register message.

Incoming register messages to an RP are passed through the configured register message filtering policy before any further processing. If the register message is rejected, the RP router sends a register-stop message to the DR. When the DR receives the register-stop message, the DR stops sending register messages for the filtered groups and sources to the RP. Two fields are used for register message filtering:

- Group multicast address
- Source address

The syntax of the existing policy statements is used to configure the filtering on these two fields. The **route-filter** statement is useful for multicast group address filtering, and the **source-address-filter** statement is useful for source address filtering. In most cases, the action is to **reject** the register messages, but more complex filtering policies are possible.

Filtering cannot be performed on other header fields, such as DR address, protocol, or port. In some configurations, an RP might not send register-stop messages when the policy action is to discard the register messages. This has no effect on the operation of the feature, but the router will continue to receive register messages.

When anycast RP is configured, register messages can be sent or received by the RP. All the RPs in the anycast RP set need to be configured with the same RP register message filtering policies. Otherwise, it might be possible to circumvent the filtering policy.



### Filtering MSDP SA Messages

Along with applying MSDP source active (SA) filters on all external MSDP sessions (in and out) to prevent SAs for groups and sources from leaking in and out of the network, you need to apply bootstrap router (BSR) filters. Applying a BSR filter to the boundary of a network prevents foreign BSR messages (which announce RP addresses) from leaking into your network. Since the routers in a PIM sparse-mode domain need to know the address of only one RP router, having more than one in the network can create issues.

If you did not use multicast scoping to create boundary filters for all customer-facing interfaces, you might want to use PIM join filters. Multicast scopes prevent the actual multicast data packets from flowing in or out of an interface. PIM join filters prevent PIM sparse-mode state from being created in the first place. Since PIM join filters apply only to the PIM sparse-mode state, it might be more beneficial to use multicast scoping to filter the actual data.



**NOTE:** When you apply firewall filters, firewall action modifiers, such as **log**, **sample**, and **count**, work only when you apply the filter on an inbound interface. The modifiers do not work on an outbound interface.

### Configuring Interface-Level PIM Neighbor Policies

You can configure a policy to filter unwanted PIM neighbors. In the following example, the PIM interface compares neighbor IP addresses with the IP address in the policy statement before any hello processing takes place. If any of the neighbor IP addresses (primary or secondary) match the IP address specified in the prefix list, PIM drops the hello packet and rejects the neighbor.

If you configure a PIM neighbor policy after PIM has already established a neighbor adjacency to an unwanted PIM neighbor, the adjacency remains intact until the neighbor hold time expires. When the unwanted neighbor sends another hello message to update its adjacency, the router recognizes the unwanted address and rejects the neighbor.

To configure a policy to filter unwanted PIM neighbors:

1. Configure the policy. The neighbor policy must be a properly structured policy statement that uses a prefix list (or a route filter) containing the neighbor primary address (or any secondary IP addresses) in a prefix list, and the **reject** option to reject the unwanted address.

```
[edit policy-options]
user@host# set prefix-list nbrGroup 1 20.20.20.1/32
user@host# set policy-statement nbr-policy from prefix-list nbrGroup1
user@host# set policy-statement nbr-policy then reject
```

2. Configure the interface globally or in the routing instance. This example shows the configuration for the routing instance.

```
[edit routing-instances PIM.master protocols pim]
user@host# set neighbor-policy nbr-policy
```

3. Verify the configuration by checking the **Hello dropped on neighbor policy** field in the output of the **show pim statistics** command.

### ***Filtering Outgoing PIM Join Messages***

When the core of your network is using MPLS, PIM join and prune messages stop at the customer edge (CE) routers and are not forwarded toward the core, because these routers do not have PIM neighbors on the core-facing interfaces. When the core of your network is using IP, PIM join and prune messages are forwarded to the upstream PIM neighbors in the core of the network.

When the core of your network is using a mix of IP and MPLS, you might want to filter certain PIM join and prune messages at the upstream egress interface of the CE routers.

You can filter PIM sparse mode (PIM-SM) join and prune messages at the egress interfaces for IPv4 and IPv6 in the upstream direction. The messages can be filtered based on the group address, source address, outgoing interface, PIM neighbor, or a combination of these values. If the filter is removed, the join is sent after the PIM periodic join timer expires.

To filter PIM sparse mode join and prune messages at the egress interfaces, create a policy rejecting the group address, source address, outgoing interface, or PIM neighbor, and then apply the policy.

The following example filters PIM join and prune messages for group addresses 224.0.1.2 and 225.1.1.1.

1. In configuration mode, create the policy.

```
user@host# set policy-options policy-statement block-groups term t1 from route-filter
224.0.1.2/32 exact
user@host# set policy-options policy-statement block-groups term t1 from route-filter
225.1.1.1/32 exact
user@host# set policy-options policy-statement block-groups term t1 then reject
user@host# set policy-options policy-statement block-groups term last then accept
```

2. Verify the policy configuration by running the **show policy-options** command.

```
user@host# show policy-options
policy-statement block-groups {
 term t1 {
 from {
 route-filter 224.0.1.2/32 exact;
 route-filter 225.1.1.1/32 exact;
 then reject;
 }
 term last {
 then accept;
 }
 }
}
```

3. Apply the PIM join and prune message filter.

```
user@host> set protocols pim export block-groups
```

4. After the configuration is committed, use the **show pim statistics** command to verify that outgoing PIM join and prune messages are being filtered.

```

user@host> show pim statistics | grep filtered
RP Filtered Source 0

Rx Joins/Prunes filtered 0

Tx Joins/Prunes filtered 254

```

The egress filter count is shown on the **Tx Joins/Prunes filtered** line.

### ***Example: Stopping Outgoing PIM Register Messages on a Designated Router***

This example shows how to stop outgoing PIM register messages on a designated router.

- [Requirements on page 3717](#)
- [Overview on page 3717](#)
- [Configuration on page 3718](#)
- [Verification on page 3719](#)

#### ***Requirements***

Before you begin:

1. Determine whether the router is directly attached to any multicast sources. Receivers must be able to locate these sources.
2. Determine whether the router is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
3. Determine whether to configure multicast to use sparse, dense, or sparse-dense mode. Each mode has different configuration considerations.
4. Determine the address of the RP if sparse or sparse-dense mode is used.
5. Determine whether to locate the RP with the static configuration, BSR, or auto-RP method.
6. Determine whether to configure multicast to use its own RPF routing table when configuring PIM in sparse, dense, or sparse-dense mode.
7. Configure the SAP and SDP protocols to listen for multicast session announcements.
8. Configure IGMP.
9. Configure the PIM static RP.
10. Filter PIM register messages from unauthorized groups and sources. See [“Example: Rejecting Incoming PIM Register Messages on RP Routers” on page 3721](#).

#### ***Overview***

In this example, you configure the group address as **224.2.2.2/32** and the source address in the group as **20.20.20.1/32**. You set the match action to not send PIM register messages for the group and source address. Then you configure the policy on the designated router to **stop-pim-register-msg-dr**.

### Configuration

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

```
set policy-options policy-statement stop-pim-register-msg-dr from route-filter
 224.2.2.2/32 exact
set policy-options policy-statement stop-pim-register-msg-dr from source-address-filter
 20.20.20.1/32 exact
set policy-options policy-statement stop-pim-register-msg-dr then reject
set protocols pim rp dr-register-policy stop-pim-register-msg-dr
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To stop outgoing PIM register messages on a designated router:

1. Configure the policy options.  

```
[edit]
user@host# edit policy-options
```
2. Set the group address.  

```
[edit policy-options]
user@host# set policy statement stop-pim-register-msg-dr from route-filter
 224.2.2.2/32 exact
```
3. Set the source address.  

```
[edit policy-options]
user@host# set policy statement stop-pim-register-msg-dr from
 source-address-filter 20.20.20.1/32 exact
```
4. Set the match action.  

```
[edit policy-options]
user@host# set policy statement stop-pim-register-msg-dr then reject
```
5. Assign the policy.  

```
[edit]
user@host# set dr-register-policy stop-pim-register-msg-dr
```

**Results** From configuration mode, confirm your configuration by entering the **show policy-options** and **show protocols** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show policy-options
policy-statement stop-pim-register-msg-dr {
 from {
 route-filter 224.2.2.2/32 exact;
 source-address-filter 20.20.20.1/32 exact;
```

```

 }
 then reject;
 }
[edit]
user@host# show protocols
pim {
 rp {
 dr-register-policy stop-pim-register-msg-dr;
 }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### **Verification**

To confirm that the configuration is working properly, perform these tasks:

- [Verifying SAP and SDP Addresses and Ports on page 3719](#)
- [Verifying the IGMP Version on page 3719](#)
- [Verifying the PIM Mode and Interface Configuration on page 3719](#)
- [Verifying the PIM RP Configuration on page 3719](#)

#### **Verifying SAP and SDP Addresses and Ports**

**Purpose** Verify that SAP and SDP are configured to listen on the correct group addresses and ports.

**Action** From operational mode, enter the **show sap listen** command.

#### **Verifying the IGMP Version**

**Purpose** Verify that IGMP version 2 is configured on all applicable interfaces.

**Action** From operational mode, enter the **show igmp interface** command.

#### **Verifying the PIM Mode and Interface Configuration**

**Purpose** Verify that PIM sparse mode is configured on all applicable interfaces.

**Action** From operational mode, enter the **show pim interfaces** command.

#### **Verifying the PIM RP Configuration**

**Purpose** Verify that the PIM RP is statically configured with the correct IP address.

**Action** From operational mode, enter the **show pim rps** command.

#### **Filtering Incoming PIM Join Messages**

Multicast scoping controls the propagation of multicast messages. Whereas multicast scoping prevents the actual multicast data packets from flowing in or out of an interface, PIM join filters prevent a state from being created in a router. A state—the (\*,G) or (S,G)

entries—is the information used for forwarding unicast or multicast packets. Using PIM join filters prevents the transport of multicast traffic across a network and the dropping of packets at a scope at the edge of the network. Also, PIM join filters reduce the potential for denial-of-service (DoS) attacks and PIM state explosion—large numbers of PIM join messages forwarded to each router on the rendezvous-point tree (RPT), resulting in memory consumption.

To use PIM join filters to efficiently restrict multicast traffic from certain source addresses, create and apply the routing policy across all routers in the network.

See [Table 278 on page 3720](#) for a list of match conditions.

**Table 278: PIM Join Filter Match Conditions**

| Match Condition              | Matches On                                                                           |
|------------------------------|--------------------------------------------------------------------------------------|
| <b>interface</b>             | Router interface or interfaces specified by name or IP address                       |
| <b>neighbor</b>              | Neighbor address (the source address in the IP header of the join and prune message) |
| <b>route-filter</b>          | Multicast group address embedded in the join and prune message                       |
| <b>source-address-filter</b> | Multicast source address embedded in the join and prune message                      |

The following example shows how to create a PIM join filter. The filter is composed of a route filter and a source address filter—**bad-groups** and **bad-sources**, respectively. the **bad-groups** filter prevents (\*G) or (S,G) join messages from being received for all groups listed. The **bad-sources** filter prevents (S,G) join messages from being received for all sources listed. The **bad-groups** filter and **bad-sources** filter are in two different terms. If route filters and source address filters are in the same term, they are logically ANDed.

To filter incoming PIM join messages:

1. Configure the policy.

```
[edit policy-statement pim-join-filter term bad-groups]
user@host# set from route-filter 224.0.1.2/32 exact
user@host# set from route-filter 239.0.0.0/8 orlonger
user@host# set then reject

[edit policy-statement pim-join-filter term bad-sources]
user@host# set from source-address-filter 10.0.0.0/8 orlonger
user@host# set from source-address-filter 127.0.0.0/8 orlonger
user@host# set then reject

[edit policy-statement pim-join-filter term last]
user@host# set then accept
```

2. Apply one or more policies to routes being imported into the routing table from PIM.

```
[edit protocols pim]
user@host# set import pim-join-filter
```

3. Verify the configuration by checking the output of the **show pim join** and **show policy** commands.

### ***Example: Rejecting Incoming PIM Register Messages on RP Routers***

This example shows how to reject incoming PIM register messages on RP routing devices.

- [Requirements on page 3721](#)
- [Overview on page 3721](#)
- [Configuration on page 3721](#)
- [Verification on page 3723](#)

### ***Requirements***

Before you begin:

1. Determine whether the routing device is directly attached to any multicast sources. Receivers must be able to locate these sources.
2. Determine whether the routing device is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
3. Determine whether to configure multicast to use sparse, dense, or sparse-dense mode. Each mode has different configuration considerations.
4. Determine the address of the RP if sparse or sparse-dense mode is used.
5. Determine whether to locate the RP with the static configuration, BSR, or auto-RP method.
6. Determine whether to configure multicast to use its own RPF routing table when configuring PIM in sparse, dense, or sparse-dense mode.
7. Configure the SAP and SDP protocols to listen for multicast session announcements. See *Configuring the Session Announcement Protocol*.
8. Configure IGMP. See [“Configuring IGMP” on page 3657](#).
9. Configure the PIM static RP. See [“Configuring Static RP” on page 3694](#).

### ***Overview***

In this example, you configure the group address as **224.1.1.1/32** and the source address in the group as **10.10.10.1/32**. You set the match action to reject PIM register messages and assign reject-pim-register-msg-rp as the policy on the RP.

### ***Configuration***

#### **CLI Quick Configuration**

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

```
set policy-options policy-statement reject-pim-register-msg-rp from route-filter
 224.1.1.1/32 exact
```

```
set policy-options policy-statement reject-pim-register-msg-rp from source-address-filter
10.10.10.1/32 exact
set policy-options policy-statement reject-pim-register-msg-rp then reject
set protocols pim rp rp-register-policy reject-pim-register-msg-rp
```

**Step-by-Step  
Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To reject the incoming PIM register messages on an RP routing device:

1. Configure the policy options.

```
[edit]
user@host# edit policy-options
```

2. Set the group address.

```
[edit policy-options]
user@host# set policy statement reject-pim-register-msg-rp from route-filter
224.1.1.1/32 exact
```

3. Set the source address.

```
[edit policy-options]
user@host# set policy statement reject-pim-register-msg-rp from
source-address-filter 10.10.10.1/32 exact
```

4. Set the match action.

```
[edit policy-options]
user@host# set policy statement reject-pim-register-msg-rp then reject
```

5. Configure the protocol.

```
[edit]
user@host# edit protocols pim rp
```

6. Assign the policy.

```
[edit]
user@host# set rp-register-policy reject-pim-register-msg-rp
```

**Results** From configuration mode, confirm your configuration by entering the **show policy-options** and **show protocols pim** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show policy-options
policy-statement reject-pim-register-msg-rp {
 from {
 route-filter 224.1.1.1/32 exact;
 source-address-filter 10.10.10.1/32 exact;
 }
 then reject;
}
[edit]
user@host# show protocols pim
rp {
```



```
rp-register-policy reject-pim-register-msg-rp;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

### **Verification**

To confirm that the configuration is working properly, perform these tasks:

- [Verifying SAP and SDP Addresses and Ports on page 3723](#)
- [Verifying the IGMP Version on page 3723](#)
- [Verifying the PIM Mode and Interface Configuration on page 3723](#)
- [Verifying the PIM Register Messages on page 3723](#)

### **Verifying SAP and SDP Addresses and Ports**

**Purpose** Verify that SAP and SDP are configured to listen on the correct group addresses and ports.

**Action** From operational mode, enter the **show sap listen** command.

### **Verifying the IGMP Version**

**Purpose** Verify that IGMP version 2 is configured on all applicable interfaces.

**Action** From operational mode, enter the **show igmp interface** command.

### **Verifying the PIM Mode and Interface Configuration**

**Purpose** Verify that PIM sparse mode is configured on all applicable interfaces.

**Action** From operational mode, enter the **show pim interfaces** command.

### **Verifying the PIM Register Messages**

**Purpose** Verify whether the rejected policy on the RP routing device is enabled.

**Action** From operational mode, enter the **show policy-options** and **show protocols pim** command.

### **Configuring Register Message Filters on a PIM RP and DR**

PIM register messages are sent to the rendezvous point (RP) by a designated router (DR). When a source for a group starts transmitting, the DR sends unicast PIM register packets to the RP.

Register messages have the following purposes:

- Notify the RP that a source is sending to a group.
- Deliver the initial multicast packets sent by the source to the RP for delivery down the shortest-path tree (SPT).

The PIM RP keeps track of all active sources in a single PIM sparse mode domain. In some cases, you want more control over which sources an RP discovers, or which sources a DR notifies other RPs about. A high degree of control over PIM register messages is provided by RP or DR register message filtering. Message filtering prevents unauthorized groups and sources from registering with an RP routing device.

You configure RP or DR register message filtering to control the number and location of multicast sources that an RP discovers. You can apply register message filters on a DR to control outgoing register messages, or apply them on an RP to control incoming register messages.

When anycast RP is configured, all RPs in the anycast RP set need to be configured with the same register message filtering policy.

You can configure message filtering globally or for a routing instance. These examples show the global configuration.

To configure an RP filter to drop the register packets for multicast group range 224.1.1.0/24 from source address 10.10.94.2:

1. On the RP, configure the policy.

```
[edit policy-options policy-statement incoming-policy-for-rp from]
user@host# set route-filter 224.1.1.0/24 orlonger
user@host# set source-address-filter 10.10.94.2/32 exact
user@host# set then reject
user@host# exit
```

2. Apply the policy to the RP.

```
[edit protocols pim rp]
user@host# set rp-register-policy incoming-policy-for-rp
user@host# set local address 10.10.10.5
user@host# exit
```

To configure a DR filter to prevent sending register packets for group range 224.1.1.0/24 and source address 10.10.10.1/32:

1. On the DR, configure the policy.

```
[edit policy-options policy-statement outgoing-policy-for-rp]
user@host# set from route-filter 224.1.1.0/24 orlonger
user@host# set from source-address-filter 10.10.10.1/32 exact
user@host# set then reject
user@host# exit
```

2. Apply the policy to the DR.

The static address is the address of the RP to which you do not want the DR to send the filtered register messages.

```
[edit protocols pim rp]
user@host# set dr-register-policy outgoing-policy-for-dr
user@host# set static 10.10.10.3
user@host# exit
```

To configure a policy expression to accept register messages for multicast group 224.1.1.5 but reject those for 224.1.1.1:

1. On the RP, configure the policies.

```
[edit policy-options policy-statement reject_224_1_1_1]
user@host# set from route-filter 224.1.1.0/24 orlonger
user@host# set from source-address-filter 10.10.94.2/32 exact
user@host# set then reject
user@host# exit

[edit policy-options policy-statement accept_224_1_1_5]
user@host# set term one from route-filter 224.1.1.5/32 exact
user@host# set term one from source-address-filter 10.10.94.2/32 exact
user@host# set term one then accept
user@host# set term two then reject
user@host# exit
```

2. Apply the policies to the RP.

```
[edit protocols pim rp]
user@host# set rp-register-policy [reject_224_1_1_1 | accept_224_1_1_5]
user@host# set local address 10.10.10.5
```

To monitor the operation of the filters, run the **show pim statistics** command. The command output contains the following fields related to filtering:

- RP Filtered Source
- Rx Joins/Prunes filtered
- Tx Joins/Prunes filtered
- Rx Register msgs filtering drop
- Tx Register msgs filtering drop

#### Related Documentation

- [Configuring PIM Auto-RP on page 3705](#)
- [Configuring PIM Bootstrap Router on page 3701](#)
- [Configuring PIM Dense Mode on page 3737](#)
- [Configuring a Designated Router for PIM on page 3692](#)
- [Example: Configuring Nonstop Active Routing for PIM](#)
- [Examples: Configuring PIM RPT and SPT Cutover](#)
- [Configuring PIM Sparse-Dense Mode on page 3740](#)
- [Configuring PIM and the Bidirectional Forwarding Detection \(BFD\) Protocol on page 3725](#)
- [Configuring Basic PIM Settings on page 3680](#)

#### Configuring PIM and the Bidirectional Forwarding Detection (BFD) Protocol

- [Understanding Bidirectional Forwarding Detection Authentication for PIM on page 3726](#)
- [Configuring BFD for PIM on page 3727](#)

- [Configuring BFD Authentication for PIM on page 3729](#)
- [Example: Configuring BFD Liveness Detection for PIM IPv6 on page 3732](#)

### ***Understanding Bidirectional Forwarding Detection Authentication for PIM***

Bidirectional Forwarding Detection (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when you run BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels.



**NOTE:** Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over PIM. BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

- [BFD Authentication Algorithms on page 3726](#)
- [Security Authentication Keychains on page 3727](#)
- [Strict Versus Loose Authentication on page 3727](#)

### ***BFD Authentication Algorithms***

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.

- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



**NOTE:** Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

### **Security Authentication Keychains**

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

### **Strict Versus Loose Authentication**

By default, strict authentication is enabled, and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

### **Configuring BFD for PIM**

The Bidirectional Forwarding Detection (BFD) Protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchanges BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the Protocol Independent Multicast (PIM) hello hold time, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

You must specify the minimum transmit and minimum receive intervals to enable BFD on PIM.

To enable failure detection:

1. Configure the interface globally or in a routing instance.

This example shows the global configuration.

```
[edit protocols pim]
user@host# edit interface fe-1/0/0.0 bfd-liveness-detection
```

2. Configure the minimum transmit interval.

This is the minimum interval after which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

```
[edit protocols pim interface fe-1/0/0.0 bfd-liveness-detection]
user@host# set transmit-interval 350
```

3. Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session.

Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set minimum-receive-interval 350
```

4. (Optional) Configure other BFD settings.

As an alternative to setting the receive and transmit intervals separately, configure one interval for both.

```
[edit protocols pim interface fe-1/0/0.0 bfd-liveness-detection]
user@host# set minimum-interval 350
```

5. Configure the threshold for the adaptation of the BFD session detection time.

When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

```
[edit protocols pim interface fe-1/0/0.0 bfd-liveness-detection]
user@host# set detection-time threshold 800
```

6. Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

```
[edit protocols pim interface fe-1/0/0.0 bfd-liveness-detection]
user@host# set multiplier 50
```

7. Configure the BFD version.

```
[edit protocols pim interface fe-1/0/0.0 bfd-liveness-detection]
user@host# set version 1
```

8. Specify that BFD sessions should not adapt to changing network conditions.

We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

```
[edit protocols pim interface fe-1/0/0.0 bfd-liveness-detection]
user@host# set no-adaptation
```

9. Verify the configuration by checking the output of the **show bfd session** command.

### ***Configuring BFD Authentication for PIM***

Beginning with Junos OS Release 9.6, you can configure authentication for Bidirectional Forwarding Detection (BFD) sessions running over Protocol Independent Multicast (PIM). Routing instances are also supported. The following steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the PIM protocol.
2. Associate the authentication keychain with the PIM protocol.
3. Configure the related security authentication keychain.

The following sections provide instructions for configuring and viewing BFD authentication on PIM:

- [Configuring BFD Authentication Parameters on page 3729](#)
- [Viewing Authentication Information for BFD Sessions on page 3730](#)

### ***Configuring BFD Authentication Parameters***

BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

To configure BFD authentication:

1. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on a PIM route or routing instance.

```
[edit protocols pim]
user@host# set interface if3-pim bfd-liveness-detection authentication algorithm
keyed-sha-1
```



**NOTE:** Nonstop active routing (NSR) is not supported with the meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

2. Specify the keychain to be used to associate BFD sessions on the specified PIM route or routing instance with the unique security authentication keychain attributes.

The keychain you specify must match the keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

```
[edit protocols pim]
user@host# set interface if3-pim bfd-liveness-detection authentication keychain
bfd-pim
```



**NOTE:** The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

3. Specify the unique security authentication information for BFD sessions:
  - The matching keychain name as specified in Step 2.
  - At least one key, a unique integer between 0 and 63. Creating multiple keys allows multiple clients to use the BFD session.
  - The secret data used to allow access to the session.
  - The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security]
user@host# set authentication-key-chains key-chain bfd-pim key 53 secret
9ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm start-time 2009-06-14.10:00:00
```

4. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit protocols pim]
user@host# set interface if3-pim bfd-liveness-detection authentication loose-check
```

5. (Optional) View your configuration by using the **show bfd session detail** or **show bfd session extensive** command.
6. Repeat these steps to configure the other end of the BFD session.

### ***Viewing Authentication Information for BFD Sessions***

You can view the existing BFD authentication configuration by using the **show bfd session detail** and **show bfd session extensive** commands.

The following example shows BFD authentication configured for the **if3-pim** BGP group. It specifies the keyed SHA-1 authentication algorithm and a keychain name of **bfd-pim**.



The authentication keychain is configured with two keys. Key 1 contains the secret data “\$9\$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm” and a start time of June 1, 2009, at 9:46:02 AM PST. Key 2 contains the secret data “\$9\$a5jiKW9L.reP38ny.TszF2/9” and a start time of June 1, 2009, at 3:29:20 PM PST.

```
[edit protocols pim]
interface if3-pim {
 bfd-liveness-detection {
 authentication {
 algorithm keyed-sha-1;
 key-chain bfd-pim;
 }
 }
}
[edit security]
authentication key-chains {
 key-chain bfd-pim {
 key 1 {
 secret "9ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm";
 start-time "2009-6-1.09:46:02 -0700";
 }
 key 2 {
 secret "9a5jiKW9L.reP38ny.TszF2/9";
 start-time "2009-6-1.15:29:20 -0700";
 }
 }
}
```

If you commit these updates to your configuration, you see output similar to the following example. In the output for the **show bfd session detail** command, **Authenticate** is displayed to indicate that BFD authentication is configured. For more information about the configuration, use the **show bfd session extensive** command. The output for this command provides the keychain name, the authentication algorithm and mode for each client in the session, and the overall BFD authentication configuration status, keychain name, and authentication algorithm and mode.

#### show bfd session detail

```
user@host# show bfd session detail
```

| Address  | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|----------|-------|------------|-------------|-------------------|------------|
| 50.0.0.2 | Up    | ge-0/1/5.0 | 0.900       | 0.300             | 3          |

Client PIM, TX interval 0.300, RX interval 0.300, **Authenticate**  
 Session up time 3d 00:34  
 Local diagnostic None, remote diagnostic NbrSignal  
 Remote state Up, version 1  
 Replicated

#### show bfd session extensive

```
user@host# show bfd session extensive
```

| Address  | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|----------|-------|------------|-------------|-------------------|------------|
| 50.0.0.2 | Up    | ge-0/1/5.0 | 0.900       | 0.300             | 3          |

Client PIM, TX interval 0.300, RX interval 0.300, **Authenticate**  
 keychain bfd-pim, algo keyed-sha-1, mode strict

```
Session up time 00:04:42
Local diagnostic None, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
Min async interval 0.300, min slow interval 1.000
Adaptive async TX interval 0.300, RX interval 0.300
Local min TX interval 0.300, minimum RX interval 0.300, multiplier 3
Remote min TX interval 0.300, min RX interval 0.300, multiplier 3
Local discriminator 2, remote discriminator 2
Echo mode disabled/inactive
Authentication enabled/active, keychain bfd-pim, algo keyed-sha-1, mode strict
```

### ***Example: Configuring BFD Liveness Detection for PIM IPv6***

This example shows how to configure Bidirectional Forwarding Detection (BFD) liveness detection for IPv6 interfaces configured for the Protocol Independent Multicast (PIM) topology. BFD is a simple hello mechanism that detects failures in a network.

The following steps are needed to configure BFD liveness detection:

1. Configure the interface.
2. Configure the related security authentication keychain.
3. Specify the BFD authentication algorithm for the PIM protocol.
4. Configure PIM, associating the authentication keychain with the desired protocol.
5. Configure BFD authentication for the routing instance.



**NOTE:** You must perform these steps on both ends of the BFD session.

---

- [Requirements on page 3732](#)
- [Overview on page 3732](#)
- [Configuration on page 3733](#)
- [Verification on page 3736](#)

### ***Requirements***

This example uses the following hardware and software components:


- Two peer routing devices.
- Junos OS 12.2 or later.

### ***Overview***

In this example, Device R1 and Device R2 are peers. Each routing device runs PIM, connected over a common medium.

[Figure 89 on page 3733](#) shows the topology used in this example.

Figure 89: BFD Liveness Detection for PIM IPv6 Topology



Assume that the routing devices initialize. No BFD session is yet established. For each routing device, PIM informs the BFD process to monitor the IPv6 address of the neighbor that is configured in the routing protocol. Addresses are not learned dynamically and must be configured.

Configure the IPv6 address and BFD liveness detection at the `[edit protocols pim]` hierarchy level for each routing device.

```
[edit protocols pim]
user@host# set interface interface-name family inet6 bfd-liveness-detection
```

Configure BFD liveness detection for the routing instance at the `[edit routing-instances instance-name protocols pim interface all family inet6]` hierarchy level (here, the *instance-name* is `instance1`):

```
[edit routing-instances instance1 protocols pim]
user@host# set bfd-liveness-detection
```

You will also configure the authentication algorithm and authentication keychain values for BFD.

In a BFD-configured network, when a client launches a BFD session with a peer, BFD begins sending slow, periodic BFD control packets that contain the interval values that you specified when you configured the BFD peers. This is known as the initialization state. BFD does not generate any up or down notifications in this state. When another BFD interface acknowledges the BFD control packets, the session moves into an up state and begins to more rapidly send periodic control packets. If a data path failure occurs and BFD does not receive a control packet within the configured amount of time, the data path is declared down and BFD notifies the BFD client. The BFD client can then perform the necessary actions to reroute traffic. This process can be different for different BFD clients.

### Configuration

#### CLI Quick Configuration

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

```
Device R1 set interfaces ge-1/1/1 unit 0 description toRouter2
 set interfaces ge-1/1/1 unit 0 family inet6
 set interfaces ge-1/1/1 unit 0 family inet6 address e80::21b:c0ff:fed5:e4dd
 set protocols pim interface ge-1/1/1 family inet6 bfd-liveness-detection authentication
 algorithm keyed-sha-1
 set protocols pim interface ge-1/1/1 family inet6 bfd-liveness-detection authentication
 key-chain bfd-pim
 set routing-instances instance1 protocols pim interface all family inet6
 bfd-liveness-detection authentication algorithm keyed-sha-1
```

```

set routing-instances instance1 protocols pim interface all family inet6
 bfd-liveness-detection authentication key-chain bfd-pim
set security authentication key-chain bfd-pim key 1 secret
 "9ggaJDmPQ6/tJgF/AtREVsyPsnCtUHM"
set security authentication key-chain bfd-pim key 1 start-time "2012-01-01.09:46:02
 -0700"
set security authentication key-chain bfd-pim key 2 secret
 "9a5jiKW9L.reP38ny.TszF2/9"
set security authentication key-chain bfd-pim key 2 start-time "2012-01-01.15:29:20
 -0700"

```

**Device R2**

```

set interfaces ge-1/1/0 unit 0 description toRouter1
set interfaces ge-1/1/0 unit 0 family inet6 address e80::21b:c0ff:fed5:e5dd
set protocols pim interface ge-1/1/0 family inet6 bfd-liveness-detection authentication
 algorithm keyed-sha-1
set protocols pim interface ge-1/1/0 family inet6 bfd-liveness-detection authentication
 key-chain bfd-pim
set routing-instances instance1 protocols pim interface all family inet6
 bfd-liveness-detection authentication algorithm keyed-sha-1
set routing-instances instance1 protocols pim interface all family inet6
 bfd-liveness-detection authentication key-chain bfd-pim
set security authentication key-chain bfd-pim key 1 secret
 "9ggaJDmPQ6/tJgF/AtREVsyPsnCtUHM"
set security authentication key-chain bfd-pim key 1 start-time "2012-01-01.09:46:02
 -0700"
set security authentication key-chain bfd-pim key 2 secret
 "9a5jiKW9L.reP38ny.TszF2/9"
set security authentication key-chain bfd-pim key 2 start-time "2012-01-01.15:29:20
 -0700"

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode”](#) on page 4704 in the *CLI User Guide*.

To configure BFD liveness detection for PIM IPv6 interfaces on Device R1:



**NOTE:** This procedure is for Device R1. Repeat this procedure for Device R2, after modifying the appropriate interface names, addresses, and any other parameters.

1. Configure the interface, using the **inet6** statement to specify that this is an IPv6 address.  

```

[edit interfaces]
user@R1# set ge-1/1/1 unit 0 description toRouter2
user@R1# set ge-1/1/1 unit 0 family inet6 address e80::21b:c0ff:fed5:e4dd

```
2. Specify the BFD authentication algorithm and keychain for the PIM protocol.  

The keychain is used to associate BFD sessions on the specified PIM route or routing instance with the unique security authentication keychain attributes. This keychain

name should match the keychain name configured at the **[edit security authentication]** hierarchy level.

```
[edit protocols]
user@R1# set pim interface ge-1/1/1.0 family inet6 bfd-liveness-detection
authentication algorithm keyed-sha-1
user@R1# set pim interface ge-1/1/1 family inet6 bfd-liveness-detection
authentication key-chain bfd-pim
```



**NOTE:** The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

3. Configure a routing instance (here, **instance1**), specifying BFD authentication and associating the security authentication algorithm and keychain.

```
[edit routing-instances]
user@R1# set instance1 protocols pim interface all family inet6
bfd-liveness-detection authentication algorithm keyed-sha-1
user@R1# set instance1 protocols pim interface all family inet6
bfd-liveness-detection authentication key-chain bfd-pim
```

4. Specify the unique security authentication information for BFD sessions:
  - The matching keychain name as specified in Step 2.
  - At least one key, a unique integer between **0** and **63**. Creating multiple keys allows multiple clients to use the BFD session.
  - The secret data used to allow access to the session.
  - The time at which the authentication key becomes active, in the format **YYYY-MM-DD.hh:mm:ss**.

```
[edit security authentication]
user@R1# set key-chain bfd-pim key 1 secret
"9ggaJDmPQ6/tJgF/AtREVsyPsnCtUHM"
user@R1# set key-chain bfd-pim key 1 start-time "2012-01-01.09:46:02 -0700"
user@R1# set key-chain bfd-pim key 2 secret "9a5jiKW9l.reP38ny.TszF2/9"
user@R1# set key-chain bfd-pim key 2 start-time "2012-01-01.15:29:20 -0700"
```

## Results

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

```
user@R1# show interfaces
ge-1/1/1 {
 unit 0 {
 description toRouter2;
 family inet6 {
 address e80::21b:c0ff:fed5:e4dd {
 }
 }
 }
}
```

```
 }
 }
}

user@R1# show protocols
pim {
 interface ge-1/1/1.0 {
 family inet6;
 bfd-liveness-detection {
 authentication {
 algorithm keyed-sha-1;
 key-chain bfd-pim;
 }
 }
 }
}

user@R1# show routing-instances
instance1 {
 protocols {
 pim {
 interface all {
 family inet 6 {
 bfd-liveness-detection {
 authentication {
 algorithm keyed-sha-1;
 key-chain bfd-pim;
 }
 }
 }
 }
 }
 }
}

user@R1# show security
authentication {
 key-chain bfd-pim {
 key 1 {
 secret "9ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm";
 start-time "2012-01-01.09:46:02 -0700";
 }
 key 2 {
 secret "9a5jiKW9l.reP38ny.TszF2/9";
 start-time "2012-01-01.15:29:20 -0700";
 }
 }
}
```

### **Verification**

Confirm that the configuration is working properly.

#### **Verifying the BFD Session**

**Purpose** Verify that BFD liveness detection is enabled.

**Action** user@R1# run `show pim neighbors detail`

Instance: PIM.master

Interface: ge-1/1/1.0

Address: fe80::21b:c0ff:fed5:e4dd, IPv6, PIM v2, Mode: Sparse, sg Join Count: 0, tsf Join Count: 0

Hello Option Holdtime: 65535 seconds

Hello Option DR Priority: 1

Hello Option Generation ID: 1417610277

Hello Option LAN Prune Delay: delay 500 ms override 2000 ms  
Join Suppression supported

Address: fe80::21b:c0ff:fedc:28dd, IPv6, PIM v2, sg Join Count: 0, tsf Join Count: 0

Secondary address: beef::2

BFD: Enabled, Operational state: Up

Hello Option Holdtime: 105 seconds 80 remaining

Hello Option DR Priority: 1

Hello Option Generation ID: 1648636754

Hello Option LAN Prune Delay: delay 500 ms override 2000 ms  
Join Suppression supported

**Meaning** The display from the `show pim neighbors detail` command shows **BFD: Enabled, Operational state: Up**, indicating that BFD is operating between the two PIM neighbors. For additional information about the BFD session (including the session ID number), use the `show bfd session extensive` command.

- Related Documentation**
- [Configuring Basic PIM Settings on page 3680](#)
  - *Example: Configuring BFD for BGP*
  - *Example: Configuring BFD Authentication for BGP*

## Configuring PIM Dense Mode

- [Understanding PIM Dense Mode on page 3737](#)
- [Configuring PIM Dense Mode Properties on page 3739](#)

### Understanding PIM Dense Mode

PIM dense mode is less sophisticated than PIM sparse mode. PIM dense mode is useful for multicast LAN applications, the main environment for all dense mode protocols.

PIM dense mode implements the same flood-and-prune mechanism that DVMRP and other dense mode routing protocols employ. The main difference between DVMRP and PIM dense mode is that PIM dense mode introduces the concept of protocol independence. PIM dense mode can use the routing table populated by any underlying unicast routing protocol to perform reverse-path-forwarding (RPF) checks.


Internet service providers (ISPs) typically appreciate the ability to use any underlying unicast routing protocol with PIM dense mode because they do not need to introduce and manage a separate routing protocol just for RPF checks. While unicast routing protocols extended as multiprotocol BGP (MBGP) and Multitopology Routing in IS-IS

(M-IS-IS) were later employed to build special tables to perform RPF checks, PIM dense mode does not require them.

PIM dense mode can use the unicast routing table populated by OSPF, IS-IS, BGP, and so on, or PIM dense mode can be configured to use a special multicast RPF table populated by MBGP or M-IS-IS when performing RPF checks.

Unlike sparse mode, in which data is forwarded only to routing devices sending an explicit request, dense mode implements a *flood-and-prune* mechanism, similar to DVMRP. In PIM dense mode, there is no RP. A routing device receives the multicast data on the interface closest to the source, then forwards the traffic to all other interfaces (see [Figure 90 on page 3738](#)).


**Figure 90: Multicast Traffic Flooded from the Source Using PIM Dense Mode**



Flooding occurs periodically. It is used to refresh state information, such as the source IP address and multicast group pair. If the routing device has no interested receivers for the data, and the OIL becomes empty, the routing device sends a prune message upstream to stop delivery of multicast traffic (see [Figure 91 on page 3739](#)).



**Figure 91: Prune Messages Sent Back to the Source to Stop Unwanted Multicast Traffic**



### Configuring PIM Dense Mode Properties

In PIM dense mode (PIM-DM), the assumption is that almost all possible subnets have at least one receiver wanting to receive the multicast traffic from a source, so the network is flooded with traffic on all possible branches, then pruned back when branches do not express an interest in receiving the packets, explicitly (by message) or implicitly (time-out silence). LANs are appropriate networks for dense-mode operation.

By default, PIM is disabled. When you enable PIM, it operates in sparse mode by default.

You can configure PIM dense mode globally or for a routing instance. This example shows how to configure the routing instance and how to specify that PIM dense mode use **inet.2** as its RPF routing table instead of **inet.0**.

To configure the routing device properties for PIM dense mode:

1. (Optional) Create an IPv4 routing table group so that interface routes are installed into two routing tables, **inet.0** and **inet.2**.

```
[edit routing-options rib-groups]
user@host# set pim-rg export-rib inet.0
user@host# set pim-rg import-rib [inet.0 inet.2]
```

2. (Optional) Associate the routing table group with a PIM routing instance.

```
[edit routing-instances PIM.dense protocols pim]
user@host# set rib-group inet pim-rg
```

3. Configure the PIM interface. If you do not specify any interfaces, PIM is enabled on all routing device interfaces. Generally, you specify interface names only if you are disabling PIM on certain interfaces.

```
[edit routing-instances PIM.dense protocols pim]
user@host# set interface fe-0/0/1.0 mode dense
```



**NOTE:** You cannot configure both PIM and Distance Vector Multicast Routing Protocol (DVMRP) in forwarding mode on the same interface. You can configure PIM on the same interface only if you configured DVMRP in unicast-routing mode.

4. Monitor the operation of PIM dense mode by running the **show pim interfaces**, **show pim join**, **show pim neighbors**, and **show pim statistics** commands.

#### Related Documentation

- [Configuring PIM Sparse-Dense Mode on page 3740](#)
- [Configuring Basic PIM Settings on page 3680](#)

### Configuring PIM Sparse-Dense Mode

- [Understanding PIM Sparse-Dense Mode on page 3740](#)
- [Mixing PIM Sparse and Dense Modes on page 3740](#)
- [Configuring PIM Sparse-Dense Mode Properties on page 3741](#)

#### Understanding PIM Sparse-Dense Mode

Sparse-dense mode, as the name implies, allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as dense is not mapped to an RP. Instead, data packets destined for that group are forwarded by means of PIM dense-mode rules. A group specified as sparse is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules.

For information about PIM sparse-mode and PIM dense-mode rules, see *Understanding PIM Sparse Mode* and “[Understanding PIM Dense Mode](#)” on page 3737.

#### Mixing PIM Sparse and Dense Modes

It is possible to mix PIM dense mode, PIM sparse mode, and PIM source-specific multicast (SSM) on the same network, the same routing device, and even the same interface. This is because modes are effectively tied to multicast groups, an IP multicast group address must be unique for a particular group's traffic, and scoping limits enforce the division between potential or actual overlaps.



**NOTE:** PIM sparse mode was capable of forming shortest-path trees (SPTs) already. Changes to PIM sparse mode to support PIM SSM mainly involved defining behavior in the SSM address range, because shared-tree behavior is prohibited for groups in the SSM address range.

A multicast routing device employing sparse-dense mode is a good example of mixing PIM modes on the same network or routing device or interface. Dense modes are easy to support because of the flooding, but scaling issues make dense modes inappropriate for Internet use beyond very restricted uses.

### ***Configuring PIM Sparse-Dense Mode Properties***

Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as “dense” is not mapped to an RP. Instead, data packets destined for that group are forwarded by means of PIM dense mode rules. A group specified as “sparse” is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules. Sparse-dense mode is useful in networks implementing auto-RP for PIM sparse mode.

By default, PIM is disabled. When you enable PIM, it operates in sparse mode by default.

You can configure PIM sparse-dense mode globally or for a routing instance. This example shows how to configure PIM sparse-dense mode globally on all interfaces, specifying that the groups 224.0.1.39 and 224.0.1.40 are using dense mode.

To configure the routing device properties for PIM sparse-dense mode:

1. Configure the dense-mode groups.

```
[protocols pim]
user@host# set dense-groups 224.0.1.39
user@host# set dense-groups 224.0.1.40
```

2. Configure all interfaces on the routing device to use sparse-dense mode. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

```
[edit protocols pim]
user@host# set interface all mode sparse-dense
user@host# set interface fxp0.0 disable
```

3. Monitor the operation of PIM sparse-dense mode by running the **show pim interfaces**, **show pim join**, **show pim neighbors**, and **show pim statistics** commands.

#### **Related Documentation**

- [Configuring PIM Dense Mode on page 3737](#)
- [Configuring Basic PIM Settings on page 3680](#)

### ***PIM Join Load Balancing on Multipath MVPN Routes Overview***

A multicast virtual private network (MPVN) is a technology to deploy the multicast service in an existing MPLS/BGP VPN.

The two main MVPN services are:

- Dual PIM MVPNs (also referred to as Draft-Rosen)
- Multiprotocol BGP-based MVPNs (also referred to as next-generation)

Next-generation MVPNs constitute the next evolution after the Draft-Rosen MVPN and provide a simpler solution for administrators who want to configure multicast over Layer 3 VPNs. A Draft-Rosen MVPN uses Protocol Independent Multicast (PIM) for customer multicast (C-multicast) signaling, and a next-generation MVPN uses BGP for C-multicast signaling.

Multipath routing in an MVPN is applied to make data forwarding more robust against network failures and to minimize shared backup capacities when resilience against network failures is required.

By default, PIM join messages are sent toward a source based on the reverse path forwarding (RPF) routing table check. If there is more than one equal-cost path toward the source [S, G] or rendezvous point (RP) [\*; G], then one upstream interface is used to send the join messages. The upstream path can be:

- A single active external BGP (EBGP) path when both EBGP and internal BGP (IBGP) paths are present.
- A single active IBGP path when there is no EBGP path present.

With the introduction of the multipath PIM join load-balancing feature, customer PIM (C-PIM) join messages are load-balanced in the following ways:

- In the case of a Draft-Rosen MVPN, unequal EBGP and IBGP paths are utilized.
- In the case of next-generation MVPN:
  - Available IBGP paths are utilized when no EBGP path is present.
  - Available EBGP paths are utilized when both EBGP and IBGP paths are present.

This feature is applicable to IPv4 C-PIM join messages over the Layer 3 MVPN service.


By default, a customer source (C-S) or a customer RP (C-RP) is considered remote if the active **rt\_entry** is a secondary route and the primary route is present in a different routing instance. Such determination is being done without taking into consideration the (C-\*;G) or (C-S,G) state for which the check is being performed. The multipath PIM join load-balancing feature determines if a source (or RP) is remote by taking into account the associated (C-\*;G) or (C-S,G) state.

When the provider network does not have provider edge (PE) routing devices with the multipath PIM join load-balancing feature enabled, hash-based join load balancing is used. Although the decision to configure this feature does not impact PIM or overall system performance, network performance can be affected temporarily, if the feature is not enabled.

With hash-based join load balancing, adding new PE routing devices to the candidate upstream toward the C-S or C-RP results in C-PIM join messages being redistributed to new upstream paths. If the number of join messages is large, network performance is impacted because of join messages being sent to the new RPF neighbor and prune messages being sent to the old RPF neighbor. In next-generation MVPN, this results in BGP C-multicast data messages being withdrawn from old upstream paths and advertised on new upstream paths, impacting network performance.

In [Figure 92 on page 3743](#), PE1 and PE2 are the upstream PE routing devices. Router PE1 learns route Source from EBGP and IBGP peers—the customer edge CE1 router and the PE2 router, respectively.

**Figure 92: PIM Join Load Balancing**



- If the PE routing devices run the Draft-Rosen MVPN, the PE1 router distributes C-PIM join messages between the EBGP path to the CE1 router and the IBGP path to the PE2 router. The join messages on the IBGP path are sent over a multicast tunnel interface through which the PE routing devices establish C-PIM adjacency with each other.

If a PE routing device loses one or all EBGP paths toward the source (or RP), the C-PIM join messages that were previously using the EBGP path are moved to a multicast tunnel interface, and the RPF neighbor on the multicast tunnel interface is selected based on a hash mechanism.

On discovering the first EBGP path toward the source (or RP), only new join messages get load-balanced across EBGP and IBGP paths, whereas the existing join messages on the multicast tunnel interface remain unaffected.

- If the PE routing devices run the next-generation MVPN, the PE1 router sends C-PIM join messages directly to the CE1 router over the EBGP path. There is no C-PIM adjacency between the PE1 and PE2 routers. Router PE3 distributes the C-PIM join messages between the two IBGP paths to PE1 and PE2. The Bitwise-XOR hash algorithm is used to send the C-multicast data according to Internet draft [draft-ietf-l3vpn-2547bis-mcast-bgp](#), *BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs*.

Because the multipath PIM join load-balancing feature in a Draft-Rosen MVPN utilizes unequal EGBP and IBGP paths to the destination, loops can be created when forwarding unicast packets to the destination. To avoid or break such loops:

- Traffic arriving from a core or master instance should not be forwarded back to the core facing interfaces.
- A single multicast tunnel interface should either be selected as the upstream interface or the downstream interface.
- An upstream or downstream multicast tunnel interface should point to a non-multicast tunnel interface.

As a result of the loop avoidance mechanism, join messages arriving from an EGBP path get load-balanced across EIBGP paths as expected, whereas join messages from an IBGP path are constrained to choose the EGBP path only.

In [Figure 92 on page 3743](#), if the CE2 host sends unicast data traffic to the CE1 host, the PE1 router could send the multicast flow to the PE2 router over the MPLS core due to traffic load balancing. A data forwarding loop is prevented by ensuring that PE2 does not forward traffic back on the MPLS core because of the load-balancing algorithm.

In the case of C-PIM join messages, assuming that both the CE2 host and the CE3 host are interested in receiving traffic from the source (S, G), and if both PE1 and PE2 choose each other as the RPF neighbor toward the source, then a multicast tree cannot be formed completely. This feature implements mechanisms to prevent such join loops in the multicast control plane in a Draft-Rosen MVPN scenario.

**NOTE:**

Disruption of multicast traffic or creation of join loops can occur, resulting in a multicast distribution tree (MDT) not being formed properly due to one of the following reasons:

- During a graceful Routing Engine switchover (GRES), the EIBGP path selection for C-PIM join messages can vary, because the upstream interface selection is performed again for the new Routing Engine based on the join messages it receives from the CE and PE neighbors. This can lead to disruption of multicast traffic depending on the number of join messages received and the load on the network at the time of the graceful restart. However, nonstop active routing (NSR) is not supported and has no impact on the multicast traffic in a Draft-Rosen MVPN scenario.
  - Any PE routing device in the provider network is running another vendor's implementation that does not apply the same hashing algorithm implemented in this feature.
  - The multipath PIM join load-balancing feature has not been configured properly.
- 

**Related Documentation**

- *Use Case for PIM Join Load Balancing*

- [Example: Configuring PIM Join Load Balancing on Draft-Rosen Multicast VPN on page 5182](#)
- [Example: Configuring PIM Join Load Balancing On Next-Generation Multicast VPN on page 5190](#)

## PIM Snooping for VPLS

- [Understanding PIM Snooping for VPLS on page 3745](#)
- [Example: Configuring PIM Snooping for VPLS on page 3746](#)

### Understanding PIM Snooping for VPLS

There are two ways to direct PIM control packets:

- By the use of PIM snooping
- By the use of PIM proxying

PIM snooping configures a device to examine and operate only on PIM hello and join/prune packets. A PIM snooping device snoops PIM hello and join/prune packets on each interface to find interested multicast receivers and populates the multicast forwarding tree with this information. PIM snooping differs from PIM proxying in that both PIM hello and join/prune packets are transparently flooded in the VPLS as opposed to the flooding of only hello packets in the case of PIM proxying. PIM snooping is configured on PE routers connected through pseudowires. PIM snooping ensures that no new PIM packets are generated in the VPLS, with the exception of PIM messages sent through LDP on pseudowires.



**NOTE:** In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

A device that supports PIM snooping snoops hello packets received on attachment circuits. It does not introduce latency in the VPLS core when it forwards PIM join/prune packets.

To configure PIM snooping on a PE router, use the **pim-snooping** statement at the **[edit routing-instances *instance-name* protocols]** hierarchy level:

```
routing-instances {
 customer {
 instance-type vpls;
 ...
 protocols {
 pim-snooping{
 traceoptions {
 file pim.log size 10m;
 flag all;
 flag timer disable;
 }
 }
 }
 }
}
```

```
}
```

“[Example: Configuring PIM Snooping for VPLS](#)” on [page 3746](#) explains the PIM snooping method. The use of the PIM proxying method is not discussed here and is outside the scope of this document. For more information about PIM proxying, see [PIM Snooping over VPLS](#).

### **Example: Configuring PIM Snooping for VPLS**

This example shows how to configure PIM snooping in a virtual private LAN service (VPLS) to restrict multicast traffic to interested devices.

- [Requirements on page 3746](#)
- [Overview on page 3746](#)
- [Configuration on page 3747](#)
- [Verification on page 3753](#)

### **Requirements**

This example uses the following hardware and software components:

- MX Series 3D Universal Edge Routers (except MX80)
- Junos OS Release 12.3 or later

### **Overview**

The following example shows how to configure PIM snooping to restrict multicast traffic to interested devices in a VPLS.



**NOTE:** This example demonstrates PIM snooping by the use of a PIM snooping device to restrict multicast traffic. The use of the PIM proxying method to achieve PIM snooping is out of the scope of this document and is yet to be implemented in Junos OS.

### **Topology**


In this example, two PE routing devices are connected to each other through a pseudowire connection. Router PE1 is connected to Routers CE1 and CE2. A multicast receiver is attached to Router CE2. Router PE2 is connected to Routers CE3 and CE4. A multicast source is connected to Router CE3, and a second multicast receiver is attached to Router CE4.

PIM snooping is configured on Routers PE1 and PE2. Hence, data sent from the multicast source is received only by members of the multicast group.

[Figure 93 on page 3747](#) shows the topology used in this example.



Figure 93: PIM Snooping for VPLS



### Configuration

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

```
Router PE1
set multicast-snooping-options traceoptions file snoop.log size 10m
set interfaces ge-2/0/0 encapsulation ethernet-vpls
set interfaces ge-2/0/0 unit 0 description toCE1
set interfaces ge-2/0/1 encapsulation ethernet-vpls
set interfaces ge-2/0/1 unit 0 description toCE2
set interfaces ge-2/0/2 unit 0 description toPE2
set interfaces ge-2/0/2 unit 0 family inet address 10.0.0.1/30
set interfaces ge-2/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.1.1/32
set routing-options router-id 10.255.1.1
set protocols mpls interface ge-2/0/1.0
set protocols bgp group toPE2 type internal
set protocols bgp group toPE2 local-address 10.255.1.1
set protocols bgp group toPE2 family l2vpn signaling
set protocols bgp group toPE2 neighbor 10.255.7.7
set protocols ospf area 0.0.0.0 interface ge-2/0/2.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-2/0/2.0
set protocols ldp interface lo0.0
set routing-instances titanium instance-type vpls
set routing-instances titanium vlan-id none
set routing-instances titanium interface ge-2/0/0.0
set routing-instances titanium interface ge-2/0/1.0
set routing-instances titanium route-distinguisher 101:101
set routing-instances titanium vrf-target target:201:201
```

```
set routing-instances titanium protocols vpls vpls-id 15
set routing-instances titanium protocols vpls site pe1 site-identifier 1
set routing-instances titanium protocols pim-snooping

Router CE1 set interfaces ge-2/0/0 unit 0 description toPE1
 set interfaces ge-2/0/0 unit 0 family inet address 10.0.0.10/30
 set interfaces lo0 unit 0 family inet address 10.255.2.2/32
 set routing-options router-id 10.255.2.2
 set protocols ospf area 0.0.0.0 interface all
 set protocols ospf area 0.0.0.0 interface lo0.0 passive
 set protocols pim rp static address 10.255.3.3
 set protocols pim interface all

Router CE2 set interfaces ge-2/0/0 unit 0 description toPE1
 set interfaces ge-2/0/0 unit 0 family inet address 10.0.0.6/30
 set interfaces ge-2/0/1 unit 0 description toReceiver1
 set interfaces ge-2/0/1 unit 0 family inet address 10.0.0.13/30
 set interfaces lo0 unit 0 family inet address 10.255.2.2
 set routing-options router-id 10.255.2.2
 set protocols ospf area 0.0.0.0 interface all
 set protocols ospf area 0.0.0.0 interface lo0.0 passive
 set protocols pim rp static address 10.255.3.3
 set protocols pim interface all

Router PE2 set multicast-snooping-options traceoptions file snoop.log size 10m
 set interfaces ge-2/0/0 encapsulation ethernet-vpls
 set interfaces ge-2/0/0 unit 0 description toCE3
 set interfaces ge-2/0/1 encapsulation ethernet-vpls
 set interfaces ge-2/0/1 unit 0 description toCE4
 set interfaces ge-2/0/2 unit 0 description toPE1
 set interfaces ge-2/0/2 unit 0 family inet address 10.0.0.2/30
 set interfaces ge-2/0/2 unit 0 family mpls
 set interfaces lo0 unit 0 family inet address 10.255.7.7/32
 set routing-options router-id 10.255.7.7
 set protocols mpls interface ge-2/0/2.0
 set protocols bgp group toPE1 type internal
 set protocols bgp group toPE1 local-address 10.255.7.7
 set protocols bgp group toPE1 family l2vpn signaling
 set protocols bgp group toPE1 neighbor 10.255.1.1
 set protocols ospf area 0.0.0.0 interface ge-2/0/2.0
 set protocols ospf area 0.0.0.0 interface lo0.0
 set protocols ldp interface ge-2/0/2.0
 set protocols ldp interface lo0.0
 set routing-instances titanium instance-type vpls
 set routing-instances titanium vlan-id none
 set routing-instances titanium interface ge-2/0/0.0
 set routing-instances titanium interface ge-2/0/1.0
 set routing-instances titanium route-distinguisher 101:101
 set routing-instances titanium vrf-target target:201:201
 set routing-instances titanium protocols vpls vpls-id 15
 set routing-instances titanium protocols vpls site pe2 site-identifier 2
 set routing-instances titanium protocols pim-snooping

Router CE3 (RP) set interfaces ge-2/0/0 unit 0 description toPE2
 set interfaces ge-2/0/0 unit 0 family inet address 10.0.0.18/30
```

```

set interfaces ge-2/0/1 unit 0 description toSource
set interfaces ge-2/0/1 unit 0 family inet address 10.0.0.29/30
set interfaces lo0 unit 0 family inet address 10.255.3.3/32
set routing-options router-id 10.255.3.3
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols pim rp local address 10.255.3.3
set protocols pim interface all

```

**Router CE4**

```

set interfaces ge-2/0/0 unit 0 description toPE2
set interfaces ge-2/0/0 unit 0 family inet address 10.0.0.22/30
set interfaces ge-2/0/1 unit 0 description toReceiver2
set interfaces ge-2/0/1 unit 0 family inet address 10.0.0.25/30
set interfaces lo0 unit 0 family inet address 10.255.4.4/32
set routing-options router-id 10.255.4.4
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols pim rp static address 10.255.3.3
set protocols pim interface all

```

**Configuring PIM Snooping for VPLS****Step-by-Step  
Procedure**

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.



**NOTE:** This section includes a step-by-step configuration procedure for one or more routing devices in the topology. For comprehensive configurations for all routing devices, see [“CLI Quick Configuration” on page 3747](#).

To configure PIM snooping for VPLS:

1. Configure the router interfaces forming the links between the routing devices.

**Router PE2**

[edit interfaces]

```

user@PE2# set ge-2/0/0 encapsulation ethernet-vpls
user@PE2# set ge-2/0/0 unit 0 description toCE3
user@PE2# set ge-2/0/1 encapsulation ethernet-vpls
user@PE2# set ge-2/0/1 unit 0 description toCE4
user@PE2# set ge-2/0/2 unit 0 description toPE1
user@PE2# set ge-2/0/2 unit 0 family mpls
user@PE2# set ge-2/0/2 unit 0 family inet address 10.0.0.2/30
user@PE2# set lo0 unit 0 family inet address 10.255.7.7/32

```



**NOTE:** ge-2/0/0.0 and ge-2/0/1.0 are configured as VPLS interfaces and connect to Routers CE3 and CE4. See *Configuring VPLS Encapsulation on CE-Facing Interfaces* for more details.

**Router CE3**

```
[edit interfaces]
user@CE3# set ge-2/0/0 unit 0 description toPE2
user@CE3# set ge-2/0/0 unit 0 family inet address 10.0.0.18/30
user@CE3# set ge-2/0/1 unit 0 description toSource
user@CE3# set ge-2/0/1 unit 0 family inet address 10.0.0.29/30
user@CE3# set lo0 unit 0 family inet address 10.255.3.3/32
```



**NOTE:** The ge-2/0/1.0 interface on Router CE3 connects to the multicast source.

```
Router CE4
[edit interfaces]
user@CE4# set ge-2/0/0 unit 0 description toPE2
user@CE4# set ge-2/0/0 unit 0 family inet address 10.0.0.22/30
user@CE4# set ge-2/0/1 unit 0 description toReceiver2
user@CE4# set ge-2/0/1 unit 0 family inet address 10.0.0.25/30
user@CE4# set lo0 unit 0 family inet address 10.255.4.4/32
```



**NOTE:** The ge-2/0/1.0 interface on Router CE4 connects to a multicast receiver.

Similarly, configure Routers PE1, CE1, and CE2.

2. Configure the router IDs of all routers.

```
Router PE2
[edit routing-options]
user@PE2# set router-id 10.255.7.7
```

Similarly, configure other routers.

3. Configure an IGP on interfaces of all routers.

```
Router PE2
[edit protocols ospf area 0.0.0.0]
user@PE2# set interface ge-2/0/2.0
user@PE2# set interface lo0.0
```

Similarly, configure other routers.

4. Configure the LDP, MPLS, and BGP protocols on the PE routers.

```
Router PE2
[edit protocols]
user@PE2# set ldp interface lo0.0
user@PE2# set mpls interface ge-2/0/2.0
user@PE2# set bgp group toPE1 type internal
user@PE2# set bgp group toPE1 local-address 10.255.7.7
user@PE2# set bgp group toPE1 family l2vpn signaling
user@PE2# set bgp group toPE1 neighbor 10.255.1.1
user@PE2# set ldp interface ge-2/0/2.0
```

The BGP group is required for interfacing with the other PE router. Similarly, configure Router PE1.

5. Configure PIM on all CE routers.

Ensure that Router CE3 is configured as the rendezvous point (RP) and that the RP address is configured on other CE routers.

```
Router CE3
[edit protocols pim]
user@CE3# set rp local address 10.255.3.3
user@CE3# set interface all
```

```
Router CE4
[edit protocols pim]
user@CE4# set rp static address 10.255.3.3
user@CE4# set interface all
```

Similarly, configure Routers CE1 and CE2.

6. Configure multicast snooping options on the PE routers.

```
Router PE2
[edit multicast-snooping-options traceoptions]
user@PE2# set file snoop.log size 10m
```

Similarly, configure Router PE1.

7. Create a routing instance (**titanium**), and configure the VPLS on the PE routers.

```
Router PE2
[edit routing-instances titanium]
user@PE2# set instance-type vpls
user@PE2# set vlan-id none
user@PE2# set interface ge-2/0/0.0
user@PE2# set interface ge-2/0/1.0
user@PE2# set route-distinguisher 101:101
user@PE2# set vrf-target target:201:201
user@PE2# set protocols vpls vpls-id 15
user@PE2# set protocols vpls site pe2 site-identifier 2
```

Similarly, configure Router PE1.

8. Configure PIM snooping on the PE routers.

```
Router PE2
[edit routing-instances titanium]
user@PE2# set protocols pim-snooping
```

Similarly, configure Router PE1.

### Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show protocols**, **show multicast-snooping-options**, and **show routing-instances** commands.

If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE2# show interfaces
ge-2/0/2 {
 unit 0 {
 description toPE1
 family inet {
 address 10.0.0.2/30;
 }
 family mpls;
 }
}
ge-2/0/0 {
 encapsulation ethernet-vpls;
 unit 0 {
 description toCE3;
 }
}
ge-2/0/1 {
 encapsulation ethernet-vpls;
 unit 0 {
 description toCE4;
 }
}
lo0 {
 unit 0 {
 family inet {
 address 10.255.7.7/32;
 }
 }
}
```

```
user@PE2# show routing-options
router-id 10.255.7.7;
```

```
user@PE2# show protocols
mpls {
 interface ge-2/0/2.0;
}
ospf {
 area 0.0.0.0 {
 interface ge-2/0/2.0;
 interface lo0.0;
 }
}
ldp {
 interface ge-2/0/2.0;
 interface lo0.0;
}
bgp {
 group toPE1 {
 type internal;
 local-address 10.255.7.7;
 family l2vpn {
 signaling;
 }
 }
}
```

```

 neighbor 10.255.1.1;
 }

user@PE2# show multicast-snooping-options
traceoptions {
 file snoop.log size 10m;
}

user@PE2# show routing-instances
titanium {
 instance-type vpls;
 vlan-id none;
 interface ge-2/0/0.0;
 interface ge-2/0/1.0;
 route-distinguisher 101:101;
 vrf-target target:201:201;
 protocols {
 vpls {
 site pe2 {
 site-identifier 2;
 }
 vpls-id 15;
 }
 pim-snooping;
 }
}

```

Similarly, confirm the configuration on all other routing devices. If you are done configuring the routing devices, enter **commit** from configuration mode.



**NOTE:** Use the **show protocols** command on the CE routers to verify the configuration for the PIM RP.

### Verification

Confirm that the configuration is working properly.

- [Verifying PIM Snooping for VPLS on page 3753](#)

### Verifying PIM Snooping for VPLS

**Purpose** Verify that PIM Snooping is operational in the network.

**Action** To verify that PIM snooping is working as desired, use the following commands:

- ***show pim snooping interfaces***
- ***show pim snooping neighbors detail***
- ***show pim snooping statistics***
- ***show pim snooping join***

- ***show pim snooping join extensive***
- ***show multicast snooping route extensive*** instance <instance-name> group <group-name>

1. From operational mode on Router PE2, run the **show pim snooping interfaces** command.

```
user@PE2> show pim snooping interfaces
Instance: titanium
```

```
Learning-Domain: default
```

| Name       | State | IP | NbrCnt |
|------------|-------|----|--------|
| ge-2/0/0.0 | Up    | 4  | 1      |
| ge-2/0/1.0 | Up    | 4  | 1      |

```
DR address: 10.0.0.22
```

```
DR flooding is ON
```

The output verifies that PIM snooping is configured on the two interfaces connecting Router PE2 to Routers CE3 and CE4.

Similarly, check the PIM snooping interfaces on Router PE1.

2. From operational mode on Router PE2, run the **show pim snooping neighbors detail** command.

```
user@PE2> show pim snooping neighbors detail
Instance: titanium
Learning-Domain: default
```

```
Interface: ge-2/0/0.0
```

```
Address: 10.0.0.18
```

```
Uptime: 00:17:06
```

```
Hello Option Holdtime: 105 seconds 99 remaining
```

```
Hello Option DR Priority: 1
```

```
Hello Option Generation ID: 552495559
```

```
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

```
Tracking is supported
```

```
Interface: ge-2/0/1.0
```

```
Address: 10.0.0.22
```

```
Uptime: 00:15:16
```

```
Hello Option Holdtime: 105 seconds 103 remaining
```

```
Hello Option DR Priority: 1
```

```
Hello Option Generation ID: 1131703485
```

```
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

```
Tracking is supported
```

The output verifies that Router PE2 can detect the IP addresses of its PIM snooping neighbors (10.0.0.18 on CE3 and 10.0.0.22 on CE4).

Similarly, check the PIM snooping neighbors on Router PE1.

3. From operational mode on Router PE2, run the **show pim snooping statistics** command.

```
user@PE2> show pim snooping statistics
```



Instance: titanium

Learning-Domain: default

|                              |      |
|------------------------------|------|
| Tx J/P messages              | 0    |
| Rx J/P messages              | 246  |
| Rx J/P messages -- seen      | 0    |
| Rx J/P messages -- received  | 246  |
| Rx Hello messages            | 1036 |
| Rx Version Unknown           | 0    |
| Rx Neighbor Unknown          | 0    |
| Rx Upstream Neighbor Unknown | 0    |
| Rx J/P Busy Drop             | 0    |
| Rx J/P Group Aggregate       | 0    |
| Rx Malformed Packet          | 0    |
|                              |      |
| Rx No PIM Interface          | 0    |
| Rx Bad Length                | 0    |
| Rx Unknown Hello Option      | 0    |
| Rx Unknown Packet Type       | 0    |
| Rx Bad TTL                   | 0    |
| Rx Bad Destination Address   | 0    |
| Rx Bad Checksum              | 0    |
| Rx Unknown Version           | 0    |

The output shows the number of hello and join/prune messages received by Router PE2. This verifies that PIM sparse mode is operational in the network.

4. Send multicast traffic from the source terminal attached to Router CE3, for the multicast group 224.1.1.1.
5. From operational mode on Router PE2, run the **show pim snooping join**, **show pim snooping join extensive**, and **show multicast snooping route extensive instance <instance-name> group <group-name>** commands to verify PIM snooping.

```
user@PE2> show pim snooping join
```

```
Instance: titanium
Learning-Domain: default
```

```
Group: 224.1.1.1
Source: *
Flags: sparse,rptree,wildcard
Upstream neighbor: 10.0.0.18, Port: ge-2/0/0.0
```

```
Group: 224.1.1.1
Source: 10.0.0.30
Flags: sparse
Upstream neighbor: 10.0.0.18, Port: ge-2/0/0.0
```

```
user@PE2> show pim snooping join extensive
```

```
Instance: titanium
Learning-Domain: default
```

```
Group: 224.1.1.1
Source: *
Flags: sparse,rptree,wildcard
Upstream neighbor: 10.0.0.18, Port: ge-2/0/0.0
Downstream port: ge-2/0/1.0
Downstream neighbors:
10.0.0.22 State: Join Flags: SRW Timeout: 180
```

```
Group: 224.1.1.1
Source: 10.0.0.30
Flags: sparse
Upstream neighbor: 10.0.0.18, Port: ge-2/0/0.0
Downstream port: ge-2/0/1.0
Downstream neighbors:
10.0.0.22 State: Join Flags: S Timeout: 180
```

The outputs show that multicast traffic sent for the group 224.1.1.1 is sent to Receiver 2 through Router CE4 and also display the upstream and downstream neighbor details.

```
user@PE2> show multicast snooping route extensive instance titanium group 224.1.1.1
Next-hop Bulking: OFF
```

Family: INET

```
Group: 224.1.1.1/32
Bridge-domain: titanium
Mesh-group: __all_ces__
Downstream interface list:
ge-2/0/1.0 -(1072)
Statistics: 0 kbps, 0 pps, 0 packets
Next-hop ID: 1048577
Route state: Active
Forwarding state: Forwarding
```

```
Group: 224.1.1.1/32
Source: 10.0.0.8
Bridge-domain: titanium
Mesh-group: __all_ces__
Downstream interface list:
ge-2/0/1.0 -(1072)
Statistics: 0 kbps, 0 pps, 0 packets
Next-hop ID: 1048577
Route state: Active
Forwarding state: Forwarding
```

**Meaning** PIM snooping is operational in the network.

## Configuration Statements: IGMP

## igmp

---

**Syntax**    `igmp {  
    accounting;  
    interface interface-name {  
        disable;  
        (accounting | no-accounting);  
        group-limit limit;  
        group-policy [ policy-names ];  
        group-threshold  
        immediate-leave;  
        log-interval  
        oif-map map-name;  
        passive;  
        promiscuous-mode;  
        ssm-map ssm-map-name;  
        ssm-map-policy ssm-map-policy-name;  
        static {  
            group multicast-group-address {  
                exclude;  
                group-count number;  
                group-increment increment;  
                source ip-address {  
                    source-count number;  
                    source-increment increment;  
                }  
            }  
        }  
        version version;  
    }  
    query-interval seconds;  
    query-last-member-interval seconds;  
    query-response-interval seconds;  
    robust-count number;  
    traceoptions {  
        file filename <files number> <size size> <world-readable | no-world-readable>;  
        flag flag <flag-modifier> <disable>;  
    }  
}`

**Hierarchy Level**    `[edit logical-systems logical-system-name protocols],`  
                          `[edit protocols]`

**Release Information**    Statement introduced before Junos OS Release 7.4.  
                              Statement introduced in Junos OS Release 12.1 for the QFX Series.  
                              Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

**Description**    Enable IGMP on the router or switch. IGMP must be enabled for the router or switch to receive multicast packets.

The remaining statements are explained separately.

|                                 |                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>                  | IGMP is disabled on the router or switch. IGMP is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP). |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling IGMP on page 3661</a></li> </ul>                                                                                                                     |

## accounting (Protocols IGMP Interface)

---

|                                 |                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (accounting   no-accounting);                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp interface interface-name</a> ],<br>[edit protocols <a href="#">igmp interface interface-name</a> ]             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Enable or disable the collection of IGMP join and leave event statistics for an interface.                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Recording IGMP Join and Leave Events on page 3674</a></li> </ul>                                                                      |

## accounting (Protocols IGMP)

---

|                                 |                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | accounting;                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp</a> ],<br>[edit protocols <a href="#">igmp</a> ]                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Enable the collection of IGMP join and leave event statistics on the system.                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Recording IGMP Join and Leave Events on page 3674</a></li> </ul>                                                                      |

## disable (Protocols IGMP)

---


|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> ],<br>[edit protocols <b>igmp interface</b> <i>interface-name</i> ]                     |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Disable IGMP on the system.                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Disabling IGMP on page 3679</a></li></ul>                                                                                                  |

## exclude (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | exclude;                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> <b>static group</b> <i>mcast-group-address</i> ],<br>[edit protocols <b>igmp interface</b> <i>interface-name</i> <b>static group</b> <i>mcast-group-address</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3.                                                                                                                                                                                                                            |
| <b>Description</b>              | Configure the static group to operate in exclude mode. In exclude mode all sources except the address configured are accepted for the group. If this statement is not included, the group operates in include mode.                                                      |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling IGMP Static Group Membership on page 3668</a></li></ul>                                                                                                                                                     |

## group (Protocols IGMP)

|                                                                                                                                                                                                                                 |                                                                                                                                                                                                |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                   | <pre>group multicast-group-address {   exclude;   group-count number;   group-increment increment;   source ip-address {     source-count number;     source-increment increment;   } }</pre>  |
| <b>Hierarchy Level</b>                                                                                                                                                                                                          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface interface-name static</b> ],<br>[edit protocols <b>igmp interface interface-name static</b> ]                     |
| <b>Release Information</b>                                                                                                                                                                                                      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>                                                                                                                                                                                                              | Specify the IGMP multicast group address and (optionally) the source address for the multicast group being statically configured on an interface.                                              |
| <div>  <p><b>NOTE:</b> You must specify a unique address for each group.</p> </div> <p>The remaining statements are explained separately.</p> |                                                                                                                                                                                                |
| <b>Required Privilege Level</b>                                                                                                                                                                                                 | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                            |
| <b>Related Documentation</b>                                                                                                                                                                                                    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling IGMP Static Group Membership on page 3668</a></li> </ul>                                                                         |

## group-count (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-count <i>number</i>;</code>                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                 |
| <b>Description</b>              | Specify the number of static groups to be created.                                                                                                                                                                                                                 |
| <b>Options</b>                  | <i>number</i> —Number of static groups.<br><b>Default:</b><br><b>Range:</b> 1 through 512                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling IGMP Static Group Membership on page 3668</a></li></ul>                                                                                                                                               |

## group-increment (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-increment <i>increment</i>;</code>                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                 |
| <b>Description</b>              | Configure the number of times the address should be incremented for each static group created. The increment is specified in dotted decimal notation similar to an IPv4 address.                                                                                   |
| <b>Options</b>                  | <i>increment</i> —Number of times the address should be incremented.<br><b>Default:</b> 0.0.0.1<br><b>Range:</b> 0.0.0.1 through 255.255.255.255                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling IGMP Static Group Membership on page 3668</a></li></ul>                                                                                                                                               |



## group-limit

---

|                                 |                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-limit <i>limit</i>;</code>                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp interface interface-name</a> ],<br>[edit protocols <a href="#">igmp interface interface-name</a> ]                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                 |
| <b>Description</b>              | <p>Configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) allowed on an interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface.</p> <p>To confirm the configured group limit on the interface, use the <a href="#">show igmp interface</a> command.</p> |
| <b>Default</b>                  | By default, there is no limit to the number of multicast groups that can join the interface.                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><i>limit</i>—group limit value for the interface.</p> <p><b>Range:</b> 1 through 32767</p>                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces on page 3676</a></li> <li>• <a href="#">group-threshold on page 3765</a></li> <li>• <a href="#">log-interval on page 3768</a></li> </ul>                                                                               |

## group-policy (Protocols IGMP)


---

|                                 |                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-policy [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> ]                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                |
| <b>Description</b>              | When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), after the routing device receives an IGMP report, the routing device compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report). |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Filtering Unwanted IGMP Reports at the IGMP Interface Level on page 3664</a></li></ul>                                                                                                                                                                        |

## group-threshold (Protocols IGMP Interface)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-threshold value;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp interface interface-name</a> ],<br>[edit protocols <a href="#">igmp interface interface-name</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | <p>Specify the threshold at which a warning message is logged for the multicast groups received on a logical interface. The threshold is a percentage of the maximum number of multicast groups allowed on a logical interface.</p> <p>For example, if you configure a maximum number of 1,000 incoming multicast groups, and you configure a threshold value of 90 percent, warning messages are logged in the system log when the interface receives 900 groups.</p> <p>To confirm the configured group threshold on the interface, use the <a href="#">show igmp interface</a> command.</p> |
| <b>Default</b>                  | By default, there is no configured threshold value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>value</b>—Percentage of the maximum number of multicast groups allowed on the interface that starts triggering the warning. You configure a percentage of the <b>group-limit</b> value that starts triggering the warnings. You must explicitly configure the <a href="#">group-limit</a> to configure a threshold value.</p> <p><b>Range:</b> 1 through 100</p>                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces on page 3676</a></li> <li>• <a href="#">group-limit on page 3763</a></li> <li>• <a href="#">log-interval on page 3768</a></li> </ul>                                                                                                                                                                                                                                                                                                                              |

## immediate-leave (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>immediate-leave;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.</p> <p>Starting in Junos OS Release 9.3, both IGMP version 2 and IGMP version 3 do host tracking when the <b>immediate-leave</b> statement is configured. This means that the multicast group leaves only when the last host leaves. The routing device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.</p> <p>When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.</p> <p>When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.</p> <div style="margin-top: 20px;">  <p><b>NOTE:</b> Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the routing device only knows about the one interested host and does not know about the others.</p> </div> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

**Related Documentation** • [Specifying Immediate-Leave Host Removal for IGMP on page 3664](#)

## interface (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> interface <i>interface-name</i> {     disable;     (accounting   no-accounting);     group-limit <i>limit</i>;     group-policy [ <i>policy-names</i> ];     immediate-leave;     oif-map <i>map-name</i>;     passive;     promiscuous-mode;     ssm-map <i>ssm-map-name</i>;     ssm-map-policy <i>ssm-map-policy-name</i>;     static {         group <i>mcast-group-address</i> {             exclude;             group-count <i>number</i>;             group-increment <i>increment</i>;             source <i>ip-address</i> {                 source-count <i>number</i>;                 source-increment <i>increment</i>;             }         }     }     version <i>version</i>; } </pre> |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b>],</p> <p>[edit protocols <b>igmp</b>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Enable IGMP on an interface and configure interface-specific properties.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b><i>interface-name</i></b>—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify <b>all</b>.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <p>• <a href="#">Enabling IGMP on page 3661</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## log-interval (Protocols IGMP Interface)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | log-interval <i>seconds</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> ],<br>[edit protocols <b>igmp interface</b> <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2.                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Specify the minimum time interval (in seconds) between sending consecutive log messages to the system log for multicast groups. To configure the time interval, you must specify the maximum number of multicast groups allowed on the interface. You must configure the <b>group-limit</b> statement before you configure the <b>log-interval</b> statement.</p> <p>To confirm the configured log interval on the interface, use the <b>show igmp interface</b> command.</p> |
| <b>Default</b>                  | By default, there is no configured time interval.                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b>seconds</b>—Minimum time interval (in seconds) between log messages. You must explicitly configure the <b>group-limit</b> to configure a time interval to send log messages.</p> <p><b>Range:</b> 6 through 32,767 seconds</p>                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces on page 3676</a></li><li>• <a href="#">group-limit on page 3763</a></li><li>• <a href="#">group-threshold on page 3765</a></li></ul>                                                                                                                                                                                                                 |

## maximum-transmit-rate (Protocols IGMP)


|                                 |                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>maximum-transmit-rate <i>packets-per-second</i>;</code>                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols igmp],<br>[edit protocols igmp]                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                         |
| <b>Description</b>              | Limit the transmission rate of IGMP packets                                                                                                                                |
| <b>Options</b>                  | <b>packets-per-second</b> —Maximum number of IGMP packets transmitted in one second by the routing device.<br><b>Range:</b> 1 through 10000<br><b>Default:</b> 500 packets |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Limiting the Maximum IGMP Message Rate on page 3668</a></li> </ul>                                                    |

## oif-map

|                                 |                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>oif-map <i>map-name</i>;</code>                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> ],<br>[edit protocols <b>igmp interface</b> <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                         |
| <b>Description</b>              | Associates an outgoing interface (OIF) map to the IGMP interface. The OIF map is a routing policy statement that can contain multiple terms.                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Multicast with Subscriber VLANs</a></li> </ul>                                                   |

## passive (IGMP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>passive &lt;allow-receive&gt; &lt;send-general-query&gt; &lt;send-group-query&gt;;</code>                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> ]                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br><b>allow-receive</b> , <b>send-general-query</b> , and <b>send-group-query</b> options were added in Junos OS Release 10.0.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                     |
| <b>Description</b>              | Specify that IGMP run on the interface and either not send and receive control traffic or selectively send and receive control traffic such as IGMP reports, queries, and leaves.                                                                                                                                                                                     |
|                                 | <div><b>NOTE:</b> You can selectively activate up to two out of the three available options for the <b>passive</b> statement while keeping the other functions passive (inactive). Activating all three options would be equivalent to not using the <b>passive</b> statement.</div> |
| <b>Options</b>                  | <b>allow-receive</b> —Enables IGMP to receive control traffic on the interface.<br><br><b>send-general-query</b> —Enables IGMP to send general queries on the interface.<br><br><b>send-group-query</b> —Enables IGMP to send group-specific and group-source-specific queries on the interface.                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Multicast with Subscriber VLANs</i></li><li>• <a href="#">Enabling IGMP on page 3661</a></li></ul>                                                                                                                                                                                                    |



## promiscuous-mode (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>promiscuous-mode;</code>                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit dynamic-profiles <i>profile-name</i> protocols igmp interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> ]                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 9.2 for dynamic profiles.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                      |
| <b>Description</b>              | Specify that the interface accepts IGMP reports from hosts on any subnetwork. Note that when enabling promiscuous-mode, all routing devices on the ethernet segment must be configured with the promiscuous mode statement. Otherwise, only the interface configured with lowest IPv4 address acts as the querier for IGMP for this Ethernet segment. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring a Dynamic Profile for Client Access</a></li> <li>• <a href="#">Accepting IGMP Messages from Remote Subnetworks on page 3665</a></li> </ul>                                                                                                                                           |

## query-interval (Protocols IGMP)

|                                 |                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>query-interval <i>seconds</i>;</code>                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> ],<br>[edit protocols <b>igmp</b> ]                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                       |
| <b>Description</b>              | Specify how often the querier routing device sends general host-query messages.                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b>seconds</b> —Time interval.<br><b>Range:</b> 1 through 1024<br><b>Default:</b> 125 seconds                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Modifying the IGMP Host-Query Message Interval on page 3662</a></li> <li>• <a href="#">query-last-member-interval (Protocols IGMP) on page 3772</a></li> <li>• <a href="#">query-response-interval (Protocols IGMP) on page 3773</a></li> </ul> |

## query-last-member-interval (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | query-last-member-interval <i>seconds</i> ;                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp</a> ],<br>[edit protocols <a href="#">igmp</a> ]                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                      |
| <b>Description</b>              | Specify how often the querier routing device sends group-specific query messages.                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>seconds</b> —Time interval, in fractions of a second or seconds.<br><b>Range:</b> 0.1 through 0.9, then in 1-second intervals 1 through 999999<br><b>Default:</b> 1 second                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Modifying the IGMP Last-Member Query Interval on page 3666</a></li><li>• <a href="#">query-interval (Protocols IGMP) on page 3771</a></li><li>• <a href="#">query-response-interval (Protocols IGMP) on page 3773</a></li></ul> |

---

## query-response-interval (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | query-response-interval <i>seconds</i> ;                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp</a> ],<br>[edit protocols <a href="#">igmp</a> ]                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                      |
| <b>Description</b>              | Specify how long the querier routing device waits to receive a response to a host-query message from a host.                                                                                                                                                                        |
| <b>Options</b>                  | <b>seconds</b> —The query response interval must be less than the query interval.<br><b>Range:</b> 1 through 1024<br><b>Default:</b> 10 seconds                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Modifying the IGMP Query Response Interval on page 3663</a></li><li>• <a href="#">query-interval (Protocols IGMP) on page 3771</a></li><li>• <a href="#">query-last-member-interval (Protocols IGMP) on page 3772</a></li></ul> |

## robust-count (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>robust-count <i>number</i>;</code>                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp</a> ],<br>[edit protocols <a href="#">igmp</a> ]                                                                   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Tune the expected packet loss on a subnet. This factor is used to calculate the group member interval, other querier present interval, and last-member query count.                            |
| <b>Options</b>                  | <i>number</i> —Robustness variable.<br><b>Range:</b> 2 through 10<br><b>Default:</b> 2                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Modifying the IGMP Robustness Variable on page 3667</a></li></ul>                                                                          |

## source (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>source ip-address {     source-count number;     source-increment increment; }</pre>                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> <b>static group</b> <i>multicast-group-address</i> ],<br>[edit protocols <b>igmp interface</b> <i>interface-name</i> <b>static group</b> <i>multicast-group-address</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                   |
| <b>Description</b>              | Specify the IP version 4 (IPv4) unicast source address for the multicast group being statically configured on an interface.                                                                                                                                                      |
| <b>Options</b>                  | <p><b><i>ip-address</i></b>—IPv4 unicast address.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling IGMP Static Group Membership on page 3668</a></li> </ul>                                                                                                                                                           |

## source-count (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source-count <i>number</i>;</code>                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> <b>source</b> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> <b>source</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                             |
| <b>Description</b>              | Configure the number of multicast source addresses that should be accepted for each static group created.                                                                                                                                                                                      |
| <b>Options</b>                  | <b><i>number</i></b> —Number of source addresses.<br><b>Default:</b> 1<br><b>Range:</b> 1 through 1024                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling IGMP Static Group Membership on page 3668</a></li></ul>                                                                                                                                                                           |

## source-increment (Protocols IGMP)

---

|                                 |                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source-increment <i>number</i>;</code>                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> <b>source</b> ],<br>[edit protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> <b>source</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                             |
| <b>Description</b>              | Configure the number of times the multicast source address should be incremented for each static group created. The increment is specified in dotted decimal notation similar to an IPv4 address.                                                                                              |
| <b>Options</b>                  | <b><i>increment</i></b> —Number of times the source address should be incremented.<br><b>Default:</b> 0.0.0.1<br><b>Range:</b> 0.0.0.1 through 255.255.255.255                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling IGMP Static Group Membership on page 3668</a></li></ul>                                                                                                                                                                           |

## ssm-map (Protocols IGMP)

---

|                                 |                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ssm-map <i>ssm-map-name</i>;</code>                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp interface interface-name</a> ],<br>[edit protocols <a href="#">igmp interface interface-name</a> ]             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Apply an SSM map to an IGMP interface.                                                                                                                                                     |
| <b>Options</b>                  | <i>ssm-map-name</i> —Name of SSM map.                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring SSM Mapping</i></li> </ul>                                                                                                |

## ssm-map-policy (IGMP)

---

|                                 |                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ssm-map-policy <i>ssm-map-policy-name</i>;</code>                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp interface interface-name</a> ],<br>[edit protocols <a href="#">igmp interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                            |
| <b>Description</b>              | Apply an SSM map policy to an IGMP interface.                                                                                                                                  |
| <b>Options</b>                  | <i>ssm-map-policy-name</i> —Name of SSM map policy.                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring SSM Maps for Different Groups to Different Sources</i></li> </ul>                                             |

## static (Protocols IGMP)

---

**Syntax**    static {  
              group *multicast-group-address* {  
                  exclude;  
                  group-count *number*;  
                  group-increment *increment*;  
                  source *ip-address* {  
                      source-count *number*;  
                      source-increment *increment*;  
                  }  
              }  
          }  
      }

**Hierarchy Level**    [edit logical-systems *logical-system-name* protocols **igmp** interface *interface-name*],  
                          [edit protocols **igmp** interface *interface-name*]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
                              Statement introduced in Junos OS Release 9.0 for EX Series switches.  
                              Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description**    Test multicast forwarding on an interface without a receiver host.

The **static** statement simulates IGMP joins on a routing device statically on an interface without any IGMP hosts. It is supported for both IGMPv2 and IGMPv3 joins. This statement is especially useful for testing multicast forwarding on an interface without a receiver host.



.....  
**NOTE:** To prevent joining too many groups accidentally, the **static** statement is not supported with the **interface all** statement.  
.....

The remaining statements are explained separately.

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

**Related Documentation**    • [Enabling IGMP Static Group Membership on page 3668](#)



## traceoptions (Protocols IGMP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> ],<br>[edit protocols <b>igmp</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>         | <p>Configure IGMP tracing options.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p>To trace the paths of multicast packets, use the <b>mtrace</b> command.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>             | The default IGMP trace options are those inherited from the routing protocols <b>traceoptions</b> statement included at the [edit routing-options] hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place tracing output in the file <b>igmp-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also include the <b>size</b> statement to specify the maximum file size.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 2 files</p> <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p><b>IGMP Tracing Flags</b></p> <ul style="list-style-type: none"> <li><b>leave</b>—Leave group messages (for IGMP version 2 only).</li> <li><b>mtrace</b>—Mtrace packets. Use the <b>mtrace</b> command to troubleshoot the software.</li> <li><b>packets</b>—All IGMP packets.</li> </ul> |

- **query**—IGMP membership query messages, including general and group-specific queries.
- **report**—Membership report messages.

#### Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

**no-stamp**—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

**no-world-readable**—(Optional) Do not allow users to read the log file.

**replace**—(Optional) Replace an existing trace file if there is one.

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

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

**Related Documentation** • [Tracing IGMP Protocol Traffic on page 3677](#)

## version (Protocols IGMP)

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>version <i>version</i>;</code>                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> ],<br>[edit protocols <b>igmp interface</b> <i>interface-name</i> ]                     |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>              | Specify the version of IGMP.                                                                                                                                                                   |
| <b>Options</b>                  | <b>version</b> —IGMP version number.<br><b>Range:</b> 1, 2, or 3<br><b>Default:</b> IGMP version 2                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                            |
| <b>Related Documentation</b>    | • <a href="#">Changing the IGMP Version on page 3668</a>                                                                                                                                       |

## Configuration Statements: IGMP Snooping

### igmp-snooping

---

```
Syntax igmp-snooping {
 vlan vlan-id {
 immediate-leave;
 interface interface-name {
 group-limit limit;
 host-only-interface;
 immediate-leave;
 multicast-router-interface;
 static {
 group ip-address {
 source ip-address;
 }
 }
 }
 }
 proxy {
 source-address ip-address;
 }
 query-interval seconds;
 query-last-member-interval seconds;
 query-response-interval seconds;
 robust-count number;
 }
```

**Hierarchy Level** [edit protocols]

**Release Information** Statement introduced in Junos OS Release 8.5.

**Description** Enable IGMP snooping on the router or switch.

**Default** IGMP snooping is disabled on the router or switch.

**Options** The statements are explained separately.

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

**Related Documentation**

- *Understanding IGMP Snooping*
- *IGMP Snooping in MC-LAG Active-Active on MX Series Routers Overview*

## group (Bridge Domains)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group <i>ip-address</i> {<br/>    source-address <i>ip-address</i>;<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <b>interface</b> <i>interface-name</i> <b>static</b> ],<br>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <b>vlan</b> <i>vlan-id</i> <b>interface</b> <i>interface-name</i> <b>static</b> ],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <b>interface</b> <i>interface-name</i> <b>static</b> ],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols <b>vlan</b> <i>vlan-id</i> igmp-snooping <b>interface</b> <i>interface-name</i> <b>static</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Configure the IGMP multicast group address that receives data on an interface and (optionally) a source address for certain packets.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <i>ip-address</i> —Group address.<br><br>The remaining statement is explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IGMP Snooping</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## group-limit

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-limit <i>limit</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i> ],<br>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">vlan</a> <i>vlan-id</i> <a href="#">interface</a> <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols <a href="#">vlan</a> <i>vlan-id</i> igmp-snooping <a href="#">interface</a> <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) allowed on an interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Default</b>                  | By default, there is no limit to the number of multicast groups joining an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <i>limit</i> —a 32-bit number for the limit on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IGMP Snooping</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## host-only-interface

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | host-only-interface;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name</a>],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">vlan vlan-id interface interface-name</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols <a href="#">vlan vlan-id</a> igmp-snooping <a href="#">interface interface-name</a>]</p> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure an interface as a host-facing interface. IGMP queries received on these interfaces are dropped.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>                  | The interface can either be a host-side or multicast-routing device interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IGMP Snooping</i></li> <li>• <a href="#">multicast-router-interface on page 3788</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## immediate-leave (Bridge Domains)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>immediate-leave;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>     | <pre>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <i>interface</i> <i>interface-name</i>], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <i>vlan</i> <i>vlan-id</i> <i>interface</i> <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <i>interface</i> <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols <i>vlan</i> <i>vlan-id</i> igmp-snooping <i>interface</i> <i>interface-name</i>]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>         | <p>The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.</p> <p>The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.</p> <p>When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.</p> <p>When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.</p> |



**NOTE:** Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the routing device only knows about the one interested host and does not know about the others.



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

**Related Documentation** • *Example: Configuring IGMP Snooping*

## interface (Bridge Domains)

**Syntax**

```
interface interface-name {
 group-limit limit;
 host-only-interface;
 multicast-router-interface;
 static {
 group ip-address {
 source ip-address;
 }
 }
}
```

**Hierarchy Level** [edit bridge-domains *bridge-domain-name* protocols igmp-snooping],  
[edit bridge-domains *bridge-domain-name* protocols igmp-snooping **vlan** *vlan-id*],  
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols  
igmp-snooping],  
[edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols  
**vlan** *vlan-id* igmp-snooping]

**Release Information** Statement introduced in Junos OS Release 8.5.

**Description** Enable IGMP snooping on an interface and configure interface-specific properties.

**Options** *interface-name*—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**.

The remaining statements are explained separately.

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

**Related Documentation** • *Example: Configuring IGMP Snooping*

## multicast-router-interface (IGMP Snooping)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | multicast-router-interface;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">vlan</a> <i>vlan-id</i> <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit protocols igmp-snooping vlan (all   <i>vlan-name</i>) interface (all   <i>interface-name</i>)]</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols <a href="#">vlan</a> <i>vlan-id</i> igmp-snooping <a href="#">interface</a> <i>interface-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.1 for EX Series switches.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | <p>Statically configure the interface as an IGMP snooping multicast-router interface—that is, an interface that faces toward a multicast router or other IGMP querier.</p> <div data-bbox="474 911 542 978" data-label="Image"> </div> <p><b>NOTE:</b> If the specified interface is a trunk port, the interface becomes a multicast-routing device interface for all VLANs configured on the trunk port. In addition, all unregistered multicast packets, whether they are IPv4 or IPv6 packets, are forwarded to the multicast routing device interface, even if the interface is configured as a multicast routing device interface only for IGMP snooping.</p> <p>Configure an interface as a bridge interface toward other multicast routing devices.</p>                                                   |
| <b>Default</b>                  | The interface can either be a host-side or multicast-routing device interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IGMP Snooping on EX Series Switches</i></li> <li>• <i>Example: Configuring IGMP Snooping</i></li> <li>• <i>Configuring IGMP Snooping (CLI Procedure)</i></li> <li>• <i>IGMP Snooping in MC-LAG Active-Active on MX Series Routers Overview</i></li> <li>• <a href="#">host-only-interface on page 3785</a></li> <li>• <i>show igmp-snooping membership</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                           |

## proxy (Bridge Domains)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>proxy {   source-address ip-address; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping],<br/>         [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <b>vlan</b> <i>vlan-id</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping],<br/>         [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols <b>vlan</b> <i>vlan-id</i> igmp-snooping]</p> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>Configure proxy mode and options, including source address. All the queries generated by IGMP snooping are sent using 0.0.0.0 as the source address in order to avoid participating in IGMP querier election. Also, all reports generated by IGMP snooping are sent with 0.0.0.0 as the source address unless there is a configured source address to use.</p>                                                                                                                                    |
| <b>Default</b>                  | <p>By default, IGMP snooping does not employ proxy mode.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.<br/>         routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IGMP Snooping</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                        |

## query-interval (Bridge Domains)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | query-interval <i>seconds</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name</a> ],<br>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">vlan vlan-id interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">vlan vlan-id interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Configure the interval for host-query message timeouts.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <b><i>seconds</i></b> —Time interval.<br><b>Range:</b> 1 through 1024<br><b>Default:</b> 125 seconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IGMP Snooping</a></li><li>• <a href="#">query-last-member-interval (Bridge Domains) on page 3791</a></li><li>• <a href="#">query-response-interval (Bridge Domains) on page 3792</a></li></ul>                                                                                                                                                                                                                                                                                                                                     |

## query-last-member-interval (Bridge Domains)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>query-last-member-interval <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">vlan</a> <i>vlan-id</i> <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">vlan</a> <i>vlan-id</i> <a href="#">interface</a> <i>interface-name</i>]</p> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Configure the interval for group-specific query timeouts.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b><i>seconds</i></b>—Time interval, in fractions of a second or seconds.</p> <p><b>Range:</b> 0.1 through 0.9, then in 1-second intervals 1 through 1024</p> <p><b>Default:</b> 1 second</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IGMP Snooping</a></li> <li>• <a href="#">query-interval on page 3790</a></li> <li>• <a href="#">query-response-interval on page 3792</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## query-response-interval (Bridge Domains)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>query-response-interval seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <code>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name</a>],</code><br><code>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">vlan vlan-id interface interface-name</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">vlan vlan-id interface interface-name</a>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Specify how long to wait to receive a response to a specific query message from a host.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <b>seconds</b> —Time interval. This interval must be less than the host-query interval.<br><b>Range:</b> 1 through 1024<br><b>Default:</b> 10 seconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IGMP Snooping</a></li><li>• <a href="#">query-interval (Bridge Domains) on page 3790</a></li><li>• <a href="#">query-last-member-interval (Bridge Domains) on page 3791</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                              |

## robust-count (Bridge Domains)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>robust-count <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">vlan</a> <i>vlan-id</i> <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">vlan</a> <i>vlan-id</i> <a href="#">interface</a> <i>interface-name</i>]</p> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Provide fine-tuning to allow for expected packet loss on a subnet. You can wait more intervals if subnet packet loss is high and IGMP report messages might be lost.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><i>number</i>—Robust interval.</p> <p><b>Range:</b> 2 through 10</p> <p><b>Default:</b> 2</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IGMP Snooping</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## source (Bridge Domains)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source ip-address;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name static group</a> ],<br>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name static group</a> ],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name static group</a> ],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols <a href="#">vlan vlan-id</a> igmp-snooping <a href="#">interface interface-name static group</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Statically define multicast group source addresses on an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <i>ip-address</i> —IP address to use as the source for the group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IGMP Snooping</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## source-address

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source-address ip-address;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">proxy</a> ],<br>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">vlan vlan-id proxy</a> ],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">proxy</a> ],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">vlan vlan-id proxy</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Specify the IP address to use as the source for IGMP snooping reports in proxy mode. Reports are sent with address 0.0.0.0 as the source address unless there is a source address configured.                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <i>ip-address</i> —IP address to use as the source for proxy-mode IGMP snooping reports.                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IGMP Snooping</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                     |



## static (Bridge Domains)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>static {   group multicast-group-address {     source ip-address;   } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">vlan</a> <i>vlan-id</i> <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">vlan</a> <i>vlan-id</i> <a href="#">interface</a> <i>interface-name</i>]</p> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>Define static multicast groups on an interface.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IGMP Snooping</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## traceoptions (Protocols IGMP Snooping)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt; ;<br/>    flag <i>flag</i> (detail   disable   receive   send);<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>     | <pre>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>domain-name</i> protocols<br/>  igmp-snooping],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> bridge-domains<br/>  <i>domain-name</i> protocols igmp-snooping],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols<br/>  igmp-snooping],<br/>[edit bridge-domains <i>domain-name</i> protocols igmp-snooping],<br/>[edit routing-instances <i>instance-name</i> bridge-domains <i>domain-name</i> protocols<br/>  igmp-snooping],<br/>[edit routing-instances <i>instance-name</i> protocols igmp-snooping]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>         | Define tracing operations for IGMP snooping.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>             | The <b>traceoptions</b> feature is disabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>             | <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. All files are placed in the directory <b>/var/log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached (<b>xk</b> to specify KB, <b>xm</b> to specify MB, or <b>xg</b> to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 3 files</p> <p><b>flag <i>flag</i></b> —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"><li>• <b>all</b>—All tracing operations.</li><li>• <b>client-notification</b>—Trace notifications.</li><li>• <b>general</b>—Trace general IGMP snooping protocol events.</li><li>• <b>group</b>—Trace group operations.</li><li>• <b>host-notification</b>—Trace host notifications.</li><li>• <b>leave</b>—Trace leave group messages (IGMPv2 only).</li><li>• <b>normal</b>—Trace normal IGMP snooping protocol events.</li><li>• <b>packets</b>—Trace all IGMP packets.</li></ul> |

- **policy**—Trace policy processing.
- **query**—Trace IGMP membership query messages.
- **report**—Trace membership report messages.
- **route**—Trace routing information.
- **state**—Trace IGMP state transitions.
- **task**—Trace routing protocol task processing.
- **timer**—Trace routing protocol timer processing.

**no-world-readable**—(Optional) Restrict file access to the user who created the file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify gigabytes

**Range:** 10 KB through 1 gigabytes

**Default:** 128 KB

**world-readable**—(Optional) Enable unrestricted file access.

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

|                              |                                                                                                                                                   |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <i>Configuring IGMP Snooping Trace Operations</i></li> <li>• <i>Configuring IGMP Snooping</i></li> </ul> |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|

## vlan (Bridge Domains)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>vlan <i>vlan-id</i> {<br/>    immediate-leave;<br/>    interface <i>interface-name</i> {<br/>        group-limit <i>limit</i>;<br/>        host-only-interface;<br/>        multicast-router-interface;<br/>        static {<br/>            group <i>multicast-group-address</i> {<br/>                source <i>ip-address</i>;<br/>            }<br/>        }<br/>    }<br/>    proxy {<br/>        source-address <i>ip-address</i>;<br/>    }<br/>    query-interval <i>seconds</i>;<br/>    query-last-member-interval <i>seconds</i>;<br/>    query-response-interval <i>seconds</i>;<br/>    robust-count <i>number</i>;<br/>}</pre> |
| Hierarchy Level          | [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping]                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Release Information      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Description              | Configure IGMP snooping parameters for a particular VLAN.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Default                  | By default, IGMP snooping options apply to all VLANs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Options                  | <i>vlan-id</i> —Apply the parameters to this VLAN.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Related Documentation    | <ul style="list-style-type: none"><li>Configuring VLAN-Specific IGMP Snooping Parameters</li><li>igmp-snooping</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## Configuration Statements: MLD

## mld

```
Syntax mld {
 accounting;
 interface interface-name {
 (accounting | no-accounting);
 disable;
 group-limit limit;
 group-policy [policy-names];
 immediate-leave;
 oif-map [map-names];
 passive;
 ssm-map ssm-map-name;
 ssm-map-policy ssm-map-policy-name;
 static {
 group multicast-group-address {
 exclude;
 group-count number;
 group-increment increment;
 source ip-address {
 source-count number;
 source-increment increment;
 }
 }
 }
 version version;
 }
 maximum-transmit-rate packets-per-second;
 query-interval seconds;
 query-last-member-interval seconds;
 query-response-interval seconds;
 robust-count number;
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols],  
[edit protocols]

**Release Information** Statement introduced before Junos OS Release 7.4.

**Description** Enable MLD on the routing device. MLD must be enabled for the routing device to receive multicast packets.

**Default** MLD is disabled on the routing device. MLD is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP).

**Options** The statements are explained separately.

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

**Related Documentation**

- *Enabling MLD*

## accounting (Protocols MLD Interface)

---

**Syntax** (accounting | no-accounting);

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols **mld interface** *interface-name*],  
[edit protocols **mld interface** *interface-name*]

**Release Information** Statement introduced in Junos OS Release 9.1.

**Description** Enable or disable the collection of MLD join and leave event statistics for an interface.

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

**Related Documentation**

- *Example: Recording MLD Join and Leave Events*

## accounting (Protocols MLD)

---

**Syntax** accounting;

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols **mld**],  
[edit protocols **mld**]

**Release Information** Statement introduced in Junos OS Release 9.1.

**Description** Enable the collection of MLD join and leave event statistics on the system.

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

**Related Documentation**

- *Example: Recording MLD Join and Leave Events*

## disable (Protocols MLD)

---

|                                 |                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>mld interface</b> <i>interface-name</i> ],<br>[edit protocols <b>mld interface</b> <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                        |
| <b>Description</b>              | Disable MLD on the system.                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Disabling MLD</i></li></ul>                                                                                                   |


## exclude (Protocols MLD)

---

|                                 |                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | exclude;                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>mld interface</b> <i>interface-name</i> <b>static group</b> <i>multicast-group-address</i> ],<br>[edit protocols <b>mld interface</b> <i>interface-name</i> <b>static group</b> <i>multicast-group-address</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3.                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the static group to operate in exclude mode. In exclude mode all sources except the address configured are accepted for the group. By default, the group operates in include mode.                                                                                   |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Enabling MLD Static Group Membership</i></li></ul>                                                                                                                                                                                  |



## group (Protocols MLD)

|                                                                                                                                                                                                                                 |                                                                                                                                                                                               |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                   | <pre>group multicast-group-address {   exclude;   group-count number;   group-increment increment;   source ip-address {     source-count number;     source-increment increment;   } }</pre> |
| <b>Hierarchy Level</b>                                                                                                                                                                                                          | [edit logical-systems <i>logical-system-name</i> protocols <b>mld interface interface-name static</b> ],<br>[edit protocols <b>mld interface interface-name static</b> ]                      |
| <b>Release Information</b>                                                                                                                                                                                                      | Statement introduced before Junos OS Release 7.4.                                                                                                                                             |
| <b>Description</b>                                                                                                                                                                                                              | The MLD multicast group address and (optionally) the source address for the multicast group being statically configured on an interface.                                                      |
| <b>Options</b>                                                                                                                                                                                                                  | <b>multicast-group-address</b> —Address of the group.                                                                                                                                         |
| <div>  <p><b>NOTE:</b> You must specify a unique address for each group.</p> </div> <p>The remaining statements are explained separately.</p> |                                                                                                                                                                                               |
| <b>Required Privilege Level</b>                                                                                                                                                                                                 | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                           |
| <b>Related Documentation</b>                                                                                                                                                                                                    | <ul style="list-style-type: none"> <li>• <i>Enabling MLD Static Group Membership</i></li> </ul>                                                                                               |

## group-count (Protocols MLD)

---

|                                 |                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-count <i>number</i>;</code>                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <code>mld interface interface-name static group multicast-group-address</code> ],<br>[edit protocols <code>mld interface interface-name static group multicast-group-address</code> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.                                                                                                                                                                                                    |
| <b>Description</b>              | Configure the number of static groups to be created.                                                                                                                                                                                             |
| <b>Options</b>                  | <i>number</i> —Number of static groups.<br><b>Default:</b> 1<br><b>Range:</b> 1 through 512                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Enabling MLD Static Group Membership</i></li></ul>                                                                                                                                                    |

## group-increment (Protocols MLD)

---

|                                 |                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-increment <i>increment</i>;</code>                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <code>mld interface interface-name static group multicast-group-address</code> ],<br>[edit protocols <code>mld interface interface-name static group multicast-group-address</code> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.                                                                                                                                                                                                    |
| <b>Description</b>              | Configure the number of times the address should be incremented for each static group created. The increment is specified in a format similar to an IPv6 address.                                                                                |
| <b>Options</b>                  | <i>increment</i> —Number of times the address should be incremented.<br><b>Default:</b> ::1<br><b>Range:</b> ::1 through ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff:                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Enabling MLD Static Group Membership</i></li></ul>                                                                                                                                                    |

## group-limit

|                                 |                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-limit <i>limit</i>;</code>                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">mld interface interface-name</a> ],<br>[edit protocols <a href="#">mld interface interface-name</a> ]                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.4.                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Configure a limit for the number of multicast groups (or [S,G] channels in MLDv2) allowed on a logical interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface.</p> <p>To confirm the configured group limit on the interface, use the <a href="#">show mld interface</a> command.</p> |
| <b>Default</b>                  | By default, there is no limit to the number of multicast groups that can join the interface.                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><i>limit</i>—group value limit for the interface.</p> <p><b>Range:</b> 1 through 32767</p>                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring the Number of MLD Multicast Group Joins on Logical Interfaces</i></li> </ul>                                                                                                                                                                                                                       |

## group-policy (Protocols MLD)


|                                 |                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-policy [ <i>policy-names</i> ];</code>                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">mld interface interface-name</a> ],<br>[edit protocols <a href="#">mld interface interface-name</a> ]                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1.                                                                                                                                                                                                                    |
| <b>Description</b>              | When a routing device running MLD version 1 or version 2 (MLDv1 or MLDv2), receives an MLD report, the routing device compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report). |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Filtering Unwanted MLD Reports at the MLD Interface Level</i></li> </ul>                                                                                                                                               |

## group-threshold (Protocols MLD Interface)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>group-threshold value;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">mld interface interface-name</a> ],<br>[edit protocols <a href="#">mld interface interface-name</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Specify the threshold at which a warning message is logged for the multicast groups received on a logical interface. The threshold is a percentage of the maximum number of multicast groups allowed on a logical interface.</p> <p>For example, if you configure a maximum number of 1,000 incoming multicast groups, and you configure a threshold value of 90 percent, warning messages are logged in the system log when the interface receives 900 groups.</p> <p>To confirm the configured group threshold on the interface, use the <a href="#">show mld interface</a> command.</p> |
| <b>Default</b>                  | By default, there is no configured threshold value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b>value</b>—Percentage of the maximum number of multicast groups allowed on the interface that starts triggering the warning. You configure a percentage of the <b>group-limit</b> value that starts triggering the warnings. You must explicitly configure the <a href="#">group-limit</a> to configure a threshold value.</p> <p><b>Range:</b> 1 through 100</p>                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Number of MLD Multicast Group Joins on Logical Interfaces</i></li><li>• <a href="#">group-limit on page 3805</a></li><li>• <a href="#">log-interval on page 3809</a></li></ul>                                                                                                                                                                                                                                                                                                                                                     |

## immediate-leave (Protocols MLD)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | immediate-leave;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>mld interface</b> <i>interface-name</i> ],<br>[edit protocols <b>mld interface</b> <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>The immediate leave setting is useful for minimizing the leave latency of MLD memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.</p> <p>The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows MLD to determine when the last host sends a leave message for the multicast group.</p> <p>When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending MLD group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the MLD leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.</p> <p>When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both MLD version 1 and MLD version 2.</p> <div>  <p><b>NOTE:</b> Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the routing device only knows about the one interested host and does not know about the others.</p> </div> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Specifying Immediate-Leave Host Removal for MLD</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## interface (Protocols MLD)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>interface <i>interface-name</i> {<br/>    disable;<br/>    (accounting   no-accounting);<br/>    group-limit <i>limit</i>;<br/>    group-policy [ <i>policy-names</i> ];<br/>    group-threshold <i>value</i>;<br/>    immediate-leave;<br/>    log-interval <i>seconds</i>;<br/>    oif-map [ <i>map-names</i> ];<br/>    passive;<br/>    ssm-map <i>ssm-map-name</i>;<br/>    ssm-map-policy <i>ssm-map-policy-name</i>;<br/>    static {<br/>        group <i>multicast-group-address</i> {<br/>            exclude;<br/>            group-count <i>number</i><br/>            group-increment <i>increment</i><br/>            source <i>ip-address</i> {<br/>                source-count <i>number</i>;<br/>                source-increment <i>increment</i>;<br/>            }<br/>        }<br/>    }<br/>    version <i>version</i>;<br/>}</pre> |
| Hierarchy Level          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">mld</a> ],<br>[edit protocols <a href="#">mld</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Release Information      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Description              | Enable MLD on an interface and configure interface-specific properties.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Options                  | <p><b><i>interface-name</i></b>—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify <b>all</b>.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Enabling MLD</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## log-interval (Protocols MLD Interface)

|                                 |                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | log-interval <i>seconds</i> ;                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">mld interface interface-name</a> ],<br>[edit protocols <a href="#">mld interface interface-name</a> ]                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2.                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Specify the minimum time interval (in seconds) between sending consecutive log messages to the system log for multicast groups. To configure the time interval, you must specify the maximum number of multicast groups allowed on the interface.</p> <p>To confirm the configured log interval on the interface, use the <a href="#">show mld interface</a> command.</p> |
| <b>Default</b>                  | By default, there is no configured time interval.                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>seconds</b>—Minimum time interval (in seconds) between log messages. You must explicitly configure the <b>group-limit</b> to configure a time interval to send log messages.</p> <p><b>Range:</b> 6 through 32,767 seconds</p>                                                                                                                                         |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Number of MLD Multicast Group Joins on Logical Interfaces</i></li> <li>• <a href="#">group-limit on page 3805</a></li> <li>• <a href="#">group-threshold on page 3806</a></li> </ul>                                                                                                                             |

## maximum-transmit-rate (Protocols MLD)

---

|                                 |                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>maximum-transmit-rate <i>packets-per-second</i>;</code>                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">mld</a> ],<br>[edit protocols <a href="#">mld</a> ]                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3.                                                                                                                             |
| <b>Description</b>              | Limit the transmission rate of MLD packets.                                                                                                                               |
| <b>Options</b>                  | <b>packets-per-second</b> —Maximum number of MLD packets transmitted in one second by the routing device.<br><b>Range:</b> 1 through 10000<br><b>Default:</b> 500 packets |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Limiting the Maximum MLD Message Rate</i></li></ul>                                                                            |

## oif-map

---

|                                 |                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>oif-map <i>map-name</i>;</code>                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">mld interface</a> <i>interface-name</i> ],<br>[edit protocols <a href="#">mld interface</a> <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.                                                                                                                                              |
| <b>Description</b>              | Associate an outgoing interface (OIF) map to an MLD logical interface. The OIF map is a routing policy statement that can contain multiple terms.                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Multicast with Subscriber VLANs</i></li></ul>                                                                              |



## passive (MLD)

|                            |                                                                                                                                                                                           |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>passive &lt;allow-receive&gt; &lt;send-general-query&gt; &lt;send-group-query&gt;;</code>                                                                                           |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <code>mld interface interface-name</code> ],<br>[edit protocols <code>mld interface interface-name</code> ]                    |
| <b>Release Information</b> | Statement introduced in Junos OS Release 9.6.<br><code>allow-receive</code> , <code>send-general-query</code> , and <code>send-group-query</code> options added in Junos OS Release 10.0. |
| <b>Description</b>         | Specify that MLD run on the interface and either not send and receive control traffic or selectively send and receive control traffic such as MLD reports, queries, and leaves.           |



**NOTE:** You can selectively activate up to two out of the three available options for the `passive` statement while keeping the other functions `passive` (inactive). Activating all three options is equivalent to not using the `passive` statement.

|                                 |                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <p><code>allow-receive</code>—Enables MLD to receive control traffic on the interface.</p> <p><code>send-general-query</code>—Enables MLD to send general queries on the interface.</p> <p><code>send-group-query</code>—Enables MLD to send group-specific and group-source-specific queries on the interface.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Example: Configuring Multicast with Subscriber VLANs</i></li> </ul>                                                                                                                                                                                                       |

## query-interval (Protocols MLD)

---

|                                 |                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | query-interval <i>seconds</i> ;                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">mld</a> ],<br>[edit protocols <a href="#">mld</a> ]                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                       |
| <b>Description</b>              | Specify how often the querier routing device sends general host-query messages.                                                                                                                                                                                         |
| <b>Options</b>                  | <b>seconds</b> —Time interval.<br><b>Range:</b> 1 through 1024<br><b>Default:</b> 125 seconds                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Modifying the MLD Host-Query Message Interval</i></li><li>• <a href="#">query-last-member-interval (Protocols MLD) on page 3812</a></li><li>• <a href="#">query-response-interval (Protocols MLD) on page 3813</a></li></ul> |

## query-last-member-interval (Protocols MLD)

---

|                                 |                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | query-last-member-interval <i>seconds</i> ;                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">mld</a> ],<br>[edit protocols <a href="#">mld</a> ]                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                          |
| <b>Description</b>              | Specify how often the querier routing device sends group-specific query messages.                                                                                                                                                                          |
| <b>Options</b>                  | <b>seconds</b> —Time interval, in fractions of a second or seconds.<br><b>Range:</b> 0.1 through 0.9, then in 1-second intervals from 1 through 1024<br><b>Default:</b> 1 second                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Modifying the MLD Last-Member Query Interval</i></li><li>• <a href="#">query-interval (Protocols MLD) on page 3812</a></li><li>• <a href="#">query-response-interval (Protocols MLD) on page 3813</a></li></ul> |

## query-response-interval (Protocols MLD)

|                                 |                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>query-response-interval <i>seconds</i>;</code>                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">mld</a> ],<br>[edit protocols <a href="#">mld</a> ]                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                      |
| <b>Description</b>              | Specify how long the querier routing device waits to receive a response to a host-query message from a host.                                                                                                                                                   |
| <b>Options</b>                  | <b><i>seconds</i></b> —Time interval. This interval must be less than the interval between general host-query messages.<br><b>Range:</b> 1 through 1024<br><b>Default:</b> 10 seconds                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Modifying the MLD Query Response Interval</i></li> <li>• <a href="#">query-interval (Protocols MLD) on page 3812</a></li> <li>• <a href="#">query-last-member-interval (Protocols MLD) on page 3812</a></li> </ul> |

## robust-count (Protocols MLD)

|                                 |                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>robust-count <i>number</i>;</code>                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">mld</a> ],<br>[edit protocols <a href="#">mld</a> ]                                                        |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                 |
| <b>Description</b>              | Tune for the expected packet loss on a subnet.                                                                                                                                    |
| <b>Options</b>                  | <b><i>number</i></b> —Time interval. This interval must be less than the interval between general host-query messages.<br><b>Range:</b> 2 through 10<br><b>Default:</b> 2 seconds |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Modifying the MLD Robustness Variable</i></li> </ul>                                                                         |

## source (Protocols MLD)

---

|                                 |                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source ip-address {<br/>    source-count number;<br/>    source-increment increment;<br/>}</code>                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <code>mld interface interface-name static group multicast-group-address</code> ],<br>[edit protocols <code>mld interface interface-name static group multicast-group-address</code> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                |
| <b>Description</b>              | IP version 6 (IPv6) unicast source address for the multicast group being statically configured on an interface.                                                                                                                                  |
| <b>Options</b>                  | <i>ip-address</i> —One or more IPv6 unicast addresses.                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Enabling MLD Static Group Membership</i></li></ul>                                                                                                                                                    |

## source-count (Protocols MLD)

---

|                                 |                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source-count number;</code>                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <code>mld interface interface-name static group multicast-group-address source</code> ],<br>[edit protocols <code>mld interface interface-name static group multicast-group-address source</code> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the number of multicast source addresses that should be accepted for each static group created.                                                                                                                                                      |
| <b>Options</b>                  | <i>number</i> —Number of source addresses.<br><b>Default:</b> 1<br><b>Range:</b> 1 through 1024                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Enabling MLD Static Group Membership</i></li></ul>                                                                                                                                                                  |

## source-increment (Protocols MLD)

|                                 |                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | source-increment <i>number</i> ;                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>mld interface</b> <i>interface-name</i> <b>static group</b> <i>mcastcast-group-address</i> <b>source</b> ],<br>[edit protocols <b>mld interface</b> <i>interface-name</i> <b>static group</b> <i>mcastcast-group-address</i> <b>source</b> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure the number of times the address should be incremented for each static group created. The increment is specified in a format similar to an IPv6 address.                                                                                                                                          |
| <b>Options</b>                  | <b>increment</b> —Number of times the source address should be incremented.<br><b>Default:</b> ::1<br><b>Range:</b> ::1 through ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff;                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Enabling MLD Static Group Membership</i></li> </ul>                                                                                                                                                                                                            |

## ssm-map (Protocols MLD)

|                                 |                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ssm-map <i>ssm-map-name</i> ;                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>mld interface</b> <i>interface-name</i> ],<br>[edit protocols <b>mld interface</b> <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.                                                                                                                            |
| <b>Description</b>              | Apply an SSM map to an MLD interface.                                                                                                                                    |
| <b>Options</b>                  | <b>ssm-map-name</b> —Name of SSM map.                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring SSM Mapping</i></li> </ul>                                                                              |

## ssm-map-policy (MLD)

---

|                                 |                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ssm-map-policy <i>ssm-map-policy-name</i>;</code>                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>mld interface</b> <i>interface-name</i> ],<br>[edit protocols <b>mld interface</b> <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.4.                                                                                                                           |
| <b>Description</b>              | Apply an SSM map policy to an MLD interface.                                                                                                                             |
| <b>Options</b>                  | <i>ssm-map-policy-name</i> —Name of SSM map policy.                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring SSM Maps for Different Groups to Different Sources</i></li></ul>                                         |

## static (Protocols MLD)

**Syntax**

```
static {
 group multicast-group-address {
 exclude;
 group-count number;
 group-increment increment;
 source ip-address {
 source-count number;
 source-increment increment;
 }
 }
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols **mld interface** *interface-name*],  
[edit protocols **mld interface** *interface-name*]

**Release Information** Statement introduced before Junos OS Release 7.4.

**Description** Test multicast forwarding on an interface.

The **static** statement simulates MLD joins on a routing device statically on an interface without any MLD hosts. It is supported for both MLDv1 and MLDv2 joins. This statement is especially useful for testing multicast forwarding on an interface without a receiver host.



**NOTE:** To prevent joining too many groups accidentally, the **static** statement is not supported with the **interface all** statement.

The remaining statements are explained separately.

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

**Related Documentation**

- *Enabling MLD Static Group Membership*

## traceoptions (Protocols MLD)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;<br/>    flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <b>mld</b> ],<br>[edit protocols <b>mld</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>         | <p>Configure MLD tracing options.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p>To trace the paths of multicast packets, use the <b>mtrace</b> command.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Default</b>             | The default MLD trace options are those inherited from the <b>traceoptions</b> statement included at the [edit routing-options] hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place tracing output in the file <b>mld-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also include the <b>size</b> statement to specify the maximum file size.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 2 files</p> <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p><b>MLD Tracing Flags</b></p> <ul style="list-style-type: none"><li>• <b>leave</b>—Leave group messages.</li><li>• <b>mtrace</b>—Mtrace packets. Use the <b>mtrace</b> command to troubleshoot the software.</li><li>• <b>packets</b>—All MLD packets.</li><li>• <b>query</b>—MLD membership query messages, including general and group-specific queries.</li></ul> |



- **report**—Membership report messages.

#### Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—Traces errors and significant events during normal packet processing

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

**no-stamp**—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

**no-world-readable**—(Optional) Do not allow users to read the log file.

**replace**—(Optional) Replace an existing trace file if there is one.

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing and trace—To view this statement in the configuration.<br>routing-control and trace-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Tracing MLD Protocol Traffic</i></li></ul>                                                           |

---

## version (Protocols MLD)

---

|                                 |                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>version version;</code>                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <code>mld interface interface-name</code> ],<br>[edit protocols <code>mld interface interface-name</code> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                      |
| <b>Description</b>              | Configure the MLD version explicitly. MLD version 2 (MLDv2) is used only to support source-specific multicast (SSM).                                                   |
| <b>Options</b>                  | <b>version</b> —MLD version to run on the interface.<br><b>Range:</b> 1 or 2<br><b>Default:</b> 1 (MLDv1)                                                              |
| <b>Required Privilege Level</b> | routing and trace—To view this statement in the configuration.<br>routing-control and trace-control—To add this statement to the configuration.                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Modifying the MLD Version</i></li></ul>                                                                                     |

## Configuration Statements: MSDP

## msdp

```
Syntax msdp {
 disable;
 active-source-limit {
 log-interval seconds;
 log-warning value;
 maximum number;
 threshold number;
 }
 data-encapsulation (disable | enable);
 export [policy-names];
 group group-name {
 ... group-configuration ...
 }
 hold-time seconds;
 import [policy-names];
 local-address address;
 keep-alive seconds;
 peer address {
 ... peer-configuration ...
 }
 rib-group group-name;
 source ip-prefix </prefix-length> {
 active-source-limit {
 maximum number;
 threshold number;
 }
 }
 sa-hold-time seconds;
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
 group group-name {
 disable;
 export [policy-names];
 import [policy-names];
 local-address address;
 mode (mesh-group | standard);
 peer address {
 ... same statements as at the [edit protocols msdp peer address] hierarchy level shown
 just following ...
 }
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
 }
 peer address {
 disable;
 active-source-limit {
 maximum number;
 threshold number;
 }
 }
 }
```

```

 }
 authentication-key peer-key;
 default-peer;
 export [policy-names];
 import [policy-names];
 local-address address;
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
}
}

```

|                                 |                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],<br>[edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.4 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                      |
| <b>Description</b>              | Enable MSDP on the router or switch. You must also configure at least one peer for MSDP to function.                                                                                                                                                                |
| <b>Default</b>                  | MSDP is disabled on the router or switch.                                                                                                                                                                                                                           |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring MSDP in a Routing Instance</i></li> </ul>                                                                                                                                                          |

## active-source-limit

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>active-source-limit {<br/>    log-interval seconds;<br/>    log-warning value;<br/>    maximum number;<br/>    threshold number;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems logical-system-name protocols msdp],<br/>[edit logical-systems logical-system-name protocols msdp group group-name peer address],<br/>[edit logical-systems logical-system-name protocols msdp peer address],<br/>[edit logical-systems logical-system-name protocols msdp source ip-address/prefix-length],<br/>[edit logical-systems logical-system-name routing-instances instance-name protocols msdp],<br/>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols<br/>    msdp group group-name peer address],<br/>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols<br/>    msdp peer address],<br/>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols<br/>    msdp source ip-address/prefix-length],<br/>[edit protocols msdp],<br/>[edit protocols msdp group group-name peer address],<br/>[edit protocols msdp peer address],<br/>[edit protocols msdp source ip-address/prefix-length],<br/>[edit routing-instances routing-instance-name protocols msdp],<br/>[edit routing-instances routing-instance-name protocols msdp group group-name<br/>    peer address],<br/>[edit routing-instances routing-instance-name protocols msdp peer address],<br/>[edit routing-instances routing-instance-name protocols msdp source<br/>    ip-address/prefix-length]</pre> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Limit the number of active source messages the routing device accepts.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Default</b>                  | If you do not include this statement, the router accepts any number of MSDP active source messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | The options are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring MSDP with Active Source Limits and Mesh Groups</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## authentication-key

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authentication-key peer-key;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <code>msdp group group-name peer address</code>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <code>msdp peer address</code>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>msdp group group-name peer address</code>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>msdp peer address</code>],</p> <p>[edit protocols <code>msdp group group-name peer address</code>],</p> <p>[edit protocols <code>msdp peer address</code>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <code>msdp group group-name peer address</code>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <code>msdp peer address</code>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Associate a Message Digest 5 (MD5) signature option authentication key with an MSDP peering session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>                  | If you do not include this statement, the routing device accepts any valid MSDP messages from the peer address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <b>peer-key</b> —MD5 authentication key. The peer key can be a text string up to 16 letters and digits long. Strings can include any ASCII characters with the exception of ( , ) , & , and [ . If you include spaces in an MSDP authentication key, enclose all characters in quotation marks ( " " ).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring MSDP in a Routing Instance</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## data-encapsulation

---

|                                 |                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>data-encapsulation (disable   enable);</code>                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a> ],<br>[edit protocols <a href="#">msdp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure a rendezvous point (RP) using MSDP to encapsulate multicast data received in MSDP register messages inside forwarded MSDP source-active messages.                                                                                                                                                                                                 |
| <b>Default</b>                  | If you do not include this statement, the RP encapsulates multicast data.                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>disable</b> —(Optional) Do not use MSDP data encapsulation.<br><b>enable</b> —Use MSDP data encapsulation.<br><b>Default:</b> <b>enable</b>                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring MSDP with Active Source Limits and Mesh Groups</i></li></ul>                                                                                                                                                                                                                                |



## default-peer

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | default-peer;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp peer</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>],</p> <p>[edit protocols <b>msdp</b>],</p> <p>[edit protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit protocols <b>msdp peer</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Establish this peer as the default MSDP peer and accept source-active messages from the peer without the usual peer-reverse-path-forwarding (peer-RPF) check.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Example: Configuring MSDP with Active Source Limits and Mesh Groups</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## disable (Protocols MSDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b>], [edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>], [edit logical-systems <i>logical-system-name</i> protocols <b>msdp peer</b> <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>], [edit protocols <b>msdp</b>], [edit protocols <b>msdp group</b> <i>group-name</i>], [edit protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>], [edit protocols <b>msdp peer</b> <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>], [edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>]</pre> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Explicitly disable MSDP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>Disabling MSDP</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## export (Protocols MSDP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp peer</a> <i>address</i>],</p> <p>[edit protocols <a href="#">msdp</a>],</p> <p>[edit protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit protocols <a href="#">msdp peer</a> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp peer</a> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Apply one or more policies to routes being exported from the routing table into MSDP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring MSDP in a Routing Instance</i></li> <li>• <a href="#">import on page 3832</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## group

```
Syntax group group-name {
 disable;
 export [policy-names];
 import [policy-names];
 local-address address;
 mode (mesh-group | standard);
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
 peer address; {
 disable;
 active-source-limit {
 maximum number;
 threshold number;
 }
 authentication-key peer-key;
 default-peer;
 export [policy-names];
 import [policy-names];
 local-address address;
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
 }
 }
 }
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols [msdp](#)],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
[msdp](#)],  
 [edit protocols [msdp](#)],  
 [edit routing-instances *routing-instance-name* protocols [msdp](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
 Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Define an MSDP peer group. MSDP peers within groups share common tracing options, if present and not overridden for an individual peer with the [peer](#) statement. To configure multiple MSDP groups, include multiple **group** statements.

By default, the group's options are identical to the global MSDP options. To override the global options, include group-specific options within the **group** statement.

The group must contain at least one peer.

**Options** *group-name*—Name of the MSDP group.

The remaining statements are explained separately.

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

**Related Documentation**

- *Example: Configuring MSDP in a Routing Instance*

## hold-time (Protocols MSDP)

**Syntax** hold-time *seconds*;

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols msdp],  
[edit logical-systems *logical-system-name* protocols msdp group *group-name* peer *address*],  
[edit logical-systems *logical-system-name* protocols msdp peer *address*],  
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols msdp],  
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols msdp  
group *group-name* peer *address*],  
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols msdp  
peer *address*],  
[edit protocols msdp],  
[edit protocols msdp group *group-name* peer *address*],  
[edit protocols msdp peer *address*],  
[edit routing-instances *instance-name* protocols msdp],  
[edit routing-instances *instance-name* protocols msdp group *group-name* peer *address*]  
[edit routing-instances *instance-name* protocols msdp peer *address*],

**Release Information** Statement introduced in Junos OS Release 12.3.

**Description** Specify the hold-time period to use when maintaining a connection with the MSDP peer. If a keepalive message is not received for the hold-time period, the MSDP peer connection is terminated. According to the RFC 3618, *Multicast Source Discovery Protocol (MSDP)*, the recommended value for the hold-time period is 75 seconds.

The hold-time period must be longer than the keepalive interval.

You might want to change the hold-time period and keepalive timer for consistency in a multi-vendor environment.

**Default** In Junos OS, the default hold-time period is 75 seconds, and the default keepalive interval is 60 seconds.

**Options** *seconds*—Hold time.  
**Range:** 15 through 150 seconds  
**Default:** 75 seconds

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

**Related Documentation**

- *Examples: Configuring MSDP*
- [keep-alive \(Protocols MSDP\) on page 3833](#)
- [sa-hold-time \(Protocols MSDP\) on page 3841](#)

## import (Protocols MSDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp peer</a> <i>address</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <a href="#">msdp</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <a href="#">msdp group</a> <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <a href="#">msdp peer</a> <i>address</i>],</code><br><code>[edit protocols <a href="#">msdp</a>],</code><br><code>[edit protocols <a href="#">msdp group</a> <i>group-name</i>],</code><br><code>[edit protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</code><br><code>[edit protocols <a href="#">msdp peer</a> <i>address</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i></code><br><code>    <a href="#">peer</a> <i>address</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp peer</a> <i>address</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Apply one or more policies to routes being imported into the routing table from MSDP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring MSDP in a Routing Instance</i></li><li>• <a href="#">export on page 3829</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## keep-alive (Protocols MSDP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>keep-alive seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols msdp],<br/> [edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i> peer address],<br/> [edit logical-systems <i>logical-system-name</i> protocols msdp peer address],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols msdp],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols msdp group <i>group-name</i> peer address],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols msdp peer address],<br/> [edit protocols msdp],<br/> [edit protocols msdp group <i>group-name</i> peer address],<br/> [edit protocols msdp peer address],<br/> [edit routing-instances <i>instance-name</i> protocols msdp],<br/> [edit routing-instances <i>instance-name</i> protocols msdp group <i>group-name</i> peer address]<br/> [edit routing-instances <i>instance-name</i> protocols msdp peer address],</p> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Specify the keepalive interval to use when maintaining a connection with the MSDP peer. If a keepalive message is not received for the hold-time period, the MSDP peer connection is terminated. According to the RFC 3618, <i>Multicast Source Discovery Protocol (MSDP)</i>, the recommended value for the keepalive timer is 60 seconds.</p> <p>The hold-time period must be longer than the keepalive interval.</p> <p>You might want to change the keepalive interval and hold-time period for consistency in a multi-vendor environment.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Default</b>                  | In Junos OS, the default hold-time period is 75 seconds, and the default keepalive interval is 60 seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b>seconds</b>—Keepalive interval.</p> <p><b>Range:</b> 10 through 60 seconds</p> <p><b>Default:</b> 60 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Examples: Configuring MSDP</a></li> <li>• <a href="#">hold-time (Protocols MSDP) on page 3831</a></li> <li>• <a href="#">sa-hold-time (Protocols MSDP) on page 3841</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## local-address

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>local-address address;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp peer</b> <i>address</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>],</code><br><code>[edit protocols <b>msdp</b>],</code><br><code>[edit protocols <b>msdp group</b> <i>group-name</i>],</code><br><code>[edit protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</code><br><code>[edit protocols <b>msdp peer</b> <i>address</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the local end of an MSDP session. You must configure at least one peer for MSDP to function. When configuring a peer, you must include this statement. This address is used to accept incoming connections to the peer and to establish connections to the remote peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <i>address</i> —IP address of the local end of the connection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li><i>Example: Configuring MSDP in a Routing Instance</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |



## log-interval (Protocols MSDP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | log-interval <i>seconds</i> ;                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp active-source-limit</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a> ],<br>[edit protocols <a href="#">msdp active-source-limit</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Specify the minimum time interval (in seconds) between sending consecutive log messages to the system log for MSDP active source messages. To configure the time interval, you must specify the maximum number of MSDP active source messages received by the device.</p> <p>To confirm the configured log interval, use the <a href="#">show msdp source-active</a> command.</p>                                                        |
| <b>Options</b>                  | <p><b>seconds</b>—Minimum time interval (in seconds) between log messages. You must explicitly configure the <a href="#">maximum</a> value to configure a time interval to send log messages.</p> <p><b>Range:</b> 6 through 32,767 seconds</p>                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups</a></li> <li>• <a href="#">log-warning</a></li> <li>• <a href="#">maximum on page 3837</a></li> </ul>                                                                                                                                                                                                              |

## log-warning (Protocols MSDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | log-warning <i>value</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp active-source-limit</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a> ],<br>[edit protocols <a href="#">msdp active-source-limit</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Specify the threshold at which the device logs a warning message in the system log for received MSDP active source messages. This threshold is a percentage of the maximum number of MSDP active source messages received by the device.</p> <p>To confirm the configured warning threshold, use the <a href="#">show msdp source-active</a> command.</p>                                                                                |
| <b>Options</b>                  | <p><b>value</b>—Percentage of the number of active source messages that starts triggering the warnings. You must explicitly configure the <a href="#">maximum</a> value to configure a warning threshold value.</p> <p><b>Range:</b> 1 through 100</p>                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring MSDP with Active Source Limits and Mesh Groups</i></li><li>• log-interval</li><li>• <a href="#">maximum on page 3837</a></li></ul>                                                                                                                                                                                                                                          |

## maximum

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>maximum <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp active-source-limit</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a> ],<br>[edit protocols <a href="#">msdp active-source-limit</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the maximum number of MSDP active source messages the router accepts.                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <i>number</i> —Maximum number of active source messages.<br><b>Range:</b> 1 through 1,000,000<br><b>Default:</b> 25,000                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring MSDP with Active Source Limits and Mesh Groups</i></li> <li>• <a href="#">threshold on page 3843</a></li> </ul>                                                                                                                                                                                                                                                            |

## mode (Protocols MSDP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | mode (mesh-group   standard);                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> ],<br>[edit protocols <b>msdp group</b> <i>group-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure groups of peers in a full mesh topology to limit excessive flooding of source-active messages to neighboring peers. The default flooding mode is <b>standard</b> .                                                                                                                                                                                                                                            |
| <b>Default</b>                  | If you do not include this statement, default flooding is applied.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b>mesh-group</b> —Group of peers that are mesh group members.<br><br><b>standard</b> —Use standard MSDP source-active flooding rules.<br><b>Default:</b> <b>standard</b>                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring MSDP with Active Source Limits and Mesh Groups</i></li></ul>                                                                                                                                                                                                                                                                                            |

## peer (Protocols MSDP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> peer address {     disable;     active-source-limit {         maximum number;         threshold number;     }     authentication-key peer-key;     default-peer;     export [ policy-names ];     import [ policy-names ];     local-address address;     traceoptions {         file filename &lt;files number&gt; &lt;size size&gt; &lt;world-readable   no-world-readable&gt;;         flag flag &lt;flag-modifier&gt; &lt;disable&gt;;     } } </pre>                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>     | <pre> [edit logical-systems logical-system-name protocols msdp], [edit logical-systems logical-system-name protocols msdp group group-name], [edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp], [edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp group group-name], [edit protocols msdp], [edit protocols msdp group group-name], [edit routing-instances routing-instance-name protocols msdp], [edit routing-instances routing-instance-name protocols msdp group group-name] </pre>                                                                                                                                                                                                                                      |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>         | <p>Define an MSDP peering relationship. An MSDP routing device must know which routing devices are its peers. You define the peer relationships explicitly by configuring the neighboring routing devices that are the MSDP peers of the local routing device. After peer relationships are established, the MSDP peers exchange messages to advertise active multicast sources. To configure multiple MSDP peers, include multiple <b>peer</b> statements.</p> <p>By default, the peer's options are identical to the global or group-level MSDP options. To override the global or group-level options, include peer-specific options within the <b>peer (Protocols MSDP)</b> statement.</p> <p>At least one peer must be configured for MSDP to function. You must configure <b>address</b> and <b>local-address</b>.</p> |
| <b>Options</b>             | <p><b>address</b>—Name of the MSDP peer.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

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

**Related Documentation**

- *Example: Configuring MSDP in a Routing Instance*

---

## rib-group (Protocols MSDP)

---

**Syntax** `rib-group group-name;`

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols [msdp](#)],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols [msdp](#)],  
[edit protocols [msdp](#)],  
[edit routing-instances *routing-instance-name* protocols [msdp](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Associate a routing table group with MSDP.

**Options** *group-name*—Name of the routing table group. The name must be one that you defined with the **rib-groups** statement at the [edit routing-options] hierarchy level.

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

**Related Documentation**

- *Example: Configuring MSDP in a Routing Instance*

## sa-hold-time (Protocols MSDP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>sa-hold-time <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols msdp], [edit logical-systems <i>logical-system-name</i> protocols msdp group <i>group-name</i> peer address], [edit logical-systems <i>logical-system-name</i> protocols msdp peer address], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols msdp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols msdp   group <i>group-name</i> peer address], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols msdp   peer address], [edit protocols msdp], [edit protocols msdp group <i>group-name</i> peer address], [edit protocols msdp peer address], [edit routing-instances <i>instance-name</i> protocols msdp], [edit routing-instances <i>instance-name</i> protocols msdp group <i>group-name</i> peer address] [edit routing-instances <i>instance-name</i> protocols msdp peer address],</pre> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>Specify the source address (SA) message hold time to use when maintaining a connection with the MSDP peer. Each entry in an SA cache has an associated hold time. The hold timer is started when an SA message is received by an MSDP peer. The timer is reset when another SA message is received before the timer expires. If another SA message is not received during the SA message hold-time period, the SA message is removed from the cache.</p> <p>You might want to change the SA message hold time for consistency in a multi-vendor environment.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b><i>seconds</i></b>—Source address message hold time.</p> <p><b>Range:</b> 75 through 300 seconds</p> <p><b>Default:</b> 75 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Examples: Configuring MSDP</i></li> <li>• <a href="#">hold-time (Protocols MSDP) on page 3831</a></li> <li>• <a href="#">keep-alive (Protocols MSDP) on page 3833</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## source

---

|                                 |                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>source ip-address &lt;/prefix-length&gt; {<br/>    active-source-limit {<br/>        maximum number;<br/>        threshold number;<br/>    }<br/>}</pre>                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp</b> ],<br>[edit protocols <b>msdp</b> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                  |
| <b>Description</b>              | Limit the number of active source messages the routing device accepts from sources in this address range.                                                                                                                                                                                                               |
| <b>Default</b>                  | If you do not include this statement, the routing device accepts any number of MSDP active source messages.                                                                                                                                                                                                             |
| <b>Options</b>                  | The other statements are explained separately.                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring MSDP with Active Source Limits and Mesh Groups</i></li></ul>                                                                                                                                                                                            |



## threshold

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>threshold <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp active-source-limit</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a> ],<br>[edit protocols <a href="#">msdp active-source-limit</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the random early detection (RED) threshold for MSDP active source messages. This number must be less than the configured or default maximum.                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <i>number</i> —RED threshold for active source messages.<br><b>Range:</b> 1 through 1,000,000<br><b>Default:</b> 24,000                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Example: Configuring MSDP with Active Source Limits and Mesh Groups</i></li> <li><a href="#">maximum on page 3837</a></li> </ul>                                                                                                                                                                                                                                                                  |

## traceoptions (Protocols MSDP)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>     | <pre> [edit logical-systems <i>logical-system-name</i> protocols <i>msdp</i>], [edit logical-systems <i>logical-system-name</i> protocols <i>msdp</i> group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols <i>msdp</i> group <i>group-name</i> peer <i>address</i>], [edit logical-systems <i>logical-system-name</i> protocols <i>msdp</i> peer <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>msdp</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>msdp</i> group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>msdp</i> group <i>group-name</i> peer <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>msdp</i> peer <i>address</i>], [edit protocols <i>msdp</i>], [edit protocols <i>msdp</i> group <i>group-name</i>], [edit protocols <i>msdp</i> group <i>group-name</i> peer <i>address</i>], [edit protocols <i>msdp</i> peer <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols <i>msdp</i>], [edit routing-instances <i>routing-instance-name</i> protocols <i>msdp</i> group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols <i>msdp</i> group <i>group-name</i> peer <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols <i>msdp</i> peer <i>address</i>] </pre> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>         | <p>Configure MSDP tracing options.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Default</b>             | The default MSDP trace options are those inherited from the routing protocol's <b>traceoptions</b> statement included at the <b>[edit routing-options]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place tracing output in the <b>msdp-log</b> file.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b><i>trace-file</i></b> reaches its maximum size, it is renamed <b><i>trace-file.0</i></b>, then <b><i>trace-file.1</i></b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

**Range:** 2 through 1000 files

**Default:** 2 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

#### MSDP Tracing Flags

- **keepalive**—Keepalive messages
- **packets**—All MSDP packets
- **route**—MSDP changes to the routing table
- **source-active**—Source-active packets
- **source-active-request**—Source-active request packets
- **source-active-response**—Source-active response packets

#### Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

***flag-modifier***—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

**no-stamp**—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

**no-world-readable**—(Optional) Do not allow any user to read the log file.

**replace**—(Optional) Replace an existing trace file if there is one.

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing and trace—To view this statement in the configuration.<br>routing-control and trace-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Tracing MSDP Protocol Traffic</i></li></ul>                                                          |

## Configuration Statements: PIM

- [\[edit protocols pim\] Hierarchy Level on page 3846](#)

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

---

The following statement hierarchy can also be included at the [\[edit logical-systems logical-system-name\]](#) hierarchy level.

```
protocols {
 pim {
 disable;
 assert-timeout seconds;
 default-vpn-source {
 interface-name interface-name; }
 }
}
```

```

}
dense-groups {
 address <announce | reject>;
}
dr-election-on-p2p;
export [policy-names];
family (inet | inet6) {
 disable;
}
graceful-restart {
 disable;
 no-bidirectional-mode;
 restart-duration seconds;
}
import [policy-names];
interface interface-name {
 ... the interface subhierarchy appears after the main [edit protocols pim] hierarchy ...
 family (inet | inet6) {
 disable;
 }
}
join-load-balance;
join-prune-timeout seconds;
nonstop-routing {
 disable;
}
override-interval milliseconds;
propagation-delay milliseconds;
reset-tracking-bit;
rib-group {
 inet group-name;
 inet6 group-name;
}
rp {
 ... the rp subhierarchy appears after the main [edit protocols pim] hierarchy ...
}
sglimit {
 family (inet | inet6) {
 log-interval seconds;
 maximum limit;
 threshold value;
 }
}
log-interval seconds;
maximum limit;
threshold value;
}
spt-threshold {
 infinity [policy-names];
}
traceoptions {
 file filename <files number> <size maximum-file-size> <world-readable |
 no-world-readable>;
 flag flag <flag-modifier> <disable>;
}

```

```
 flag (route | state) <flag-modifier> <disable> <filter <match-on prefix>
 <policy [policy-names]>>;
 }
}

pim {
 interface interface-name {
 accept-remote-source;
 disable;
 bfd-liveness-detection {
 authentication {
 algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
 meticulous-keyed-sha-1 | simple-password);
 key-chain key-chain-name;
 loose-check;
 }
 detection-time {
 threshold milliseconds;
 }
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 version (1 | automatic);
 }
 bidirectional {
 df-election {
 backoff-period milliseconds;
 offer-period milliseconds;
 robustness-count number;
 }
 }
 family (inet | inet6) {
 disable;
 }
 hello-interval seconds;
 bidirectional-sparse | bidirectional-sparse-dense mode (bidirectional-sparse |
 bidirectional-sparse-dense | dense | sparse | sparse-dense);
 neighbor-policy [policy-names];
 override-interval milliseconds;
 priority number;
 propagation-delay milliseconds;
 reset-tracking-bit;
 version (1 | 2);
 }
}

pim {
 rp {
 auto-rp {
 (announce | discovery | mapping);
 (mapping-agent-election | no-mapping-agent-election);
 }
 }
}
```

```

}
bidirectional {
 address address {
 group-ranges {
 destination-ip-prefix</prefix-length>;
 }
 hold-time seconds;
 priority number;
 }
}
bootstrap {
 family (inet | inet6) {
 export [policy-names];
 import [policy-names];
 priority number;
 }
}
bootstrap-export [policy-names];
bootstrap-import [policy-names];
bootstrap-priority number;
dr-register-policy [policy-names];
embedded-rp {
 group-ranges {
 ip-prefix</prefix-length>;
 }
 maximum-rps limit;
}
group-rp-mapping {
 family (inet | inet6) {
 log-interval seconds;
 maximum limit;
 threshold value;
 }
}
log-interval seconds;
maximum limit;
threshold value;
}
}
local {
 ... the local subhierarchy appears after the main [edit protocols pim rp] hierarchy ...
}
register-limit {
 family (inet | inet6) {
 log-interval seconds;
 maximum limit;
 threshold value;
 }
}
log-interval seconds;
maximum limit;
threshold value;
}
}
rp-register-policy [policy-names];
static {

```

```
 address address {
 group-ranges {
 ip-prefix</prefix-length>;
 }
 override;
 version (1 | 2);
 }
 }
}

rp {
 local {
 disable;
 address address;
 family (inet | inet6) {
 disable;
 address address;
 anycast-pim {
 local-address address;
 rp-set {
 address address <forward-msdp-sa>;
 }
 }
 group-ranges {
 ip-prefix</prefix-length>;
 }
 hold-time seconds;
 override;
 priority number;
 }
 group-ranges {
 ip-prefix</prefix-length>;
 }
 hold-time seconds;
 override;
 priority number;
 }
}
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
  - *[edit protocols] Hierarchy Level*



## accept-remote-source

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | accept-remote-source;                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>pim interface <i>interface-name</i> ],<br>[edit protocols pim interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 9.6 for EX Series switches.                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Accept traffic from a remote source. A remote source is a source that is not on the same subnet as the incoming interface. This statement enables the remote source to be learned and advertised by MSDP so that receivers in other MSDP areas can join the source. You do not need to disable RPF checking, but you do need to ensure that the best path to reach the remote source is through the incoming interface.    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Interface to Accept Traffic from a Remote Source</i></li> <li>• <i>Example: Allowing MBGP MVPN Remote Sources</i></li> </ul>                                                                                                                                                                                                                                   |

## address (Anycast RPs)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>address <i>address</i> &lt;forward-msdp-sa&gt;;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <b>pim rp local</b> (inet   inet6) <b>anycast-pim rp-set</b>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim rp local</b> (inet   inet6) <b>anycast-pim rp-set</b>],</code><br><code>[edit protocols <b>pim rp local</b> (inet   inet6) <b>anycast-pim rp-set</b>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim rp local</b> (inet   inet6) <b>anycast-pim rp-set</b>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the anycast rendezvous point (RP) addresses in the RP set. Multiple addresses can be configured in an RP set. If the RP has peer Multicast Source Discovery Protocol (MSDP) connections, then the RP must forward MSDP source active (SA) messages.                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <b><i>address</i></b> —RP address in an RP set.<br><br><b><i>forward-msdp-sa</i></b> —(Optional) Forward MSDP SAs to this address.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## address (Bidirectional Rendezvous Points)

|                                 |                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> address address {   group-ranges {     destination-ip-prefix &lt;/prefix-length&gt;;   }   hold-time seconds;   priority number; } </pre>                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <pre> [edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   pim rp bidirectional], [edit protocols pim rp bidirectional], [edit routing-instances <i>routing-instance-name</i> protocols pim rp bidirectional] </pre> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Configure bidirectional rendezvous point (RP) addresses. The address can be a loopback interface address, an address of a link interface, or an address that is not assigned to an interface but belongs to a subnet that is reachable by the bidirectional PIM routing devices in the network.                                                               |
| <b>Options</b>                  | <p><b>address</b>—Bidirectional RP address.</p> <p><b>Default:</b> 232.0.0.0/8</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding Bidirectional PIM</i></li> <li>• <i>Example: Configuring Bidirectional PIM</i></li> </ul>                                                                                                                                                                                                           |

## address (Local RPs)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>address address;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <code>pim rp local family</code> (inet   inet6)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br><code>pim rp local family</code> (inet   inet6)],<br>[edit protocols <code>pim rp local family</code> (inet   inet6)],<br>[edit routing-instances <i>routing-instance-name</i> protocols <code>pim rp local family</code> (inet   inet6)] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Configure the local rendezvous point (RP) address.                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <i>address</i> —Local RP address.                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Local PIM RPs on page 3695</a></li></ul>                                                                                                                                                                                                                                                                                                                                                               |

## address (Static RPs)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> address address {   group-ranges {     destination-ip-prefix&lt;/prefix-length&gt;;   }   override;   version version; } </pre>                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp static</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp static</a>],</p> <p>[edit protocols <a href="#">pim static</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp static</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                               |
| <b>Description</b>              | <p>Configure static rendezvous point (RP) addresses. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.</p> <p>For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.</p>                                                       |
| <b>Options</b>                  | <p><b>address</b>—Static RP address.</p> <p><b>Default:</b> 224.0.0.0/4</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Static PIM RP Address on the Non-RP Routing Device on page 3699</a></li> </ul>                                                                                                                                                                                                                                                         |

## algorithm

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>algorithm <i>algorithm-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication]                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Specify the algorithm to use for BFD authentication.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><i>algorithm-name</i>—Name of algorithm to use for BFD authentication:</p> <ul style="list-style-type: none"><li>• <b>simple-password</b>—Plain-text password. One to 16 bytes of plain text. One or more passwords can be configured.</li><li>• <b>keyed-md5</b>—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive rates greater than 100 ms.</li><li>• <b>meticulous-keyed-md5</b>—Meticulous keyed Message Digest 5 hash algorithm.</li><li>• <b>keyed-sha-1</b>—Keyed Secure Hash Algorithm I for sessions with transmit and receive rates greater than 100 ms.</li><li>• <b>meticulous-keyed-sha-1</b>—Meticulous keyed Secure Hash Algorithm I.</li></ul> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Bidirectional Forwarding Detection Authentication for PIM on page 3726</a></li><li>• <a href="#">Configuring BFD Authentication for PIM on page 3729</a></li><li>• <a href="#">authentication on page 3859</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                    |

## anycast-pim

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>anycast-pim {   rp-set {     address address &lt;forward-msdp-sa&gt;;   } }</pre>                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>pim rp local family</b> (inet   inet6)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br><b>pim rp local family</b> (inet   inet6)],<br>[edit protocols <b>pim rp local family</b> (inet   inet6)],<br>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim rp local family</b> (inet   inet6)] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Configure properties for anycast RP using PIM.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring PIM Anycast With or Without MSDP</i></li> </ul>                                                                                                                                                                                                                                                                                                                               |

## assert-timeout

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>assert-timeout seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ]                                                                                                                  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Multicast routing devices running PIM sparse mode often forward the same stream of multicast packets onto the same LAN through the rendezvous-point tree (RPT) and shortest-path tree (SPT). PIM assert messages help routing devices determine which routing device forwards the traffic and prunes the RPT for this group. By default, routing devices enter an assert cycle every 180 seconds. You can configure this assert timeout to be between 5 and 210 seconds. |
| <b>Options</b>                  | <b>seconds</b> —Time for routing device to wait before another assert message cycle.<br><b>Range:</b> 5 through 210 seconds<br><b>Default:</b> 180 seconds                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring the PIM Assert Timeout</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                     |



## authentication (Protocols PIM)


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>authentication {   algorithm <i>algorithm-name</i>;   key-chain <i>key-chain-name</i>;   loose-check; }</pre>                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <pre>[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>   bfd-liveness-detection]</pre>                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Configure the algorithm, security keychain, and level of authentication for BFD sessions running on PIM interfaces.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD Authentication for PIM on page 3729</a></li> <li>• <a href="#">Configuring BFD for PIM on page 3727</a></li> <li>• <a href="#">Understanding Bidirectional Forwarding Detection Authentication for PIM on page 3726</a></li> <li>• <a href="#">bfd-liveness-detection on page 3862</a></li> <li>• <a href="#">key-chain (Protocols PIM) on page 3894</a></li> <li>• <a href="#">loose-check on page 3898</a></li> </ul> |

## auto-rp

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>auto-rp {<br/>    (announce   discovery   mapping);<br/>    (mapping-agent-election   no-mapping-agent-election);<br/>}</pre>                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ],<br>[edit protocols <a href="#">pim rp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ]                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure automatic RP announcement and discovery.                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b>announce</b>—Configure the routing device to listen only for mapping packets and also to advertise itself if it is an RP.</p> <p><b>discovery</b>—Configure the routing device to listen only for mapping packets.</p> <p><b>mapping</b>—Configures the routing device to announce, listen for and generate mapping packets, and announce that the routing device is eligible to be an RP.</p> <p>The remaining statement is explained separately.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring PIM Auto-RP on page 3705</a></li></ul>                                                                                                                                                                                                                                                                                                                                                       |

## backoff-period

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <code>backoff-period <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <p>[edit logical-systems <i>logical-system-name</i> protocols <code>pim interface interface-name</code> bidirectional df-election],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>pim interface interface-name</code> bidirectional df-election],</p> <p>[edit protocols <code>pim interface interface-name</code> bidirectional df-election],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <code>pim interface interface-name</code> bidirectional df-election]</p> |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Statement introduced in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <p>Configure the designated forwarder (DF) election backoff period for bidirectional PIM. The <b>backoff-period</b> statement configures the period that the acting DF waits between receiving a better DF Offer and sending the Pass message to transfer DF responsibility.</p>                                                                                                                                                                                                                                                                                       |
| <div>  <p><b>NOTE:</b> Junos OS checks rendezvous point (RP) unicast reachability before accepting incoming DF messages. DF messages for unreachable rendezvous points are ignored. This is needed to prevent the following example scenario. Routers A and B are downstream routing devices on the same LAN, and both are supposed to send DF election messages with an infinite metric on their upstream interfaces (reverse-path forwarding [RPF] interfaces). Router A has a higher IP address than Router B. When both routing devices lose the path to the RP, both send an Offer message with the infinite metric onto the LAN. Router A wins the election because it has a higher IP address, and Router B backs off as a result. After three Offer messages, according to RFC 5015, Router A looks up the RP and finds no path to the RP. As a result, Router A transitions to the Lose state and sends nothing. On the other hand, after backing off for an interval of 3 x the Offer period, Router B does not receive any messages, and resumes the DF election by sending a new Offer message. Hence, the pattern repeats indefinitely.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <p><b>milliseconds</b>—Period that the acting DF waits between receiving a better DF Offer and sending the Pass message to transfer DF responsibility.</p> <p><b>Range:</b> 100 through 65,535 milliseconds</p> <p><b>Default:</b> 1000</p>                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <ul style="list-style-type: none"> <li>• <i>Understanding Bidirectional PIM</i></li> <li>• <i>Example: Configuring Bidirectional PIM</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                    |

## bfd-liveness-detection (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>bfd-liveness-detection {<br/>  authentication {<br/>    algorithm <i>algorithm-name</i>;<br/>    key-chain <i>key-chain-name</i>;<br/>    loose-check;<br/>  }<br/>  detection-time {<br/>    threshold <i>milliseconds</i>;<br/>  }<br/>  minimum-interval <i>milliseconds</i>;<br/>  minimum-receive-interval <i>milliseconds</i>;<br/>  multiplier <i>number</i>;<br/>  no-adaptation;<br/>  transmit-interval {<br/>    minimum-interval <i>milliseconds</i>;<br/>    threshold <i>milliseconds</i>;<br/>  }<br/>  version (0   1   automatic);<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">pim interface interface-name</a> ],<br>[edit protocols <a href="#">pim interface interface-name family</a> (inet   inet6)],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim interface interface-name family</a> (inet   inet6)]                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br><b>authentication</b> option introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Configure bidirectional forwarding detection (BFD) timers and authentication for PIM.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 3727</a></li><li>• <a href="#">Configuring BFD Authentication for PIM on page 3729</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                 |

## bidirectional (Interface)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> bidirectional {   df-election {     backoff-period <i>milliseconds</i>;     offer-period <i>milliseconds</i>;     robustness-count <i>number</i>;   } } </pre>                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim interface interface-name</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim interface interface-name</a>],</p> <p>[edit protocols <a href="#">pim interface interface-name</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim interface interface-name</a>]</p> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | <p>Configure parameters for bidirectional PIM.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding Bidirectional PIM</i></li> <li>• <i>Example: Configuring Bidirectional PIM</i></li> </ul>                                                                                                                                                                                                                                                                                                                        |

## bidirectional (RP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>bidirectional {<br/>    address address {<br/>        group-ranges {<br/>            destination-ip-prefix&lt;/prefix-length&gt;;<br/>        }<br/>        hold-time seconds;<br/>        priority number;<br/>    }<br/>}</pre>                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ],<br>[edit protocols <a href="#">pim rp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure the routing device's rendezvous-point (RP) properties for bidirectional PIM.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Understanding Bidirectional PIM</i></li><li>• <i>Example: Configuring Bidirectional PIM</i></li></ul>                                                                                                                                                                                                                    |

## bootstrap

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>bootstrap {     family (inet   inet6) {         export [ <i>policy-names</i> ];         import [ <i>policy-names</i> ];         priority <i>number</i>;     } }</pre>                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a>],</p> <p>[edit protocols <a href="#">pim rp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                          |
| <b>Description</b>              | <p>Configure parameters to control bootstrap routers and messages.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4 on page 3701</a></li> <li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 3702</a></li> </ul>                                                                                                                                                       |

## bootstrap-export

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bootstrap-export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ],<br>[edit protocols <a href="#">pim rp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                      |
| <b>Description</b>              | Apply one or more export policies to control outgoing PIM bootstrap messages.                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4 on page 3701</a></li><li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 3702</a></li><li>• <a href="#">bootstrap-import on page 3867</a></li></ul>                                                                                   |



## bootstrap-import

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bootstrap-import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a>],</p> <p>[edit protocols <a href="#">pim rp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                      |
| <b>Description</b>              | Apply one or more import policies to control incoming PIM bootstrap messages.                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4 on page 3701</a></li> <li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 3702</a></li> <li>• <a href="#">bootstrap-export on page 3866</a></li> </ul>                                                                                              |

## bootstrap-priority

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>bootstrap-priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ],<br>[edit protocols <a href="#">pim rp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                      |
| <b>Description</b>              | Configure whether this routing device is eligible to be a bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router.                                                                                                                                                                             |
| <b>Options</b>                  | <b><i>number</i></b> —Priority for becoming the bootstrap router. A value of 0 means that the routing device is not eligible to be the bootstrap router.<br><b>Range:</b> 0 through 255<br><b>Default:</b> 0                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4 on page 3701</a></li></ul>                                                                                                                                                                                                                                        |

## dense-groups

---

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>dense-groups {<br/>    <i>addresses</i>;<br/>}</code>                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                          |
| <b>Description</b>              | Configure which groups are operating in dense mode.                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <i>addresses</i> —Address of groups operating in dense mode.                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring PIM Sparse-Dense Mode Properties on page 3741</a></li> </ul>                                                                                                                                                                                                                           |

## detection-time (BFD for PIM)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>detection-time {<br/>    threshold milliseconds;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Hierarchy Level          | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Release Information      | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for BFD authentication introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Description              | <p>Enable BFD failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the <b>clear bfd adaptation</b> command to return BFD interval timers to their configured values. The <b>clear bfd adaptation</b> command is hitless, meaning that the command does not affect traffic flow on the routing device.</p> <p>The remaining statement is explained separately.</p> |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 3727</a></li><li>• <a href="#">bfd-liveness-detection on page 3862</a></li><li>• <a href="#">threshold on page 3938</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## df-election

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | df-election {<br>backoff-period <i>milliseconds</i> ;<br>offer-period <i>milliseconds</i> ;<br>robustness-count <i>number</i> ;<br>}                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>pim interface</b> <i>interface-name</i> bidirectional],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim interface</b> <i>interface-name</i> bidirectional],<br>[edit protocols <b>pim interface</b> <i>interface-name</i> bidirectional],<br>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim interface</b> <i>interface-name</i> bidirectional] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Optionally, configure the designated forwarder (DF) election parameters for bidirectional PIM.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Understanding Bidirectional PIM</i></li> <li>• <i>Example: Configuring Bidirectional PIM</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                     |

## disable (PIM Graceful Restart)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>pim graceful-restart</b> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim graceful-restart</b> ],<br>[edit protocols <b>pim graceful-restart</b> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim graceful-restart</b> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                          |
| <b>Description</b>              | Explicitly disable PIM sparse mode graceful restart.                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Sparse Mode Graceful Restart</i></li> </ul>                                                                                                                                                                                                                                                                                 |

## disable (PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>pim</b>],<br/>[edit logical-systems <i>logical-system-name</i> protocols <b>pim family</b> (inet   inet6)],<br/>[edit logical-systems <i>logical-system-name</i> protocols <b>pim interface</b> <i>interface-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> protocols <b>pim rp local family</b> (inet   inet6)],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim</b>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim interface</b> <i>interface-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim rp local family</b> (inet   inet6)],<br/>[edit protocols <b>pim</b>],<br/>[edit protocols <b>pim family</b> (inet   inet6)],<br/>[edit protocols <b>pim interface</b> <i>interface-name</i>],<br/>[edit protocols <b>pim rp local family</b> (inet   inet6)],<br/>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim</b>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim family</b> (inet   inet6)],<br/>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim interface</b> <i>interface-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim rp local family</b> (inet   inet6)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>disable</b> statement extended to the <b>[family]</b> hierarchy level in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Explicitly disable PIM at the protocol, interface or family hierarchy levels.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Disabling PIM on page 3687</a></li><li>• <a href="#">disable (PIM Graceful Restart) on page 3871</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## dr-election-on-p2p

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>dr-election-on-p2p;</code>                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.1.<br>Statement introduced in Junos OS Release 9.1 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                              |
| <b>Description</b>              | Enable PIM designated router (DR) election on point-to-point (P2P) links.                                                                                                                                                                                                                                                                               |
| <b>Default</b>                  | No PIM DR election is performed on point-to-point links.                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring PIM Designated Router Election on Point-to-Point Links on page 3693</a></li> </ul>                                                                                                                                                                                                     |

## dr-register-policy

|                                 |                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>dr-register-policy [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ],<br>[edit protocols <a href="#">pim rp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                          |
| <b>Description</b>              | Apply one or more policies to control outgoing PIM register messages.                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Register Message Filters on a PIM RP and DR on page 3723</a></li> <li>• <a href="#">rp-register-policy on page 3930</a></li> </ul>                                                                                                                                                                 |

## embedded-rp

---

|                          |                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>embedded-rp {<br/>  group-ranges {<br/>    destination-ip-prefix &lt;/prefix-length&gt;;<br/>  }<br/>  maximum-rps limit;<br/>}</pre>                                                                                                                                                                                                                          |
| Hierarchy Level          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ],<br>[edit protocols <a href="#">pim rp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ] |
| Release Information      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                      |
| Description              | Configure properties for embedded IP version 6 (IPv6) RPs.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                 |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Configuring PIM Embedded RP for IPv6 on page 3711</a></li></ul>                                                                                                                                                                                                                                                 |



## export (Protocols PIM Bootstrap)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp bootstrap family</a> (inet   inet6)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp bootstrap family</a> (inet   inet6)],<br>[edit protocols <a href="#">pim rp bootstrap family</a> (inet   inet6)],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp bootstrap family</a> (inet   inet6)] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Apply one or more export policies to control outgoing PIM bootstrap messages.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4 on page 3701</a></li> <li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 3702</a></li> <li>• <a href="#">import (Protocols PIM Bootstrap) on page 3887</a></li> </ul>                                                                                                                                                                                           |

## export (Protocols PIM)

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>export [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                      |
| <b>Description</b>              | Apply one or more export policies to control outgoing PIM join and prune messages. PIM join and prune filters can be applied to PIM-SM and PIM-SSM messages. PIM join and prune filters cannot be applied to PIM-DM messages.                                                                                                                           |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Filtering Outgoing PIM Join Messages on page 3716</a></li> </ul>                                                                                                                                                                                                                                   |

## family (Bootstrap)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>family (inet   inet6) {<br/>    export [ <i>policy-names</i> ];<br/>    import [ <i>policy-names</i> ];<br/>    priority <i>number</i>;<br/>}</pre>                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp bootstrap</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp bootstrap</a> ],<br>[edit protocols <a href="#">pim rp bootstrap</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp bootstrap</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure which IP protocol type bootstrap properties to apply.                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>inet</b>—Apply IP version 4 (IPv4) local RP properties.</p> <p><b>inet6</b>—Apply IPv6 local RP properties.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4 on page 3701</a></li><li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 3702</a></li></ul>                                                                                                                                                                                   |

## family (Protocols PIM)

---

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                | family (inet   inet6) {<br>disable;<br>}                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>       | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> <a href="#">interface</a> <i>interface-name</i> ] |
| <b>Release Information</b>   | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>           | Enable the PIM protocol for the specified family.                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>               | <b>inet</b> —Enable the PIM protocol for the IP version 4 (IPv4) address family.<br><br><b>inet6</b> —Enable the PIM protocol for the IP version 6 (IPv6) address family.<br><br>The remaining statement is explained separately.                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <a href="#">Disabling PIM on page 3687</a></li> <li>• <a href="#">disable (PIM Graceful Restart) on page 3871</a></li> <li>• <a href="#">disable (PIM) on page 3872</a></li> </ul>                                                                                                                                                                                                                                                                      |

## family (Protocols PIM Interface)

---

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                | <pre>family (inet   inet6) {<br/>    bfd-liveness-detection {<br/>        authentication {<br/>            algorithm <i>algorithm-name</i>;<br/>            key-chain <i>key-chain-name</i>;<br/>            loose-check;<br/>        }<br/>        detection-time {<br/>            threshold <i>milliseconds</i>;<br/>        }<br/>        minimum-interval <i>milliseconds</i>;<br/>        minimum-receive-interval <i>milliseconds</i>;<br/>        multiplier <i>number</i>;<br/>        no-adaptation;<br/>        transmit-interval {<br/>            minimum-interval <i>milliseconds</i>;<br/>            threshold <i>milliseconds</i>;<br/>        }<br/>        version (0   1   automatic);<br/>    }<br/>    disable;<br/>}</pre> |
| <b>Hierarchy Level</b>       | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br><a href="#">pim interface interface-name</a> ],<br>[edit protocols <a href="#">pim interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim interface interface-name</a> ]                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>   | Statement introduced in Junos OS Release 9.6.<br>Support for the Bidirectional Forwarding Detection (BFD) Protocol statements was introduced in Junos OS Release 12.2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>           | Configure one of the following PIM protocol settings for the specified family on the specified interface: <ul style="list-style-type: none"><li>• BFD protocol settings</li><li>• Disable PIM</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>               | <b>inet</b> —Enable the PIM protocol for the IP version 4 (IPv4) address family.<br><br><b>inet6</b> —Enable the PIM protocol for the IP version 6 (IPv6) address family.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Configuring PIM and the Bidirectional Forwarding Detection (BFD) Protocol on page 3725</a></li><li>• <a href="#">Disabling PIM on page 3687</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## family (Local RP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>family (inet   inet6) {     disable;     address address;     anycast-pim {         local-address address;         rp-set {             address address &lt;forward-msdp-sa&gt;;         }     }     group-ranges {         destination-ip-prefix &lt;/prefix-length&gt;;     }     hold-time seconds;     override;     priority number; }</pre>                                                     |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp local</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp local</a>],</p> <p>[edit protocols <a href="#">pim rp local</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp local</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                              |
| <b>Description</b>              | Configure which IP protocol type local RP properties to apply.                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>inet</b>—Apply IP version 4 (IPv4) local RP properties.</p> <p><b>inet6</b>—Apply IPv6 local RP properties.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring Local PIM RPs on page 3695</a></li> </ul>                                                                                                                                                                                                                                                                                                   |

## graceful-restart (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>graceful-restart {<br/>  disable;<br/>  no-bidirectional-mode;<br/>  restart-duration seconds;<br/>}</pre>                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                          |
| <b>Description</b>              | Configure PIM sparse mode graceful restart.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring PIM Sparse Mode Graceful Restart</i></li></ul>                                                                                                                                                                                                                                                   |

## group (RPF Selection)

---


|                                 |                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> group group-address{   source source-address{     next-hop next-hop-address;   }   wildcard-source {     next-hop next-hop-address;   } } </pre> |
| <b>Hierarchy Level</b>          | [edit routing-instances <i>routing-instance-name</i> edit protocols pim rpf-selection]                                                                 |
| <b>Release Information</b>      | <p>Statement introduced in JUNOS Release 10.4.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                            |
| <b>Description</b>              | Configure the PIM group address for which you configure RPF selection <a href="#">group (RPF Selection)</a> .                                          |
| <b>Default</b>                  | By default, PIM RPF selection is not configured.                                                                                                       |
| <b>Options</b>                  | <b>group-address</b> —PIM group address for which you configure RPF selection.                                                                         |
| <b>Required Privilege Level</b> | <p>view-level—To view this statement in the configuration.</p> <p>control-level—To add this statement to the configuration.</p>                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring PIM RPF Selection</i></li> </ul>                                                      |

## group-ranges

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>group-ranges {     destination-ip-prefix&lt;/prefix-length&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp embedded-rp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp embedded-rp</a>],</p> <p>[edit protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols <a href="#">pim rp embedded-rp</a>],</p> <p>[edit protocols <a href="#">pim rp local family</a> (inet   inet6)],</p> <p>[edit protocols <a href="#">pim rp static address</a> <i>address</i>],</p> <p>[edit routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp embedded-rp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp local family</a> (inet   inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp static address</a> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional RP addresses introduced in Junos OS Release 12.1.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure the address ranges of the multicast groups for which this routing device can be a rendezvous point (RP).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>                  | The routing device is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b><i>destination-ip-prefix&lt;/prefix-length&gt;</i></b> —Addresses or address ranges for which this routing device can be an RP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Local PIM RPs on page 3695</a> in the <i>Multicast Protocols Configuration Guide</i></li> <li>• <a href="#">Configuring PIM Embedded RP for IPv6 on page 3711</a> in the <i>Multicast Protocols Configuration Guide</i></li> <li>• <a href="#">Example: Configuring Bidirectional PIM</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |



## group-rp-mapping

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>group-rp-mapping {   family (inet   inet6) {     log-interval seconds;     maximum limit;     threshold value;   }   log-interval seconds;   maximum limit;   threshold value; }</pre>                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a>],</p> <p>[edit protocols <a href="#">pim rp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a>]</p>                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Configure a limit for the number of incoming group-to-RP mappings.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                                 | <div>  <p><b>NOTE:</b> The maximum limit settings that you configure with the <code>maximum</code> and the <code>family (inet   inet6) maximum</code> statements are mutually exclusive. For example, if you configure a global maximum group-to-RP mapping limit, you cannot configure a limit at the family level for IPv4 or IPv6. If you attempt to configure a limit at both the global level and the family level, the device will not accept the configuration.</p> </div> |
| <b>Options</b>                  | <p><b>family (inet   inet6)</b>—(Optional) Specify either IPv4 or IPv6 messages to be counted towards the configured group-to-RP mapping limit.</p> <p><b>Default:</b> Both IPv4 and IPv6 messages are counted towards the configured group-to-RP limit.</p> <p>The remaining statements are described separately.</p>                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring PIM State Limits on page 5198</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## hello-interval (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | hello-interval <i>seconds</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br><a href="#">pim interface interface-name</a> ],<br>[edit protocols <a href="#">pim interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Specify how often the routing device sends PIM hello packets out of an interface.                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <b>seconds</b> —Length of time between PIM hello packets.<br><b>Range:</b> 0 through 255<br><b>Default:</b> 30 seconds                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">hold-time on page 3885</a></li><li>• <a href="#">Modifying the PIM Hello Interval on page 3683</a></li></ul>                                                                                                                                                                                                                                                                                               |

## hold-time (Protocols PIM)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>hold-time seconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols <b>pim rp local family</b> (inet   inet6)],</p> <p>[edit routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim rp local family</b> (inet   inet6)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional RP addresses introduced in Junos OS Release 12.1.</p>                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Specify the time period for which a neighbor is to consider the sending routing device (this routing device) to be operative (up).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>seconds</b>—Hold time.</p> <p><b>Range:</b> 0 through 255</p> <p><b>Default:</b> 150 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Local PIM RPs on page 3695</a> in the <i>Multicast Protocols Configuration Guide</i></li> <li>• <i>Example: Configuring Bidirectional PIM</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                   |

## idle-standby-path-switchover-delay

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>idle-standby-path-switchover-delay &lt;seconds&gt;;</code>                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ]                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2.                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Configure the time interval after which an ECMP join is moved to the standby path in the absence of traffic on the path.</p> <p>In the absence of this statement, ECMP joins are not moved to the standby path until traffic is detected on the path.</p>                                                                                                                |
| <b>Options</b>                  | <code>&lt;seconds&gt;</code> —Time interval after which an ECMP join is moved to the standby RPF path in the absence of traffic on the path.                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring PIM Make-Before-Break Join Load Balancing</i></li><li>• <i>Configuring PIM Join Load Balancing</i></li><li>• <a href="#">clear pim join-distribution on page 4018</a></li><li>• <a href="#">join-load-balance on page 3892</a></li><li>• <a href="#">standby-path-creation-delay on page 3936</a></li></ul> |

## import (Protocols PIM Bootstrap)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp bootstrap</a> (inet   inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp bootstrap</a> (inet   inet6)],</p> <p>[edit protocols <a href="#">pim rp bootstrap</a> (inet   inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp bootstrap</a> (inet   inet6)]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Apply one or more import policies to control incoming PIM bootstrap messages.                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4 on page 3701</a></li> <li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 3702</a></li> <li>• <a href="#">export (Protocols PIM Bootstrap) on page 3875</a></li> </ul>                                                                                                                                                                                  |

## import (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                          |
| <b>Description</b>              | Apply one or more policies to routes being imported into the routing table from PIM. Use the <b>import</b> statement to filter PIM join messages and prevent them from entering the network.                                                                                                                                                            |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Filtering Incoming PIM Join Messages on page 3719</a></li></ul>                                                                                                                                                                                                                                     |

## infinity

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>infinity [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim spt-threshold</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim spt-threshold</a>],</p> <p>[edit protocols <a href="#">pim spt-threshold</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim spt-threshold</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                      |
| <b>Description</b>              | Apply one or more policies to set the SPT threshold to infinity for a source-group address pair. Use the <b>infinity</b> statement to prevent the last-hop routing device from transitioning from the RPT rooted at the RP to an SPT rooted at the source for that source-group address pair.                                                                                                                                  |
| <b>Options</b>                  | <b><i>policy-names</i></b> —Name of one or more policies.                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring the PIM SPT Threshold Policy</i></li> </ul>                                                                                                                                                                                                                                                                                                                   |

## interface (Protocols PIM)

---

```
Syntax interface (Protocols PIM) (all | interface-name) {
 accept-remote-source;
 disable;
 bfd-liveness-detection {
 authentication {
 algorithm algorithm-name;
 key-chain key-chain-name;
 loose-check;
 }
 detection-time {
 threshold milliseconds;
 }
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 version (0 | 1 | automatic);
 }
 bidirectional {
 df-election {
 backoff-period milliseconds;
 offer-period milliseconds;
 robustness-count number;
 }
 }
 family (inet | inet6) {
 bfd-liveness-detection {
 authentication {
 algorithm algorithm-name;
 key-chain key-chain-name;
 loose-check;
 }
 detection-time {
 threshold milliseconds;
 }
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 version (0 | 1 | automatic);
 }
 disable;
 }
 hello-interval seconds;
```



```

mode (bidirectional-sparse | bidirectional-sparse-dense | dense | sparse | sparse-dense);
neighbor-policy [policy-names];
override-interval milliseconds;
priority number;
propagation-delay milliseconds;
reset-tracking-bit;
version version;
}

```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols [pim](#)],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
[pim](#)],  
 [edit protocols [pim](#)],  
 [edit routing-instances *routing-instance-name* protocols [pim](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
 Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Enable PIM on an interface and configure interface-specific properties.

**Options** *interface-name*—Name of the interface. Specify the full interface name, including the  
 physical and logical address components. To configure all interfaces, you can specify  
[all](#).

The remaining statements are explained separately.

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

**Related Documentation** • [PIM on Aggregated Interfaces on page 3684](#)

## join-load-balance

---

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>join-load-balance {<br/>    automatic;<br/>}</pre>                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                              |
| <b>Description</b>              | Enable load balancing of PIM join messages across interfaces and routing devices.                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>automatic</b> —Enables automatic load balancing of PIM join messages. When a new interface or neighbor is introduced into the network, ECMP joins are redistributed with minimal disruption to traffic.                                                                                                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring PIM Make-Before-Break Join Load Balancing</i></li><li>• <i>Configuring PIM Join Load Balancing</i></li><li>• <a href="#">clear pim join-distribution on page 4018</a> in the <i>Junos OS Operational Mode Commands</i></li></ul>                                                        |

---

## join-prune-timeout

---

|                                 |                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | join-prune-timeout <i>seconds</i> ;                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols pim],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],<br>[edit protocols pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                  |
| <b>Description</b>              | Configure the timeout for the join state. If the periodic join refresh message is not received before the timeout expires, the join state is removed.                                                                                                                               |
| <b>Options</b>                  | <b>seconds</b> —Number of seconds to wait for the periodic join message to arrive.<br><b>Range:</b> 210 through 240 seconds<br><b>Default:</b> 210 seconds                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Modifying the Join State Timeout</i></li></ul>                                                                                                                                                                                           |

## key-chain (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>key-chain <i>key-chain-name</i>;</code>                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> family {inet   inet6} bfd-liveness-detection authentication],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> family {inet   inet6} bfd-liveness-detection authentication]       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement modified in Junos OS Release 12.2 to include <b>family</b> in the hierarchy level.                                                                         |
| <b>Description</b>              | Specify the security keychain to use for BFD authentication.                                                                                                                                                                                                                               |
| <b>Options</b>                  | <b><i>key-chain-name</i></b> —Name of the security keychain to use for BFD authentication. The name is a unique integer between <b>0</b> and <b>63</b> . This must match one of the keychains in the <b>authentication-key-chains</b> statement at the [edit security] hierarchy level.    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD Authentication for PIM on page 3729</a></li><li>• <a href="#">Understanding Bidirectional Forwarding Detection Authentication for PIM on page 3726</a></li><li>• <a href="#">authentication on page 3859</a></li></ul> |

## local

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> local {   disable;   address address;   family (inet   inet6) {     disable;     address address;     anycast-pim {       local-address address;       rp-set {         address address &lt;forward-msdp-sa&gt;;       }     }     group-ranges {       destination-ip-prefix &lt;/prefix-length&gt;;     }     hold-time seconds;     override;     priority number;   }   group-ranges {     destination-ip-prefix &lt;/prefix-length&gt;;   }   hold-time seconds;   override;   priority number; } </pre> |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>pim rp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim rp</b>],</p> <p>[edit protocols <b>pim rp</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim rp</b>]</p>                                                                                                                                                                      |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure the routing device's RP properties.                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Configuring Local PIM RPs on page 3695</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                            |

## local-address (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>local-address address;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <b>pim rp local family</b> (inet   inet6) <b>anycast-pim</b>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim rp local family</b> (inet   inet6) <b>anycast-pim</b>],</code><br><code>[edit protocols <b>pim rp local family</b> (inet   inet6) <b>anycast-pim</b>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim rp local family</b> (inet   inet6) <b>anycast-pim</b>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the routing device local address for the anycast rendezvous point (RP). If this statement is omitted, the router ID is used as this address.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <b>address</b> —Anycast RP IPv4 or IPv6 address, depending on <b>family</b> configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring PIM Anycast With or Without MSDP</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## log-interval (PIM Entries)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | log-interval <i>value</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim sglimit],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim sglimit <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim sglimit],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim sglimit <i>family</i>],</p> <p>[edit protocols pim sglimit],</p> <p>[edit protocols pim sglimit <i>family</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim sglimit],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim sglimit <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp group-rp-mapping],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp group-rp-mapping <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp group-rp-mapping],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp group-rp-mapping <i>family</i>],</p> <p>[edit protocols pim rp group-rp-mapping],</p> <p>[edit protocols pim rp group-rp-mapping <i>family</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp group-rp-mapping],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp group-rp-mapping <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp register-limit],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp register-limit <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp register-limit],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp register-limit <i>family</i>],</p> <p>[edit protocols pim rp register-limit],</p> <p>[edit protocols pim rp register-limit <i>family</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp register-limit],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp register-limit <i>family</i>],</p> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Configure the amount of time between log messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><i>seconds</i>—Minimum time interval (in seconds) between log messages. To configure the time interval, you must explicitly configure the maximum number of entries received with the <b>maximum</b> statement. You can apply the log interval to incoming PIM join messages, PIM register messages, and group-to-RP mappings.</p> <p><b>Range:</b> 1 through 65,535</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>add new concept and example topic to related topic list.</li> <li><a href="#">clear pim join on page 4017</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## loose-check

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | loose-check;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication]                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Specify loose authentication checking on the BFD session. Use loose authentication for transitional periods only when authentication might not be configured at both ends of the BFD session.</p> <p>By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure <i>loose checking</i>. When loose checking is configured, packets are accepted without authentication being checked at each end of the session.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD Authentication for PIM on page 3729</a></li><li>• <a href="#">Understanding Bidirectional Forwarding Detection Authentication for PIM on page 3726</a></li><li>• <a href="#">authentication on page 3859</a></li></ul>                                                                                                                                                                                                                                                                                           |




## mapping-agent-election

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (mapping-agent-election   no-mapping-agent-election);                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp auto-rp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp auto-rp</a> ],<br>[edit protocols <a href="#">pim rp auto-rp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp auto-rp</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                          |
| <b>Description</b>              | Configure the routing device mapping announcements as a mapping agent.                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>mapping-agent-election</b>—Mapping agents do not announce mappings when receiving mapping messages from a higher-addressed mapping agent.</p> <p><b>no-mapping-agent-election</b>—Mapping agents always announce mappings and do not perform mapping agent election.</p> <p><b>Default:</b> mapping-agent-election</p>                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring PIM Auto-RP on page 3705</a></li> </ul>                                                                                                                                                                                                                                                                                            |

## maximum (PIM Entries)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>maximum <i>limit</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols pim sglimit],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim sglimit <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim sglimit],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim sglimit <i>family</i>],</p> <p>[edit protocols pim sglimit],</p> <p>[edit protocols pim sglimit <i>family</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim sglimit],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim sglimit <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp group-rp-mapping],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp group-rp-mapping <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp group-rp-mapping],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp group-rp-mapping <i>family</i>],</p> <p>[edit protocols pim rp group-rp-mapping],</p> <p>[edit protocols pim rp group-rp-mapping <i>family</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp group-rp-mapping],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp group-rp-mapping <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp register-limit],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp register-limit <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp register-limit],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp register-limit <i>family</i>],</p> <p>[edit protocols pim rp register-limit],</p> <p>[edit protocols pim rp register-limit <i>family</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp register-limit],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp register-limit <i>family</i>],</p> |
| <b>Release Information</b> | Statement introduced in Junos OS Release 12.2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>         | Configure the maximum number of specified PIM entries received by the device. If the device reaches the configured limit, no new entries are received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                            | <div>  <p><b>NOTE:</b> The maximum limit settings that you configure with the <b>maximum</b> and the <b>family</b> (<b>inet</b>   <b>inet6</b>) <b>maximum</b> statements are mutually exclusive. For example, if you configure a global maximum PIM join state limit, you cannot configure a limit at the family level for IPv4 or IPv6 joins. If you attempt to configure a limit at both the global level and the family level, the device will not accept the configuration.</p> </div>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>             | <p><b>limit</b>—Maximum number of PIM entries received by the device. If you configure both the <b>log-interval</b> and the <b>maximum</b> statements, a warning is triggered when the maximum limit is reached.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

Depending on your configuration, this limit specifies the maximum number of PIM joins, PIM register messages, or group-to-RP mappings received by the device.

**Range:** 1 through 65,535

|                                 |                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• add new concept and example topic to related topic list.</li> <li>• <a href="#">clear pim join on page 4017</a></li> </ul> |

## maximum-rps


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>maximum-rps limit;</code>                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp embedded-rp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp embedded-rp</a> ],<br>[edit protocols <a href="#">pim rp embedded-rp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp embedded-rp</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                      |
| <b>Description</b>              | Limit the number of RPs that the routing device acknowledges.                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <i>limit</i> —Number of RPs.<br><b>Range:</b> 1 through 500<br><b>Default:</b> 100                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring PIM Embedded RP for IPv6 on page 3711</a></li> </ul>                                                                                                                                                                                                                                                                                               |

## minimum-interval (PIM BFD Liveness Detection)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>minimum-interval <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit protocols <code>pim interface interface-name bfd-liveness-detection</code> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <code>pim interface interface-name bfd-liveness-detection</code> ]                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure the minimum interval after which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the <code>transmit-interval</code> <code>minimum-interval</code> and <code>minimum-receive-interval</code> statements. |
| <b>Options</b>                  | <i>milliseconds</i> —Minimum transmit and receive interval.<br><b>Range:</b> 1 through 255,000 milliseconds                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 3727</a></li></ul>                                                                                                                                                                                                                                                                                                                                       |

## minimum-interval (PIM BFD Transmit Interval)

|                                                                                                                                                                                                                                                                                                                    |                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                      | <code>minimum-interval <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                             | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval]                                                                                                                                                      |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                         | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for BFD authentication introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                    |
| <b>Description</b>                                                                                                                                                                                                                                                                                                 | Configure the minimum interval after which the local routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the <b>minimum-interval</b> statement at the [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection] hierarchy level. |
| <b>Options</b>                                                                                                                                                                                                                                                                                                     | <i>milliseconds</i> —Minimum transmit interval value.<br><b>Range:</b> 1 through 255,000                                                                                                                                                                                                                                                                                                            |
| <div>  <p><b>NOTE:</b> The threshold value specified in the <b>threshold</b> statement must be greater than the value specified in the <b>minimum-interval</b> statement for the <b>transmit-interval</b> statement.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                    | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                       | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 3727</a></li> <li>• <a href="#">bfd-liveness-detection on page 3862</a></li> <li>• <a href="#">minimum-interval on page 3902</a></li> <li>• <a href="#">threshold on page 3939</a></li> </ul>                                                                                                                  |

## minimum-receive-interval

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | minimum-receive-interval <i>milliseconds</i> ;                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit protocols <b>pim interface</b> <i>interface-name</i> <b>bfd-liveness-detection</b> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim interface</b> <i>interface-name</i> <b>bfd-liveness-detection</b> ]                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                       |
| <b>Description</b>              | Configure the minimum interval after which the local routing device must receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the <b>minimum-interval</b> statement at the [edit protocols <b>pim interface</b> <i>interface-name</i> <b>bfd-liveness-detection</b> ] hierarchy level. |
| <b>Options</b>                  | <b>milliseconds</b> —Minimum receive interval.<br><b>Range:</b> 1 through 255,000 milliseconds                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 3727</a></li></ul>                                                                                                                                                                                                                                                                                                           |

## mode (Protocols PIM)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | mode (bidirectional-sparse   bidirectional-sparse-dense   dense   sparse   sparse-dense);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>pim interface</b> <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim interface</b> <i>interface-name</i> ],<br>[edit protocols <b>pim interface</b> <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim interface</b> <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br><b>bidirectional-sparse</b> and <b>bidirectional-sparse-dense</b> options introduced in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Configure the PIM mode on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p>The choice of PIM mode is closely tied to controlling how groups are mapped to PIM modes, as follows:</p> <ul style="list-style-type: none"> <li>• <b>bidirectional-sparse</b>—Use if all multicast groups are operating in bidirectional, sparse, or SSM mode.</li> <li>• <b>bidirectional-sparse-dense</b>—Use if multicast groups, except those that are specified in the <b>dense-groups</b> statement, are operating in bidirectional, sparse, or SSM mode.</li> <li>• <b>dense</b>—Use if all multicast groups are operating in dense mode.</li> <li>• <b>sparse</b>—Use if all multicast groups are operating in sparse mode or SSM mode.</li> <li>• <b>sparse-dense</b>—Use if multicast groups, except those that are specified in the <b>dense-groups</b> statement, are operating in sparse mode or SSM mode.</li> </ul> <p><b>Default:</b> Sparse mode</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring PIM Dense Mode Properties on page 3739</a> in the <i>Multicast Protocols Configuration Guide</i></li> <li>• <a href="#">Configuring PIM Sparse-Dense Mode Properties on page 3741</a> in the <i>Multicast Protocols Configuration Guide</i></li> <li>• <i>Example: Configuring Bidirectional PIM</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## multiplier

---

|                                 |                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>multiplier <i>number</i>;</code>                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit protocols <code>pim interface <i>interface-name</i> bfd-liveness-detection</code> ],<br>[edit routing-instances <code><i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection</code> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                             |
| <b>Description</b>              | Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.                                                                                                            |
| <b>Options</b>                  | <b><i>number</i></b> —Number of hello packets.<br><b>Range:</b> 1 through 255<br><b>Default:</b> 3                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 3727</a></li></ul>                                                                                                                                 |

## neighbor-policy

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>neighbor-policy [ <i>policy-names</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit logical-systems <code><i>logical-system-name</i> protocols pim interface <i>interface-name</i></code> ],<br>[edit logical-systems <code><i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i></code> ],<br>[edit protocols <code>pim interface <i>interface-name</i></code> ],<br>[edit routing-instances <code><i>routing-instance-name</i> protocols pim interface <i>interface-name</i></code> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Apply a PIM interface-level policy to filter neighbor IP addresses.                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <b><i>policy-name</i></b> —Name of the policy that filters neighbor IP addresses.                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Interface-Level PIM Neighbor Policies on page 3715</a></li></ul>                                                                                                                                                                                                                                                                                                                                            |



## next-hop (PIM RPF Selection)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>next-hop <i>next-hop-address</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> source <i>source-address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> wildcard-source],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> source <i>source-address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> wildcard-source]</p> |
| <b>Release Information</b>      | <p>Statement introduced in JUNOS Release 10.4.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure the specific next-hop address for the PIM group source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <i>next-hop-address</i> —Specific next-hop address for the PIM group source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>view-level—To view this statement in the configuration.</p> <p>control-level—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring PIM RPF Selection</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## no-adaptation (PIM BFD Liveness Detection)

|                                 |                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>no-adaptation;</code>                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <p>[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]</p>                                                              |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> |
| <b>Description</b>              | Configure BFD sessions not to adapt to changing network conditions. We recommend that you <i>do not</i> disable BFD adaptation unless it is preferable to have BFD adaptation disabled in your network.                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 3727</a></li> <li>• <a href="#">bfd-liveness-detection on page 3862</a></li> </ul>                                                                                                            |

## no-bidirectional-mode

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | no-bidirectional-mode;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <b>pim</b> graceful-restart],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim</b> graceful-restart],<br>[edit protocols <b>pim</b> graceful-restart],<br>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim</b> graceful-restart]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b> | Statement introduced in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>         | <p>Disable forwarding for bidirectional PIM routes during graceful restart recovery, both in cases of a routing protocol process (rpd) restart and graceful Routing Engine switchover.</p> <p>Bidirectional PIM accepts packets for a bidirectional route on multiple interfaces. This means that some topologies might develop multicast routing loops if all PIM neighbors are not synchronized with regard to the identity of the designated forwarder (DF) on each link. If one routing device is forwarding without actively participating in DF elections, particularly after unicast routing changes, multicast routing loops might occur.</p> <p>If graceful restart for PIM is enabled and the forwarding of packets on bidirectional routes is disallowed (by including the <b>no-bidirectional-mode</b> statement in the configuration), PIM behaves conservatively to avoid multicast routing loops during the recovery period. When the routing protocol process (rpd) restarts, all bidirectional routes are deleted. After graceful restart has completed, the routes are re-added, based on the converged unicast and bidirectional PIM state. While graceful restart is active, bidirectional multicast flows drop packets.</p> |
| <b>Default</b>             | If graceful restart for PIM is enabled and the bidirectional PIM is enabled, the default graceful restart behavior is to continue forwarding packets on bidirectional routes. If the gracefully restarting routing device was serving as a DF for some interfaces to rendezvous points, the restarting routing device sends a DF Winner message with a metric of 0 on each of these RP interfaces. This ensures that a neighbor routing device does not become the DF due to unicast topology changes that might occur during the graceful restart period. Sending a DF Winner message with a metric of 0 prevents another PIM neighbor from assuming the DF role until after graceful restart completes. When graceful restart completes, the gracefully restarted routing device sends another DF Winner message with the actual converged unicast metric.                                                                                                                                                                                                                                                                                                                                                                                     |



**NOTE:** Graceful Routing Engine switchover operates independently of the graceful restart behavior. If graceful Routing Engine switchover is configured without graceful restart, all PIM routes for all modes are deleted when the rpd process restarts. If graceful Routing Engine switchover is configured with graceful restart, the behavior is the same as described here, except that the recovery happens on the Routing Engine that assumes mastership.

|                                 |                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM Sparse Mode Graceful Restart</i> in the <i>Multicast Protocols Configuration Guide</i></li> <li>• <i>Understanding Bidirectional PIM</i></li> <li>• <i>Example: Configuring Bidirectional PIM</i></li> </ul> |

## no-dr-flood (PIM Snooping)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-dr-flood;                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit routing-instances <instance-name> protocols <a href="#">pim-snooping traceoptions</a> ],<br>[edit logical-systems <logical-system-name> routing-instances <instance-name> protocols <a href="#">pim-snooping traceoptions</a> ],<br>[edit routing-instances <instance-name> protocols <a href="#">pim-snooping</a> vlan <vlan-id>],<br>[edit logical-systems <logical-system-name> routing-instances <instance-name> protocols <a href="#">pim-snooping</a> vlan <vlan-id>] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for MX Series 3D Universal Edge devices.                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Disable default flooding of multicast data on the PIM designated router port.                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                               |

## offer-period

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>offer-period milliseconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim interface interface-name</a> bidirectional df-election],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim interface interface-name</a> bidirectional df-election],<br>[edit protocols <a href="#">pim interface interface-name</a> bidirectional df-election],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim interface interface-name</a> bidirectional df-election]                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Configure the designated forwarder (DF) election offer period for bidirectional PIM. When a DF election Offer or Winner message fails to be received, the message is retransmitted. The <b>offer-period</b> statement modifies the interval between repeated DF election messages. The <a href="#">robustness-count</a> statement determines the minimum number of DF election messages that must fail to be received for DF election to fail. To prevent routing loops, all routing devices on the link must have a consistent view of the DF. When the DF election fails because DF election messages are not received, forwarding on bidirectional PIM routes is suspended.</p> <p>If a routing device receives from a neighbor a better offer than its own, the routing device stops participating in the election for a period of <a href="#">robustness-count</a> * <b>offer-period</b>. Eventually, all routing devices except the best candidate stop sending Offer messages.</p> |
| <b>Options</b>                  | <p><b>milliseconds</b>—Interval to wait before retransmitting DF Offer and Winner messages.</p> <p><b>Range:</b> 100 through 10,000 milliseconds</p> <p><b>Default:</b> 100</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Understanding Bidirectional PIM</i></li><li>• <i>Example: Configuring Bidirectional PIM</i></li><li>• <a href="#">robustness-count on page 3927</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## override (PIM static RP)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | override;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp local],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp local family inet],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp local family inet6],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp static address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim rp local],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim rp local family inet],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim rp local family inet6],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim rp static address <i>address</i>],</p> <p>[edit protocols pim rp local],</p> <p>[edit protocols pim rp local family inet],</p> <p>[edit protocols pim rp local family inet6],</p> <p>[edit protocols pim rp static address <i>address</i>],</p> <p>[edit routing-instances <i>instance-name</i> protocols pim rp local],</p> <p>[edit routing-instances <i>instance-name</i> protocols pim rp local family inet],</p> <p>[edit routing-instances <i>instance-name</i> protocols pim rp local family inet6],</p> <p>[edit routing-instances <i>instance-name</i> protocols pim rp static address <i>address</i>]</p> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | When you configure both static RP mapping and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for a given group range, and allow dynamic RP mapping for all other groups.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Static RP on page 3694</a></li> <li>• <a href="#">Configuring PIM Auto-RP on page 3705</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## override-interval

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>override-interval milliseconds;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim interface</a> <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim interface</a> <i>interface-name</i> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim interface</a> <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ]<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim interface</a> <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.1.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Set the maximum time in milliseconds to delay sending override join messages for a multicast network that has join suppression enabled. When a router or switch sees a prune message for a join it is currently suppressing, it waits for the interval specified by the override timer before it sends an override join message.                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | This is a random timer with a value in milliseconds.<br><b>Range:</b> 0 through maximum override value<br><b>Default:</b> 2000 milliseconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Enabling Join Suppression</a></li><li>• <a href="#">propagation-delay on page 3922</a></li><li>• <a href="#">reset-tracking-bit on page 3924</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## pim

```

Syntax pim {
 disable;
 assert-timeout seconds;
 dense-groups {
 addresses;
 }
 dr-election-on-p2p;
 export;
 family (inet | inet6) {
 disable;
 }
 graceful-restart {
 disable;
 no-bidirectional-mode;
 restart-duration seconds;
 }
 import [policy-names];
 interface interface-name {
 family (inet | inet6) {
 disable;
 }
 bfd-liveness-detection {
 authentication {
 algorithm algorithm-name;
 key-chain key-chain-name;
 }
 loose-check;
 detection-time {
 threshold milliseconds;
 }
 }
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 multiplier number;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 version (0 | 1 | automatic);
 }
 accept-remote-source;
 disable;
 bidirectional {
 df-election {
 backoff-period milliseconds;
 offer-period milliseconds;
 robustness-count number;
 }
 }
 family (inet | inet6) {
 disable;
 }
 hello-interval seconds;

```

```
mode (bidirectional-sparse | bidirectional-sparse-dense | dense | sparse |
sparse-dense);
neighbor-policy [policy-names];
override-interval milliseconds;
priority number;
propagation-delay milliseconds;
reset-tracking-bit;
version version;
}
join-load-balance;
join-prune-timeout;
mdt {
 data-mdt-reuse;
 group-range multicast-prefix;
 threshold {
 group group-address {
 source source-address {
 rate threshold-rate;
 }
 }
 }
 tunnel-limit limit;
}
}
mvpn {
 autodiscovery {
 inet-mdt;
 }
}
nonstop-routing;
override-interval milliseconds;
propagation-delay milliseconds;
reset-tracking-bit;
rib-group group-name;
rp {
 auto-rp {
 (announce | discovery | mapping);
 (mapping-agent-election | no-mapping-agent-election);
 }
 bidirectional {
 address address {
 group-ranges {
 destination-ip-prefix</prefix-length>;
 }
 hold-time seconds;
 priority number;
 }
 }
 bootstrap {
 family (inet | inet6) {
 export [policy-names];
 import [policy-names];
 priority number;
 }
 }
}
bootstrap-import [policy-names];
bootstrap-export [policy-names];
```



```

bootstrap-priority number;
dr-register-policy [policy-names];
embedded-rp {
 group-ranges {
 destination-ip-prefix</prefix-length>;
 }
 maximum-rps limit;
}
group-rp-mapping {
 family (inet | inet6) {
 log-interval seconds;
 maximum limit;
 threshold value;
 }
}
log-interval seconds;
maximum limit;
threshold value;
}
local {
 family (inet | inet6) {
 address address;
 anycast-pim {
 rp-set {
 address address <forward-msdp-sa>;
 }
 disable;
 local-address address;
 }
 group-ranges {
 destination-ip-prefix</prefix-length>;
 }
 hold-time seconds;
 override;
 priority number;
 }
}
register-limit {
 family (inet | inet6) {
 log-interval seconds;
 maximum limit;
 threshold value;
 }
}
log-interval seconds;
maximum limit;
threshold value;
}
rp-register-policy [policy-names];
spt-threshold {
 infinity [policy-names];
}
static {
 address address {

```

```
 override;
 version version;
 group-ranges {
 destination-ip-prefix </prefix-length>;
 }
 }
}
rpf-selection {
 group group-address {
 source source-address {
 next-hop next-hop-address;
 }
 wildcard-source {
 next-hop next-hop-address;
 }
 }
 prefix-list prefix-list-addresses {
 source source-address {
 next-hop next-hop-address;
 }
 wildcard-source {
 next-hop next-hop-address;
 }
 }
}
sglimit {
 family (inet | inet6) {
 log-interval seconds;
 maximum limit;
 threshold value;
 }
 log-interval seconds;
 maximum limit;
 threshold value;
}
}
traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
}
tunnel-devices [mt-fpc/pic/port];
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],  
[edit protocols],  
[edit routing-instances *routing-instance-name* protocols]

**Release Information** Statement introduced before Junos OS Release 7.4.  
**family** statement introduced in Junos OS Release 9.6.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.

|                                 |                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b>              | Enable PIM on the routing device.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                       |
| <b>Default</b>                  | PIM is disabled on the routing device.                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode</a></li> <li>• <a href="#">Configuring PIM Dense Mode Properties on page 3739</a></li> <li>• <a href="#">Configuring PIM Sparse-Dense Mode Properties on page 3741</a></li> </ul> |

## pim-snooping

|                                 |                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>pim-snooping {   traceoptions{     file;     flag [all   route   normal   general   state   policy   task   timer   packets   hello   join   prune]     [detail   disable   receive   send];   }   no-dr-flood;   vlan&lt;vlan-id&gt;{     no-dr-flood;   } }</pre> |
| <b>Hierarchy Level</b>          | [edit routing-instances <instance-name> protocols],<br>[edit logical-systems <logical-system-name> routing-instances <instance-name> protocols]                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for MX Series 3D Universal Edge devices.                                                                                                                                                                                   |
| <b>Description</b>              | Enable PIM snooping on the device.                                                                                                                                                                                                                                       |
| <b>Default</b>                  | PIM snooping is disabled on the device.                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b>traceoptions</b>—Configure tracing options for PIM snooping.</p> <p><b>no-dr-flood</b>—Disable default flooding of multicast data on the PIM designated router port.</p> <p><b>vlan &lt;vlan-id&gt;</b>—Configure PIM snooping parameters for a VLAN.</p>          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">PIM Snooping for VPLS on page 3745</a></li> </ul>                                                                                                                                                                   |

## prefix-list (PIM RPF Selection)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>prefix-list <i>prefix-list-addresses</i> {<br/>    source <i>source-address</i> {<br/>        next-hop <i>next-hop-address</i>;<br/>    }<br/>    wildcard-source {<br/>        next-hop <i>next-hop-address</i>;<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                        |
| Hierarchy Level          | <pre>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i><br/>    source <i>source-address</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i><br/>    wildcard-source],<br/>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list<br/>    <i>prefix-list-addresses</i> source <i>source-address</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list<br/>    <i>prefix-list-addresses</i> wildcard-source]</pre> |
| Release Information      | Statement introduced in Junos OS Release 10.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Description              | (Optional) Configure a list of prefixes (addresses) for multiple PIM groups.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Options                  | <b><i>prefix-list-addresses</i></b> —List of prefixes (addresses) for multiple PIM groups.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Required Privilege Level | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Related Documentation    | <ul style="list-style-type: none"><li>• <i>Example: Configuring PIM RPF Selection</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## priority (Bootstrap)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <code>pim rp bootstrap</code> (inet   inet6)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>pim rp bootstrap</code> (inet   inet6)],</p> <p>[edit protocols <code>pim rp bootstrap</code> (inet   inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <code>pim rp bootstrap</code> (inet   inet6)]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure the routing device's likelihood to be elected as the bootstrap router.                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b><i>number</i></b>—Routing device's priority for becoming the bootstrap router. A higher value corresponds to a higher priority.</p> <p><b>Range:</b> 0 through a 32-bit number</p> <p><b>Default:</b> 0 (The routing device has the least likelihood of becoming the bootstrap router and sends packets with a priority of 0.)</p>                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4 on page 3701</a></li> <li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 3702</a></li> <li>• <a href="#">bootstrap-priority on page 3868</a></li> </ul>                                                                                                                                                                                    |

## priority (PIM Interfaces)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <code>pim interface <i>interface-name</i></code> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>pim interface <i>interface-name</i></code> ],<br>[edit protocols <code>pim interface <i>interface-name</i></code> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <code>pim interface <i>interface-name</i></code> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure the routing device's likelihood to be elected as the designated router.                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <b><i>number</i></b> —Routing device's priority for becoming the designated router. A higher value corresponds to a higher priority.<br><b>Range:</b> 0 through a 32-bit number<br><b>Default:</b> 0 (The routing device has the least likelihood of becoming the designated router.)                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Interface Priority for PIM Designated Router Selection on page 3692</a></li></ul>                                                                                                                                                                                                                                                                                                                           |

## priority (PIM RPs)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>priority <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols <b>pim rp local family</b> (inet   inet6)],</p> <p>[edit routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim rp local family</b> (inet   inet6)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional RP addresses introduced in Junos OS Release 12.1.</p>                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>For PIM-SM, configure this routing device's priority for becoming an RP.</p> <p>For bidirectional PIM, configure this RP address' priority for becoming an RP.</p> <p>The bootstrap router uses this field when selecting the list of candidate rendezvous points to send in the bootstrap message. A smaller number increases the likelihood that the routing device or RP address becomes the RP. A priority value of 0 means that bootstrap router can override the group range being advertised by the candidate RP.</p>                                                                                                             |
| <b>Options</b>                  | <p><b><i>number</i></b>—Priority for becoming an RP. A lower value corresponds to a higher priority.</p> <p><b>Range:</b> 0 through 255</p> <p><b>Default:</b> 1</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Local PIM RPs on page 3695</a> in the <i>Multicast Protocols Configuration Guide</i></li> <li>• <i>Example: Configuring Bidirectional PIM</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                   |

## propagation-delay

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>propagation-delay <i>milliseconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | <code>[edit protocols <a href="#">pim</a>],</code><br><code>[edit protocols <a href="#">pim interface</a> <i>interface-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim interface</a> <i>interface-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim interface</a> <i>interface-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code><a href="#">pim interface</a> <i>interface-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.1.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Set a delay for implementing a PIM prune message on the upstream routing device on a multicast network for which join suppression has been enabled. The routing device waits for the prune pending period to detect whether a join message is currently being suppressed by another routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <b><i>milliseconds</i></b> —Interval for the prune pending timer, which is the sum of the <b>propagation-delay</b> value and the <b>override-interval</b> value.<br><b>Range:</b> 250 through 2000 milliseconds<br><b>Default:</b> 500 milliseconds                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Enabling Join Suppression</i></li><li>• <a href="#">override-interval on page 3912</a></li><li>• <a href="#">reset-tracking-bit on page 3924</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |



## register-limit

|                            |                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> register-limit {     family (inet   inet6) {         log-interval <i>seconds</i>;         maximum <i>limit</i>;         threshold <i>value</i>;     }     log-interval <i>seconds</i>;     maximum <i>limit</i>;     threshold <i>value</i>; } </pre>                                                                                                                        |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a>],</p> <p>[edit protocols <a href="#">pim rp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a>]</p> |
| <b>Release Information</b> | Statement introduced in Junos OS Release 12.2.                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>         | Configure a limit for the number of incoming (S,G) PIM registers.                                                                                                                                                                                                                                                                                                                  |



**NOTE:** The maximum limit settings that you configure with the **maximum** and the **family (inet | inet6) maximum** statements are mutually exclusive. For example, if you configure a global maximum PIM register message limit, you cannot configure a limit at the family level for IPv4 or IPv6. If you attempt to configure a limit at both the global level and the family level, the device will not accept the configuration.

|                                 |                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <p><b>family (inet   inet6)</b>—(Optional) Specify either IPv4 or IPv6 messages to be counted towards the configured register message limit.</p> <p><b>Default:</b> Both IPv4 and IPv6 messages are counted towards the configured register message limit.</p> <p>The remaining statements are described separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring PIM State Limits on page 5198</a></li> <li>• <a href="#">clear pim join on page 4017</a></li> <li>• <a href="#">clear pim register on page 4020</a></li> </ul>                                                                                 |

## reset-tracking-bit

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | reset-tracking-bit;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim interface</a> <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim interface</a> <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim interface</a> <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim interface</a> <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.1.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Change the value of a tracking bit (T-bit) field in the LAN prune delay hello option from the default of 1 to 0, which enables join suppression for a multicast interface. When the network starts receiving multiple identical join messages, join suppression triggers a random timer with a value of 66 through 84 milliseconds ( $1.1 \times$ periodic through $1.4 \times$ periodic, where periodic is 60 seconds). This creates an interval during which no identical join messages are sent. Eventually, only one of the identical messages is sent. Join suppression is triggered each time identical messages are sent for the same join.                                                              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Enabling Join Suppression</i></li><li>• <a href="#">override-interval on page 3912</a></li><li>• <a href="#">propagation-delay on page 3922</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## restart-duration (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>restart-duration <i>seconds</i>;</code>                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim graceful-restart</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim graceful-restart</a> ],<br>[edit protocols <a href="#">pim graceful-restart</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim graceful-restart</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure the duration of the graceful restart interval.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <b><i>seconds</i></b> —Time that the routing device waits (in seconds) to complete PIM sparse mode graceful restart.<br><b>Range:</b> 30 through 300<br><b>Default:</b> 60                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring PIM Sparse Mode Graceful Restart</i></li> </ul>                                                                                                                                                                                                                                                                                                                       |

## rib-group (Protocols PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>rib-group {<br/>    inet <i>group-name</i>;<br/>    inet6 <i>group-name</i>;<br/>}</pre>                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                          |
| <b>Description</b>              | Associate a routing table group with PIM.                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <b><i>table-name</i></b> —Name of the routing table. The name must be one that you defined with the <b>rib-groups</b> statement at the <b>[edit routing-options]</b> hierarchy level.                                                                                                                                                                   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring a Dedicated PIM RPF Routing Table</i></li></ul>                                                                                                                                                                                                                                         |

## robustness-count

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>robustness-count <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <code>pim interface interface-name</code> bidirectional df-election],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <code>pim interface interface-name</code> bidirectional df-election],</p> <p>[edit protocols <code>pim interface interface-name</code> bidirectional df-election],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <code>pim interface interface-name</code> bidirectional df-election]</p>                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Configure the designated forwarder (DF) election robustness count for bidirectional PIM. When a DF election Offer or Winner message fails to be received, the message is retransmitted. The <b>robustness-count</b> statement sets the minimum number of DF election messages that must fail to be received for DF election to fail. To prevent routing loops, all routers on the link must have a consistent view of the DF. When the DF election fails because DF election messages are not received, forwarding on bidirectional PIM routes is suspended.</p> <p>If a router receives from a neighbor a better offer than its own, the router stops participating in the election for a period of <b>robustness-count</b> * <code>offer-period</code>. Eventually, all routers except the best candidate stop sending Offer messages.</p> |
| <b>Options</b>                  | <p><i>number</i>—Number of transmission attempts for DF election messages.</p> <p><b>Range:</b> 1 through 10</p> <p><b>Default:</b> 3</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Understanding Bidirectional PIM</i></li> <li><i>Example: Configuring Bidirectional PIM</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## rp

```

Syntax rp {
 auto-rp {
 (announce | discovery | mapping);
 (mapping-agent-election | no-mapping-agent-election);
 }
 bidirectional {
 address address {
 group-ranges {
 destination-ip-prefix </prefix-length>;
 }
 hold-time seconds;
 priority number;
 }
 }
 bootstrap {
 family (inet | inet6) {
 export [policy-names];
 import [policy-names];
 priority number;
 }
 }
 bootstrap-export [policy-names];
 bootstrap-import [policy-names];
 bootstrap-priority number;
 dr-register-policy [policy-names];
 embedded-rp {
 group-ranges {
 destination-ip-prefix </prefix-length>;
 }
 maximum-rps limit;
 }
 group-rp-mapping {
 family (inet | inet6) {
 log-interval seconds;
 maximum limit;
 threshold value;
 }
 }
 log-interval seconds;
 maximum limit;
 threshold value;
}
local {
 family (inet | inet6) {
 disable;
 address address;
 anycast-pim {
 local-address address;
 address address <forward-msdp-sa>;
 rp-set {
 }
 }
 }
}

```

```

 }
 group-ranges {
 destination-ip-prefix</prefix-length>;
 }
 hold-time seconds;
 override;
 priority number;
}
}
register-limit {
 family (inet | inet6) {
 log-interval seconds;
 maximum limit;
 threshold value;
 }
}
log-interval seconds;
maximum limit;
threshold value;
}
}
register-probe-time register-probe-time;
}
rp-register-policy [policy-names];
static {
 address address {
 override;
 version version;
 group-ranges {
 destination-ip-prefix</prefix-length>;
 }
 }
}
}
}

```

|                                 |                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a>],</p> <p>[edit protocols <a href="#">pim</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                          |
| <b>Description</b>              | <p>Configure the routing device as an actual or potential RP. A routing device can be an RP for more than one group.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                     |
| <b>Default</b>                  | If you do not include the <b>rp</b> statement, the routing device can never become the RP.                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                         |

**Related Documentation**    • [Understanding PIM Sparse Mode](#)

---

## rp-register-policy

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | rp-register-policy [ <i>policy-names</i> ];                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ],<br>[edit protocols <a href="#">pim rp</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                          |
| <b>Description</b>              | Apply one or more policies to control incoming PIM register messages.                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | • <a href="#">Configuring Register Message Filters on a PIM RP and DR on page 3723</a><br>• <a href="#">dr-register-policy on page 3873</a>                                                                                                                                                                                                                         |



## rp-set


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>rp-set {   address address &lt;forward-msdp-sa&gt;; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>pim local family</b> (inet   inet6) <b>anycast-pim</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim local family</b> (inet   inet6) <b>anycast-pim</b>],</p> <p>[edit protocols <b>pim local family</b> (inet   inet6) <b>anycast-pim</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim local family</b> (inet   inet6) <b>anycast-pim</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Configure a set of rendezvous point (RP) addresses for anycast RP. You can configure up to 15 RPs.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring PIM Anycast With or Without MSDP</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                               |

## rpf-selection

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>rpf-selection {<br/>  group group-address {<br/>    source source-address {<br/>      next-hop next-hop-address;<br/>    }<br/>    wildcard-source {<br/>      next-hop next-hop-address;<br/>    }<br/>  }<br/>  prefix-list prefix-list-addresses {<br/>    source source-address {<br/>      next-hop next-hop-address;<br/>    }<br/>    wildcard-source {<br/>      next-hop next-hop-address;<br/>    }<br/>  }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit routing-instances <i>routing-instance-name</i> protocols pim]                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in JUNOS Release 10.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Configure the PIM RPF next-hop neighbor for a specific group and source for a VRF routing instance.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                          |
| <b>Default</b>                  | If you omit the <b>rpf-selection</b> statement, PIM RPF checks typically choose the best path determined by the unicast protocol for all multicast flows.                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <b>source-address</b> —Specific source address for the PIM group.                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring PIM RPF Selection</i></li></ul>                                                                                                                                                                                                                                                                                                                                        |

## sglimit

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>sglimit {     family (inet   inet6) {         log-interval seconds;         maximum limit;         threshold value;     }     log-interval seconds;     maximum limit;     threshold value; }</pre>                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a>],</p> <p>[edit protocols <a href="#">pim</a> ],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a>]</p>                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Configure a limit for the number of accepted (*G) and (S,G) PIM join states.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|                                 | <div>  <p><b>NOTE:</b> The maximum limit settings that you configure with the <code>maximum</code> and the <code>family (inet   inet6) maximum</code> statements are mutually exclusive. For example, if you configure a global maximum PIM join state limit, you cannot configure a limit at the family level for IPv4 or IPv6 joins. If you attempt to configure a limit at both the global level and the family level, the device will not accept the configuration.</p> </div> |
| <b>Options</b>                  | <p><b>family (inet   inet6)</b>—(Optional) Specify either IPv4 or IPv6 join states to be counted towards the configured join state limit.</p> <p><b>Default:</b> Both IPv4 and IPv6 join states are counted towards the configured join state limit.</p> <p>The remaining statements are described separately.</p>                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring PIM State Limits on page 5198</a></li> <li>• <a href="#">clear pim join on page 4017</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                        |

## source (PIM RPF Selection)

---

|                                 |                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source source-address {<br/>    next-hop next-hop-address;<br/>}</code>                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> ] |
| <b>Release Information</b>      | Statement introduced in JUNOS Release 10.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                              |
| <b>Description</b>              | Configure the source address for the PIM group.                                                                                                                                                                                               |
| <b>Options</b>                  | <i>source-address</i> —Specific source address for the PIM group.<br><br>The remaining statements are explained separately.                                                                                                                   |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring PIM RPF Selection</i></li></ul>                                                                                                                                               |

## spt-threshold

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | spt-threshold {<br>infinity [ <i>policy-names</i> ];<br>}                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>pim</b> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>pim</b> ],<br>[edit protocols <b>pim</b> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <b>pim</b> ]                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Set the SPT threshold to infinity for a source-group address pair. Last-hop multicast routing devices running PIM sparse mode can forward the same stream of multicast packets onto the same LAN through an RPT rooted at the RP or an SPT rooted at the source. By default, last-hop routing devices transition to a direct SPT to the source. You can configure this routing device to set the SPT transition value to infinity to prevent this transition for any source-group address pair.<br><br>The remaining statements are explained separately. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring the PIM SPT Threshold Policy</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## standby-path-creation-delay

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>standby-path-creation-delay &lt;seconds&gt;;</code>                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ]                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2.                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Configure the time interval after which a standby path is created, when a new ECMP interface or neighbor is added to the network.</p> <p>In the absence of this statement, ECMP joins are redistributed as soon as a new ECMP interface or neighbor is added to the network.</p>                                                                                                |
| <b>Options</b>                  | <code>&lt;seconds&gt;</code> —Time interval after which a standby path is created, when a new ECMP interface or neighbor is added to the network. Range is from 1 through 300.                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring PIM Make-Before-Break Join Load Balancing</i></li><li>• <i>Configuring PIM Join Load Balancing</i></li><li>• <a href="#">clear pim join-distribution on page 4018</a></li><li>• <a href="#">join-load-balance on page 3892</a></li><li>• <a href="#">idle-standby-path-switchover-delay on page 3886</a></li></ul> |

## static (Protocols PIM)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>static {   address address {     group-ranges {       destination-ip-prefix&lt;/prefix-length&gt;;     }     override;     version version;   } }</pre>                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a>],</p> <p>[edit protocols <a href="#">pim rp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp</a>]</p>                                                                                                                                      |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Configure static RP addresses. The default static RP address is 224.0.0.0/4. To configure other addresses, include one or more <b>address</b> statements. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.</p> <p>For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Static PIM RP Address on the Non-RP Routing Device on page 3699</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                     |

## threshold (PIM BFD Detection Time)

---

|                            |                                                                                                                                                                                                                                                                  |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>threshold <i>milliseconds</i>;</code>                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>     | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection detection-time],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection detection-time]                         |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Support for BFD authentication introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>         | Specify the threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.                                            |



.....

**NOTE:** The threshold value must be equal to or greater than the transmit interval.


The threshold time must be equal to or greater than the value specified in the [minimum-interval](#) or the [minimum-receive-interval](#) statement.

.....

|                                 |                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <i>milliseconds</i> —Value for the detection time adaptation threshold.<br><b>Range:</b> 1 through 255,000                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 3727</a></li><li>• <a href="#">bfd-liveness-detection on page 3862</a></li><li>• <a href="#">detection-time on page 3870</a></li><li>• <a href="#">minimum-interval on page 3902</a></li><li>• <a href="#">minimum-receive-interval on page 3904</a></li></ul> |



## threshold (PIM BFD Transmit Interval)

|                                                                                                                                                                                                                                                                                                                                    |                                                                                                                                                                                                                                                |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                      | <code>threshold <i>milliseconds</i>;</code>                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                             | [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval] |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                         | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                     |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                 | Specify the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.                                    |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                     | <i>milliseconds</i> —Value for the transmit interval adaptation threshold.<br><b>Range:</b> 0 through 4,294,967,295 ( $2^{32} - 1$ )                                                                                                           |
| <div>  <p><b>NOTE:</b> The threshold value specified in the <code>threshold</code> statement must be greater than the value specified in the <code>minimum-interval</code> statement for the <code>transmit-interval</code> statement.</p> </div> |                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                    | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                            |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                       | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 3727</a></li> <li>• <a href="#">bfd-liveness-detection on page 3862</a></li> </ul>                                                                        |

## threshold (PIM Entries)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>threshold value;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols pim sglimit],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim sglimit <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim sglimit],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim sglimit <i>family</i>],</p> <p>[edit protocols pim sglimit],</p> <p>[edit protocols pim sglimit <i>family</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim sglimit],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim sglimit <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp group-rp-mapping],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp group-rp-mapping <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp group-rp-mapping],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp group-rp-mapping <i>family</i>],</p> <p>[edit protocols pim rp group-rp-mapping],</p> <p>[edit protocols pim rp group-rp-mapping <i>family</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp group-rp-mapping],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp group-rp-mapping <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp register-limit],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim rp register-limit <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp register-limit],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp register-limit <i>family</i>],</p> <p>[edit protocols pim rp register-limit],</p> <p>[edit protocols pim rp register-limit <i>family</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp register-limit],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp register-limit <i>family</i>],</p> |
| <b>Release Information</b> | Statement introduced in Junos OS Release 12.2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>         | Configure a threshold at which a warning message is logged when a certain number of PIM entries have been received by the device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>             | <p><i>value</i>—Threshold at which a warning message is logged. This is a percentage of the maximum number of entries accepted by the device as defined with the <b>maximum</b> statement. You can apply this threshold to incoming PIM join messages, PIM register messages, and group-to-RP mappings.</p> <p>For example, if you configure a maximum number of 1,000 incoming group-to-RP mappings, and you configure a threshold value of 90 percent, warning messages are logged in the system log when the device receives 900 group-to-RP mappings. The same formula applies to incoming PIM join messages and PIM register messages if configured with both the <b>maximum limit</b> and the <b>threshold value</b> statements.</p> <p><b>Default:</b> 1 through 100</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

|                              |                                                                                                                                                                  |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege</b>    | routing—To view this statement in the configuration.                                                                                                             |
| <b>Level</b>                 | routing-control—To add this statement to the configuration.                                                                                                      |
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• add new concept and example topic to related topic list.</li><li>• <a href="#">clear pim join on page 4017</a></li></ul> |

## traceoptions (Protocols PIM)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;<br/>    flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ],<br>[edit protocols <a href="#">pim</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>         | <p>Configure PIM tracing options.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>             | The default PIM trace options are those inherited from the routing protocol's <b>traceoptions</b> statement included at the <b>[edit routing-options]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place tracing output in the <b>pim-log</b> file.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also include the <b>size</b> statement to specify the maximum file size.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 2 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p><b>PIM Tracing Flags</b></p> <ul style="list-style-type: none"><li>• <b>assert</b>—Assert messages</li><li>• <b>bidirectional-df-election</b>—Bidirectional PIM designated-forwarder (DF) election events</li></ul> |

- **bootstrap**—Bootstrap messages
- **cache**—Packets in the PIM sparse mode routing cache
- **graft**—Graft and graft acknowledgment messages
- **hello**—Hello packets
- **join**—Join messages
- **mt**—Multicast tunnel messages
- **nsr-synchronization**—Nonstop active routing (NSR) synchronization messages
- **packets**—All PIM packets
- **prune**—Prune messages
- **register**—Register and register stop messages
- **rp**—Candidate RP advertisements
- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

**no-stamp**—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

**no-world-readable**—(Optional) Do not allow users to read the log file.

**replace**—(Optional) Replace an existing trace file if there is one.

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 0 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |                                                                               |
|---------------------------------|-------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | routing and trace—To view this statement in the configuration.                |
|                                 | routing-control and trace-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | • <a href="#">Configuring PIM Trace Options on page 3685</a>                  |
|                                 | • <i>Tracing DVMRP Protocol Traffic</i>                                       |
|                                 | • <i>Tracing MSDP Protocol Traffic</i>                                        |
|                                 | • <a href="#">Configuring PIM Trace Options on page 3685</a>                  |

## traceoptions (PIM Snooping)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> traceoptions {     file;     flag [all   general   hello   join   normal   packets   policy   prune   route   state   task   timer     ]     [detail   disable   receive   send]; } </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>     | <p>[edit routing-instances &lt;instance-name&gt; protocols <a href="#">pim-snooping</a>],<br/> [edit logical-systems &lt;logical-system-name&gt; routing-instances &lt;instance-name&gt; protocols<br/> <a href="#">pim-snooping</a>]</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b> | Statement introduced in Junos OS Release 12.3 for MX Series 3D Universal Edge devices.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>         | Define tracing operations for PIM snooping.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Default</b>             | <p>The <b>traceoptions</b> feature is disabled by default.</p> <p>The default PIM trace options are those inherited from the routing protocol's <b>traceoptions</b> statement included at the <b>[edit routing-options]</b> hierarchy level.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>             | <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>.</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p><b>PIM Snooping Tracing Flags:</b></p> <ul style="list-style-type: none"> <li>• <b>all</b>—All tracing operations.</li> <li>• <b>general</b>—Trace general PIM snooping events.</li> <li>• <b>hello</b>—Trace hello packets.</li> <li>• <b>join</b>—Trace join messages.</li> <li>• <b>normal</b>—Trace normal PIM snooping events. If you do not specify this flag, only unusual or abnormal operations are traced.</li> <li>• <b>packets</b>—Trace all PIM packets.</li> <li>• <b>policy</b>—Trace policy processing.</li> <li>• <b>prune</b>—Trace prune messages.</li> <li>• <b>route</b>—Trace routing information.</li> <li>• <b>state</b>—Trace PIM state transitions.</li> <li>• <b>task</b>—Trace PIM protocol task processing.</li> <li>• <b>timer</b>—Trace PIM protocol timer processing.</li> </ul> |

***flag-modifier***—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers per flag:

- **detail**—Provide detailed trace information
- **disable**—Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

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

|                              |                                                                                                      |
|------------------------------|------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">PIM Snooping for VPLS on page 3745</a></li></ul> |
|------------------------------|------------------------------------------------------------------------------------------------------|



## transmit-interval (PIM BFD Liveness Detection)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>transmit-interval {     minimum-interval milliseconds;     threshold milliseconds; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | <pre>[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>   bfd-liveness-detection]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Specify the transmit interval for the <b>bfd-liveness-detection</b> statement. The negotiated transmit interval for a peer is the interval between the sending of BFD packets to peers. The receive interval for a peer is the minimum interval between receiving packets sent from its peer; the receive interval is not negotiated between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 3727</a></li> <li>• <a href="#">bfd-liveness-detection on page 3862</a></li> <li>• <a href="#">threshold on page 3939</a></li> <li>• <a href="#">minimum-interval on page 3903</a></li> <li>• <a href="#">minimum-receive-interval on page 3904</a></li> </ul>                                                                                                                                                                                                                                                                                              |

## tunnel-devices

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>tunnel-devices [ <i>mt-fpc/pic/port</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim],<br>[edit routing-instances <i>instance-name</i> protocols pim]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2.<br>Statement introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>List one or more tunnel-capable PICs to be used for creating multicast tunnel (<b>mt</b>) interfaces. Creating a PIC list enables you to control the load-balancing implementation.</p> <p>Tunnel-capable PICs include:</p> <ul style="list-style-type: none"><li>• Adaptive Services PIC</li><li>• Multiservices PIC or Multiservices DPC</li><li>• Tunnel Services PIC</li><li>• On MX Series routers, a PIC created with the <b>tunnel-services</b> statement at the [edit chassis fpc <i>slot-number</i> pic <i>number</i>] hierarchy level.</li></ul> <p>The physical position of the PIC in the routing device determines the multicast tunnel interface name. For example, if you have an Adaptive Services PIC installed in FPC slot 0 and PIC slot 0, the corresponding multicast tunnel interface name is <b>mt-0/0/0</b>. The same is true for Tunnel Services PICs, Multiservices PICs, and Multiservices DPCs.</p> |
| <b>Default</b>                  | Multicast tunnel interfaces are created on all available tunnel-capable PICs, based on a round-robin algorithm.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <b>mt-fpc/pic/port</b> —Interface that is automatically generated when a tunnel-capable PIC is installed in the routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Load Balancing Multicast Tunnel Interfaces Among Available PICs</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## version (BFD)

---

|                                 |                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | version (0   1   automatic);                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">piminterface</a> <i>interface-name</i> <a href="#">bfd-liveness-detection</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim interface</a> <i>interface-name</i> <a href="#">bfd-liveness-detection</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.1.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                  |
| <b>Description</b>              | Specify the bidirectional forwarding detection (BFD) protocol version that you want to detect.                                                                                                                                                                              |
| <b>Options</b>                  | Configure the BFD version to detect: <b>1</b> (BFD version 1) or <b>automatic</b> (autodetect the BFD version)<br><b>Default:</b> automatic                                                                                                                                 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 3727</a></li> </ul>                                                                                                                                                                    |

## version (PIM)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>version version;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">pim rp static address address</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br><a href="#">pim interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br><a href="#">pim rp static address address</a> ],<br>[edit protocols <a href="#">pim interface interface-name</a> ],<br>[edit protocols <a href="#">pim rp static address address</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">pim rp static address address</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Specify the version of PIM.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <b>version</b> —PIM version number.<br><b>Range:</b> 1 or 2<br><b>Default:</b> PIMv1 for rendezvous point (RP) mode (at the [edit protocols <a href="#">pim rp static address address</a> ] hierarchy level). PIMv2 for interface mode (at the [edit protocols <a href="#">pim interface interface-name</a> ] hierarchy level).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling PIM Sparse Mode</a></li><li>• <a href="#">Configuring PIM Dense Mode Properties on page 3739</a></li><li>• <a href="#">Configuring PIM Sparse-Dense Mode Properties on page 3741</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## vlan (PIM Snooping)

---

|                                 |                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>vlan &lt;vlan-id&gt;{<br/>no-dr-flood;<br/>}</code>                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit routing-instances <instance-name> protocols <a href="#">pim-snooping</a> ],<br>[edit logical-systems <logical-system-name> routing-instances <instance-name> protocols<br><a href="#">pim-snooping</a> ] |
| <b>Description</b>              | Configure PIM snooping parameters for a VLAN.                                                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">vlan</a></li> </ul>                                                                                                                                       |

## vpn-group-address

---

|                                 |                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>vpn-group-address address;</code>                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>pim],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                         |
| <b>Description</b>              | Configure the group address for the Layer 3 VPN in the service provider's network.                                                                                                        |
| <b>Options</b>                  | <b>address</b> —Address for the Layer 3 VPN in the service provider's network.                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Multicast Layer 3 VPNs</a></li> <li>• <a href="#">Multicast Protocols Configuration Guide</a></li> </ul>                 |

## wildcard-source (PIM RPF Selection)

---

|                                 |                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | wildcard-source {<br>next-hop next-hop-address;<br>}                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.4.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                           |
| <b>Description</b>              | Use a wildcard for the multicast source instead of (or in addition to) a specific multicast source.<br><br>The remaining statements are explained separately.                                                                                 |
| <b>Required Privilege Level</b> | view-level—To view this statement in the configuration.<br>control-level—To add this statement to the configuration.                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring PIM RPF Selection</i></li></ul>                                                                                                                                               |

## Administration

---

- [Operational Commands: IGMP on page 3952](#)
- [Operational Commands: IGMP Snooping on page 3964](#)
- [Operational Commands: MLD on page 3976](#)
- [Operational Commands: MSDP on page 3991](#)
- [Operational Commands: PIM on page 4016](#)

## Operational Commands: IGMP

## clear igmp statistics

|                                    |                                                                                                                                                                                                                                                                                                                                          |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | clear igmp statistics<br><interface <i>interface-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switches)</b> | clear igmp statistics<br><interface <i>interface-name</i> >                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                 |
| <b>Description</b>                 | Clear Internet Group Management Protocol (IGMP) statistics.                                                                                                                                                                                                                                                                              |
| <b>Options</b>                     | <p><b>none</b>—Clear IGMP statistics on all interfaces.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear IGMP statistics for the specified interface only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | clear                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"> <li><a href="#">show igmp statistics</a></li> </ul>                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>       | <a href="#">clear igmp statistics on page 3953</a>                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>               | See <a href="#">show igmp statistics</a> for an explanation of output fields.                                                                                                                                                                                                                                                            |

## Sample Output

### clear igmp statistics

The following sample output displays IGMP statistics information before and after the **clear igmp statistics** command is entered:

```

user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type Received Sent Rx errors
Membership Query 8883 459 0
V1 Membership Report 0 0 0
DVMRP 19784 35476 0
PIM V1 18310 0 0
Cisco Trace 0 0 0
V2 Membership Report 0 0 0
Group Leave 0 0 0
Mtrace Response 0 0 0
Mtrace Request 0 0 0
Domain Wide Report 0 0 0
V3 Membership Report 0 0 0

```

|                                     |   |
|-------------------------------------|---|
| Other Unknown types                 | 0 |
| IGMP v3 unsupported type            | 0 |
| IGMP v3 source required for SSM     | 0 |
| IGMP v3 mode not applicable for SSM | 0 |

|                        |      |
|------------------------|------|
| IGMP Global Statistics |      |
| Bad Length             | 0    |
| Bad Checksum           | 0    |
| Bad Receive If         | 0    |
| Rx non-local           | 1227 |

user@host> clear igmp statistics

user@host> show igmp statistics

IGMP packet statistics for all interfaces

| IGMP Message type                   | Received | Sent | Rx errors |
|-------------------------------------|----------|------|-----------|
| Membership Query                    | 0        | 0    | 0         |
| V1 Membership Report                | 0        | 0    | 0         |
| DVMRP                               | 0        | 0    | 0         |
| PIM V1                              | 0        | 0    | 0         |
| Cisco Trace                         | 0        | 0    | 0         |
| V2 Membership Report                | 0        | 0    | 0         |
| Group Leave                         | 0        | 0    | 0         |
| Mtrace Response                     | 0        | 0    | 0         |
| Mtrace Request                      | 0        | 0    | 0         |
| Domain Wide Report                  | 0        | 0    | 0         |
| V3 Membership Report                | 0        | 0    | 0         |
| Other Unknown types                 |          |      | 0         |
| IGMP v3 unsupported type            |          |      | 0         |
| IGMP v3 source required for SSM     |          |      | 0         |
| IGMP v3 mode not applicable for SSM |          |      | 0         |
| IGMP Global Statistics              |          |      |           |
| Bad Length                          | 0        |      |           |
| Bad Checksum                        | 0        |      |           |
| Bad Receive If                      | 0        |      |           |
| Rx non-local                        | 0        |      |           |



## show igmp group

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                       | show igmp group<br><brief   detail><br><group-name><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | show igmp group<br><brief   detail><br><group-name>                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                         |
| <b>Description</b>                                  | Display Internet Group Management Protocol (IGMP) group membership information.                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about membership for all IGMP groups.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>group-name</b>—(Optional) Display group membership for the specified IP address only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <i>clear igmp membership</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>                        | <a href="#">show igmp group (Include Mode) on page 3956</a><br><a href="#">show igmp group (Exclude Mode) on page 3957</a><br><a href="#">show igmp group brief on page 3957</a><br><a href="#">show igmp group detail on page 3957</a>                                                                                                                                                                                          |
| <b>Output Fields</b>                                | Table 279 on page 3955 describes the output fields for the <b>show igmp group</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                               |

**Table 279: show igmp group Output Fields**

| Field Name        | Field Description                                                                                                                                       | Level of Output |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Interface</b>  | Name of the interface that received the IGMP membership report. A name of <b>local</b> indicates that the local routing device joined the group itself. | All levels      |
| <b>Group</b>      | Group address.                                                                                                                                          | All levels      |
| <b>Group Mode</b> | Mode the SSM group is operating in: <b>Include</b> or <b>Exclude</b> .                                                                                  | All levels      |
| <b>Source</b>     | Source address.                                                                                                                                         | All levels      |

Table 279: show igmp group Output Fields (*continued*)

| Field Name       | Field Description                                                                                                                                                                                                         | Level of Output |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Source timeout   | Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.      | detail          |
| Last reported by | Address of the host that last reported membership in this group.                                                                                                                                                          | All levels      |
| Timeout          | Time remaining until the group membership is removed.                                                                                                                                                                     | brief none      |
| Group timeout    | Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer. | detail          |
| Type             | Type of group membership: <ul style="list-style-type: none"> <li>• <b>Dynamic</b>—Host reported the membership.</li> <li>• <b>Static</b>—Membership is configured.</li> </ul>                                             | All levels      |

## Sample Output

### show igmp group (Include Mode)

```

user@host> show igmp group
Interface: t1-0/1/0.0
 Group: 232.1.1.1
 Group mode: Include
 Source: 10.0.0.2
 Last reported by: 10.9.5.2
 Timeout: 24 Type: Dynamic
 Group: 232.1.1.1
 Group mode: Include
 Source: 10.0.0.3
 Last reported by: 10.9.5.2
 Timeout: 24 Type: Dynamic
 Group: 232.1.1.1
 Group mode: Include
 Source: 10.0.0.4
 Last reported by: 10.9.5.2
 Timeout: 24 Type: Dynamic
 Group: 232.1.1.2
 Group mode: Include
 Source: 10.0.0.4
 Last reported by: 10.9.5.2
 Timeout: 24 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
 Group: 224.0.0.2
 Source: 0.0.0.0
 Last reported by: Local
 Timeout: 0 Type: Dynamic
 Group: 224.0.0.22
 Source: 0.0.0.0

```

```

Last reported by: Local
Timeout: 0 Type: Dynamic

```

### show igmp group (Exclude Mode)

```

user@host> show igmp group
Interface: t1-0/1/0.0
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
 Group: 224.0.0.2
 Source: 0.0.0.0
 Last reported by: Local
 Timeout: 0 Type: Dynamic
 Group: 224.0.0.22
 Source: 0.0.0.0
 Last reported by: Local
 Timeout: 0 Type: Dynamic

```

### show igmp group brief

The output for the **show igmp group brief** command is identical to that for the **show igmp group** command.

### show igmp group detail

```

user@host> show igmp group detail
Interface: t1-0/1/0.0
 Group: 232.1.1.1
 Group mode: Include
 Source: 10.0.0.2
 Source timeout: 12
 Last reported by: 10.9.5.2
 Group timeout: 0 Type: Dynamic
 Group: 232.1.1.1
 Group mode: Include
 Source: 10.0.0.3
 Source timeout: 12
 Last reported by: 10.9.5.2
 Group timeout: 0 Type: Dynamic
 Group: 232.1.1.1
 Group mode: Include
 Source: 10.0.0.4
 Source timeout: 12
 Last reported by: 10.9.5.2
 Group timeout: 0 Type: Dynamic
 Group: 232.1.1.2
 Group mode: Include
 Source: 10.0.0.4
 Source timeout: 12
 Last reported by: 10.9.5.2
 Group timeout: 0 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
 Group: 224.0.0.2
 Group mode: Exclude
 Source: 0.0.0.0
 Source timeout: 0

```

```
 Last reported by: Local
 Group timeout: 0 Type: Dynamic
Group: 224.0.0.22
 Group mode: Exclude
 Source: 0.0.0.0
 Source timeout: 0
 Last reported by: Local
 Group timeout: 0 Type: Dynamic
```

## show igmp interface

|                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                         | show igmp interface<br><brief   detail><br><interface-name><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                        |
| <b>Syntax (EX Series Switches and the QFX Series)</b> | show igmp interface<br><brief   detail><br><interface-name>                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                            | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                   |
| <b>Description</b>                                    | Display information about Internet Group Management Protocol (IGMP)-enabled interfaces.                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                        | <p><b>none</b>—Display standard information about all IGMP-enabled interfaces.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface-name</b>—(Optional) Display information about the specified IGMP-enabled interface only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                       | view                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>                          | <ul style="list-style-type: none"> <li>• <i>clear igmp membership</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>                          | <a href="#">show igmp interface on page 3961</a><br><a href="#">show igmp interface brief on page 3961</a><br><a href="#">show igmp interface detail on page 3962</a><br><a href="#">show igmp interface &lt;interface-name&gt; on page 3962</a>                                                                                                                                                                                           |
| <b>Output Fields</b>                                  | <a href="#">Table 280 on page 3959</a> describes the output fields for the <b>show igmp interface</b> command.<br>Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                  |

**Table 280: show igmp interface Output Fields**

| Field Name | Field Description                                                               | Level of Output |
|------------|---------------------------------------------------------------------------------|-----------------|
| Interface  | Name of the interface.                                                          | All levels      |
| Querier    | Address of the routing device that has been elected to send membership queries. | All levels      |
| State      | State of the interface: <b>Up</b> or <b>Down</b> .                              | All levels      |

Table 280: show igmp interface Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>SSM Map Policy</b>     | Name of the source-specific multicast (SSM) map policy that has been applied to the IGMP interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels      |
| <b>Timeout</b>            | How long until the IGMP querier is declared to be unreachable, in seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels      |
| <b>Version</b>            | IGMP version being used on the interface: 1, 2, or 3.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | All levels      |
| <b>Groups</b>             | Number of groups on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels      |
| <b>Group limit</b>        | Maximum number of groups allowed on the interface. Any joins requested after the limit is reached are rejected.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | All levels      |
| <b>Group threshold</b>    | Configured threshold at which a warning message is generated.<br><br>This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | All levels      |
| <b>Group log-interval</b> | Time (in seconds) between consecutive log messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels      |
| <b>Immediate Leave</b>    | State of the immediate leave option: <ul style="list-style-type: none"> <li>• <b>On</b>—Indicates that the router removes a host from the multicast group as soon as the router receives a leave group message from a host associated with the interface.</li> <li>• <b>Off</b>—Indicates that after receiving a leave group message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels      |
| <b>Promiscuous Mode</b>   | State of the promiscuous mode option: <ul style="list-style-type: none"> <li>• <b>On</b>—Indicates that the router can accept IGMP reports from subnetworks that are not associated with its interfaces.</li> <li>• <b>Off</b>—Indicates that the router can accept IGMP reports only from subnetworks that are associated with its interfaces.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels      |
| <b>Passive</b>            | State of the passive mode option: <ul style="list-style-type: none"> <li>• <b>On</b>—Indicates that the router can run IGMP on the interface but not send or receive control traffic such as IGMP reports, queries, and leaves.</li> <li>• <b>Off</b>—Indicates that the router can run IGMP on the interface and send or receive control traffic such as IGMP reports, queries, and leaves.</li> </ul> <p>The <b>passive</b> statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the <b>on</b> state declaration:</p> <ul style="list-style-type: none"> <li>• <b>send-general-query</b>—The interface sends general queries.</li> <li>• <b>send-group-query</b>—The interface sends group-specific and group-source-specific queries.</li> <li>• <b>allow-receive</b>—The interface receives control traffic.</li> </ul> | All levels      |
| <b>OIF map</b>            | Name of the OIF map (if configured) associated with the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels      |

Table 280: show igmp interface Output Fields (*continued*)

| Field Name            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| SSM map               | Name of the source-specific multicast (SSM) map (if configured) used on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | All levels      |
| Configured Parameters | Information configured by the user: <ul style="list-style-type: none"> <li><b>IGMP Query Interval</b>—Interval (in seconds) at which this router sends membership queries when it is the querier.</li> <li><b>IGMP Query Response Interval</b>—Time (in seconds) that the router waits for a report in response to a general query.</li> <li><b>IGMP Last Member Query Interval</b>—Time (in seconds) that the router waits for a report in response to a group-specific query.</li> <li><b>IGMP Robustness Count</b>—Number of times the router retries a query.</li> </ul> | All levels      |
| Derived Parameters    | Derived information: <ul style="list-style-type: none"> <li><b>IGMP Membership Timeout</b>—Timeout period (in seconds) for group membership. If no report is received for these groups before the timeout expires, the group membership is removed.</li> <li><b>IGMP Other Querier Present Timeout</b>—Time (in seconds) that the router waits for the IGMP querier to send a query.</li> </ul>                                                                                                                                                                              | All levels      |

## Sample Output

### show igmp interface

```

user@host> show igmp interface
Interface: at-0/3/1.0
 Querier: 10.111.30.1
 State: Up Timeout: None Version: 2 Groups: 4
 SSM Map Policy: ssm-policy-A
Interface: so-1/0/0.0
 Querier: 10.111.10.1
 State: Up Timeout: None Version: 2 Groups: 2
 SSM Map Policy: ssm-policy-B
Interface: so-1/0/1.0
 Querier: 10.111.20.1
 State: Up Timeout: None Version: 2 Groups: 4
 SSM Map Policy: ssm-policy-C
Immediate Leave: On
Promiscuous Mode: Off

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0

```

### show igmp interface brief

The output for the **show igmp interface brief** command is identical to that for the **show igmp interface** command. For sample output, see [show igmp interface on page 3961](#).

### show igmp interface detail

The output for the **show igmp interface detail** command is identical to that for the **show igmp interface** command. For sample output, see [show igmp interface on page 3961](#).

### show igmp interface <interface-name>

```
user@host# show igmp interface ge-3/2/0.0
Interface: ge-3/2/0.0
Querier: 20.1.1.1
State: Up Timeout: None Version: 3 Groups: 1
Group limit: 8
Group threshold: 60
Group log-interval: 10
Immediate leave: Off
Promiscuous mode: Off
```



## show multicast pim-to-igmp-proxy

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                       | show multicast pim-to-igmp-proxy<br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                        |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | show multicast pim-to-igmp-proxy<br><instance <i>instance-name</i> >                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                          | Command introduced in Junos OS Release 9.6.<br>Command introduced in Junos OS Release 9.6 for EX Series switches.<br><b>instance</b> option introduced in Junos OS Release 10.3.<br><b>instance</b> option introduced in Junos OS Release 10.3 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                           |
| <b>Description</b>                                  | Display configuration information about PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy.                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                                      | <p><b>none</b>—Display configuration information about PIM-to-IGMP message translation for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display configuration information about PIM-to-IGMP message translation for a specific multicast instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <i>Configuring PIM-to-IGMP and PIM-to-MLD Message Translation</i></li> </ul>                                                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>                        | <a href="#">show multicast pim-to-igmp-proxy on page 3964</a><br><a href="#">show multicast pim-to-igmp-proxy instance on page 3964</a>                                                                                                                                                                                                                                                                                                             |
| <b>Output Fields</b>                                | <a href="#">Table 281 on page 3963</a> describes the output fields for the <b>show multicast pim-to-igmp-proxy</b> command. Output fields are listed in the order in which they appear.                                                                                                                                                                                                                                                             |

**Table 281: show multicast pim-to-igmp-proxy Output Fields**

| Field Name                   | Field Description                                                                                                                                     |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Instance</b>              | Routing instance. Default instance is <b>master</b> (inet.0 routing table).                                                                           |
| <b>Proxy state</b>           | State of PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy, on the configured upstream interfaces: <b>enabled</b> or <b>disabled</b> . |
| <b><i>interface-name</i></b> | Name of upstream interface (no more than two allowed) on which PIM-to-IGMP message translation is configured.                                         |

## Sample Output

### show multicast pim-to-igmp-proxy

```
user@host> show multicast pim-to-igmp-proxy
Instance: master Proxy state: enabled
ge-0/1/0.1
ge-0/1/0.2
```

### show multicast pim-to-igmp-proxy instance

```
user@host> show multicast pim-to-igmp-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/1/0.1
```

## Operational Commands: IGMP Snooping

## clear igmp snooping membership

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear igmp snooping membership<br><group   source <i>address</i> ><br><instance <i>instance-name</i> ><br><interface <i>interface-name</i> ><br><learning-domain <i>learning-domain-name</i> ><br><vlan-id <i>vlan-identifier</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Command introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Clear IP IGMP snooping membership information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>none</b>—Clear IGMP snooping membership for all supported address families on all interfaces.</p> <p><b>group   source <i>address</i></b>—(Optional) Clear IGMP snooping membership for the specified multicast group or source address.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear IGMP snooping membership for the specified instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear IGMP snooping membership on a specific interface.</p> <p><b>learning-domain <i>learning-domain-name</i></b>—(Optional) Perform this operation on all learning domains or on a particular learning domain.</p> <p><b>vlan-id <i>vlan-identifier</i></b>—(Optional) Perform this operation on a particular VLAN.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show igmp snooping membership on page 3970</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>    | <a href="#">clear igmp snooping membership on page 3965</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## Sample Output

### clear igmp snooping membership

```
user@host> clear igmp snooping membership
```

## clear igmp snooping statistics

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear igmp snooping statistics</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;interface <i>interface-name</i>&gt;</code><br><code>&lt;learning-domain (all   <i>learning-domain-name</i>)&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Command introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Clear IP IGMP snooping statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>none</b>—Clear IGMP snooping statistics for all supported address families on all interfaces.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear IGMP snooping statistics for the specified instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear IGMP snooping statistics on a specific interface.</p> <p><b>learning-domain (all   <i>learning-domain-name</i>)</b>—(Optional) Perform this operation on all learning domains or on a particular learning domain.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show igmp snooping statistics on page 3974</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>    | <a href="#">clear igmp snooping statistics on page 3966</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Sample Output

### clear igmp snooping statistics

```
user@host> clear igmp snooping statistics
```

## show igmp snooping interface

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show igmp snooping interface <i>interface-name</i><br><brief   detail><br><bridge-domain <i>bridge-domain-name</i> ><br><virtual-switch <i>virtual-switch-name</i> ><br><vlan-id <i>vlan-identifier</i> >                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Command introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Display IGMP snooping interface information.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>none</b>—Display detailed information.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>bridge-domain <i>bridge-domain-name</i></b>—(Optional) Display information about a particular bridge domain.</p> <p><b>virtual-switch <i>virtual-switch-name</i></b>—(Optional) Display information about a particular virtual switch.</p> <p><b>vlan-id <i>vlan-identifier</i></b>—(Optional) Display information about a particular VLAN.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show igmp snooping membership on page 3970</a></li> <li>• <a href="#">show igmp snooping statistics on page 3974</a></li> </ul>                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <p><a href="#">show igmp snooping interface on page 3968</a></p> <p><a href="#">show igmp snooping interface (Group Limit Configured) on page 3969</a></p>                                                                                                                                                                                                                                                                                                                            |
| <b>Output Fields</b>            | Table 282 on page 3967 lists the output fields for the <b>show igmp snooping interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                           |

Table 282: show igmp snooping interface Output Fields

| Field Name                   | Field Description                                                                              | Level of Output |
|------------------------------|------------------------------------------------------------------------------------------------|-----------------|
| Routing-instance             | Routing instance for IGMP snooping.                                                            | All levels      |
| Learning Domain              | Learning domain for snooping.                                                                  | All levels      |
| IGMP Query Interval          | Frequency (in seconds) with which this router sends membership queries when it is the querier. | detail          |
| IGMP Query Response Interval | Time (in seconds) that the router waits for a response to a general query.                     | detail          |

Table 282: show igmp snooping interface Output Fields (*continued*)

| Field Name                         | Field Description                                                                                                                    | Level of Output |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| IGMP Last Member Query Interval    | Time (in seconds) that the router waits for a report in response to a group-specific query.                                          | detail          |
| IGMP Robustness Count              | Number of times the router retries a query.                                                                                          | detail          |
| immediate-leave                    | State of immediate leave: <b>On</b> or <b>Off</b> .                                                                                  | All levels      |
| router-interface                   | Router interfaces that are part of this learning domain.                                                                             | All levels      |
| Group limit                        | Maximum number of (source,group) pairs allowed per interface. When a group limit is not configured, this field is not shown.         | All levels      |
| interface                          | Interfaces that are being snooped in this learning domain.                                                                           | All levels      |
| Groups                             | Number of groups on the interface.                                                                                                   | none            |
| State                              | State of the interface: <b>Up</b> or <b>Down</b> .                                                                                   | none            |
| Up Groups                          | Number of active multicast groups attached to the logical interface.                                                                 | All levels      |
| IGMP Membership Timeout            | Timeout for group membership. If no report is received for these groups before the timeout expires, the group membership is removed. | none            |
| IGMP Other Querier Present Timeout | Time that the router waits for the IGMP querier to send a query.                                                                     | none            |

## Sample Output

### show igmp snooping interface

```

user@host> show igmp snooping interface
Instance: bridge-domain bar

Learning-Domain: default
Interface: ge-0/1/0.200
 State: Up Groups: 0
 Immediate leave: Off
 Router interface: yes
Interface: ge-0/1/2.200
 State: Up Groups: 2
 Immediate leave: On
 Router interface: no
Interface: ge-0/1/3.200
 State: Up Groups: 1
 Immediate leave: Off
 Router interface: no

Configured Parameters:
IGMP Query Interval: 130.0
IGMP Query Response Interval: 15.0

```

```
IGMP Last Member Query Interval: 2.0
IGMP Robustness Count: 3

Derived Parameters:
IGMP Membership Timeout: 405.0
IGMP Other Querier Present Timeout: 397.500
```

## Sample Output

### show igmp snooping interface (Group Limit Configured)

```
user@host> show igmp snooping interface instance vpls1
Instance: vpls1

Learning-Domain: default
Interface: ge-1/3/9.0
 State: Up Groups: 0
 Immediate leave: Off
 Router interface: yes
Interface: ge-1/3/8.0
 State: Up Groups: 0
 Immediate leave: Off
 Router interface: yes
 Group limit: 1000

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2
```

## show igmp snooping membership

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show igmp snooping membership<br><brief   detail><br><bridge-domain <i>bridge-domain-name</i> ><br><group <i>group-name</i> ><br><virtual-switch <i>virtual-switch-name</i> ><br><vlan-id <i>vlan-identifier</i> >                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Display IGMP snooping membership information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <p><b>none</b>—Display detailed information.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>bridge-domain <i>bridge-domain-name</i></b>—(Optional) Display information about a particular bridge domain.</p> <p><b>group <i>group-name</i></b> —(Optional) Display information about this group address.</p> <p><b>virtual-switch <i>virtual-switch-name</i></b>—(Optional) Display information about a particular virtual switch.</p> <p><b>vlan-id <i>vlan-identifier</i></b>—(Optional) Display information about a particular VLAN.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show igmp snooping interface on page 3967</a></li> <li>• <a href="#">show igmp snooping statistics on page 3974</a></li> <li>• <a href="#">clear igmp snooping membership on page 3965</a></li> </ul>                                                                                                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>    | <a href="#">show igmp snooping membership on page 3971</a><br><a href="#">show igmp snooping membership (Exclude Mode) on page 3972</a><br><a href="#">show igmp snooping membership interface ge-0/1/2.200 on page 3972</a><br><a href="#">show igmp snooping membership vlan-id 1 on page 3972</a>                                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>            | Table 283 on page 3970 lists the output fields for the <b>show igmp snooping membership</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                          |

**Table 283: show igmp snooping membership Output Fields**

| Field Name      | Field Description                   | Level of Output |
|-----------------|-------------------------------------|-----------------|
| Instance        | Routing instance for IGMP snooping. | All levels      |
| Learning Domain | Learning domain for snooping.       | All levels      |



Table 283: show igmp snooping membership Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                    | Level of Output |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Interface</b>        | Interface on which this router is a proxy.                                                                                                                                                                                                           | <b>detail</b>   |
| <b>Up Groups</b>        | Number of active multicast groups attached to the logical interface.                                                                                                                                                                                 | All levels      |
| <b>Group</b>            | Multicast group address in the membership database.                                                                                                                                                                                                  | All levels      |
| <b>Group Mode</b>       | Mode the SSM group is operating in: <b>Include</b> or <b>Exclude</b> .                                                                                                                                                                               | All levels      |
| <b>Source</b>           | Source address used on queries.                                                                                                                                                                                                                      | <b>detail</b>   |
| <b>Last reported by</b> | Address of source last replying to the query.                                                                                                                                                                                                        | <b>detail</b>   |
| <b>Group Timeout</b>    | Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.                            | All levels      |
| <b>Timeout</b>          | Length of time (in seconds) left until the entry is purged.                                                                                                                                                                                          | <b>detail</b>   |
| <b>Type</b>             | Way that the group membership information was learned: <ul style="list-style-type: none"> <li>• <b>Dynamic</b>—Group membership was learned by the IGMP protocol.</li> <li>• <b>Static</b>—Group membership was learned by configuration.</li> </ul> | <b>detail</b>   |
| <b>Include receiver</b> | Source address of receiver included in membership with timeout (in seconds).                                                                                                                                                                         | <b>detail</b>   |

## Sample Output

### show igmp snooping membership

```

user@host> show igmp snooping membership
Instance: vpls2

Learning-Domain: vlan-id 2
Interface: ge-3/0/0.2
Up Groups: 0
Interface: ge-3/1/0.2
Up Groups: 0
Interface: ge-3/1/5.2
Up Groups: 0

Instance: vpls1

Learning-Domain: vlan-id 1
Interface: ge-3/0/0.1
Up Groups: 0
Interface: ge-3/1/0.1
Up Groups: 0
Interface: ge-3/1/5.1
Up Groups: 1
 Group: 225.10.10.1
 Group mode: Exclude
 Source: 0.0.0.0

```

```
Last reported by: 100.6.85.2
Group timeout: 173 Type: Dynamic
```

#### show igmp snooping membership (Exclude Mode)

```
user@host> show igmp snooping membership
Instance: vpls2

Learning-Domain: vlan-id 2
Interface: ge-3/0/0.2
Up Groups: 0
Interface: ge-3/1/0.2
Up Groups: 0
Interface: ge-3/1/5.2
Up Groups: 0

Instance: vpls1

Learning-Domain: vlan-id 1
Interface: ge-3/0/0.1
Up Groups: 0
Interface: ge-3/1/0.1
Up Groups: 0
Interface: ge-3/1/5.1
Up Groups: 1
 Group: 225.10.10.1
 Group mode: Exclude
 Source: 0.0.0.0
 Last reported by: 100.6.85.2
 Group timeout: 173 Type: Dynamic
```

#### show igmp snooping membership interface ge-0/1/2.200

```
user@host> show igmp snooping membership interface ge-0/1/2.200
Instance: bridge-domain bar

Learning-Domain: default
Interface: ge-0/1/2.200
 Group: 225.1.1.1
 Source: 0.0.0.0
 Timeout: 391 Type: Static
 Group: 232.1.1.1
 Source: 192.168.1.1
 Timeout: 0 Type: Static
```

#### show igmp snooping membership vlan-id 1

```
user@host> show igmp snooping membership vlan-id 1
Instance: vpls2

Instance: vpls1

Learning-Domain: vlan-id 1
Interface: ge-3/0/0.1
Up Groups: 0
Interface: ge-3/1/0.1
Up Groups: 0
Interface: ge-3/1/5.1
Up Groups: 1
 Group: 225.10.10.1
 Group mode: Exclude
 Source: 0.0.0.0
```

Last reported by: 100.6.85.2  
Group timeout: 209 Type: Dynamic

## show igmp snooping statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show igmp snooping statistics<br><brief   detail><br><bridge-domain <i>bridge-domain-name</i> ><br><virtual-switch <i>virtual-switch-name</i> ><br><vlan-id <i>vlan-identifier</i> >                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Command introduced in Junos OS Release 8.5.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Display IGMP snooping statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b>none</b>—(Optional) Display detailed information.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>bridge-domain <i>bridge-domain-name</i></b>—(Optional) Display information about a particular bridge domain.</p> <p><b>virtual-switch <i>virtual-switch-name</i></b>—(Optional) Display information about a particular virtual switch.</p> <p><b>vlan-id <i>vlan-identifier</i></b>—(Optional) Display information about a particular VLAN.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show igmp snooping interface on page 3967</a></li> <li>• <a href="#">show igmp snooping membership on page 3970</a></li> <li>• <a href="#">clear igmp snooping statistics on page 3966</a></li> </ul>                                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>    | <a href="#">show igmp snooping statistics on page 3975</a>                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>            | Table 284 on page 3974 lists the output fields for the <b>show igmp snooping statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                     |

**Table 284: show igmp snooping statistics Output Fields**

| Field Name             | Field Description                                                                       | Level of Output |
|------------------------|-----------------------------------------------------------------------------------------|-----------------|
| Routing-instance       | Routing instance for IGMP snooping.                                                     | All levels      |
| IGMP packet statistics | Heading for IGMP snooping statistics for all interfaces or for the specified interface. | All levels      |
| learning-domain        | Appears at end of “IGMP packets statistics” line.                                       | All levels      |

Table 284: show igmp snooping statistics Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>IGMP Message type</b>      | Summary of IGMP statistics: <ul style="list-style-type: none"> <li>• <b>Membership Query</b>—Number of membership queries sent and received.</li> <li>• <b>V1 Membership Report</b>—Number of version 1 membership reports sent and received.</li> <li>• <b>DVMRP</b>—Number of DVMRP messages sent or received.</li> <li>• <b>PIM V1</b>—Number of PIM version 1 messages sent or received.</li> <li>• <b>Cisco Trace</b>—Number of Cisco trace messages sent or received.</li> <li>• <b>V2 Membership Report</b>—Number of version 2 membership reports sent or received.</li> <li>• <b>Group Leave</b>—Number of group leave messages sent or received.</li> <li>• <b>Domain Wide Report</b>—Number of domain-wide reports sent or received.</li> <li>• <b>V3 Membership Report</b>—Number of version 3 membership reports sent or received.</li> <li>• <b>Other Unknown types</b>—Number of unknown message types received.</li> <li>• <b>IGMP v3 unsupported type</b>—Number of messages received with unknown and unsupported IGMP version 3 message types.</li> <li>• <b>IGMP v3 source required for SSM</b>—Number of IGMP version 3 messages received that contained no source.</li> <li>• <b>IGMP v3 mode not applicable for SSM</b>—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM).</li> </ul> | All levels      |
| <b>Received</b>               | Number of messages received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | All levels      |
| <b>Sent</b>                   | Number of messages sent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels      |
| <b>Rx errors</b>              | Number of received packets that contained errors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | All levels      |
| <b>IGMP Global Statistics</b> | Summary of IGMP snooping statistics for all interfaces. <ul style="list-style-type: none"> <li>• <b>Bad Length</b>—Number of messages received with length errors so severe that further classification could not occur.</li> <li>• <b>Bad Checksum</b>—Number of messages received with a bad IP checksum. No further classification was performed.</li> <li>• <b>Rx non-local</b>—Number of messages received from senders that are not local.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels      |

## Sample Output

### show igmp snooping statistics

```
user@host> show igmp snooping statistics
Routing-instance foo
```

```
IGMP packet statistics for all interfaces in learning-domain vlan-100
```

| IGMP Message type    | Received | Sent | Rx errors |
|----------------------|----------|------|-----------|
| Membership Query     | 89       | 51   | 0         |
| V1 Membership Report | 0        | 0    | 0         |
| DVMRP                | 0        | 0    | 0         |

|                                     |     |   |    |
|-------------------------------------|-----|---|----|
| PIM V1                              | 0   | 0 | 0  |
| Cisco Trace                         | 0   | 0 | 0  |
| V2 Membership Report                | 139 | 0 | 0  |
| Group Leave                         | 0   | 0 | 0  |
| Domain Wide Report                  | 0   | 0 | 0  |
| V3 Membership Report                | 136 | 0 | 0  |
| Other Unknown types                 |     |   | 0  |
| IGMP v3 unsupported type            |     |   | 0  |
| IGMP v3 source required for SSM     |     |   | 23 |
| IGMP v3 mode not applicable for SSM |     |   | 0  |

#### IGMP Global Statistics

|              |   |
|--------------|---|
| Bad Length   | 0 |
| Bad Checksum | 0 |
| Rx non-local | 0 |

#### Routing-instance bar

IGMP packet statistics for all interfaces in learning-domain vlan-100

| IGMP Message type                   | Received | Sent | Rx errors |
|-------------------------------------|----------|------|-----------|
| Membership Query                    | 89       | 51   | 0         |
| V1 Membership Report                | 0        | 0    | 0         |
| DVMRP                               | 0        | 0    | 0         |
| PIM V1                              | 0        | 0    | 0         |
| Cisco Trace                         | 0        | 0    | 0         |
| V2 Membership Report                | 139      | 0    | 0         |
| Group Leave                         | 0        | 0    | 0         |
| Domain Wide Report                  | 0        | 0    | 0         |
| V3 Membership Report                | 136      | 0    | 0         |
| Other Unknown types                 |          |      | 0         |
| IGMP v3 unsupported type            |          |      | 0         |
| IGMP v3 source required for SSM     |          |      | 23        |
| IGMP v3 mode not applicable for SSM |          |      | 0         |

#### IGMP Global Statistics

|              |   |
|--------------|---|
| Bad Length   | 0 |
| Bad Checksum | 0 |
| Rx non-local | 0 |

## Operational Commands: MLD

## clear mld membership

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear mld membership<br><group <i>group-name</i> >   <interface <i>interface-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Clear Multicast Listener Discovery (MLD) group membership.                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b>none</b>—Clear all MLD memberships.</p> <p><b>group <i>group-name</i></b>—(Optional) Clear MLD membership for the specified group.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear MLD group membership for the specified interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show mld group on page 3979</a></li> </ul>                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>    | <a href="#">clear mld membership on page 3977</a>                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                     |

### Sample Output

#### clear mld membership

```
user@host> clear mld membership
```

## clear mld statistics

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear mld statistics</code><br><code>&lt;interface <i>interface-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Clear Multicast Listener Discovery (MLD) statistics.                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <b>none</b> —(Same as <b>logical-system all</b> ) Clear MLD statistics for all interfaces.<br><br><b>interface <i>interface-name</i></b> —(Optional) Clear MLD statistics for the specified interface.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show mld statistics on page 3987</a></li></ul>                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <a href="#">clear mld statistics on page 3978</a>                                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                               |

## Sample Output

### clear mld statistics

```
user@host> clear mld statistics
```



## show mld group

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show mld group<br><brief   detail><br><group-name><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Display information about Multicast Listener Discovery (MLD) group membership.                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>none</b>—Display standard information about all MLD groups.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>group-name</b>—(Optional) Display MLD information about the specified group.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">clear mld membership on page 3977</a></li> </ul>                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>    | <p><a href="#">show mld group (Include Mode) on page 3980</a></p> <p><a href="#">show mld group (Exclude Mode) on page 3981</a></p> <p><a href="#">show mld group brief on page 3981</a></p> <p><a href="#">show mld group detail (Include Mode) on page 3981</a></p> <p><a href="#">show mld group detail (Exclude Mode) on page 3982</a></p>                                                          |
| <b>Output Fields</b>            | <p><a href="#">Table 285 on page 3979</a> describes the output fields for the <b>show mld group</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                |

**Table 285: show mld group Output Fields**

| Field Name              | Field Description                                                                                                                | Level of Output |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Interface</b>        | Name of the interface that received the MLD membership report; <b>local</b> means that the local router joined the group itself. | All levels      |
| <b>Group</b>            | Group address.                                                                                                                   | All levels      |
| <b>Source</b>           | Source address.                                                                                                                  | All levels      |
| <b>Group Mode</b>       | Mode the SSM group is operating in: <b>Include</b> or <b>Exclude</b> .                                                           | All levels      |
| <b>Last reported by</b> | Address of the host that last reported membership in this group.                                                                 | All levels      |

Table 285: show mld group Output Fields (*continued*)

| Field Name     | Field Description                                                                                                                                                                                                         | Level of Output |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Source timeout | Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.      | detail          |
| Timeout        | Time remaining until the group membership is removed.                                                                                                                                                                     | brief none      |
| Group timeout  | Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer. | detail          |
| Type           | Type of group membership: <ul style="list-style-type: none"> <li>• <b>Dynamic</b>—Host reported the membership.</li> <li>• <b>Static</b>—Membership is configured.</li> </ul>                                             | All levels      |

## Sample Output

### show mld group (Include Mode)

```

user@host> show mld group
Interface: fe-0/1/2.0
 Group: ff02::1:ff05:1a67
 Group mode: Include
 Source: ::
 Last reported by: fe80::2e0:81ff:fe05:1a67
 Timeout: 245 Type: Dynamic
 Group: ff02::1:ffa8:c35e
 Group mode: Include
 Source: ::
 Last reported by: fe80::2e0:81ff:fe05:1a67
 Timeout: 241 Type: Dynamic
 Group: ff02::2:43e:d7f6
 Group mode: Include
 Source: ::
 Last reported by: fe80::2e0:81ff:fe05:1a67
 Timeout: 244 Type: Dynamic
 Group: ff05::2
 Group mode: Include
 Source: ::
 Last reported by: fe80::2e0:81ff:fe05:1a67
 Timeout: 244 Type: Dynamic
Interface: local
 Group: ff02::2
 Source: ::
 Last reported by: Local
 Timeout: 0 Type: Dynamic
 Group: ff02::16
 Source: ::
 Last reported by: Local
 Timeout: 0 Type: Dynamic

```

### show mld group (Exclude Mode)

```

user@host> show mld group
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
 Group: ff02::6
 Source: ::
 Last reported by: fe80::21f:12ff:feb6:4b3a
 Timeout: 245 Type: Dynamic
 Group: ff02::16
 Source: ::
 Last reported by: fe80::21f:12ff:feb6:4b3a
 Timeout: 28 Type: Dynamic
Interface: local
 Group: ff02::2
 Source: ::
 Last reported by: Local
 Timeout: 0 Type: Dynamic
 Group: ff02::16
 Source: ::
 Last reported by: Local
 Timeout: 0 Type: Dynamic

```

### show mld group brief

The output for the **show mld group brief** command is identical to that for the **show mld group** command. For sample output, see [show mld group \(Include Mode\) on page 3980](#) and [show mld group \(Exclude Mode\) on page 3981](#).

### show mld group detail (Include Mode)

```

user@host> show mld group detail
Interface: fe-0/1/2.0
 Group: ff02::1:ff05:1a67
 Group mode: Include
 Source: ::
 Last reported by: fe80::2e0:81ff:fe05:1a67
 Timeout: 224 Type: Dynamic
 Group: ff02::1:ffa8:c35e
 Group mode: Include
 Source: ::
 Last reported by: fe80::2e0:81ff:fe05:1a67
 Timeout: 220 Type: Dynamic
 Group: ff02::2:43e:d7f6
 Group mode: Include
 Source: ::
 Last reported by: fe80::2e0:81ff:fe05:1a67
 Timeout: 223 Type: Dynamic
 Group: ff05::2
 Group mode: Include
 Source: ::
 Last reported by: fe80::2e0:81ff:fe05:1a67
 Timeout: 223 Type: Dynamic
Interface: so-1/0/1.0
 Group: ff02::2
 Group mode: Include
 Source: ::
 Last reported by: fe80::280:42ff:fe15:f445
 Timeout: 258 Type: Dynamic
Interface: local

```

```
Group: ff02::2
 Group mode: Include
 Source: ::
 Last reported by: Local
 Timeout: 0 Type: Dynamic
Group: ff02::16
 Source: ::
 Last reported by: Local
 Timeout: 0 Type: Dynamic
```

#### show mld group detail (Exclude Mode)

```
user@host> show mld group detail
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
 Group: ff02::6
 Group mode: Exclude
 Source: ::
 Source timeout: 0
 Last reported by: fe80::21f:12ff:feb6:4b3a
 Group timeout: 226 Type: Dynamic
 Group: ff02::16
 Group mode: Exclude
 Source: ::
 Source timeout: 0
 Last reported by: fe80::21f:12ff:feb6:4b3a
 Group timeout: 246 Type: Dynamic
Interface: local
 Group: ff02::2
 Group mode: Exclude
 Source: ::
 Source timeout: 0
 Last reported by: Local
 Group timeout: 0 Type: Dynamic
 Group: ff02::16
 Group mode: Exclude
 Source: ::
 Source timeout: 0
 Last reported by: Local
 Group timeout: 0 Type: Dynamic
```

## show mld interface

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show mld interface<br><brief   detail><br><interface-name><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Display information about Multicast Listener Discovery (MLD)-enabled interfaces.                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>none</b>—Display standard information about all MLD-enabled interfaces.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface-name</b>—(Optional) Display information about the specified interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear mld membership on page 3977</a></li> </ul>                                                                                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">show mld interface on page 3985</a><br><a href="#">show mld interface brief on page 3985</a><br><a href="#">show mld interface detail on page 3986</a><br><a href="#">show mld interface &lt;interface-name&gt; on page 3986</a>                                                                                                                                                                            |
| <b>Output Fields</b>            | <p><a href="#">Table 286 on page 3983</a> describes the output fields for the <b>show mld interface</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                            |

**Table 286: show mld interface Output Fields**

| Field Name            | Field Description                                                                              | Level of Output |
|-----------------------|------------------------------------------------------------------------------------------------|-----------------|
| <b>Interface</b>      | Name of the interface.                                                                         | All levels      |
| <b>Querier</b>        | Address of the router that has been elected to send membership queries.                        | All levels      |
| <b>State</b>          | State of the interface: <b>Up</b> or <b>Down</b> .                                             | All levels      |
| <b>SSM Map Policy</b> | Name of the source-specific multicast (SSM) map policy that has been applied to the interface. | All levels      |
| <b>SSM Map Policy</b> | Name of the source-specific multicast (SSM) map policy at the MLD interface.                   | All levels      |
| <b>Timeout</b>        | How long until the MLD querier is declared to be unreachable, in seconds.                      | All levels      |
| <b>Version</b>        | MLD version being used on the interface: 1 or 2.                                               | All levels      |

Table 286: show mld interface Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Groups</b>             | Number of groups on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | All levels      |
| <b>Passive</b>            | <p>State of the passive mode option:</p> <ul style="list-style-type: none"> <li>• <b>On</b>—Indicates that the router can run IGMP or MLD on the interface but not send or receive control traffic such as IGMP or MLD reports, queries, and leaves.</li> <li>• <b>Off</b>—Indicates that the router can run IGMP or MLD on the interface and send or receive control traffic such as IGMP or MLD reports, queries, and leaves.</li> </ul> <p>The <b>passive</b> statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the <b>on</b> state declaration:</p> <ul style="list-style-type: none"> <li>• <b>send-general-query</b>—The interface sends general queries.</li> <li>• <b>send-group-query</b>—The interface sends group-specific and group-source-specific queries.</li> <li>• <b>allow-receive</b>—The interface receives control traffic</li> </ul> | All levels      |
| <b>OIF map</b>            | Name of the OIF map associated to the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels      |
| <b>SSM map</b>            | Name of the source-specific multicast (SSM) map used on the interface, if configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels      |
| <b>Group limit</b>        | Maximum number of groups allowed on the interface. Any memberships requested after the limit is reached are rejected.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels      |
| <b>Group threshold</b>    | <p>Configured threshold at which a warning message is generated.</p> <p>This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | All levels      |
| <b>Group log-interval</b> | Time (in seconds) between consecutive log messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | All levels      |
| <b>Immediate Leave</b>    | <p>State of the immediate leave option:</p> <ul style="list-style-type: none"> <li>• <b>On</b>—Indicates that the router removes a host from the multicast group as soon as the router receives a multicast listener done message from a host associated with the interface.</li> <li>• <b>Off</b>—Indicates that after receiving a multicast listener done message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                            | All levels      |

Table 286: show mld interface Output Fields (*continued*)

| Field Name                   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Configured Parameters</b> | Information configured by the user. <ul style="list-style-type: none"> <li>• <b>MLD Query Interval (.1 secs)</b>—Interval at which this router sends membership queries when it is the querier.</li> <li>• <b>MLD Query Response Interval (.1 secs)</b>—Time that the router waits for a report in response to a general query.</li> <li>• <b>MLD Last Member Query Interval (.1 secs)</b>—Time that the router waits for a report in response to a group-specific query.</li> <li>• <b>MLD Robustness Count</b>—Number of times the router retries a query.</li> </ul> | All levels      |
| <b>Derived Parameters</b>    | Derived information. <ul style="list-style-type: none"> <li>• <b>MLD Membership Timeout (.1 secs)</b>—Timeout period for group membership. If no report is received for these groups before the timeout expires, the group membership will be removed.</li> <li>• <b>MLD Other Querier Present Timeout (.1 secs)</b>—Time that the router waits for the IGMP querier to send a query.</li> </ul>                                                                                                                                                                        | All levels      |

## Sample Output

### show mld interface

```

user@host> show mld interface
Interface: fe-0/0/0
 Querier: None
 State: Up Timeout: 0 Version: 1 Groups: 0
 SSM Map Policy: ssm-policy-A
Interface: at-0/3/1.0
 Querier: 8038::c0a8:c345
 State: Up Timeout: None Version: 1 Groups: 0
 SSM Map Policy: ssm-policy-B
Interface: fe-1/0/1.0
 Querier: ::192.168.195.73
 State: Up Timeout: None Version: 1 Groups: 3
 SSM Map Policy: ssm-policy-C
 SSM map: ipv6map1
Immediate Leave: On

Configured Parameters:
MLD Query Interval (.1 secs): 1250
MLD Query Response Interval (.1 secs): 100
MLD Last Member Query Interval (.1 secs): 10
MLD Robustness Count: 2

Derived Parameters:
MLD Membership Timeout (.1secs): 2600
MLD Other Querier Present Timeout (.1 secs): 2550

```

### show mld interface brief

The output for the **show mld interface brief** command is identical to that for the **show mld interface** command. For sample output, see [show mld interface on page 3985](#).

### show mld interface detail

The output for the **show mld interface detail** command is identical to that for the **show mld interface** command. For sample output, see [show mld interface on page 3985](#).

### show mld interface <interface-name>

```
user@host# show mld interface ge-3/2/0.0
Interface: ge-3/2/0.0
Querier: 20.1.1.1
State: Up Timeout: None Version: 3 Groups: 1
Group limit: 8
Group threshold: 60
Group log-interval: 10
Immediate leave: Off
Promiscuous mode: Off
```



## show mld statistics

|                                 |                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show mld statistics<br><interface <i>interface-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Display information about Multicast Listener Discovery (MLD) statistics.                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>none</b>—Display MLD statistics for all interfaces.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display statistics about the specified interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear mld statistics on page 3978</a></li> </ul>                                                                                                                                                                                                                                |
| <b>List of Sample Output</b>    | <a href="#">show mld statistics on page 3988</a><br><a href="#">show mld statistics interface on page 3989</a>                                                                                                                                                                                                                       |
| <b>Output Fields</b>            | <p><a href="#">Table 287 on page 3987</a> describes the output fields for the <b>show mld statistics</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                        |

**Table 287: show mld statistics Output Fields**

| Field Name | Field Description                                 |
|------------|---------------------------------------------------|
| Received   | Number of received packets.                       |
| Sent       | Number of transmitted packets.                    |
| Rx errors  | Number of received packets that contained errors. |

Table 287: show mld statistics Output Fields (*continued*)

| Field Name                   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>MLD Message type</b>      | Summary of MLD statistics. <ul style="list-style-type: none"> <li>• <b>Listener Query (v1/v2)</b>—Number of membership queries sent and received.</li> <li>• <b>Listener Report (v1)</b>—Number of version 1 membership reports sent and received.</li> <li>• <b>Listener Done (v1/v2)</b>—Number of Listener Done messages sent and received.</li> <li>• <b>Listener Report (v2)</b>—Number of version 2 membership reports sent and received.</li> <li>• <b>Other Unknown types</b>—Number of unknown message types received.</li> <li>• <b>MLD v2 source required for SSM</b>—Number of MLD version 2 messages received that contained no source.</li> <li>• <b>MLD v2 mode not applicable for SSM</b>—Number of MLD version 2 messages received that did not contain a mode applicable for source-specific multicast (SSM).</li> </ul>          |
| <b>MLD Global Statistics</b> | Summary of MLD statistics for all interfaces. <ul style="list-style-type: none"> <li>• <b>Bad Length</b>—Number of messages received with length errors so severe that further classification could not occur.</li> <li>• <b>Bad Checksum</b>—Number of messages received with an invalid IP checksum. No further classification was performed.</li> <li>• <b>Bad Receive If</b>—Number of messages received on an interface not enabled for MLD.</li> <li>• <b>Rx non-local</b>—Number of messages received from nonlocal senders.</li> <li>• <b>Timed out</b>—Number of groups that timed out as a result of not receiving an explicit leave message.</li> <li>• <b>Rejected Report</b>—Number of reports dropped because of the MLD group policy.</li> <li>• <b>Total Interfaces</b>—Number of interfaces configured to support IGMP.</li> </ul> |

## Sample Output

### show mld statistics

```

user@host> show mld statistics
MLD packet statistics for all interfaces
MLD Message type Received Sent Rx errors
Listener Query (v1/v2) 0 2 0
Listener Report (v1) 0 0 0
Listener Done (v1/v2) 0 0 0
Listener Report (v2) 0 0 0
Other Unknown types 0 0 0
MLD v2 source required for SSM 2
MLD v2 mode not applicable for SSM 0

MLD Global Statistics
Bad Length 0
Bad Checksum 0
Bad Receive If 0
Rx non-local 0
Timed out 0

```

|                  |   |
|------------------|---|
| Rejected Report  | 0 |
| Total Interfaces | 2 |

#### show mld statistics interface

```
user@host> show mld statistics interface fe-1/0/1.0
MLD interface packet statistics for fe-1/0/1.0
MLD Message type Received Sent Rx errors
Listener Query (v1/v2) 0 2 0
Listener Report (v1) 0 0 0
Listener Done (v1/v2) 0 0 0
Listener Report (v2) 0 0 0
Other Unknown types 0 0
MLD v2 source required for SSM 2
MLD v2 mode not applicable for SSM 0

MLD Global Statistics
Bad Length 0
Bad Checksum 0
Bad Receive If 0
Rx non-local 0
Timed out 0
Rejected Report 0
Total Interfaces 2
```

## show multicast pim-to-mld-proxy

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                       | show multicast pim-to-mld-proxy<br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                   |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | show multicast pim-to-mld-proxy<br><instance <i>instance-name</i> >                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>                          | Command introduced in Junos OS Release 9.6.<br>Command introduced in Junos OS Release 9.6 for EX Series switches.<br><b>instance</b> option introduced in Junos OS Release 10.3.<br><b>instance</b> option introduced in Junos OS Release 10.3 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                     |
| <b>Description</b>                                  | Display configuration information about PIM-to-MLD message translation, also known as PIM-to-MLD proxy.                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                                      | <b>none</b> —Display configuration information about PIM-to-MLD message translation for all routing instances.<br><br><b>instance <i>instance-name</i></b> —(Optional) Display configuration information about PIM-to-MLD message translation for a specific multicast instance.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>                        | <a href="#">show multicast pim-to-mld-proxy on page 3991</a><br><a href="#">show multicast pim-to-mld-proxy instance on page 3991</a>                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>                                | <a href="#">Table 288 on page 3990</a> describes the output fields for the <b>show multicast pim-to-mld-proxy</b> command. Output fields are listed in the order in which they appear.                                                                                                                                                                                                                                                        |

**Table 288: show multicast pim-to-mld-proxy Output Fields**

| Field Name            | Field Description                                                                                                                                   |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| Proxy state           | State of PIM-to-MLD message translation, also known as PIM-to-MLD proxy, on the configured upstream interfaces: <b>enabled</b> or <b>disabled</b> . |
| <i>interface-name</i> | Name of upstream interface (no more than two allowed) on which PIM-to-MLD message translation is configured.                                        |

## Sample Output

### show multicast pim-to-mld-proxy

```
user@host> show multicast pim-to-mld-proxy
Instance: master Proxy state: enabled
ge-0/5/0.1
ge-0/5/0.2
```

### show multicast pim-to-mld-proxy instance

```
user@host> show multicast pim-to-mld-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/5/0.1
```

## Operational Commands: MSDP

## show msdp

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show msdp<br><brief   detail><br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )><br><peer <i>peer-address</i> >                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Display Multicast Source Discovery Protocol (MSDP) information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>none</b>—Display standard MSDP information for all routing instances.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information for the specified instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>peer <i>peer-address</i></b>—(Optional) Display information about the specified peer only.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show msdp source on page 3994</a></li> <li>• <a href="#">show msdp source-active on page 3996</a></li> <li>• <a href="#">show msdp statistics on page 3999</a></li> </ul>                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>    | <a href="#">show msdp on page 3993</a><br><a href="#">show msdp brief on page 3993</a><br><a href="#">show msdp detail on page 3993</a>                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>            | <a href="#">Table 289 on page 3992</a> describes the output fields for the <b>show msdp</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                |

**Table 289: show msdp Output Fields**

| Field Name    | Field Description                                                                        | Level of Output |
|---------------|------------------------------------------------------------------------------------------|-----------------|
| Peer address  | IP address of the peer.                                                                  | All levels      |
| Local address | Local address of the peer.                                                               | All levels      |
| State         | Status of the MSDP connection: <b>Listen</b> , <b>Established</b> , or <b>Inactive</b> . | All levels      |
| Last up/down  | Time at which the most recent peer-state change occurred.                                | All levels      |

Table 289: show msdp Output Fields (*continued*)

| Field Name           | Field Description                                                                                                                                                                   | Level of Output |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Peer-Group           | Peer group name.                                                                                                                                                                    | All levels      |
| SA Count             | Number of source-active cache entries advertised by each peer that were accepted, compared to the number that were received, in the format <i>number-accepted/number-received</i> . | All levels      |
| Peer Connect Retries | Number of peer connection retries.                                                                                                                                                  | detail          |
| State timer expires  | Number of seconds before another message is sent to a peer.                                                                                                                         | detail          |
| Peer Times out       | Number of seconds to wait for a response from the peer before the peer is declared unavailable.                                                                                     | detail          |
| SA accepted          | Number of entries in the source-active cache accepted from the peer.                                                                                                                | detail          |
| SA received          | Number of entries in the source-active cache received by the peer.                                                                                                                  | detail          |

## Sample Output

### show msdp

```

user@host> show msdp
Peer address Local address State Last up/down Peer-Group SA Count
198.32.8.193 198.32.8.195 Established 5d 19:25:44 North23 120/150
198.32.8.194 198.32.8.195 Established 3d 19:27:27 North23 300/345
198.32.8.196 198.32.8.195 Established 5d 19:39:36 North23 10/13
198.32.8.197 198.32.8.195 Established 5d 19:32:27 North23 5/6
198.32.8.198 198.32.8.195 Established 3d 19:33:04 North23 2305/3000

```

### show msdp brief

The output for the **show msdp brief** command is identical to that for the **show msdp** command. For sample output, see [show msdp on page 3993](#).

### show msdp detail

```

user@host> show msdp detail
Peer: 10.255.70.15
Local address: 10.255.70.19
State: Established
Peer Connect Retries: 0
State timer expires: 22
Peer Times out: 49
SA accepted: 0
SA received: 0

```

## show msdp source

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show msdp source</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;source-address&gt;</code>                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Display multicast sources learned from Multicast Source Discovery Protocol (MSDP).                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <b>none</b> —Display standard MSDP source information for all routing instances.<br><br><b>instance <i>instance-name</i></b> —(Optional) Display information for the specified instance only.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b>source-address</b> —(Optional) IP address and optional prefix length. Display information for the specified source address only. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show msdp on page 3992</a></li><li>• <a href="#">show msdp source-active on page 3996</a></li><li>• <a href="#">show msdp statistics on page 3999</a></li></ul>                                                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">show msdp source on page 3995</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |



**Output Fields** Table 290 on page 3995 describes the output fields for the **show msdp source** command. Output fields are listed in the approximate order in which they appear.

**Table 290: show msdp source Output Fields**

| Field Name     | Field Description                                                                                                                                                                                                                                                       |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Source address | IP address of the source.                                                                                                                                                                                                                                               |
| /Len           | Length of the prefix for this IP address.                                                                                                                                                                                                                               |
| Type           | Discovery method for this multicast source: <ul style="list-style-type: none"> <li>• <b>Configured</b>—Source-active limit explicitly configured for this source.</li> <li>• <b>Dynamic</b>—Source-active limit established when this source was discovered.</li> </ul> |
| Maximum        | Source-active limit applied to this source.                                                                                                                                                                                                                             |
| Threshold      | Source-active threshold applied to this source.                                                                                                                                                                                                                         |
| Exceeded       | Number of source-active messages received from this source exceeding the established maximum.                                                                                                                                                                           |

## Sample Output

**show msdp source**

```

user@host> show msdp source
Source address /Len Type Maximum Threshold Exceeded
0.0.0.0 /0 Configured 5 none 0
10.1.0.0 /16 Configured 500 none 0
10.1.1.1 /32 Configured 10000 none 0
10.1.1.2 /32 Dynamic 6936 none 0
10.1.5.5 /32 Dynamic 500 none 123
10.2.1.1 /32 Dynamic 2 none 0

```

## show msdp source-active

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>show msdp source-active</code><br><code>&lt;brief   detail&gt;</code><br><code>&lt;group <i>group</i>&gt;</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;local&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;originator <i>originator</i>&gt;</code><br><code>&lt;peer <i>peer-address</i>&gt;</code><br><code>&lt;source <i>source-address</i>&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Display the Multicast Source Discovery Protocol (MSDP) source-active cache.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <b>none</b> —Display standard MSDP source-active cache information for all routing instances.<br><br><b>brief   detail</b> —(Optional) Display the specified level of output.<br><br><b>group <i>group</i></b> —(Optional) Display source-active cache information for the specified group.<br><br><b>instance <i>instance-name</i></b> —(Optional) Display information for the specified instance.<br><br><b>local</b> —(Optional) Display all source-active caches originated by this router.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b>originator <i>originator</i></b> —(Optional) Display information about the peer that originated the source-active cache entries.<br><br><b>peer <i>peer-address</i></b> —(Optional) Display the source-active cache of the specified peer.<br><br><b>source <i>source-address</i></b> —(Optional) Display the source-active cache of the specified source. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show msdp on page 3992</a></li><li>• <a href="#">show msdp source on page 3994</a></li><li>• <a href="#">show msdp statistics on page 3999</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">show msdp source-active on page 3997</a><br><a href="#">show msdp source-active brief on page 3997</a><br><a href="#">show msdp source-active detail on page 3998</a><br><a href="#">show msdp source-active source on page 3998</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Output Fields</b>            | <a href="#">Table 291 on page 3997</a> describes the output fields for the <b>show msdp source-active</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

Table 291: show msdp source-active Output Fields

| Field Name                              | Field Description                                                                                                                                          |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Global active source limit exceeded     | Number of times all peers have exceeded configured active source limits.                                                                                   |
| Global active source limit maximum      | Configured number of active source messages accepted by the device.                                                                                        |
| Global active source limit threshold    | Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.                                         |
| Global active source limit log-warning  | Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device).                                |
| Global active source limit log interval | Time (in seconds) between consecutive log messages.                                                                                                        |
| Group address                           | Multicast address of the group.                                                                                                                            |
| Source address                          | IP address of the source.                                                                                                                                  |
| Peer address                            | IP address of the peer.                                                                                                                                    |
| Originator                              | Router ID configured on the source of the rendezvous point (RP) that originated the message, or the loopback address when the router ID is not configured. |
| Flags                                   | Flags: Accept, Reject, or Filtered.                                                                                                                        |

## Sample Output

### show msdp source-active

```

user@host> show msdp source-active
Group address Source address Peer address Originator Flags
230.0.0.0 192.168.195.46 local 10.255.14.30 Accept
230.0.0.1 192.168.195.46 local 10.255.14.30 Accept
230.0.0.2 192.168.195.46 local 10.255.14.30 Accept
230.0.0.3 192.168.195.46 local 10.255.14.30 Accept
230.0.0.4 192.168.195.46 local 10.255.14.30 Accept

```

### show msdp source-active brief

The output for the **show msdp source-active brief** command is identical to that for the **show msdp source-active** command. For sample output, see [show msdp source-active on page 3997](#).

### show msdp source-active detail

The output for the **show msdp source-active detail** command is identical to that for the **show msdp source-active** command. For sample output, see [show msdp source-active on page 3997](#).

### show msdp source-active source

```
user@host> show msdp source-active source 192.168.215.246
```

```
Global active source limit exceeded: 0
```

```
Global active source limit maximum: 25000
```

```
Global active source limit threshold: 24000
```

```
Global active source limit log-warning: 100
```

```
Global active source limit log interval: 0
```

| Group address | Source address  | Peer address   | Originator     | Flags  |
|---------------|-----------------|----------------|----------------|--------|
| 226.2.2.1     | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.3     | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.4     | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.5     | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.7     | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.10    | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.11    | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.13    | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.14    | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |
| 226.2.2.15    | 192.168.215.246 | 10.255.182.140 | 10.255.182.140 | Accept |

## show msdp statistics

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show msdp statistics<br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )><br><peer <i>peer-address</i> >                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Display statistics about Multicast Source Discovery Protocol (MSDP) peers.                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>none</b>—Display statistics about all MSDP peers for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display statistics about a specific MSDP instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>peer <i>peer-address</i></b>—(Optional) Display statistics about a particular MSDP peer.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear msdp statistics</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">show msdp statistics on page 4001</a><br><a href="#">show msdp statistics peer on page 4001</a>                                                                                                                                                                                                                                                                                                                                                    |
| <b>Output Fields</b>            | <a href="#">Table 292 on page 3999</a> describes the output fields for the <b>show msdp statistics</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                        |

**Table 292: show msdp statistics Output Fields**

| Field Name                              | Field Description                                                                                                           |
|-----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| Global active source limit exceeded     | Number of times all peers have exceeded configured active source limits.                                                    |
| Global active source limit maximum      | Configured number of active source messages accepted by the device.                                                         |
| Global active source limit threshold    | Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.          |
| Global active source limit log-warning  | Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device). |
| Global active source limit log interval | Time (in seconds) between consecutive log messages.                                                                         |
| Peer                                    | Address of peer.                                                                                                            |

Table 292: show msdp statistics Output Fields (*continued*)

| Field Name                          | Field Description                                                                                                                                   |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| Last State Change                   | How long ago the peer state changed.                                                                                                                |
| Last message received from the peer | How long ago the last message was received from the peer.                                                                                           |
| RPF Failures                        | Number of reverse path forwarding (RPF) failures.                                                                                                   |
| Remote Closes                       | Number of times the remote peer closed.                                                                                                             |
| Peer Timeouts                       | Number of peer timeouts.                                                                                                                            |
| SA messages sent                    | Number of source-active messages sent.                                                                                                              |
| SA messages received                | Number of source-active messages received.                                                                                                          |
| SA request messages sent            | Number of source-active request messages sent.                                                                                                      |
| SA request messages received        | Number of source-active request messages received.                                                                                                  |
| SA response messages sent           | Number of source-active response messages sent.                                                                                                     |
| SA response messages received       | Number of source-active response messages received.                                                                                                 |
| Active source exceeded              | Number of times this peer has exceeded configured source-active limits.                                                                             |
| Active source Maximum               | Configured number of active source messages accepted by this peer.                                                                                  |
| Active source threshold             | Configured threshold on this peer for applying random early discard (RED) to drop some but not all MSDP active source messages.                     |
| Active source log-warning           | Configured threshold on this peer at which a warning message is logged (percentage of the number of active source messages accepted by the device). |
| Active source log-interval          | Time (in seconds) between consecutive log messages on this peer.                                                                                    |
| Keepalive messages sent             | Number of keepalive messages sent.                                                                                                                  |
| Keepalive messages received         | Number of keepalive messages received.                                                                                                              |
| Unknown messages received           | Number of unknown messages received.                                                                                                                |

Table 292: show msdp statistics Output Fields (*continued*)

| Field Name              | Field Description                  |
|-------------------------|------------------------------------|
| Error messages received | Number of error messages received. |

## Sample Output

### show msdp statistics

```

user@host> show msdp statistics
Global active source limit exceeded: 0
Global active source limit maximum: 10
Global active source limit threshold: 8
Global active source limit log-warning: 60
Global active source limit log interval: 60

Peer: 10.255.245.39
Last State Change: 11:54:49 (00:24:59)
Last message received from peer: 11:53:32 (00:26:16)
RPF Failures: 0
Remote Closes: 0
Peer Timeouts: 0
SA messages sent: 376
SA messages received: 459
SA request messages sent: 0
SA request messages received: 0
SA response messages sent: 0
SA response messages received: 0
Active source exceeded: 0
Active source Maximum: 10
Active source threshold: 8
Active source log-warning: 60
Active source log-interval 120
Keepalive messages sent: 17
Keepalive messages received: 19
Unknown messages received: 0
Error messages received: 0

```

### show msdp statistics peer

```

user@host> show msdp statistics peer 10.255.182.140
Peer: 10.255.182.140
 Last State Change: 8:19:23 (00:01:08)
 Last message received from peer: 8:20:05 (00:00:26)
 RPF Failures: 0
 Remote Closes: 0
 Peer Timeouts: 0
 SA messages sent: 17
 SA messages received: 16
 SA request messages sent: 0
 SA request messages received: 0
 SA response messages sent: 0
 SA response messages received: 0
 Active source exceeded: 20
 Active source Maximum: 10
 Active source threshold: 8
 Active source log-warning: 60
 Active source log-interval: 120
 Keepalive messages sent: 0

```

Keepalive messages received: 0  
Unknown messages received: 0  
Error messages received: 0



## show multicast usage

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                       | <pre>show multicast usage &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show multicast usage &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                                  | Display usage information about the 10 most active Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM) groups.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                      | <p><b>none</b>—Display multicast usage information for all supported address families for all routing instances.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>inet   inet6</b>—(Optional) Display usage information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about the most active DVMRP or PIM groups for a specific multicast instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>                        | <p><a href="#">show multicast usage on page 4004</a></p> <p><a href="#">show multicast usage brief on page 4004</a></p> <p><a href="#">show multicast usage instance on page 4004</a></p> <p><a href="#">show multicast usage detail on page 4005</a></p>                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>                                | <p>Table 293 on page 4003 describes the output fields for the <b>show multicast usage</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

**Table 293: show multicast usage Output Fields**

| Field Name      | Field Description                                                                                 |
|-----------------|---------------------------------------------------------------------------------------------------|
| <b>Instance</b> | Name of the routing instance. (Displayed when multicast is configured within a routing instance.) |

Table 293: show multicast usage Output Fields (*continued*)

| Field Name     | Field Description                                                                                                                                                                        |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Group</b>   | Group address.                                                                                                                                                                           |
| <b>Sources</b> | Number of sources.                                                                                                                                                                       |
| <b>Packets</b> | Number of packets that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the packets field displays <b>unavailable</b> . |
| <b>Bytes</b>   | Number of bytes that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the bytes field displays <b>unavailable</b> .     |
| <b>Prefix</b>  | IP address.                                                                                                                                                                              |
| <b>/len</b>    | Prefix length.                                                                                                                                                                           |
| <b>Groups</b>  | Number of multicast groups.                                                                                                                                                              |

## Sample Output

### show multicast usage

```

user@host> show multicast usage
Group Sources Packets Bytes
228.0.0.0 1 52847 4439148
239.1.1.1 2 13450 1125530

Prefix /len Groups Packets Bytes
10.255.14.144 /32 2 66254 5561304
10.255.70.15 /32 1 43 3374...
```

### show multicast usage brief

The output for the **show multicast usage brief** command is identical to that for the **show multicast usage** command. For sample output, see [show multicast usage on page 4004](#).

### show multicast usage instance

```

user@host> show multicast usage instance VPN-A
Group Sources Packets Bytes
224.2.127.254 1 5538 509496
224.0.1.39 1 13 624
224.0.1.40 1 13 624

Prefix /len Groups Packets Bytes
192.168.195.34 /32 1 5538 509496
10.255.14.30 /32 1 13 624
10.255.245.91 /32 1 13 624
...
```

### show multicast usage detail

```
user@host> show multicast usage detail
```

| Group                                                   | Sources | Packets | Bytes   |
|---------------------------------------------------------|---------|---------|---------|
| 228.0.0.0                                               | 1       | 53159   | 4465356 |
| Source: 10.255.14.144 /32 Packets: 53159 Bytes: 4465356 |         |         |         |
| 239.1.1.1                                               | 2       | 13450   | 1125530 |
| Source: 10.255.14.144 /32 Packets: 13407 Bytes: 1122156 |         |         |         |
| Source: 10.255.70.15 /32 Packets: 43 Bytes: 3374        |         |         |         |

| Prefix           | /len | Groups         | Packets        | Bytes   |
|------------------|------|----------------|----------------|---------|
| 10.255.14.144    | /32  | 2              | 66566          | 5587512 |
| Group: 228.0.0.0 |      | Packets: 53159 | Bytes: 4465356 |         |
| Group: 239.1.1.1 |      | Packets: 13407 | Bytes: 1122156 |         |
| 10.255.70.15     | /32  | 1              | 43             | 3374    |
| Group: 239.1.1.1 |      | Packets: 43    | Bytes: 3374    |         |

## show route table

---

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | <code>show route table <i>routing-table-name</i></code><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Syntax (EX Series Switches)</b> | <code>show route table <i>routing-table-name</i></code><br><brief   detail   extensive   terse>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                 | Display the route entries in a particular routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b><i>routing-table-name</i></b> —Display route entries for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the <b>show route table inet</b> command).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"><li>• <a href="#">show route summary</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>       | <a href="#">show route table bgp.l2.vpn on page 4007</a><br><a href="#">show route table bgp.l3vpn.0 on page 4007</a><br><a href="#">show route table bgp.l3vpn.0 detail on page 4007</a><br><a href="#">show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured) on page 4009</a><br><a href="#">show route table inet.0 on page 4009</a><br><a href="#">show route table inet6.0 on page 4009</a><br><a href="#">show route table inet6.3 on page 4010</a><br><a href="#">show route table inetflow detail on page 4010</a><br><a href="#">show route table l2circuit.0 on page 4010</a><br><a href="#">show route table mpls on page 4011</a><br><a href="#">show route table mpls extensive on page 4011</a><br><a href="#">show route table mpls.0 on page 4011</a><br><a href="#">show route table mpls.0 (RSVP Route—Transit LSP) on page 4012</a><br><a href="#">show route table vpls_1 detail on page 4012</a><br><a href="#">show route table vpn-a on page 4012</a><br><a href="#">show route table vpn-a.mdt.0 on page 4013</a><br><a href="#">show route table VPN-A detail on page 4013</a><br><a href="#">show route table VPN-AB.inet.0 on page 4013</a><br><a href="#">show route table VPN_blue.mvpn-inet6.0 on page 4014</a><br><a href="#">show route table VPN-A detail on page 4014</a> |

[show route table inetflow detail on page 4015](#)

**Output Fields** For information about output fields, see the output field tables for the [show route](#) command, the [show route detail](#) command, the [show route extensive](#) command, or the [show route terse](#) command.

## Sample Output

### show route table bgp.l2vpn

```
user@host> show route table bgp.l2vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
 *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
 AS path: I
 > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

### show route table bgp.l3vpn.0

```
user@host> show route table bgp.l3vpn.0
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.71.15:100:10.255.71.17/32
 *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
 AS path: I
 > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.15:200:10.255.71.18/32
 *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
 AS path: I
 > via so-2/1/0.0, Push 100021, Push 100011(top)
```

### show route table bgp.l3vpn.0 detail

```
user@host> show route table bgp.l3vpn.0 detail
bgp.l3vpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

10.255.245.12:1:4.0.0.0/8 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 10.255.245.12:1
 Source: 10.255.245.12
 Next hop: 192.168.208.66 via fe-0/0/0.0, selected
 Label operation: Push 182449
 Protocol next hop: 10.255.245.12
 Push 182449
 Indirect next hop: 863a630 297
 State: <Active Int Ext>
 Local AS: 35 Peer AS: 35
 Age: 12:19 Metric2: 1
 Task: BGP_35.10.255.245.12+179
 Announcement bits (1): 0-BGP.0.0.0.0+179
 AS path: 30 10458 14203 2914 3356 I (Atomic) Aggregator: 3356 4.68.0.11

 Communities: 2914:420 target:11111:1 origin:56:78
 VPN Label: 182449
 Localpref: 100
```

```
Router ID: 10.255.245.12

10.255.245.12:1:4.17.225.0/24 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 10.255.245.12:1
 Source: 10.255.245.12
 Next hop: 192.168.208.66 via fe-0/0/0.0, selected
 Label operation: Push 182465
 Protocol next hop: 10.255.245.12
 Push 182465
 Indirect next hop: 863a8f0 305
 State: <Active Int Ext>
 Local AS: 35 Peer AS: 35
 Age: 12:19 Metric2: 1
 Task: BGP_35.10.255.245.12+179
 Announcement bits (1): 0-BGP.0.0.0.0+179
 AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496 6496 I
 Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
 VPN Label: 182465
 Localpref: 100
 Router ID: 10.255.245.12

10.255.245.12:1:4.17.226.0/23 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 10.255.245.12:1
 Source: 10.255.245.12
 Next hop: 192.168.208.66 via fe-0/0/0.0, selected
 Label operation: Push 182465
 Protocol next hop: 10.255.245.12
 Push 182465
 Indirect next hop: 86bd210 330
 State: <Active Int Ext>
 Local AS: 35 Peer AS: 35
 Age: 12:19 Metric2: 1
 Task: BGP_35.10.255.245.12+179
 Announcement bits (1): 0-BGP.0.0.0.0+179
 AS path: 30 10458 14203 2914 11853 11853 11853 11853 6496 6496 6496 6496 6496
 6496 I
 Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
 VPN Label: 182465
 Localpref: 100
 Router ID: 10.255.245.12

10.255.245.12:1:4.17.251.0/24 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 10.255.245.12:1
 Source: 10.255.245.12
 Next hop: 192.168.208.66 via fe-0/0/0.0, selected
 Label operation: Push 182465
 Protocol next hop: 10.255.245.12
 Push 182465
 Indirect next hop: 86bd210 330
 State: <Active Int Ext>
 Local AS: 35 Peer AS: 35
 Age: 12:19 Metric2: 1
 Task: BGP_35.10.255.245.12+179
 Announcement bits (1): 0-BGP.0.0.0.0+179
 AS path: 30 10458 14203 2914 11853 11853 11853 11853 6496 6496 6496 6496 6496
 6496 I
```

```

Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100

```

### show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured)

```

user@host> show route table bgp.rtarget.0
bgp.rtarget.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100:100:100/96
 *[RTarget/5] 00:03:14
 Type Proxy
 for 10.255.165.103
 for 10.255.166.124
 Local

```

### show route table inet.0

```

user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0 *[Static/5] 00:51:57
 > to 111.222.5.254 via fxp0.0
1.0.0.1/32 *[Direct/0] 00:51:58
 > via at-5/3/0.0
1.0.0.2/32 *[Local/0] 00:51:58
 Local
12.12.12.21/32 *[Local/0] 00:51:57
 Reject
13.13.13.13/32 *[Direct/0] 00:51:58
 > via t3-5/2/1.0
13.13.13.14/32 *[Local/0] 00:51:58
 Local
13.13.13.21/32 *[Local/0] 00:51:58
 Local
13.13.13.22/32 *[Direct/0] 00:33:59
 > via t3-5/2/0.0
127.0.0.1/32 [Direct/0] 00:51:58
 > via lo0.0
111.222.5.0/24 *[Direct/0] 00:51:58
 > via fxp0.0
111.222.5.81/32 *[Local/0] 00:51:58
 Local

```

### show route table inet6.0

```

user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64 *[Direct/0] 00:01:34
>via fe-0/1/0.0

fec0:0:0:3::/128 *[Local/0] 00:01:34
>Local

fec0:0:0:4::/64 *[Static/5] 00:01:34
>to fec0:0:0:3::ffff via fe-0/1/0.0

```

### show route table inet6.3

```
user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
 *[LDP/9] 00:00:22, metric 1
 > via so-1/0/0.0
::10.255.245.196/128
 *[LDP/9] 00:00:08, metric 1
 > via so-1/0/0.0, Push 100008
```

### show route table inetflow detail

```
user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Next-hop reference count: 2
 State: **Active Ext>
 Local AS: 65002 Peer AS: 65000
 Age: 4
 Task: BGP_65000.10.12.99.5+3792
 Announcement bits (1): 0-Flow
 AS path: 65000 I
 Communities: traffic-rate:0:0
 Validation state: Accept, Originator: 10.12.99.5
 Via: 10.12.44.0/24, Active
 Localpref: 100
 Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
 *Flow Preference: 5
 Next-hop reference count: 2
 State: **Active>
 Local AS: 65002
 Age: 6:30
 Task: RT Flow
 Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
 AS path: I
 Communities: 1:1
```

### show route table l2circuit.0

```
user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
 *[L2CKT/7] 00:50:47
 > via so-0/1/2.0, Push 100049
 > via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
 *[LDP/9] 00:50:14
 Discard
10.1.1.195:CtrlWord:1:2:Local/96
 *[L2CKT/7] 00:50:47
 > via so-0/1/2.0, Push 100049
 > via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
```



```
*[LDP/9] 00:50:14
Discard
```

### show route table mpls

```
user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0 *[MPLS/0] 00:13:55, metric 1
 Receive
1 *[MPLS/0] 00:13:55, metric 1
 Receive
2 *[MPLS/0] 00:13:55, metric 1
 Receive
1024 *[VPN/0] 00:04:18
 to table red.inet.0, Pop
```

### show route table mpls extensive

```
user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}
 *LDP Preference: 9
 Next hop: via so-1/0/0.0, selected
 Pop
 State: <Active Int>
 Age: 29:50 Metric: 1
 Task: LDP
 Announcement bits (1): 0-KRT
 AS path: I
 Prefixes bound to route: 10.0.0.194/32
```

### show route table mpls.0

```
user@host> show route table mpls.0
mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0 *[MPLS/0] 00:45:09, metric 1
 Receive
1 *[MPLS/0] 00:45:09, metric 1
 Receive
2 *[MPLS/0] 00:45:09, metric 1
 Receive
100000 *[L2VPN/7] 00:43:04
 > via so-0/1/0.1, Pop
100001 *[L2VPN/7] 00:43:03
 > via so-0/1/0.2, Pop Offset: 4
100002 *[LDP/9] 00:43:22, metric 1
 via so-0/1/2.0, Pop
 > via so-0/1/3.0, Pop
100002(S=0) *[LDP/9] 00:43:22, metric 1
 via so-0/1/2.0, Pop
 > via so-0/1/3.0, Pop
100003 *[LDP/9] 00:43:22, metric 1
 > via so-0/1/2.0, Swap 100002
 via so-0/1/3.0, Swap 100002
100004 *[LDP/9] 00:43:16, metric 1
 via so-0/1/2.0, Swap 100049
 > via so-0/1/3.0, Swap 100049
```

```

so-0/1/0.1 *[L2VPN/7] 00:43:04
 > via so-0/1/2.0, Push 100001, Push 100049(top)
 via so-0/1/3.0, Push 100001, Push 100049(top)
so-0/1/0.2 *[L2VPN/7] 00:43:03
 via so-0/1/2.0, Push 100000, Push 100049(top) Offset: -4
 > via so-0/1/3.0, Push 100000, Push 100049(top) Offset: -4

```

#### show route table mpls.0 (RSVP Route—Transit LSP)

```
user@host> show route table mpls.0
```

```

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

0 *[MPLS/0] 00:37:31, metric 1
 Receive
1 *[MPLS/0] 00:37:31, metric 1
 Receive
2 *[MPLS/0] 00:37:31, metric 1
 Receive
13 *[MPLS/0] 00:37:31, metric 1
 Receive
300352 *[RSVP/7/1] 00:08:00, metric 1
 > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300352(S=0) *[RSVP/7/1] 00:08:00, metric 1
 > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384 *[RSVP/7/2] 00:05:20, metric 1
 > to 8.64.1.106 via ge-1/0/0.0, Pop
300384(S=0) *[RSVP/7/2] 00:05:20, metric 1
 > to 8.64.1.106 via ge-1/0/0.0, Pop

```

#### show route table vpls\_1 detail

```
user@host> show route table vpls_1 detail
```

```

vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

```

```

1.1.1.11:1000:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Receive table: vpls_1.l2vpn.0
Next-hop reference count: 2
State: <Active Int Ext>
Age: 4:29:47 Metric2: 1
Task: vpls_1-l2vpn
Announcement bits (1): 1-BGP.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF

```

#### show route table vpn-a

```
user@host> show route table vpn-a
```

```
vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

```

192.168.16.1:1:1/96
 *[VPN/7] 05:48:27
 Discard
192.168.24.1:1:2:1/96
 *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
 AS path: I
 > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am

```

```

192.168.24.1:1:3:1/96
 *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
 AS path: I
 > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am

```

#### show route table vpn-a.mdt.0

```

user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
 *[MVPN/70] 01:23:05, metric2 1
 Indirect
1:1:1:10.255.14.218:232.1.1.1/144
 *[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
 AS path: I
 > via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
 *[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
 AS path: I
 > via so-0/0/1.0, label-switched-path r0-to-r2

```

#### show route table VPN-A detail

```

user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 10.255.179.13:200
 Next hop type: Indirect
 Next-hop reference count: 5
 Source: 10.255.179.13
 Next hop type: Router, Next hop index: 732
 Next hop: 10.39.1.14 via fe-0/3/0.0, selected
 Label operation: Push 299824, Push 299824(top)
 Protocol next hop: 10.255.179.13
 Push 299824
 Indirect next hop: 8f275a0 1048574
 State: (Secondary Active Int Ext)
 Local AS: 1 Peer AS: 1
 Age: 3:41:06 Metric: 1 Metric2: 1
 Task: BGP_1.10.255.179.13+64309
 Announcement bits (2): 0-KRT 1-BGP RT Background
 AS path: I
 Communities: target:1:200 rte-type:0.0.0.0:1:0
 Import Accepted
 VPN Label: 299824 TTL Action: vrf-ttl-propagate
 Localpref: 100
 Router ID: 10.255.179.13
 Primary Routing Table bgp.13vpn.0

```

#### show route table VPN-AB.inet.0

```

user@host> show route table VPN-AB.inet.0
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.0/30 *[OSPF/10] 00:07:24, metric 1
 > via so-7/3/1.0
10.39.1.4/30 *[Direct/0] 00:08:42
 > via so-5/1/0.0

```

```

10.39.1.6/32 *[Local/0] 00:08:46
 Local
10.255.71.16/32 *[Static/5] 00:07:24
 > via so-2/0/0.0
10.255.71.17/32 *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
 AS path: I
 > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.18/32 *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
 AS path: I
 > via so-2/1/0.0, Push 100021, Push 100011(top)
10.255.245.245/32 *[BGP/170] 00:08:35, localpref 100
 AS path: 2 I
 > to 10.39.1.5 via so-5/1/0.0
10.255.245.246/32 *[OSPF/10] 00:07:24, metric 1
 > via so-7/3/1.0

```

#### show route table VPN\_blue.mvpn-inet6.0

```

user@host> show route table VPN_blue.mvpn-inet6.0
vpn_blue.mvpn-inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.255.2.202:65535:10.255.2.202/432
 *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
 AS path: I
 > via so-0/1/3.0
1:10.255.2.203:65535:10.255.2.203/432
 *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
 AS path: I
 > via so-0/1/0.0
1:10.255.2.204:65535:10.255.2.204/432
 *[MVPN/70] 00:57:23, metric2 1
 Indirect
5:10.255.2.202:65535:128::192.168.90.2:128:ffff::1/432
 *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
 AS path: I
 > via so-0/1/3.0
6:10.255.2.203:65535:65000:128::10.12.53.12:128:ffff::1/432
 *[PIM/105] 00:02:37
 Multicast (IPv6)
7:10.255.2.202:65535:65000:128::192.168.90.2:128:ffff::1/432
 *[MVPN/70] 00:02:37, metric2 1
 Indirect

```

#### show route table VPN-A detail

```

user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Route Distinguisher: 10.255.179.13:200
 Next hop type: Indirect
 Next-hop reference count: 5
 Source: 10.255.179.13
 Next hop type: Router, Next hop index: 732
 Next hop: 10.39.1.14 via fe-0/3/0.0, selected
 Label operation: Push 299824, Push 299824(top)
 Protocol next hop: 10.255.179.13
 Push 299824

```

```

Indirect next hop: 8f275a0 1048574
State: (Secondary Active Int Ext)
Local AS: 1 Peer AS: 1
Age: 3:41:06 Metric: 1 Metric2: 1
Task: BGP_1.10.255.179.13+64309
Announcement bits (2): 0-KRT 1-BGP RT Background
AS path: I
Communities: target:1:200 rte-type:0.0.0.0:1:0
Import Accepted
VPN Label: 299824 TTL Action: vrf-ttl-propagate
Localpref: 100
Router ID: 10.255.179.13
Primary Routing Table bgp.13vpn.0

```

### show route table inetflow detail

```

user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
 *BGP Preference: 170/-101
 Next-hop reference count: 2
 State: **Active Ext>
 Local AS: 65002 Peer AS: 65000
 Age: 4
 Task: BGP_65000.10.12.99.5+3792
 Announcement bits (1): 0-Flow
 AS path: 65000 I
 Communities: traffic-rate:0:0
 Validation state: Accept, Originator: 10.12.99.5
 Via: 10.12.44.0/24, Active
 Localpref: 100
 Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
 *Flow Preference: 5
 Next-hop reference count: 2
 State: **Active>
 Local AS: 65002
 Age: 6:30
 Task: RT Flow
 Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
 AS path: I
 Communities: 1:1

user@PE1> show route table green.l2vpn.0 (VPLS Multihoming with FEC 129)
green.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.2:100:1.1.1.2/96 AD
 *[VPLS/170] 1d 03:11:03, metric2 1
 Indirect
1.1.1.4:100:1.1.1.4/96 AD
 *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
 AS path: I, validation-state: unverified
 > via ge-1/2/1.5
1.1.1.2:100:1:0/96 MH
 *[VPLS/170] 1d 03:11:03, metric2 1
 Indirect
1.1.1.4:100:1:0/96 MH
 *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
 AS path: I, validation-state: unverified

```

```

> via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.2:1.1.1.4/176
*[VPLS/7] 1d 03:11:02, metric2 1
> via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.4:1.1.1.2/176
*[LDP/9] 1d 03:11:02
Discard
```

## Operational Commands: PIM

## clear pim join

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                       | clear pim join<br><group-address><br><inet   inet6><br><instance instance-name><br><logical-system (all   logical-system-name)>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | clear pim join<br><group-address><br><inet   inet6><br><instance instance-name>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>                                  | Clear the Protocol Independent Multicast (PIM) join and prune states.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                                      | <p><b>none</b>—Clear the PIM join and prune states for all groups, family addresses, and instances.</p> <p><b>group-address</b>—(Optional) Clear the PIM join and prune states for a group address.</p> <p><b>inet   inet6</b>—(Optional) Clear the PIM join and prune states for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance instance-name</b>—(Optional) Clear the join and prune states for a specific PIM-enabled routing instance.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Additional Information</b>                       | The <b>clear pim join</b> command cannot be used to clear the PIM join and prune state on a backup Routing Engine when nonstop active routing is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">show pim join on page 4036</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>                        | <a href="#">clear pim join on page 4017</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Output Fields</b>                                | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## Sample Output

### clear pim join

```
user@host> clear pim join
```

## clear pim join-distribution

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear pim join-distribution</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Redistribute the Protocol Independent Multicast (PIM) join states.</p> <p>You can find out if there are multiple paths available for a source (for example, an RP) with the output of the <b>show pim source</b> command.</p> <p>When you include the <b>join-load-balance</b> statement in the configuration, the PIM join states are distributed evenly on available equal-cost multipath links. When an upstream neighbor link fails, Junos OS redistributes the PIM join states to the remaining links. However, when new links are added or the failed link is restored, the existing PIM joins are not redistributed to the new link. New flows will be distributed to the new links. However, in a network without new joins and prunes, the new link is not used for multicast traffic. The <b>clear pim join-distribution</b> command redistributes the existing flows to the new upstream neighbors. Redistributing the existing flows causes traffic to be disrupted, so we recommend that you run the <b>clear pim join-distribution</b> command during a maintenance window.</p> |
| <b>Options</b>                  | <p><b>none</b>—Redistribute the PIM join states for the default master instance.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Redistribute the join states for a specific PIM-enabled routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Additional Information</b>   | The <b>clear pim join-distribution</b> command cannot be used to redistribute the PIM join states on a backup Routing Engine when nonstop active routing is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show pim neighbors on page 4045</a></li><li>• <a href="#">show pim join on page 4036</a></li><li>• <a href="#">join-load-balance on page 3892</a> in the <i>Multicast Protocols Configuration Guide</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>    | <a href="#">clear pim join-distribution on page 4019</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Output Fields</b>            | When you enter this command, you are provided no feedback on the status of your request. You can enter the <b>show pim join</b> command before and after distributing the join state to verify the operation.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |



## Sample Output

clear pim join-distribution

```
user@host> clear pim join-distribution
```

## clear pim register

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                       | clear pim register<br><inet   inet6><br><instance <i>instance-name</i> ><br><interface <i>interface-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | clear pim register<br><inet   inet6><br><instance <i>instance-name</i> ><br><interface <i>interface-name</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax (PTX Series)</b>                          | clear pim register<br><inet   inet6><br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>                          | Command introduced in Junos OS Release 7.6.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>                                  | Clear Protocol Independent Multicast (PIM) register message counters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                      | <p>none—Clear PIM register message counters for all family addresses, instances, and interfaces.</p> <p><b>inet   inet6</b>—(Optional) Clear PIM register message counters for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear register message counters for a specific PIM-enabled routing instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear PIM register message counters for a specific interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Additional Information</b>                       | The <b>clear pim register</b> command cannot be used to clear the PIM register state on a backup Routing Engine when nonstop active routing is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"><li>• <a href="#">show pim statistics on page 4059</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>List of Sample Output</b>                        | <a href="#">clear pim register on page 4021</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Output Fields</b>                                | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Sample Output

clear pim register

```
user@host> clear pim register
```

## clear pim statistics

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                       | clear pim statistics<br><inet   inet6><br><instance <i>instance-name</i> ><br><interface <i>interface-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | clear pim statistics<br><inet   inet6><br><instance <i>instance-name</i> ><br><interface <i>interface-name</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                |
| <b>Description</b>                                  | Clear Protocol Independent Multicast (PIM) statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                                      | <b>none</b> —Clear PIM statistics for all family addresses, instances, and interfaces.<br><br><b>inet   inet6</b> —(Optional) Clear PIM statistics for IPv4 or IPv6 family addresses, respectively.<br><br><b>instance <i>instance-name</i></b> —(Optional) Clear statistics for a specific PIM-enabled routing instance.<br><br><b>interface <i>interface-name</i></b> —(Optional) Clear PIM statistics for a specific interface.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Additional Information</b>                       | The <b>clear pim statistics</b> command cannot be used to clear the PIM statistics on a backup Routing Engine when nonstop active routing is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b>                     | clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"><li>• <a href="#">show pim statistics on page 4059</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>                        | <a href="#">clear pim statistics on page 4022</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>                                | See <a href="#">show pim statistics</a> for an explanation of output fields.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## Sample Output

### clear pim statistics

The following sample output displays PIM statistics before and after the **clear pim statistics** command is entered:

```

user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type Received Sent Rx errors
Hello 0 0 0
Register 0 0 0
Register Stop 0 0 0
Join Prune 0 0 0
Bootstrap 0 0 0
Assert 0 0 0
Graft 0 0 0
Graft Ack 0 0 0
Candidate RP 0 0 0
V1 Query 2111 4222 0
V1 Register 0 0 0
V1 Register Stop 0 0 0
V1 Join Prune 14200 13115 0
V1 RP Reachability 0 0 0
V1 Assert 0 0 0
V1 Graft 0 0 0
V1 Graft Ack 0 0 0
PIM statistics summary for all interfaces:
Unknown type 0
V1 Unknown type 0
Unknown Version 0
Neighbor unknown 0
Bad Length 0
Bad Checksum 0
Bad Receive If 0
Rx Intf disabled 2007
Rx V1 Require V2 0
Rx Register not RP 0
RP Filtered Source 0
Unknown Reg Stop 0
Rx Join/Prune no state 1040
Rx Graft/Graft Ack no state 0
...

```

```

user@host> clear pim statistics
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type Received Sent Rx errors
Hello 0 0 0
Register 0 0 0
Register Stop 0 0 0
Join Prune 0 0 0
Bootstrap 0 0 0
Assert 0 0 0
Graft 0 0 0
Graft Ack 0 0 0
Candidate RP 0 0 0
V1 Query 1 0 0
V1 Register 0 0 0
...

```

## request pim multicast-tunnel rebalance

---

|                                    |                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | request pim multicast-tunnel rebalance<br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                     |
| <b>Syntax (EX Series Switches)</b> | request pim multicast-tunnel rebalance<br><instance <i>instance-name</i> >                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>         | Command introduced in Junos OS Release 10.2.<br>Command introduced in Junos OS Release 10.2 for EX Series switches.                                                                                                                                                                                                                                                    |
| <b>Description</b>                 | Rebalance the assignment of multicast tunnel encapsulation interfaces across available tunnel-capable PICs or across a configured list of tunnel-capable PICs. You can determine whether a rebalance is necessary by running the <b>show pim interfaces instance <i>instance-name</i></b> command.                                                                     |
| <b>Options</b>                     | <b>none</b> —Re-create and rebalance all tunnel interfaces for all routing instances.<br><br><b>instance <i>instance-name</i></b> —Re-create and rebalance all tunnel interfaces for a specific instance.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | maintenance                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"><li>• <a href="#">show pim interfaces on page 4033</a></li><li>• <i>Load Balancing Multicast Tunnel Interfaces Among Available PICs</i> in the <i>Junos Multicast Protocols Configuration Guide</i></li></ul>                                                                                                                        |
| <b>Output Fields</b>               | This command produces no output. To verify the operation of the command, run the <b>show pim interface instance <i>instance-name</i></b> before and after running the <b>request pim multicast-tunnel rebalance</b> command.                                                                                                                                           |

## show pim bidirectional df-election

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show pim bidirectional df-election &lt;brief   detail &gt; &lt;inet   inet6&gt; &lt;instance <i>instance name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;rpa <i>address</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | For bidirectional PIM, display the designated forwarder (DF) election results for each interface grouped by the rendezvous point addresses (RPAs).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>none</b>—Display standard information about all interfaces.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>inet   inet6</b>—(Optional) Display DF election results for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display DF election results for a specific routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>rpa <i>address</i></b>—(Optional) Display the DF election results for an RP address.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">show pim bidirectional df-election on page 4026</a><br><a href="#">show pim bidirectional df-election brief on page 4026</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>            | Table 294 on page 4025 describes the output fields for the <b>show pim bidirectional df-election</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

**Table 294: show pim bidirectional df-election Output Fields**

| Field Name          | Field Description                                                            | Level of Output |
|---------------------|------------------------------------------------------------------------------|-----------------|
| <b>Family</b>       | IPv4 address family ( <b>INET</b> ) or IPv6 address family ( <b>INET6</b> ). | All levels      |
| <b>Instance</b>     | Name of the routing instance.                                                | All levels      |
| <b>RPA</b>          | RP address.                                                                  | All levels      |
| <b>Group ranges</b> | Address ranges of the multicast groups mapped to this RP address.            | All levels      |

Table 294: show pim bidirectional df-election Output Fields (*continued*)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                           | Level of Output                                                      |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------|
| <b>Interfaces</b> | Bidirectional PIM interfaces on this routing device. An interface can win the DF election ( <b>Win</b> ), lose the DF election ( <b>Lose</b> ), or be the RP link ( <b>RPL</b> ). The RP link is the interface directly connected to a subnet that contains a phantom RP address. A phantom RP address is an RP address that is not assigned to a routing device interface. | All levels<br><br><b>brief</b> displays the DF election winner only. |
| <b>DF</b>         | IP address of the designated forwarder.                                                                                                                                                                                                                                                                                                                                     | All levels                                                           |

## Sample Output

### show pim bidirectional df-election

```

user@host> show pim bidirectional df-election
Instance: PIM.master Family: INET

RPA: 10.10.1.3
Group ranges: 224.1.3.0/24, 225.1.3.0/24
Interfaces:
 ge-0/0/1.0 (RPL) DF: none
 lo0.0 (Win) DF: 10.255.179.246
 xe-4/1/0.0 (Win) DF: 10.10.2.1

RPA: 10.10.13.2
Group ranges: 224.1.1.0/24, 225.1.1.0/24
Interfaces:
 ge-0/0/1.0 (Lose) DF: 10.10.1.2
 lo0.0 (Win) DF: 10.255.179.246
 xe-4/1/0.0 (Lose) DF: 10.10.2.2

Instance: PIM.master Family: INET6

RPA: fec0::10:10:1:3
Group ranges: ff00::/8
Interfaces:
 ge-0/0/1.0 (Lose) DF: fe80::b2c6:9aff:fe95:86fa
 lo0.0 (Win) DF: fe80::2a0:a50f:fc64:e661
 xe-4/1/0.0 (Win) DF: fe80::226:88ff:fec5:3c37

RPA: fec0::10:10:13:2
Group ranges: ff00::/8
Interfaces:
 ge-0/0/1.0 (Lose) DF: fe80::b2c6:9aff:fe95:86fa
 lo0.0 (Win) DF: fe80::2a0:a50f:fc64:e661
 xe-4/1/0.0 (Win) DF: fe80::226:88ff:fec5:3c37

```

### show pim bidirectional df-election brief

```

user@host> show pim bidirectional df-election brief
Instance: PIM.master Family: INET

RPA: 10.10.1.3
Group ranges: 224.1.3.0/24, 225.1.3.0/24
Interfaces:
 lo0.0 (Win) DF: 10.255.179.246
 xe-4/1/0.0 (Win) DF: 10.10.2.1

```



```
RPA: 10.10.13.2
Group ranges: 224.1.1.0/24, 225.1.1.0/24
Interfaces:
 lo0.0 (Win) DF: 10.255.179.246

Instance: PIM.master Family: INET6

RPA: fec0::10:10:1:3
Group ranges: ff00::/8
Interfaces:
 lo0.0 (Win) DF: fe80::2a0:a50f:fc64:e661
 xe-4/1/0.0 (Win) DF: fe80::226:88ff:fec5:3c37

RPA: fec0::10:10:13:2
Group ranges: ff00::/8
Interfaces:
 lo0.0 (Win) DF: fe80::2a0:a50f:fc64:e661
 xe-4/1/0.0 (Win) DF: fe80::226:88ff:fec5:3c37
```

## show pim bidirectional df-election interface

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show pim bidirectional df-election interface<br><inet   inet6><br><instance <i>instance name</i> ><br><interface-name><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | For bidirectional PIM, display the default and the configured designated forwarder (DF) election parameters for each interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>none</b>—Display standard information about all interfaces.</p> <p><b>inet   inet6</b>—(Optional) Display DF election parameters for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display DF election parameters for a specific routing instance.</p> <p><b>interface-name</b>—(Optional) Display DF election parameters for a specific interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>    | <a href="#">show pim bidirectional df-election interface on page 4029</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>            | <a href="#">Table 295 on page 4028</a> describes the output fields for the <b>show pim bidirectional df-election interface</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                 |

**Table 295: show pim bidirectional df-election interface Output Fields**

| Field Name             | Field Description                                                                                                               |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| <b>Instance</b>        | Name of the routing instance.                                                                                                   |
| <b>Family</b>          | IPv4 address family ( <b>INET</b> ) or IPv6 address family ( <b>INET6</b> ).                                                    |
| <b>Interface</b>       | Name of the bidirectional PIM interface.                                                                                        |
| <b>Robustnes Count</b> | Minimum number of DF election messages that must fail to be received for DF election to fail.                                   |
| <b>Offer Period</b>    | Interval between repeated DF election messages.                                                                                 |
| <b>Backoff Period</b>  | Period that the acting DF waits between receiving a better DF Offer and sending the Pass message to transfer DF responsibility. |

Table 295: show pim bidirectional df-election interface Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                            |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RPA        | RP address.                                                                                                                                                  |
| State      | For each RP address, state of each interface with respect to the DF election: <b>Offer</b> (when the election is in progress), <b>Win</b> , or <b>Lose</b> . |
| DF         | IP address of the designated forwarder.                                                                                                                      |

## Sample Output

### show pim bidirectional df-election interface

```

user@host> show pim bidirectional df-election interface
Instance: PIM.master Family: INET

Interface: ge-0/0/1.0
 Robustness Count: 3
 Offer Period: 100 ms
 Backoff Period: 1000 ms

 RPA State DF
 10.10.1.3 Offer none
 10.10.13.2 Lose 10.10.1.2

Interface: lo0.0
 Robustness Count: 3
 Offer Period: 100 ms
 Backoff Period: 1000 ms

 RPA State DF
 10.10.1.3 Win 10.255.179.246
 10.10.13.2 Win 10.255.179.246

Interface: xe-4/1/0.0
 Robustness Count: 3
 Offer Period: 100 ms
 Backoff Period: 1000 ms

 RPA State DF
 10.10.1.3 Win 10.10.2.1
 10.10.13.2 Lose 10.10.2.2

Instance: PIM.master Family: INET6

Interface: ge-0/0/1.0
 Robustness Count: 3
 Offer Period: 100 ms
 Backoff Period: 1000 ms

 RPA State DF
 fec0::10:10:1:3 Lose fe80::b2c6:9aff:fe95:86fa
 fec0::10:10:13:2 Lose fe80::b2c6:9aff:fe95:86fa

Interface: lo0.0

```

Robustness Count: 3  
Offer Period: 100 ms  
Backoff Period: 1000 ms

|                  |       |                          |
|------------------|-------|--------------------------|
| RPA              | State | DF                       |
| fec0::10:10:1:3  | Win   | fe80::2a0:a50f:fc64:e661 |
| fec0::10:10:13:2 | Win   | fe80::2a0:a50f:fc64:e661 |

Interface: xe-4/1/0.0  
Robustness Count: 3  
Offer Period: 100 ms  
Backoff Period: 1000 ms

|                  |       |                          |
|------------------|-------|--------------------------|
| RPA              | State | DF                       |
| fec0::10:10:1:3  | Win   | fe80::226:88ff:fec5:3c37 |
| fec0::10:10:13:2 | Win   | fe80::226:88ff:fec5:3c37 |

## show pim bootstrap

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                       | show pim bootstrap<br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                     |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | show pim bootstrap<br><instance <i>instance-name</i> >                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>instance</b> option introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                     |
| <b>Description</b>                                  | For sparse mode only, display information about Protocol Independent Multicast (PIM) bootstrap routers.                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                                      | <p><b>none</b>—Display PIM bootstrap router information for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about bootstrap routers for a specific PIM-enabled routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>                        | <a href="#">show pim bootstrap on page 4032</a><br><a href="#">show pim bootstrap instance on page 4032</a>                                                                                                                                                                                                                                                                                        |
| <b>Output Fields</b>                                | Table 296 on page 4031 describes the output fields for the <b>show pim bootstrap</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                              |

**Table 296: show pim bootstrap Output Fields**

| Field Name           | Field Description                                                                               |
|----------------------|-------------------------------------------------------------------------------------------------|
| <b>Instance</b>      | Name of the routing instance.                                                                   |
| <b>BSR</b>           | Bootstrap router.                                                                               |
| <b>Pri</b>           | Priority of the routing device as elected to be the bootstrap router.                           |
| <b>Local address</b> | Local routing device address.                                                                   |
| <b>Pri</b>           | Local routing device address priority to be elected as the bootstrap router.                    |
| <b>State</b>         | Local routing device election state: <b>Candidate</b> , <b>Elected</b> , or <b>Ineligible</b> . |

Table 296: show pim bootstrap Output Fields (*continued*)

| Field Name | Field Description                                                                                    |
|------------|------------------------------------------------------------------------------------------------------|
| Timeout    | How long until the local routing device declares the bootstrap router to be unreachable, in seconds. |

## Sample Output

### show pim bootstrap

```
user@host> show pim bootstrap
Instance: PIM.master
```

| BSR                     | Pri | Local address           | Pri | State      | Timeout |
|-------------------------|-----|-------------------------|-----|------------|---------|
| None                    | 0   | 10.255.71.46            | 0   | InEligible | 0       |
| feco:1:1:1:1:0:aff:785c | 34  | feco:1:1:1:1:0:aff:7c12 | 0   | InEligible | 0       |

### show pim bootstrap instance

```
user@host> show pim bootstrap instance VPN-A
Instance: PIM.VPN-A
```

| BSR  | Pri | Local address   | Pri | State      | Timeout |
|------|-----|-----------------|-----|------------|---------|
| None | 0   | 192.168.196.105 | 0   | InEligible | 0       |

## show pim interfaces

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                       | show pim interfaces<br><inet   inet6><br><instance ( <i>instance-name</i>   all)><br><logical-system (all   <i>logical-system-name</i> )>                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | show pim interfaces<br><inet   inet6><br><instance ( <i>instance-name</i>   all)>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.<br>Support for bidirectional PIM added in Junos OS Release 12.1.<br>Support for the <b>instance all</b> option added in Junos OS Release 12.1.                                                                                                                 |
| <b>Description</b>                                  | Display information about the interfaces on which Protocol Independent Multicast (PIM) is configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                                      | <p><b>none</b>—Display interface information for all family addresses for the main instance.</p> <p><b>inet   inet6</b>—(Optional) Display interface information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance (<i>instance-name</i>   all)</b>—(Optional) Display information about interfaces for a specific PIM-enabled routing instance or for all routing instances.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>                        | <a href="#">show pim interfaces on page 4034</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Output Fields</b>                                | <a href="#">Table 297 on page 4033</a> describes the output fields for the <b>show pim interfaces</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                          |

**Table 297: show pim interfaces Output Fields**

| Field Name      | Field Description                                                                          |
|-----------------|--------------------------------------------------------------------------------------------|
| <b>Instance</b> | Name of the routing instance.                                                              |
| <b>Name</b>     | Interface name.                                                                            |
| <b>State</b>    | State of the interface. The state also is displayed in the <b>show interfaces</b> command. |

Table 297: show pim interfaces Output Fields (*continued*)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Mode</b>         | <p>PIM mode running on the interface:</p> <ul style="list-style-type: none"> <li>• <b>B</b>—In bidirectional mode, multicast groups are carried across the network over bidirectional shared trees. This type of tree minimizes PIM routing state, which is especially important in networks with numerous and dispersed senders and receivers.</li> <li>• <b>S</b>—In sparse mode, routing devices must join and leave multicast groups explicitly. Upstream routing devices do not forward multicast traffic to this routing device unless this device has sent an explicit request (using a join message) to receive multicast traffic.</li> <li>• <b>Dense</b>—Unlike sparse mode, where data is forwarded only to routing devices sending an explicit request, dense mode implements a flood-and-prune mechanism, similar to DVMRP (the first multicast protocol used to support the multicast backbone). (Not supported on QFX Series.)</li> <li>• <b>Sparse-Dense</b>—Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as <b>dense</b> is not mapped to a rendezvous point (RP). Instead, data packets destined for that group are forwarded using PIM-Dense Mode (PIM-DM) rules. A group specified as <b>sparse</b> is mapped to an RP, and data packets are forwarded using PIM-Sparse Mode (PIM-SM) rules. (Not supported on QFX Series.)</li> </ul> <p>When sparse-dense mode is configured, the output includes both <b>S</b> and <b>D</b>. When bidirectional-sparse mode is configured, the output includes <b>S</b> and <b>B</b>. When bidirectional-sparse-dense mode is configured, the output includes <b>B</b>, <b>S</b>, and <b>D</b>.</p> |
| <b>IP</b>           | Version number of the address family on the interface: <b>4</b> (IPv4) or <b>6</b> (IPv6).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>V</b>            | PIM version running on the interface: 1 or 2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>State</b>        | <p>State of PIM on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Active</b>—Bidirectional mode is enabled on the interface and on all PIM neighbors.</li> <li>• <b>DR</b>—Designated router.</li> <li>• <b>NotCap</b>—Bidirectional mode is not enabled on the interface. This can happen when bidirectional PIM is not configured locally, when one of the neighbors is not configured for bidirectional PIM, or when one of the neighbors has not implemented the bidirectional PIM protocol.</li> <li>• <b>NotDR</b>—Not the designated router.</li> <li>• <b>P2P</b>—Point to point.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>NbrCnt</b>       | Number of neighbors that have been seen on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>JoinCnt(sg)</b>  | Number of (s,g) join messages that have been seen on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>JointCnt(*g)</b> | Number of (*g) join messages that have been seen on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>DR address</b>   | Address of the designated router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## Sample Output

### show pim interfaces

```

user@host> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,

```



Active = Bidirectional is active, NotCap = Not Bidirectional Capable

| Name           | Stat | Mode | IP | V | State        | NbrCnt | JoinCnt(sg/*g) | DR address |
|----------------|------|------|----|---|--------------|--------|----------------|------------|
| ge-0/3/0.0     | Up   | S    | 4  | 2 | NotDR,NotCap | 1      | 0/0            | 40.0.0.3   |
| ge-0/3/3.50    | Up   | S    | 4  | 2 | DR,NotCap    | 1      | 9901/100       | 50.0.0.2   |
| ge-0/3/3.51    | Up   | S    | 4  | 2 | DR,NotCap    | 1      | 0/0            | 51.0.0.2   |
| pe-1/2/0.32769 | Up   | S    | 4  | 2 | P2P,NotCap   | 0      | 0/0            |            |

## show pim join

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                       | <code>show pim join</code><br><code>&lt;brief   detail   extensive   summary&gt;</code><br><code>&lt;inet   inet6&gt;</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;range&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <code>show pim join</code><br><code>&lt;brief   detail   extensive   summary&gt;</code><br><code>&lt;inet   inet6&gt;</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;range&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>summary</b> option introduced in Junos OS Release 9.6.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Support for bidirectional PIM added in Junos OS Release 12.1.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>                                  | Display information about Protocol Independent Multicast (PIM) groups for all PIM modes.<br><br>For bidirectional PIM, display information about PIM group ranges (*G-range) for each active bidirectional RP group range, in addition to each of the joined (*G) routes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                                      | <b>none</b> —Display the standard information about PIM groups for all supported family addresses for all routing instances.<br><br><b>brief   detail   extensive   summary</b> —(Optional) Display the specified level of output.<br><br><b>inet   inet6</b> —(Optional) Display PIM group information for IPv4 or IPv6 family addresses, respectively.<br><br><b>instance <i>instance-name</i></b> —(Optional) Display information about groups for the specified PIM-enabled routing instance only.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b>range</b> —(Optional) Address range of the group, specified as <i>prefix/prefix-length</i> . |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"><li>• <a href="#">clear pim join on page 4017</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>                        | <a href="#">show pim join summary on page 4040</a><br><a href="#">show pim join (PIM Sparse Mode) on page 4040</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

[show pim join \(Bidirectional PIM\) on page 4040](#)  
[show pim join instance <instance-name> on page 4041](#)  
[show pim join detail on page 4041](#)  
[show pim join extensive \(PIM Sparse Mode\) on page 4042](#)  
[show pim join extensive \(Bidirectional PIM\) on page 4043](#)  
[show pim join extensive \(Bidirectional PIM with a Directly Connected Phantom RP\) on page 4044](#)  
[show pim join instance <instance-name> extensive on page 4044](#)

**Output Fields** [Table 298 on page 4037](#) describes the output fields for the **show pim join** command. Output fields are listed in the approximate order in which they appear.

**Table 298: show pim join Output Fields**

| Field Name                               | Field Description                                                                                                                                      | Level of Output                            |
|------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------|
| <b>Instance</b>                          | Name of the routing instance.                                                                                                                          | <b>brief detail extensive summary none</b> |
| <b>Family</b>                            | Name of the address family: <b>inet</b> (IPv4) or <b>inet6</b> (IPv6).                                                                                 | <b>brief detail extensive summary none</b> |
| <b>Route type</b>                        | Type of multicast route: (S,G) or (*G).                                                                                                                | <b>summary</b>                             |
| <b>Route count</b>                       | Number of (S,G) routes and number of (*G) routes.                                                                                                      | <b>summary</b>                             |
| <b>R</b>                                 | Rendezvous Point Tree.                                                                                                                                 | <b>brief detail extensive none</b>         |
| <b>S</b>                                 | Sparse.                                                                                                                                                | <b>brief detail extensive none</b>         |
| <b>W</b>                                 | Wildcard.                                                                                                                                              | <b>brief detail extensive none</b>         |
| <b>Group</b>                             | Group address.                                                                                                                                         | <b>brief detail extensive none</b>         |
| <b>Bidirectional group prefix length</b> | For bidirectional PIM, length of the IP prefix for RP group ranges.                                                                                    | All levels                                 |
| <b>Source</b>                            | Multicast source: <ul style="list-style-type: none"> <li>• * (wildcard value)</li> <li>• <i>ipv4-address</i></li> <li>• <i>ipv6-address</i></li> </ul> | <b>brief detail extensive none</b>         |
| <b>RP</b>                                | Rendezvous point for the PIM group.                                                                                                                    | <b>brief detail extensive none</b>         |

Table 298: show pim join Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output                    |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------|
| <b>Flags</b>              | PIM flags: <ul style="list-style-type: none"> <li>• <b>bidirectional</b>—Bidirectional mode entry.</li> <li>• <b>dense</b>—Dense mode entry.</li> <li>• <b>rptree</b>—Entry is on the rendezvous point tree.</li> <li>• <b>sparse</b>—Sparse mode entry.</li> <li>• <b>spt</b>—Entry is on the shortest-path tree for the source.</li> <li>• <b>wildcard</b>—Entry is on the shared tree.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>brief detail extensive none</b> |
| <b>Upstream interface</b> | RPF interface toward the source address for the source-specific state (S,G) or toward the rendezvous point (RP) address for the non-source-specific state (*,G).<br><br>For bidirectional PIM, <b>RP Link</b> means that the interface is directly connected to a subnet that contains a phantom RP address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>brief detail extensive none</b> |
| <b>Upstream neighbor</b>  | Information about the upstream neighbor: <b>Direct</b> , <b>Local</b> , <b>Unknown</b> , or a specific IP address.<br><br>For bidirectional PIM, <b>Direct</b> means that the interface is directly connected to a subnet that contains a phantom RP address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>extensive</b>                   |
| <b>Upstream state</b>     | Information about the upstream interface: <ul style="list-style-type: none"> <li>• <b>Join to RP</b>—Sending a join to the rendezvous point.</li> <li>• <b>Join to Source</b>—Sending a join to the source.</li> <li>• <b>Local RP</b>—Sending neither join messages nor prune messages toward the RP, because this routing device is the rendezvous point.</li> <li>• <b>Local Source</b>—Sending neither join messages nor prune messages toward the source, because the source is locally attached to this routing device.</li> <li>• <b>Prune to RP</b>—Sending a prune to the rendezvous point.</li> <li>• <b>Prune to Source</b>—Sending a prune to the source.</li> </ul> <p><b>NOTE:</b> RP group range entries have <b>None</b> in the <b>Upstream state</b> field because RP group ranges do not trigger actual PIM join messages between routing devices.</p> | <b>extensive</b>                   |

Table 298: show pim join Output Fields (*continued*)

| Field Name                                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output  |
|-------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Downstream neighbors</b>               | <p>Information about downstream interfaces:</p> <ul style="list-style-type: none"> <li>• <b>Interface</b>—Interface name for the downstream neighbor.</li> </ul> <p><b>NOTE:</b> A pseudo PIM-SM interface appears for all IGMP-only interfaces.</p> <ul style="list-style-type: none"> <li>• <b>Interface address</b>—Address of the downstream neighbor.</li> <li>• <b>State</b>—Information about the downstream neighbor: <b>join</b> or <b>prune</b>.</li> <li>• <b>Flags</b>—PIM join flags: <b>R (RPtree)</b>, <b>S (Sparse)</b>, <b>W (Wildcard)</b>, or <b>zero</b>.</li> <li>• <b>Uptime</b>—Time since the downstream interface joined the group.</li> <li>• <b>Time since last Join</b>—Time since the last join message was received from the downstream interface.</li> <li>• <b>Time since last Prune</b>—Time since the last prune message was received from the downstream interface.</li> </ul> | <b>extensive</b> |
| <b>Number of downstream interfaces</b>    | Total number of outgoing interfaces for each (S,G) entry.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>extensive</b> |
| <b>Assert Timeout</b>                     | Length of time between assert cycles on the downstream interface. Not displayed if the assert timer is null.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>extensive</b> |
| <b>Keepalive timeout</b>                  | Time remaining until the downstream join state is updated (in seconds). If the downstream join state is not updated before this keepalive timer reaches zero, the entry is deleted. If there is a directly connected host, <b>Keepalive timeout</b> is <b>Infinity</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>extensive</b> |
| <b>Uptime</b>                             | Time since the creation of (S,G) or (*,G) state. The uptime is not refreshed every time a PIM join message is received for an existing (S,G) or (*,G) state.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>extensive</b> |
| <b>Bidirectional accepting interfaces</b> | <p>Interfaces on the routing device that forward bidirectional PIM traffic.</p> <p>The reasons for forwarding bidirectional PIM traffic are that the interface is the winner of the designated forwarder election (<b>DF Winner</b>), or the interface is the reverse path forwarding (RPF) interface toward the RP (<b>RPF</b>).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>extensive</b> |

## Sample Output

### show pim join summary

```
user@host> show pim join summary
Instance: PIM.master Family: INET

Route type Route count
(s,g) 2
(*,g) 1

Instance: PIM.master Family: INET6
```

### show pim join (PIM Sparse Mode)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

### show pim join (Bidirectional PIM)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.13.2
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0

Group: 224.1.3.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.1.3
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0 (RP Link)

Group: 225.1.1.0
Bidirectional group prefix length: 24
Source: *
```

```

RP: 10.10.13.2
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0

Group: 225.1.3.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.1.3
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0 (RP Link)

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

#### show pim join instance <instance-name>

```

user@host> show pim join instance VPN-A
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
Source: *
RP: 10.10.47.100
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 235.1.1.2
Source: 192.168.195.74
Flags: sparse,spt
Upstream interface: at-0/3/1.0

Group: 235.1.1.2
Source: 192.168.195.169
Flags: sparse
Upstream interface: so-1/0/1.0

Instance: PIM.VPN-A Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

#### show pim join detail

```

user@host> show pim join detail
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0

```

Instance: PIM.master Family: INET6  
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

### show pim join extensive (PIM Sparse Mode)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local RP
Uptime: 00:03:49
Downstream neighbors:
 Interface: so-1/0/0.0
 10.111.10.2 State: Join Flags: SRW Timeout: 174
 Uptime: 00:03:49 Time since last Join: 00:01:49
 Interface: mt-1/1/0.32768
 10.10.47.100 State: Join Flags: SRW Timeout: Infinity
 Uptime: 00:03:49 Time since last Join: 00:01:49
Number of downstream interfaces: 2

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local Source, Local RP
Keepalive timeout: 344
Uptime: 00:03:49
Downstream neighbors:
 Interface: so-1/0/0.0
 10.111.10.2 State: Join Flags: S Timeout: 174
 Uptime: 00:03:49 Time since last Prune: 00:01:49
 Interface: mt-1/1/0.32768
 10.10.47.100 State: Join Flags: S Timeout: Infinity
 Uptime: 00:03:49 Time since last Prune: 00:01:49
Number of downstream interfaces: 2

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0
Upstream neighbor: 10.111.10.2
Upstream state: Local RP, Join to Source
Keepalive timeout: 344
Uptime: 00:03:49
Downstream neighbors:
 Interface: Pseudo-GMP
 fe-0/0/0.0 fe-0/0/1.0 fe-0/0/3.0
 Interface: so-1/0/0.0 (pruned)
 10.111.10.2 State: Prune Flags: SR Timeout: 174
 Uptime: 00:03:49 Time since last Prune: 00:01:49
 Interface: mt-1/1/0.32768
 10.10.47.100 State: Join Flags: S Timeout: Infinity
 Uptime: 00:03:49 Time since last Prune: 00:01:49
Number of downstream interfaces: 3
```



Instance: PIM.master Family: INET6  
 R = Rendezvous Point Tree, S = Sparse, W = Wildcard

### show pim join extensive (Bidirectional PIM)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.13.2
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0
 Upstream neighbor: 10.10.1.2
 Upstream state: None
 Uptime: 00:03:49
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)
 Number of downstream interfaces: 0

Group: 225.1.1.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.13.2
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0
 Upstream neighbor: 10.10.1.2
 Upstream state: None
 Uptime: 00:03:49
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)
 Downstream neighbors:
 Interface: lt-1/0/10.24
 10.0.24.4 State: Join RW Timeout: 185
 Interface: lt-1/0/10.23
 10.0.23.3 State: Join RW Timeout: 184
 Number of downstream interfaces: 2

Group: 225.1.3.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.1.3
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0 (RP Link)
 Upstream neighbor: Direct
 Upstream state: Local RP
 Uptime: 00:03:49
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)
 Interface: xe-4/1/0.0 (DF Winner)
 Number of downstream interfaces: 0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

**show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP)**

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.3.0
 Bidirectional group prefix length: 24
 Source: *
 RP: 10.10.1.3
 Flags: bidirectional,rptree,wildcard
 Upstream interface: ge-0/0/1.0 (RP Link)
 Upstream neighbor: Direct
 Upstream state: Local RP
 Uptime: 00:03:49
 Bidirectional accepting interfaces:
 Interface: ge-0/0/1.0 (RPF)
 Interface: lo0.0 (DF Winner)
 Interface: xe-4/1/0.0 (DF Winner)
 Number of downstream interfaces: 0
```

**show pim join instance <instance-name> extensive**

```
user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
 Source: *
 RP: 10.10.47.100
 Flags: sparse,rptree,wildcard
 Upstream interface: Local
 Upstream neighbor: Local
 Upstream state: Local RP
 Uptime: 00:03:49
 Downstream neighbors:
 Interface: mt-1/1/0.32768
 10.10.47.101 State: Join Flags: SRW Timeout: 156
 Uptime: 00:03:49 Time since last Join: 00:01:49
 Number of downstream interfaces: 1

Group: 235.1.1.2
 Source: 192.168.195.74
 Flags: sparse,spt
 Upstream interface: at-0/3/1.0
 Upstream neighbor: 10.111.30.2
 Upstream state: Local RP, Join to Source
 Keepalive timeout: 156
 Uptime: 00:14:52

Group: 235.1.1.2
 Source: 192.168.195.169
 Flags: sparse
 Upstream interface: so-1/0/1.0
 Upstream neighbor: 10.111.20.2
 Upstream state: Local RP, Join to Source
 Keepalive timeout: 156
 Uptime: 00:14:52
```

## show pim neighbors

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                       | <pre>show pim neighbors &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance (instance-name   all)&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim neighbors &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance (instance-name   all)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Support for the <b>instance all</b> option added in Junos OS Release 12.1.</p>                                                                                                                                                                                                                       |
| <b>Description</b>                                  | Display information about Protocol Independent Multicast (PIM) neighbors.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                                      | <p><b>none</b>—(Same as <b>brief</b>) Display standard information about PIM neighbors for all supported family addresses for the main instance.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>inet   inet6</b>—(Optional) Display information about PIM neighbors for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance (instance-name   all)</b>—(Optional) Display information about neighbors for the specified PIM-enabled routing instance or for all routing instances.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>                        | <p><a href="#">show pim neighbors on page 4047</a></p> <p><a href="#">show pim neighbors brief on page 4047</a></p> <p><a href="#">show pim neighbors instance on page 4047</a></p> <p><a href="#">show pim neighbors detail on page 4047</a></p> <p><a href="#">show pim neighbors detail (With BFD) on page 4048</a></p>                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>                                | <p><a href="#">Table 299 on page 4046</a> describes the output fields for the <b>show pim neighbors</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

Table 299: show pim neighbors Output Fields

| Field Name                                       | Field Description                                                                                                                                                                                                                                                                                              | Level of Output   |
|--------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| <b>Instance</b>                                  | Name of the routing instance.                                                                                                                                                                                                                                                                                  | All levels        |
| <b>Interface</b>                                 | Interface through which the neighbor is reachable.                                                                                                                                                                                                                                                             | All levels        |
| <b>Neighbor addr</b>                             | Address of the neighboring PIM routing device.                                                                                                                                                                                                                                                                 | All levels        |
| <b>IP</b>                                        | IP version: 4 or 6.                                                                                                                                                                                                                                                                                            | All levels        |
| <b>V</b>                                         | PIM version running on the neighbor: 1 or 2.                                                                                                                                                                                                                                                                   | All levels        |
| <b>Mode</b>                                      | PIM mode of the neighbor: <b>Sparse</b> , <b>Dense</b> , <b>SparseDense</b> , or <b>Unknown</b> . When the neighbor is running PIM version 2, this mode is always <b>Unknown</b> .                                                                                                                             | All levels        |
| <b>Option</b>                                    | Can be one or more of the following: <ul style="list-style-type: none"> <li>• <b>B</b>—Bidirectional Capable.</li> <li>• <b>H</b>—Hello Option Holdtime.</li> <li>• <b>G</b>—Generation Identifier.</li> <li>• <b>P</b>—Hello Option DR Priority.</li> <li>• <b>L</b>—Hello Option LAN Prune Delay.</li> </ul> | <b>brief</b> none |
| <b>Uptime</b>                                    | Time the neighbor has been operational since the PIM process was last initialized, in the format <b>dd:hh:mm:ss ago</b> for less than a week and <b>nwnd:hh:mm:ss ago</b> for more than a week.                                                                                                                | All levels        |
| <b>Address</b>                                   | Address of the neighboring PIM routing device.                                                                                                                                                                                                                                                                 | <b>detail</b>     |
| <b>BFD</b>                                       | Status and operational state of the Bidirectional Forwarding Detection (BFD) protocol on the interface: <b>Enabled</b> , <b>Operational state is up</b> , or <b>Disabled</b> .                                                                                                                                 | <b>detail</b>     |
| <b>Hello Option Holdtime</b>                     | Time for which the neighbor is available, in seconds. The range of values is 0 through 65,535.                                                                                                                                                                                                                 | <b>detail</b>     |
| <b>Hello Default Holdtime</b>                    | Default holdtime and the time remaining if the <b>holdtime</b> option is not in the received hello message.                                                                                                                                                                                                    | <b>detail</b>     |
| <b>Hello Option DR Priority</b>                  | Designated router election priority. The range of values is 0 through 255.                                                                                                                                                                                                                                     | <b>detail</b>     |
| <b>Hello Option Generation ID</b>                | 9-digit or 10-digit number used to tag hello messages.                                                                                                                                                                                                                                                         | <b>detail</b>     |
| <b>Hello Option Bi-Directional PIM supported</b> | Neighbor can process bidirectional PIM messages.                                                                                                                                                                                                                                                               | <b>detail</b>     |
| <b>Hello Option LAN Prune Delay</b>              | Time to wait before the neighbor receives prune messages, in the format <b>delay nnn ms override nnnn ms</b> .                                                                                                                                                                                                 | <b>detail</b>     |

Table 299: show pim neighbors Output Fields (*continued*)

| Field Name                 | Field Description                                                                                                                                                                                                                                                                   | Level of Output |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Join Suppression supported | Neighbor is capable of join suppression.                                                                                                                                                                                                                                            | detail          |
| Rx Join                    | Information about joins received from the neighbor. <ul style="list-style-type: none"> <li><b>Group</b>—Group addresses in the join message.</li> <li><b>Source</b>—Address of the source in the join message.</li> <li><b>Timeout</b>—Time for which the join is valid.</li> </ul> | detail          |

## Sample Output

### show pim neighbors

```

user@host> show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority

Interface IP V Mode Option Uptime Neighbor addr
so-1/0/0.0 4 2 HPLG 00:07:10 10.111.10.2

```

### show pim neighbors brief

The output for the **show pim neighbors brief** command is identical to that for the **show pim neighbors** command. For sample output, see [show pim neighbors on page 4047](#).

### show pim neighbors instance

```

user@host> show pim neighbors instance VPN-A
Instance: PIM.VPN-A
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority

Interface IP V Mode Option Uptime Neighbor addr
at-0/3/1.0 4 2 HPLG 00:07:54 10.111.30.2
mt-1/1/0.32768 4 2 HPLG 00:07:22 10.10.47.101
so-1/0/1.0 4 2 HPLG 00:07:50 10.111.20.2

```

### show pim neighbors detail

```

user@host> show pim neighbors detail
Instance: PIM.master
Interface: ge-0/0/1.0

Address: 10.10.1.1, IPv4, PIM v2, Mode: SparseDense, sg Join Count: 0, tsf
Join Count: 2
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option Generation ID: 2053759302
Hello Option Bi-Directional PIM supported
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported

```

```
Address: 10.10.1.2, IPv4, PIM v2, sg Join Count: 0, tsg Join Count: 2
 BFD: Disabled
 Hello Option Holdtime: 105 seconds 93 remaining
 Hello Option DR Priority: 1
 Hello Option Generation ID: 1734018161
 Hello Option Bi-Directional PIM supported
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
 Join Suppression supported
```

Interface: lo0.0

```
Address: 10.255.179.246, IPv4, PIM v2, Mode: SparseDense, sg Join Count:
0, tsg Join Count: 0
 Hello Option Holdtime: 65535 seconds
 Hello Option DR Priority: 1
 Hello Option Generation ID: 1997462267
 Hello Option Bi-Directional PIM supported
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
 Join Suppression supported
```

#### show pim neighbors detail (With BFD)

```
user@host> show pim neighbors detail
```

Instance: PIM.master

Interface: fe-1/0/0.0

```
Address: 192.168.11.1, IPv4, PIM v2, Mode: Sparse
 Hello Option Holdtime: 65535 seconds
 Hello Option DR Priority: 1
 Hello Option Generation ID: 836607909
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

```
Address: 192.168.11.2, IPv4, PIM v2
 BFD: Enabled, Operational state is up
 Hello Default Holdtime: 105 seconds 104 remaining
 Hello Option DR Priority: 1
 Hello Option Generation ID: 1907549685
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

Interface: fe-1/0/1.0

```
Address: 192.168.12.1, IPv4, PIM v2
 BFD: Disabled
 Hello Default Holdtime: 105 seconds 80 remaining
 Hello Option DR Priority: 1
 Hello Option Generation ID: 1971554705
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

## show pim rps

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                       | <pre>show pim rps &lt;brief   detail   extensive&gt; &lt;group-address&gt; &lt;inet   inet6&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim rps &lt;brief   detail   extensive&gt; &lt;group-address&gt; &lt;inet   inet6&gt; &lt;instance instance-name&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p>                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>                                  | Display information about Protocol Independent Multicast (PIM) rendezvous points (RPs).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                      | <p><b>none</b>—Display standard information about PIM RPs for all groups and family addresses for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>group-address</b>—(Optional) Display the RPs for a particular group. If you specify a group address, the output lists the routing device that is the RP for that group.</p> <p><b>inet   inet6</b>—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance instance-name</b>—(Optional) Display information about RPs for a specific PIM-enabled routing instance.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Bidirectional PIM</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>List of Sample Output</b>                        | <p><a href="#">show pim rps on page 4052</a></p> <p><a href="#">show pim rps brief on page 4052</a></p> <p><a href="#">show pim rps &lt;group-address&gt; (Bidirectional PIM) on page 4052</a></p> <p><a href="#">show pim rps &lt;group-address&gt; (PIM Dense Mode) on page 4052</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

[show pim rps <group-address> \(SSM Range Without asm-override-ssm Configured\) on page 4052](#)  
[show pim rps <group-address> \(SSM Range With asm-override-ssm Configured and a Sparse-Mode RP\) on page 4053](#)  
[show pim rps <group-address> \(SSM Range With asm-override-ssm Configured and a Bidirectional RP\) on page 4053](#)  
[show pim rps instance on page 4053](#)  
[show pim rps extensive \(PIM Sparse Mode\) on page 4053](#)  
[show pim rps extensive \(Bidirectional PIM\) on page 4054](#)  
[show pim rps extensive \(PIM Anycast RP in Use\) on page 4054](#)

**Output Fields** [Table 300 on page 4050](#) describes the output fields for the **show pim rps** command. Output fields are listed in the approximate order in which they appear.

**Table 300: show pim rps Output Fields**

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                          | Level of Output         |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Instance</b>                 | Name of the routing instance.                                                                                                                                                                                                                                                                                                                                                              | All levels              |
| <b>Family or Address family</b> | Name of the address family: <b>inet</b> (IPv4) or <b>inet6</b> (IPv6).                                                                                                                                                                                                                                                                                                                     | All levels              |
| <b>RP address</b>               | Address of the rendezvous point.                                                                                                                                                                                                                                                                                                                                                           | All levels              |
| <b>Type</b>                     | Type of RP: <ul style="list-style-type: none"> <li><b>auto-rp</b>—Address of the RP known through the Auto-RP protocol.</li> <li><b>bootstrap</b>—Address of the RP known through the bootstrap router protocol (BSR).</li> <li><b>embedded</b>—Address of the RP known through an embedded RP (IPv6).</li> <li><b>static</b>—Address of RP known through static configuration.</li> </ul> | <b>brief none</b>       |
| <b>Holdtime</b>                 | How long to keep the RP active, with time remaining, in seconds.                                                                                                                                                                                                                                                                                                                           | All levels              |
| <b>Timeout</b>                  | How long until the local routing device determines the RP to be unreachable, in seconds.                                                                                                                                                                                                                                                                                                   | All levels              |
| <b>Groups</b>                   | Number of groups currently using this RP.                                                                                                                                                                                                                                                                                                                                                  | All levels              |
| <b>Group prefixes</b>           | Addresses of groups that this RP can span.                                                                                                                                                                                                                                                                                                                                                 | <b>brief none</b>       |
| <b>Learned via</b>              | Address and method by which the RP was learned.                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Mode</b>                     | The PIM mode of the RP: bidirectional or sparse.<br><br>If a sparse and bidirectional RPs are configured with the same RP address, they appear as separate entries in both formats.                                                                                                                                                                                                        | All levels              |
| <b>Time Active</b>              | How long the RP has been active, in the format <b>hh:mm:ss</b> .                                                                                                                                                                                                                                                                                                                           | <b>detail extensive</b> |



Table 300: show pim rps Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output                                     |
|---------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------|
| <b>Device Index</b>                   | Index value of the order in which Junos OS finds and initializes the interface.<br><br>For bidirectional RPs, the <b>Device Index</b> output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail extensive</b>                             |
| <b>Subunit</b>                        | Logical unit number of the interface.<br><br>For bidirectional RPs, the <b>Subunit</b> output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b>                             |
| <b>Interface</b>                      | Either the encapsulation or the de-encapsulation logical interface, depending on whether this routing device is a designated router (DR) facing an RP router, or is the local RP, respectively.<br><br>For bidirectional RPs, the <b>Interface</b> output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b>                             |
| <b>Group Ranges</b>                   | Addresses of groups that this RP spans.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b><br><br><i>group-address</i> |
| <b>Active groups using RP</b>         | Number of groups currently using this RP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive</b>                             |
| <b>total</b>                          | Total number of active groups for this RP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b>                             |
| <b>Register State for RP</b>          | Current register state for each group:<br><ul style="list-style-type: none"> <li>• <b>Group</b>—Multicast group address.</li> <li>• <b>Source</b>—Multicast source address for which the PIM register is sent or received, depending on whether this router is a designated router facing an RP router, or is the local RP, respectively:</li> <li>• <b>First Hop</b>—PIM-designated routing device that sent the Register message (the source address in the IP header).</li> <li>• <b>RP Address</b>—RP to which the Register message was sent (the destination address in the IP header).</li> <li>• <b>State</b>:<br/>On the designated router: <ul style="list-style-type: none"> <li>• <b>Send</b>—Sending Register messages.</li> <li>• <b>Probe</b>—Sent a null register. If a Register-Stop message does not arrive in 5 seconds, the designated router resumes sending Register messages.</li> <li>• <b>Suppress</b>—Received a Register-Stop message. The designated router is waiting for the timer to resume before changing to <b>Probe</b> state.</li> </ul> </li> <li>• On the RP: <ul style="list-style-type: none"> <li>• <b>Receive</b>—Receiving Register messages.</li> </ul> </li> </ul> | <b>extensive</b>                                    |
| <b>Anycast-PIM rpset</b>              | If anycast RP is configured, the addresses of the RPs in the set.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>extensive</b>                                    |
| <b>Anycast-PIM local address used</b> | If anycast RP is configured, the local address used by the RP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>extensive</b>                                    |

Table 300: show pim rps Output Fields (*continued*)

| Field Name                        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output      |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|
| <b>Anycast-PIM Register State</b> | <p>If anycast RP is configured, the current register state for each group:</p> <ul style="list-style-type: none"> <li>• <b>Group</b>—Multicast group address.</li> <li>• <b>Source</b>—Multicast source address for which the PIM register is sent or received, depending on whether this routing device is a designated router facing an RP router, or is the local RP, respectively.</li> <li>• <b>Origin</b>—How the information was obtained: <ul style="list-style-type: none"> <li>• <b>DIRECT</b>—From a local attachment</li> <li>• <b>MSDP</b>—From the Multicast Source Discovery Protocol (MSDP)</li> <li>• <b>DR</b>—From the designated router</li> </ul> </li> </ul> | <b>extensive</b>     |
| <b>RP selected</b>                | For sparse mode and bidirectional mode, the identity of the RP for the specified group address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <i>group-address</i> |

## Sample Output

### show pim rps

```

user@host> show pim rps
Instance: PIM.master
Address family INET
RP address Type Mode Holdtime Timeout Groups Group prefixes
10.10.1.3 static bidir 150 None 2 224.1.3.0/24
 225.1.3.0/24
10.10.13.2 static bidir 150 None 2 224.1.1.0/24
 225.1.1.0/24

```

### show pim rps brief

The output for the **show pim rps brief** command is identical to that for the **show pim rps** command. For sample output, see [show pim rps on page 4052](#).

### show pim rps <group-address> (Bidirectional PIM)

```

user@host> show pim rps 224.1.1.1
Instance: PIM.master

224.1.0.0/16
 11.4.12.75 (Bidirectional)

RP selected: 11.4.12.75

```

### show pim rps <group-address> (PIM Dense Mode)

```

user@host> show pim rps 224.1.1.1
Instance: PIM.master

Dense Mode active for group 224.1.1.1

```

### show pim rps <group-address> (SSM Range Without asm-override-ssm Configured)

```

user@host> show pim rps 224.1.1.1

```

Instance: PIM.master

Source-specific Mode (SSM) active for group 224.1.1.1

#### show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Sparse-Mode RP)

user@host> show pim rps 224.1.1.1

Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16  
11.4.12.75

RP selected: 11.4.12.75

#### show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Bidirectional RP)

user@host> show pim rps 224.1.1.1

Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16  
11.4.12.75 (Bidirectional)

RP selected: (null)

#### show pim rps instance

user@host> show pim rps instance VPN-A

Instance: PIM.VPN-A

Address family INET

| RP address   | Type   | Holdtime | Timeout | Groups | Group prefixes |
|--------------|--------|----------|---------|--------|----------------|
| 10.10.47.100 | static | 0        | None    | 1      | 224.0.0.0/4    |

Address family INET6

#### show pim rps extensive (PIM Sparse Mode)

user@host> show pim rps extensive

Instance: PIM.master

Family: INET

RP: 10.255.245.91

Learned via: static configuration

Time Active: 00:05:48

Holdtime: 45 with 36 remaining

Device Index: 122

Subunit: 32768

Interface: pd-6/0/0.32768

Group Ranges:

224.0.0.0/4, 36s remaining

Active groups using RP:

225.1.1.1

total 1 groups active

Register State for RP:

| Group     | Source         | FirstHop      | RP Address    | State   | Timeout |
|-----------|----------------|---------------|---------------|---------|---------|
| 225.1.1.1 | 192.168.195.78 | 10.255.14.132 | 10.255.245.91 | Receive | 0       |

#### show pim rps extensive (Bidirectional PIM)

```
user@host> show pim rps extensive
Instance: PIM.master
Address family INET
```

```
RP: 10.10.1.3
Learned via: static configuration
Mode: Bidirectional
Time Active: 01:58:07
Holdtime: 150
Group Ranges:
 224.1.3.0/24
 225.1.3.0/24
```

```
RP: 10.10.13.2
Learned via: static configuration
Mode: Bidirectional
Time Active: 01:58:07
Holdtime: 150
Group Ranges:
 224.1.1.0/24
 225.1.1.0/24
```

#### show pim rps extensive (PIM Anycast RP in Use)

```
user@host> show pim rps extensive
Instance: PIM.master
```

```
Family: INET
RP: 10.10.10.2
Learned via: static configuration
Time Active: 00:54:52
Holdtime: 0
Device Index: 130
Subunit: 32769
Interface: pimd.32769
Group Ranges:
 224.0.0.0/4
Active groups using RP:
 224.10.10.10
```

total 1 groups active

```
Anycast-PIM rpset:
 10.100.111.34
 10.100.111.17
 10.100.111.55
```

Anycast-PIM local address used: 10.100.111.1

Anycast-PIM Register State:

| Group        | Source     | Origin |
|--------------|------------|--------|
| 224.1.1.1    | 10.10.95.2 | DIRECT |
| 224.1.1.2    | 10.10.95.2 | DIRECT |
| 224.10.10.10 | 10.10.70.1 | MSDP   |
| 224.10.10.11 | 10.10.70.1 | MSDP   |
| 224.20.20.1  | 10.10.71.1 | DR     |

Address family INET6

Anycast-PIM rpset:

ab::1

ab::2

Anycast-PIM local address used: cd::1

Anycast-PIM Register State:

| Group         | Source       | Origin |
|---------------|--------------|--------|
| ::224.1.1.1   | ::10.10.95.2 | DIRECT |
| ::224.1.1.2   | ::10.10.95.2 | DIRECT |
| ::224.20.20.1 | ::10.10.71.1 | DR     |

## show pim source

---

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                       | <code>show pim source</code><br><code>&lt;brief   detail&gt;</code><br><code>&lt;inet   inet6&gt;</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;source-prefix&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <code>show pim source</code><br><code>&lt;brief   detail&gt;</code><br><code>&lt;inet   inet6&gt;</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;source-prefix&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>                          | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>                                  | Display information about the Protocol Independent Multicast (PIM) source reverse path forwarding (RPF) state.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                                      | <b>none</b> —Display standard information about the PIM RPF state for all supported family addresses for all routing instances.<br><br><b>brief   detail</b> —(Optional) Display the specified level of output.<br><br><b>inet   inet6</b> —(Optional) Display information for IPv4 or IPv6 family addresses, respectively.<br><br><b>instance <i>instance-name</i></b> —(Optional) Display information about the RPF state for a specific PIM-enabled routing instance.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b>source-prefix</b> —(Optional) Display the state for source RPF states in the given range. |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>                        | <a href="#">show pim source on page 4057</a><br><a href="#">show pim source brief on page 4057</a><br><a href="#">show pim source detail on page 4057</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>                                | <a href="#">Table 301 on page 4057</a> describes the output fields for the <b>show pim source</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

Table 301: show pim source Output Fields

| Field Name         | Field Description                                                     |
|--------------------|-----------------------------------------------------------------------|
| Instance           | Name of the routing instance.                                         |
| Source             | Address of the source or reverse path.                                |
| Prefix/length      | Prefix and prefix length for the route used to reach the RPF address. |
| Upstream interface | RPF interface toward the source address.                              |
| Upstream Neighbor  | Address of the RPF neighbor used to reach the source address.         |

## Sample Output

### show pim source

```

user@host> show pim source
Instance: PIM.master Family: INET

Source 10.255.14.144
 Prefix 10.255.14.144/32
 Upstream interface Local
 Upstream neighbor Local

Source 10.255.70.15
 Prefix 10.255.70.15/32
 Upstream interface so-1/0/0.0
 Upstream neighbor 10.111.10.2

Instance: PIM.master Family: INET6

```

### show pim source brief

The output for the **show pim source brief** command is identical to that for the **show pim source** command. For sample output, see [show pim source on page 4057](#).

### show pim source detail

```

user@host> show pim source detail
Instance: PIM.master Family: INET

Source 10.255.14.144
 Prefix 10.255.14.144/32
 Upstream interface Local
 Upstream neighbor Local
 Active groups:228.0.0.0
 239.1.1.1
 239.1.1.1

Source 10.255.70.15
 Prefix 10.255.70.15/32
 Upstream interface so-1/0/0.0
 Upstream neighbor 10.111.10.2
 Active groups:239.1.1.1

```

Instance: PIM.master Family: INET6



## show pim statistics

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                       | <pre>show pim statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (EX Series Switch and the QFX Series)</b> | <pre>show pim statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                          | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p>                                                                                                                                                                              |
| <b>Description</b>                                  | Display Protocol Independent Multicast (PIM) statistics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                      | <p><b>none</b>—Display PIM statistics.</p> <p><b>inet   inet6</b>—(Optional) Display IPv4 or IPv6 PIM statistics, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display statistics for a specific routing instance enabled by Protocol Independent Multicast (PIM).</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display statistics about the specified interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                     | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>                        | <ul style="list-style-type: none"> <li>• <a href="#">clear pim statistics on page 4022</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>                        | <p><a href="#">show pim statistics on page 4066</a></p> <p><a href="#">show pim statistics inet interface &lt;interface-name&gt; on page 4068</a></p> <p><a href="#">show pim statistics inet6 interface &lt;interface-name&gt; on page 4068</a></p> <p><a href="#">show pim statistics instance &lt;instance-name&gt; on page 4069</a></p> <p><a href="#">show pim statistics interface &lt;interface-name&gt; on page 4070</a></p>                                                                                                                                  |
| <b>Output Fields</b>                                | <p><a href="#">Table 302 on page 4060</a> describes the output fields for the <b>show pim statistics</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                         |

Table 302: show pim statistics Output Fields

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Instance</b>         | <p>Name of the routing instance.</p> <p>This field only appears if you specify an interface, for example:</p> <ul style="list-style-type: none"> <li>• <b>inet interface <i>interface-name</i></b></li> <li>• <b>inet6 interface <i>interface-name</i></b></li> <li>• <b>interface <i>interface-name</i></b></li> </ul>                                                                                                 |
| <b>Family</b>           | <p>Output is for IPv4 or IPv6 PIM statistics. <b>INET</b> indicates IPv4 statistics, and <b>INET6</b> indicates IPv6 statistics.</p> <p>This field only appears if you specify an interface, for example:</p> <ul style="list-style-type: none"> <li>• <b>inet interface <i>interface-name</i></b></li> <li>• <b>inet6 interface <i>interface-name</i></b></li> <li>• <b>interface <i>interface-name</i></b></li> </ul> |
| <b>PIM statistics</b>   | PIM statistics for all interfaces or for the specified interface.                                                                                                                                                                                                                                                                                                                                                       |
| <b>PIM message type</b> | Message type for which statistics are displayed.                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Received</b>         | Number of received statistics.                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Sent</b>             | Number of messages sent of a certain type.                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Rx errors</b>        | Number of received packets that contained errors.                                                                                                                                                                                                                                                                                                                                                                       |
| <b>V2 Hello</b>         | PIM version 2 hello packets.                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>V2 Register</b>      | PIM version 2 register packets.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>V2 Register Stop</b> | PIM version 2 register stop packets.                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>V2 Join Prune</b>    | PIM version 2 join and prune packets.                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>V2 Bootstrap</b>     | PIM version 2 bootstrap packets.                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>V2 Assert</b>        | PIM version 2 assert packets.                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>V2 Graft</b>         | PIM version 2 graft packets.                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>V2 Graft Ack</b>     | PIM version 2 graft acknowledgment packets.                                                                                                                                                                                                                                                                                                                                                                             |
| <b>V2 Candidate RP</b>  | PIM version 2 candidate RP packets.                                                                                                                                                                                                                                                                                                                                                                                     |

Table 302: show pim statistics Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                              |
|-----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>V2 State Refresh</b>                 | PIM version 2 control messages related to PIM dense mode (PIM-DM) state refresh.<br><br>State refresh is an extension to PIM-DM. It not supported in Junos OS. |
| <b>V2 DF Election</b>                   | PIM version 2 send and receive messages associated with bidirectional PIM designated forwarder election.                                                       |
| <b>V1 Query</b>                         | PIM version 1 query packets.                                                                                                                                   |
| <b>V1 Register</b>                      | PIM version 1 register packets.                                                                                                                                |
| <b>V1 Register Stop</b>                 | PIM version 1 register stop packets.                                                                                                                           |
| <b>V1 Join Prune</b>                    | PIM version 1 join and prune packets.                                                                                                                          |
| <b>V1 RP Reachability</b>               | PIM version 1 RP reachability packets.                                                                                                                         |
| <b>V1 Assert</b>                        | PIM version 1 assert packets.                                                                                                                                  |
| <b>V1 Graft</b>                         | PIM version 1 graft packets.                                                                                                                                   |
| <b>V1 Graft Ack</b>                     | PIM version 1 graft acknowledgment packets.                                                                                                                    |
| <b>AutoRP Announce</b>                  | Auto-RP announce packets.                                                                                                                                      |
| <b>AutoRP Mapping</b>                   | Auto-RP mapping packets.                                                                                                                                       |
| <b>AutoRP Unknown type</b>              | Auto-RP packets with an unknown type.                                                                                                                          |
| <b>Anycast Register</b>                 | Auto-RP announce packets.                                                                                                                                      |
| <b>Anycast Register Stop</b>            | Auto-RP announce packets.                                                                                                                                      |
| <b>Global Statistics</b>                | Summary of PIM statistics for all interfaces.                                                                                                                  |
| <b>Hello dropped on neighbor policy</b> | Number of hello packets dropped because of a configured neighbor policy.                                                                                       |
| <b>Unknown type</b>                     | Number of PIM control packets received with an unknown type.                                                                                                   |
| <b>V1 Unknown type</b>                  | Number of PIM version 1 control packets received with an unknown type.                                                                                         |
| <b>Unknown Version</b>                  | Number of PIM control packets received with an unknown version. The version is not version 1 or version 2.                                                     |

Table 302: show pim statistics Output Fields (*continued*)

| Field Name                             | Field Description                                                                                                         |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Neighbor unknown</b>                | Number of PIM control packets received (excluding PIM hello) without first receiving the hello packet.                    |
| <b>Bad Length</b>                      | Number of PIM control packets received for which the packet size does not match the PIM length field in the packet.       |
| <b>Bad Checksum</b>                    | Number of PIM control packets received for which the calculated checksum does not match the checksum field in the packet. |
| <b>Bad Receive If</b>                  | Number of PIM control packets received on an interface that does not have PIM configured.                                 |
| <b>Rx Bad Data</b>                     | Number of PIM control packets received that contain data for TCP Bad register packets.                                    |
| <b>Rx Intf disabled</b>                | Number of PIM control packets received on an interface that has PIM disabled.                                             |
| <b>Rx V1 Require V2</b>                | Number of PIM version 1 control packets received on an interface configured for PIM version 2.                            |
| <b>Rx V2 Require V1</b>                | Number of PIM version 2 control packets received on an interface configured for PIM version 1.                            |
| <b>Rx Register not RP</b>              | Number of PIM register packets received when the routing device is not the RP for the group.                              |
| <b>Rx Register no route</b>            | Number of PIM register packets received when the RP does not have a unicast route back to the source.                     |
| <b>Rx Register no decap if</b>         | Number of PIM register packets received when the RP does not have a de-encapsulation interface.                           |
| <b>Null Register Timeout</b>           | Number of NULL register timeout packets.                                                                                  |
| <b>RP Filtered Source</b>              | Number of PIM packets received when the routing device has a source address filter configured for the RP.                 |
| <b>Rx Unknown Reg Stop</b>             | Number of register stop messages received with an unknown type.                                                           |
| <b>Rx Join/Prune no state</b>          | Number of join and prune messages received for which the routing device has no state.                                     |
| <b>Rx Join/Prune on upstream if</b>    | Number of join and prune messages received on the interface used to reach the upstream routing device, toward the RP.     |
| <b>Rx Join/Prune for invalid group</b> | Number of join or prune messages received for invalid multicast group addresses.                                          |

Table 302: show pim statistics Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                           |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Rx Join/Prune messages dropped</b> | Number of join and prune messages received and dropped.                                                                                                                     |
| <b>Rx sparse join for dense group</b> | Number of PIM sparse mode join messages received for a group that is configured for dense mode.                                                                             |
| <b>Rx Graft/Graft Ack no state</b>    | Number of graft and graft acknowledgment messages received for which the router or switch has no state.                                                                     |
| <b>Rx Graft on upstream if</b>        | Number of graft messages received on the interface used to reach the upstream routing device, toward the RP.                                                                |
| <b>Rx CRP not BSR</b>                 | Number of BSR messages received in which the PIM message type is Candidate-RP-Advertisement, not Bootstrap.                                                                 |
| <b>Rx BSR when BSR</b>                | Number of BSR messages received in which the PIM message type is Bootstrap.                                                                                                 |
| <b>Rx BSR not RPF if</b>              | Number of BSR messages received on an interface that is not the RPF interface.                                                                                              |
| <b>Rx unknown hello opt</b>           | Number of PIM hello packets received with options that Junos OS does not support.                                                                                           |
| <b>Rx data no state</b>               | Number of PIM control packets received for which the routing device has no state for the data type.                                                                         |
| <b>Rx RP no state</b>                 | Number of PIM control packets received for which the routing device has no state for the RP.                                                                                |
| <b>Rx aggregate</b>                   | Number of PIM aggregate MDT packets received.                                                                                                                               |
| <b>Rx malformed packet</b>            | Number of PIM control packets received with a malformed IP unicast or multicast address family.                                                                             |
| <b>No RP</b>                          | Number of PIM control packets received with no RP address.                                                                                                                  |
| <b>No register encaps if</b>          | Number of PIM register packets received when the first-hop routing device does not have an encapsulation interface.                                                         |
| <b>No route upstream</b>              | Number of PIM control packets received when the routing device does not have a unicast route to the the interface used to reach the upstream routing device, toward the RP. |
| <b>Nexthop Unusable</b>               | Number of PIM control packets with an unusable nexthop. A path can be unusable if the route is hidden or the link is down.                                                  |
| <b>RP mismatch</b>                    | Number of PIM control packets received for which the routing device has an RP mismatch.                                                                                     |

Table 302: show pim statistics Output Fields (*continued*)

| Field Name                                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>RP mode mismatch</b>                    | RP mode (sparse or bidirectional) mismatches encountered when processing join and prune messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>RPF neighbor unknown</b>                | Number of PIM control packets received for which the routing device has an unknown RPF neighbor for the source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Rx Joins/Prunes filtered</b>            | The number of join and prune messages filtered because of configured route filters and source address filters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Tx Joins/Prunes filtered</b>            | The number of join and prune messages filtered because of configured route filters and source address filters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Embedded-RP invalid addr</b>            | Number of packets received with an invalid embedded RP address in PIM join messages and other types of messages sent between routing domains.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Embedded-RP limit exceed</b>            | Number of times the limit configured with the <b>maximum-rps</b> statement is exceeded. The <b>maximum-rps</b> statement limits the number of embedded RPs created in a specific routing instance. The range is from 1 through 500. The default is 100.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Embedded-RP added</b>                   | <p>Number of packets in which the embedded RP for IPv6 is added.</p> <p>The following receive events trigger extraction of an IPv6 embedded RP address on the routing device:</p> <ul style="list-style-type: none"> <li>• Multicast Listener Discovery (MLD) report for an embedded RP multicast group address</li> <li>• PIM join message with an embedded RP multicast group address</li> <li>• Static embedded RP multicast group address associated with an interface</li> <li>• Packets sent to an embedded RP multicast group address received on the DR</li> </ul> <p>An embedded RP node discovered through these receive events is added if it does not already exist on the routing platform.</p> |
| <b>Embedded-RP removed</b>                 | Number of packets in which the embedded RP for IPv6 is removed. The embedded RP is removed whenever all PIM join states using this RP are removed or the configuration changes to remove the embedded RP feature.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Rx Register msgs filtering drop</b>     | Number of received register messages dropped because of a filter configured for PIM register messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Tx Register msgs filtering drop</b>     | Number of register messages dropped because of a filter configured for PIM register messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Rx Bidir Join/Prune on non-Bidir if</b> | Error counter for join and prune messages received on non-bidirectional PIM interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

Table 302: show pim statistics Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                                                   |
|-----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Rx Bidir Join/Prune on non-DF if</b> | Error counter for join and prune messages received on non-designated forwarder interfaces.                                                                                          |
| <b>V4 (S,G) Maximum</b>                 | Maximum number of (S,G) IPv4 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted. |
| <b>V4 (S,G) Accepted</b>                | Number of accepted (S,G) IPv4 multicast routes.                                                                                                                                     |
| <b>V4 (S,G) Threshold</b>               | Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv4 multicast routes accepted by the device).                                            |
| <b>V4 (S,G) Log Interval</b>            | Time (in seconds) between consecutive log messages.                                                                                                                                 |
| <b>V6 (S,G) Maximum</b>                 | Maximum number of (S,G) IPv6 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted. |
| <b>V6 (S,G) Accepted</b>                | Number of accepted (S,G) IPv6 multicast routes.                                                                                                                                     |
| <b>V6 (S,G) Threshold</b>               | Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv6 multicast routes accepted by the device).                                            |
| <b>V6 (S,G) Log Interval</b>            | Time (in seconds) between consecutive log messages.                                                                                                                                 |
| <b>V4 (grp-prefix, RP) Maximum</b>      | Maximum number of group-to-rendezvous point (RP) IPv4 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.        |
| <b>V4 (grp-prefix, RP) Accepted</b>     | Number of accepted group-to-RP IPv4 multicast mappings.                                                                                                                             |
| <b>V4 (grp-prefix, RP) Threshold</b>    | Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv4 multicast mappings accepted by the device).                                    |
| <b>V4 (grp-prefix, RP) Log Interval</b> | Time (in seconds) between consecutive log messages.                                                                                                                                 |
| <b>V6 (grp-prefix, RP) Maximum</b>      | Maximum number of group-to RP IPv6 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.                           |
| <b>V6 (grp-prefix, RP) Accepted</b>     | Number of accepted group-to-RP IPv6 multicast mappings.                                                                                                                             |

Table 302: show pim statistics Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                                                                 |
|-----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>V6 (grp-prefix, RP) Threshold</b>    | Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv6 multicast mappings accepted by the device).                                                  |
| <b>V6 (grp-prefix, RP) Log Interval</b> | Time (in seconds) between consecutive log messages.                                                                                                                                               |
| <b>V4 Register Maximum</b>              | Maximum number of IPv4 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted.<br><br>You configure the register limits on the RP. |
| <b>V4 Register Accepted</b>             | Number of accepted IPv4 PIM registers.                                                                                                                                                            |
| <b>V4 Register Threshold</b>            | Threshold at which a warning message is logged (percentage of the maximum number of IPv4 PIM registers accepted by the device).                                                                   |
| <b>V4 Register Log Interval</b>         | Time (in seconds) between consecutive log messages.                                                                                                                                               |
| <b>V6 Register Maximum</b>              | Maximum number of IPv6 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted.<br><br>You configure the register limits on the RP. |
| <b>V6 Register Accepted</b>             | Number of accepted IPv6 PIM registers.                                                                                                                                                            |
| <b>V6 Register Threshold</b>            | Threshold at which a warning message is logged (percentage of the maximum number of IPv6 PIM registers accepted by the device).                                                                   |
| <b>V6 Register Log Interval</b>         | Time (in seconds) between consecutive log messages.                                                                                                                                               |

## Sample Output

### show pim statistics

```

user@host> show pim statistics
PIM Message type Received Sent Rx errors
V2 Hello 15 32 0
V2 Register 0 362 0
V2 Register Stop 483 0 0
V2 Join Prune 18 518 0
V2 Bootstrap 0 0 0
V2 Assert 0 0 0
V2 Graft 0 0 0
V2 Graft Ack 0 0 0
V2 Candidate RP 0 0 0
V2 State Refresh 0 0 0
V2 DF Election 0 0 0
V1 Query 0 0 0
V1 Register 0 0 0

```



|                       |   |   |   |
|-----------------------|---|---|---|
| V1 Register Stop      | 0 | 0 | 0 |
| V1 Join Prune         | 0 | 0 | 0 |
| V1 RP Reachability    | 0 | 0 | 0 |
| V1 Assert             | 0 | 0 | 0 |
| V1 Graft              | 0 | 0 | 0 |
| V1 Graft Ack          | 0 | 0 | 0 |
| AutoRP Announce       | 0 | 0 | 0 |
| AutoRP Mapping        | 0 | 0 | 0 |
| AutoRP Unknown type   | 0 |   |   |
| Anycast Register      | 0 | 0 | 0 |
| Anycast Register Stop | 0 | 0 | 0 |

## Global Statistics

|                                  |   |
|----------------------------------|---|
| Hello dropped on neighbor policy | 0 |
| Unknown type                     | 0 |
| V1 Unknown type                  | 0 |
| Unknown Version                  | 0 |
| Neighbor unknown                 | 0 |
| Bad Length                       | 0 |
| Bad Checksum                     | 0 |
| Bad Receive If                   | 0 |
| Rx Bad Data                      | 0 |
| Rx Intf disabled                 | 0 |
| Rx V1 Require V2                 | 0 |
| Rx V2 Require V1                 | 0 |
| Rx Register not RP               | 0 |
| Rx Register no route             | 0 |
| Rx Register no decap if          | 0 |
| Null Register Timeout            | 0 |
| RP Filtered Source               | 0 |
| Rx Unknown Reg Stop              | 0 |
| Rx Join/Prune no state           | 0 |
| Rx Join/Prune on upstream if     | 0 |
| Rx Join/Prune for invalid group  | 5 |
| Rx Join/Prune messages dropped   | 0 |
| Rx sparse join for dense group   | 0 |
| Rx Graft/Graft Ack no state      | 0 |
| Rx Graft on upstream if          | 0 |
| Rx CRP not BSR                   | 0 |
| Rx BSR when BSR                  | 0 |
| Rx BSR not RPF if                | 0 |
| Rx unknown hello opt             | 0 |
| Rx data no state                 | 0 |
| Rx RP no state                   | 0 |
| Rx aggregate                     | 0 |
| Rx malformed packet              | 0 |
| Rx illegal TTL                   | 0 |
| Rx illegal destination address   | 0 |
| No RP                            | 0 |
| No register encap if             | 0 |
| No route upstream                | 0 |
| Nexthop Unusable                 | 0 |
| RP mismatch                      | 0 |
| RP mode mismatch                 | 0 |
| RPF neighbor unknown             | 0 |
| Rx Joins/Prunes filtered         | 0 |
| Tx Joins/Prunes filtered         | 0 |
| Embedded-RP invalid addr         | 0 |
| Embedded-RP limit exceed         | 0 |
| Embedded-RP added                | 0 |

```
Embedded-RP removed 0
Rx Register msgs filtering drop 0
Tx Register msgs filtering drop 0
Rx Bidir Join/Prune on non-Bidir if 0
Rx Bidir Join/Prune on non-DF if 0
```

## Sample Output

**show pim statistics inet interface <interface-name>**

```
user@host> show pim statistics inet interface ge-0/3/0.0
Instance: PIM.master Family: INET
```

PIM Interface statistics for ge-0/3/0.0

| PIM Message type      | Received | Sent | Rx errors |
|-----------------------|----------|------|-----------|
| V2 Hello              | 0        | 4    | 0         |
| V2 Register           | 0        | 0    | 0         |
| V2 Register Stop      | 0        | 0    | 0         |
| V2 Join Prune         | 0        | 0    | 0         |
| V2 Bootstrap          | 0        | 0    | 0         |
| V2 Assert             | 0        | 0    | 0         |
| V2 Graft              | 0        | 0    | 0         |
| V2 Graft Ack          | 0        | 0    | 0         |
| V2 Candidate RP       | 0        | 0    | 0         |
| V1 Query              | 0        | 0    | 0         |
| V1 Register           | 0        | 0    | 0         |
| V1 Register Stop      | 0        | 0    | 0         |
| V1 Join Prune         | 0        | 0    | 0         |
| V1 RP Reachability    | 0        | 0    | 0         |
| V1 Assert             | 0        | 0    | 0         |
| V1 Graft              | 0        | 0    | 0         |
| V1 Graft Ack          | 0        | 0    | 0         |
| AutoRP Announce       | 0        | 0    | 0         |
| AutoRP Mapping        | 0        | 0    | 0         |
| AutoRP Unknown type   | 0        |      |           |
| Anycast Register      | 0        | 0    | 0         |
| Anycast Register Stop | 0        | 0    | 0         |

## Sample Output

**show pim statistics inet6 interface <interface-name>**

```
user@host> show pim statistics inet6 interface ge-0/3/0.0
Instance: PIM.master Family: INET6
```

PIM Interface statistics for ge-0/3/0.0

| PIM Message type      | Received | Sent | Rx errors |
|-----------------------|----------|------|-----------|
| V2 Hello              | 0        | 4    | 0         |
| V2 Register           | 0        | 0    | 0         |
| V2 Register Stop      | 0        | 0    | 0         |
| V2 Join Prune         | 0        | 0    | 0         |
| V2 Bootstrap          | 0        | 0    | 0         |
| V2 Assert             | 0        | 0    | 0         |
| V2 Graft              | 0        | 0    | 0         |
| V2 Graft Ack          | 0        | 0    | 0         |
| V2 Candidate RP       | 0        | 0    | 0         |
| Anycast Register      | 0        | 0    | 0         |
| Anycast Register Stop | 0        | 0    | 0         |

**show pim statistics instance <instance-name>**

```

user@host> show pim statistics instance VPN-A
PIM Message type Received Sent Rx errors
V2 Hello 31 37 0
V2 Register 0 0 0
V2 Register Stop 0 0 0
V2 Join Prune 0 16 0
V2 Bootstrap 0 0 0
V2 Assert 0 0 0
V2 Graft 0 0 0
V2 Graft Ack 0 0 0
V2 Candidate RP 0 0 0
V2 State Refresh 0 0 0
V2 DF Election 0 0 0
V1 Query 0 0 0
V1 Register 0 0 0
V1 Register Stop 0 0 0
V1 Join Prune 0 0 0
V1 RP Reachability 0 0 0
V1 Assert 0 0 0
V1 Graft 0 0 0
V1 Graft Ack 0 0 0
AutoRP Announce 0 0 0
AutoRP Mapping 0 0 0
AutoRP Unknown type 0 0 0
Anycast Register 0 0 0
Anycast Register Stop 0 0 0

```

**Global Statistics**

```

Hello dropped on neighbor policy 0
Unknown type 0
V1 Unknown type 0
Unknown Version 0
Neighbor unknown 0
Bad Length 0
Bad Checksum 0
Bad Receive If 0
Rx Bad Data 0
Rx Intf disabled 0
Rx V1 Require V2 0
Rx V2 Require V1 0
Rx Register not RP 0
Rx Register no route 0
Rx Register no decap if 0
Null Register Timeout 0
RP Filtered Source 0
Rx Unknown Reg Stop 0
Rx Join/Prune no state 0
Rx Join/Prune on upstream if 0
Rx Join/Prune for invalid group 0
Rx Join/Prune messages dropped 0
Rx sparse join for dense group 0
Rx Graft/Graft Ack no state 0
Rx Graft on upstream if 0
Rx CRP not BSR 0
Rx BSR when BSR 0
Rx BSR not RPF if 0
Rx unknown hello opt 0
Rx data no state 0

```

|                                     |     |
|-------------------------------------|-----|
| Rx RP no state                      | 0   |
| Rx aggregate                        | 0   |
| Rx malformed packet                 | 0   |
| Rx illegal TTL                      | 0   |
| Rx illegal destination address      | 0   |
| No RP                               | 0   |
| No register encap if                | 0   |
| No route upstream                   | 28  |
| Nexthop Unusable                    | 0   |
| RP mismatch                         | 0   |
| RP mode mismatch                    | 0   |
| RPF neighbor unknown                | 0   |
| Rx Joins/Prunes filtered            | 0   |
| Tx Joins/Prunes filtered            | 0   |
| Embedded-RP invalid addr            | 0   |
| Embedded-RP limit exceed            | 0   |
| Embedded-RP added                   | 0   |
| Embedded-RP removed                 | 0   |
| Rx Register msgs filtering drop     | 0   |
| Tx Register msgs filtering drop     | 0   |
| Rx Bidir Join/Prune on non-Bidir if | 0   |
| Rx Bidir Join/Prune on non-DF if    | 0   |
| V4 (S,G) Maximum                    | 10  |
| V4 (S,G) Accepted                   | 9   |
| V4 (S,G) Threshold                  | 80  |
| V4 (S,G) Log Interval               | 80  |
| V6 (S,G) Maximum                    | 8   |
| V6 (S,G) Accepted                   | 8   |
| V6 (S,G) Threshold                  | 50  |
| V6 (S,G) Log Interval               | 100 |
| V4 (grp-prefix, RP) Maximum         | 100 |
| V4 (grp-prefix, RP) Accepted        | 5   |
| V4 (grp-prefix, RP) Threshold       | 80  |
| V4 (grp-prefix, RP) Log Interval    | 10  |
| V6 (grp-prefix, RP) Maximum         | 20  |
| V6 (grp-prefix, RP) Accepted        | 0   |
| V6 (grp-prefix, RP) Threshold       | 90  |
| V6 (grp-prefix, RP) Log Interval    | 20  |
| V4 Register Maximum                 | 100 |
| V4 Register Accepted                | 10  |
| V4 Register Threshold               | 80  |
| V4 Register Log Interval            | 10  |
| V6 Register Maximum                 | 20  |
| V6 Register Accepted                | 0   |
| V6 Register Threshold               | 90  |
| V6 Register Log Interval            | 20  |

## Sample Output

**show pim statistics interface <interface-name>**

```
user@host> show pim statistics interface ge-0/3/0.0
Instance: PIM.master Family: INET
```

PIM Interface statistics for ge-0/3/0.0

| PIM Message type | Received | Sent | Rx errors |
|------------------|----------|------|-----------|
| V2 Hello         | 0        | 3    | 0         |
| V2 Register      | 0        | 0    | 0         |
| V2 Register Stop | 0        | 0    | 0         |
| V2 Join Prune    | 0        | 0    | 0         |

|                       |   |   |   |
|-----------------------|---|---|---|
| V2 Bootstrap          | 0 | 0 | 0 |
| V2 Assert             | 0 | 0 | 0 |
| V2 Graft              | 0 | 0 | 0 |
| V2 Graft Ack          | 0 | 0 | 0 |
| V2 Candidate RP       | 0 | 0 | 0 |
| V1 Query              | 0 | 0 | 0 |
| V1 Register           | 0 | 0 | 0 |
| V1 Register Stop      | 0 | 0 | 0 |
| V1 Join Prune         | 0 | 0 | 0 |
| V1 RP Reachability    | 0 | 0 | 0 |
| V1 Assert             | 0 | 0 | 0 |
| V1 Graft              | 0 | 0 | 0 |
| V1 Graft Ack          | 0 | 0 | 0 |
| AutoRP Announce       | 0 | 0 | 0 |
| AutoRP Mapping        | 0 | 0 | 0 |
| AutoRP Unknown type   | 0 |   |   |
| Anycast Register      | 0 | 0 | 0 |
| Anycast Register Stop | 0 | 0 | 0 |

Instance: PIM.master Family: INET6

PIM Interface statistics for ge-0/3/0.0

| PIM Message type      | Received | Sent | Rx errors |
|-----------------------|----------|------|-----------|
| V2 Hello              | 0        | 3    | 0         |
| V2 Register           | 0        | 0    | 0         |
| V2 Register Stop      | 0        | 0    | 0         |
| V2 Join Prune         | 0        | 0    | 0         |
| V2 Bootstrap          | 0        | 0    | 0         |
| V2 Assert             | 0        | 0    | 0         |
| V2 Graft              | 0        | 0    | 0         |
| V2 Graft Ack          | 0        | 0    | 0         |
| V2 Candidate RP       | 0        | 0    | 0         |
| Anycast Register      | 0        | 0    | 0         |
| Anycast Register Stop | 0        | 0    | 0         |



# Network Management and Monitoring

- [Overview on page 4073](#)
- [Configuration on page 4074](#)
- [Administration on page 4116](#)

## Overview

---

- [Ethernet OAM Link Fault Management on page 4073](#)

## Ethernet OAM Link Fault Management

- [IEEE 802.3ah OAM Link-Fault Management Overview on page 4073](#)

### [IEEE 802.3ah OAM Link-Fault Management Overview](#)

---

Ethernet interfaces capable of running at 100 Mbps or faster on EX Series switches, MX Series, M Series (except M5 and M10 routers), and T Series routers support the IEEE 802.3ah standard for Operation, Administration, and Management (OAM). You can configure IEEE 802.3ah OAM on Ethernet point-to-point direct links or links across Ethernet repeaters. The IEEE 802.3ah standard meets the requirement for OAM capabilities as Ethernet moves from being solely an enterprise technology to being a WAN and access technology, as well as being backward-compatible with existing Ethernet technology. Junos OS supports IEEE 802.3ah link-fault management.

The features of link-fault management are:

- Discovery
- Link monitoring
- Remote fault detection
- Remote loopback

The following features are not supported:

- Ethernet running on top of a Layer 2 protocol, such as Ethernet over ATM, is not supported in OAM configurations.
- Remote loopback is not supported on the 10-Gigabit Ethernet LAN/WAN PIC with SFP+.

- The remote loopback feature mentioned in section 57.2.11 of IEEE 802.3ah is not supported on T4000 routers.



**NOTE:** Aggregated Ethernet member links will now use the physical MAC address as the source MAC address in 802.3ah OAM packets.

#### Related Documentation

- [Configuring IEEE 802.3ah OAM Link-Fault Management on page 4075](#)
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- [Configuring Link Discovery on page 4077](#)
- [Configuring the OAM PDU Interval on page 4078](#)
- [Configuring the OAM PDU Threshold on page 4079](#)
- [Configuring Threshold Values for Local Fault Events on an Interface on page 4079](#)
- [Disabling the Sending of Link Event TLVs on page 4080](#)
- [Detecting Remote Faults on page 4081](#)
- [Configuring an OAM Action Profile on page 4082](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083](#)
- [Monitoring the Loss of Link Adjacency on page 4085](#)
- [Monitoring Protocol Status on page 4086](#)
- [Configuring Threshold Values for Fault Events in an Action Profile on page 4087](#)
- [Applying an Action Profile on page 4088](#)
- [Setting a Remote Interface into Loopback Mode on page 4088](#)
- [Enabling Remote Loopback Support on the Local Interface on page 4089](#)
- [Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 4090](#)
- [\*Junos® OS Ethernet Interfaces\*](#)

## Configuration

---

- [Configuration: Ethernet OAM Link Fault Management on page 4074](#)
- [Configuration Statements on page 4092](#)

### Configuration: Ethernet OAM Link Fault Management

- [Configuring IEEE 802.3ah OAM Link-Fault Management on page 4075](#)
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- [Configuring Link Discovery on page 4077](#)
- [Configuring the OAM PDU Interval on page 4078](#)
- [Configuring the OAM PDU Threshold on page 4079](#)



- [Configuring Threshold Values for Local Fault Events on an Interface on page 4079](#)
- [Disabling the Sending of Link Event TLVs on page 4080](#)
- [Detecting Remote Faults on page 4081](#)
- [Configuring an OAM Action Profile on page 4082](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083](#)
- [Monitoring the Loss of Link Adjacency on page 4085](#)
- [Monitoring Protocol Status on page 4086](#)
- [Configuring Threshold Values for Fault Events in an Action Profile on page 4087](#)
- [Applying an Action Profile on page 4088](#)
- [Setting a Remote Interface into Loopback Mode on page 4088](#)
- [Enabling Remote Loopback Support on the Local Interface on page 4089](#)
- [Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 4090](#)

### [Configuring IEEE 802.3ah OAM Link-Fault Management](#)

You can configure threshold values for fault events that trigger the sending of link event TLVs when the values exceed the threshold. To set threshold values for fault events on an interface, include the **event-thresholds** statement at the **[edit protocols oam ethernet link-fault-management interface]** hierarchy level.

You can also configure OAM threshold values within an action profile and apply the action profile to multiple interfaces. To create an action profile, include the **action-profile** statement at the **[edit protocols oam ethernet link-fault-management]** hierarchy level.

You can configure Ethernet OAM either on an aggregate interface or on each of its member links. However, we recommend that you configure Ethernet OAM on the aggregate interface, and this will internally enable Ethernet OAM on the member links.

To view OAM statistics, use the **show oam ethernet link-fault-management** operational mode command. To clear OAM statistics, use the **clear oam ethernet link-fault-management statistics** operational mode command. To clear link-fault management state information and restart the link discovery process on Ethernet interfaces, use the **clear oam ethernet link-fault-management state** operational mode command. For more information about these commands, see the *Junos OS Operational Mode Commands*.

#### Related Documentation

- [event-thresholds on page 4101](#)
- [action-profile](#)
- [IEEE 802.3ah OAM Link-Fault Management Overview on page 4073](#)
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- [Configuring Link Discovery on page 4077](#)
- [Configuring the OAM PDU Interval on page 4078](#)
- [Configuring the OAM PDU Threshold on page 4079](#)
- [Configuring Threshold Values for Local Fault Events on an Interface on page 4079](#)

- [Disabling the Sending of Link Event TLVs on page 4080](#)
- [Detecting Remote Faults on page 4081](#)
- [Configuring an OAM Action Profile on page 4082](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083](#)
- [Monitoring the Loss of Link Adjacency on page 4085](#)
- [Monitoring Protocol Status on page 4086](#)
- [Configuring Threshold Values for Fault Events in an Action Profile on page 4087](#)
- [Applying an Action Profile on page 4088](#)
- [Setting a Remote Interface into Loopback Mode on page 4088](#)
- [Enabling Remote Loopback Support on the Local Interface on page 4089](#)
- [Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 4090](#)
- [\*Junos® OS Ethernet Interfaces\*](#)

---

### Enabling IEEE 802.3ah OAM Support

To enable IEEE 802.3ah OAM support, include the **interface** statement at the **[edit protocols oam ethernet link-fault-management]** hierarchy level:

**[edit protocols oam ethernet link-fault-management interface interface-name]**

When you enable IEEE 802.3ah OAM on a physical interface, the discovery process is automatically triggered.

#### Related Documentation

- [link-fault-management on page 4093](#)
- [IEEE 802.3ah OAM Link-Fault Management Overview on page 4073](#)
- [Configuring IEEE 802.3ah OAM Link-Fault Management on page 4075](#)
- [Configuring Link Discovery on page 4077](#)
- [Configuring the OAM PDU Interval on page 4078](#)
- [Configuring the OAM PDU Threshold on page 4079](#)
- [Configuring Threshold Values for Local Fault Events on an Interface on page 4079](#)
- [Disabling the Sending of Link Event TLVs on page 4080](#)
- [Detecting Remote Faults on page 4081](#)
- [Configuring an OAM Action Profile on page 4082](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083](#)
- [Monitoring the Loss of Link Adjacency on page 4085](#)
- [Monitoring Protocol Status on page 4086](#)
- [Configuring Threshold Values for Fault Events in an Action Profile on page 4087](#)
- [Applying an Action Profile on page 4088](#)

- [Setting a Remote Interface into Loopback Mode on page 4088](#)
- [Enabling Remote Loopback Support on the Local Interface on page 4089](#)
- [Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 4090](#)
- *Junos® OS Ethernet Interfaces*

### Configuring Link Discovery

When the IEEE 802.3ah OAM protocol is enabled on a physical interface, the discovery process is automatically triggered. The discovery process permits Ethernet interfaces to discover and monitor the peer on the link if it also supports the IEEE 802.3ah standard.

You can specify the discovery mode used for IEEE 802.3ah OAM support. The discovery process is triggered automatically when OAM IEEE 802.3ah functionality is enabled on a port. Link monitoring is done when the interface sends periodic OAM PDUs.

To configure the discovery mode, include the **link-discovery** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
 link-discovery (active | passive);
```

In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. After the discovery process has been initiated, both sides participate in discovery.

#### Related Documentation

- [link-discovery on page 4108](#)
- [IEEE 802.3ah OAM Link-Fault Management Overview on page 4073](#)
- [Configuring IEEE 802.3ah OAM Link-Fault Management on page 4075](#)
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- [Configuring the OAM PDU Interval on page 4078](#)
- [Configuring the OAM PDU Threshold on page 4079](#)
- [Configuring Threshold Values for Local Fault Events on an Interface on page 4079](#)
- [Disabling the Sending of Link Event TLVs on page 4080](#)
- [Detecting Remote Faults on page 4081](#)
- [Configuring an OAM Action Profile on page 4082](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083](#)
- [Monitoring the Loss of Link Adjacency on page 4085](#)
- [Monitoring Protocol Status on page 4086](#)
- [Configuring Threshold Values for Fault Events in an Action Profile on page 4087](#)
- [Applying an Action Profile on page 4088](#)
- [Setting a Remote Interface into Loopback Mode on page 4088](#)

- [Enabling Remote Loopback Support on the Local Interface on page 4089](#)
- [Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 4090](#)
- [Junos® OS Ethernet Interfaces](#)

### [Configuring the OAM PDU Interval](#)

---

Periodic OAM PDUs are sent to perform link monitoring.

You can specify the periodic OAM PDU sending interval for fault detection.

To configure the sending interval, include the **pdu-interval** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
 pdu-interval interval;
```

The periodic OAM PDU interval range is from 100 through 1000 milliseconds. The default sending interval is 1000 milliseconds.

#### Related Documentation

- [pdu-interval on page 4112](#)
- [IEEE 802.3ah OAM Link-Fault Management Overview on page 4073](#)
- [Configuring IEEE 802.3ah OAM Link-Fault Management on page 4075](#)
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- [Configuring Link Discovery on page 4077](#)
- [Configuring the OAM PDU Threshold on page 4079](#)
- [Configuring Threshold Values for Local Fault Events on an Interface on page 4079](#)
- [Disabling the Sending of Link Event TLVs on page 4080](#)
- [Detecting Remote Faults on page 4081](#)
- [Configuring an OAM Action Profile on page 4082](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083](#)
- [Monitoring the Loss of Link Adjacency on page 4085](#)
- [Monitoring Protocol Status on page 4086](#)
- [Configuring Threshold Values for Fault Events in an Action Profile on page 4087](#)
- [Applying an Action Profile on page 4088](#)
- [Setting a Remote Interface into Loopback Mode on page 4088](#)
- [Enabling Remote Loopback Support on the Local Interface on page 4089](#)
- [Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 4090](#)
- [Junos® OS Ethernet Interfaces](#)

### Configuring the OAM PDU Threshold

You can specify the number of OAM PDUs that an interface can miss before the link between peers is considered down.

To configure the number of PDUs that can be missed from the peer, include the **pdu-threshold** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
 pdu-threshold threshold-value;
```

The threshold value range is from 3 through 10. The default is three PDUs.

#### Related Documentation

- [pdu-threshold on page 4113](#)
- [IEEE 802.3ah OAM Link-Fault Management Overview on page 4073](#)
- [Configuring IEEE 802.3ah OAM Link-Fault Management on page 4075](#)
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- [Configuring Link Discovery on page 4077](#)
- [Configuring the OAM PDU Interval on page 4078](#)
- [Configuring Threshold Values for Local Fault Events on an Interface on page 4079](#)
- [Disabling the Sending of Link Event TLVs on page 4080](#)
- [Detecting Remote Faults on page 4081](#)
- [Configuring an OAM Action Profile on page 4082](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083](#)
- [Monitoring the Loss of Link Adjacency on page 4085](#)
- [Monitoring Protocol Status on page 4086](#)
- [Configuring Threshold Values for Fault Events in an Action Profile on page 4087](#)
- [Applying an Action Profile on page 4088](#)
- [Setting a Remote Interface into Loopback Mode on page 4088](#)
- [Enabling Remote Loopback Support on the Local Interface on page 4089](#)
- [Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 4090](#)
- *Junos® OS Ethernet Interfaces*

### Configuring Threshold Values for Local Fault Events on an Interface

You can configure threshold values on an interface for the local errors that trigger the sending of link event TLVs.

To set the error threshold values for sending event TLVs, include the **frame-error**, **frame-period**, **frame-period-summary**, and **symbol-period** statements at the **[edit protocols**

**oam ethernet link-fault-management interface *interface-name* event-thresholds]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
event-thresholds {
 frame-error count;
 frame-period count;
 frame-period-summary count;
 symbol-period count;
}
```

**Related  
Documentation**

- [event-thresholds on page 4101](#)
- [frame-error on page 4103](#)
- [frame-period on page 4104](#)
- [frame-period-summary on page 4105](#)
- [symbol-period on page 4115](#)
- [IEEE 802.3ah OAM Link-Fault Management Overview on page 4073](#)
- [Configuring IEEE 802.3ah OAM Link-Fault Management on page 4075](#)
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- [Configuring Link Discovery on page 4077](#)
- [Configuring the OAM PDU Interval on page 4078](#)
- [Configuring the OAM PDU Threshold on page 4079](#)
- [Disabling the Sending of Link Event TLVs on page 4080](#)
- [Detecting Remote Faults on page 4081](#)
- [Configuring an OAM Action Profile on page 4082](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083](#)
- [Monitoring the Loss of Link Adjacency on page 4085](#)
- [Monitoring Protocol Status on page 4086](#)
- [Configuring Threshold Values for Fault Events in an Action Profile on page 4087](#)
- [Applying an Action Profile on page 4088](#)
- [Setting a Remote Interface into Loopback Mode on page 4088](#)
- [Enabling Remote Loopback Support on the Local Interface on page 4089](#)
- [Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 4090](#)
- [Junos® OS Ethernet Interfaces](#)

---

### Disabling the Sending of Link Event TLVs

You can disable the sending of link event TLVs.

To disable the monitoring and sending of PDUs containing link event TLVs in periodic PDUs, include the **no-allow-link-events** statement at the **[edit protocols oam ethernet link-fault-management interface *interface-name* negotiation-options]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name
 negotiation-options]
no-allow-link-events;
```

#### Related Documentation

- [no-allow-link-events on page 4109](#)
- [IEEE 802.3ah OAM Link-Fault Management Overview on page 4073](#)
- [Configuring IEEE 802.3ah OAM Link-Fault Management on page 4075](#)
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- [Configuring Link Discovery on page 4077](#)
- [Configuring the OAM PDU Interval on page 4078](#)
- [Configuring the OAM PDU Threshold on page 4079](#)
- [Configuring Threshold Values for Local Fault Events on an Interface on page 4079](#)
- [Detecting Remote Faults on page 4081](#)
- [Configuring an OAM Action Profile on page 4082](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083](#)
- [Monitoring the Loss of Link Adjacency on page 4085](#)
- [Monitoring Protocol Status on page 4086](#)
- [Configuring Threshold Values for Fault Events in an Action Profile on page 4087](#)
- [Applying an Action Profile on page 4088](#)
- [Setting a Remote Interface into Loopback Mode on page 4088](#)
- [Enabling Remote Loopback Support on the Local Interface on page 4089](#)
- [Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 4090](#)
- *Junos® OS Ethernet Interfaces*

#### Detecting Remote Faults

Fault detection is either based on flags or fault event type, length, and values (TLVs) received in OAM protocol data units (PDUs). Flags that trigger a link fault are:

- Critical Event
- Dying Gasp
- Link Fault

The link event TLVs are sent by the remote DTE by means of event notification PDUs. Link event TLVs are:

- Errored Symbol Period Event
- Errored Frame Event
- Errored Frame Period Event
- Errored Frame Seconds Summary Event

**Related Documentation**

- [IEEE 802.3ah OAM Link-Fault Management Overview on page 4073](#)
- [Configuring IEEE 802.3ah OAM Link-Fault Management on page 4075](#)
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- [Configuring Link Discovery on page 4077](#)
- [Configuring the OAM PDU Interval on page 4078](#)
- [Configuring the OAM PDU Threshold on page 4079](#)
- [Configuring Threshold Values for Local Fault Events on an Interface on page 4079](#)
- [Disabling the Sending of Link Event TLVs on page 4080](#)
- [Configuring an OAM Action Profile on page 4082](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083](#)
- [Monitoring the Loss of Link Adjacency on page 4085](#)
- [Monitoring Protocol Status on page 4086](#)
- [Configuring Threshold Values for Fault Events in an Action Profile on page 4087](#)
- [Applying an Action Profile on page 4088](#)
- [Setting a Remote Interface into Loopback Mode on page 4088](#)
- [Enabling Remote Loopback Support on the Local Interface on page 4089](#)
- [Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 4090](#)
- *Junos® OS Ethernet Interfaces*

---

### Configuring an OAM Action Profile

You can create an action profile to define event fault flags and thresholds and the action to be taken. You can then apply the action profile to one or more interfaces.

To configure an action profile, include the **action-profile** statement at the **[edit protocols oam ethernet link-fault-management]** hierarchy level:

```
action-profile profile-name {
 action {
 syslog;
 link-down;
 send-critical-event;
 }
 event {
 link-adjacency-loss;
 link-event-rate {
 10;
 }
 }
}
```



```

 frame-error count;
 frame-period count;
 frame-period-summary count;
 symbol-period count;
 }
 protocol-down;
}

```

#### Related Documentation

- [action-profile on page 4096](#)
- [IEEE 802.3ah OAM Link-Fault Management Overview on page 4073](#)
- [Configuring IEEE 802.3ah OAM Link-Fault Management on page 4075](#)
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- [Configuring Link Discovery on page 4077](#)
- [Configuring the OAM PDU Interval on page 4078](#)
- [Configuring the OAM PDU Threshold on page 4079](#)
- [Configuring Threshold Values for Local Fault Events on an Interface on page 4079](#)
- [Disabling the Sending of Link Event TLVs on page 4080](#)
- [Detecting Remote Faults on page 4081](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083](#)
- [Monitoring the Loss of Link Adjacency on page 4085](#)
- [Monitoring Protocol Status on page 4086](#)
- [Configuring Threshold Values for Fault Events in an Action Profile on page 4087](#)
- [Applying an Action Profile on page 4088](#)
- [Setting a Remote Interface into Loopback Mode on page 4088](#)
- [Enabling Remote Loopback Support on the Local Interface on page 4089](#)
- [Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 4090](#)
- *Junos® OS Ethernet Interfaces*

#### Specifying the Actions to Be Taken for Link-Fault Management Events

You can specify the action to be taken by the system when the configured link-fault event occurs. Multiple action profiles can be applied to a single interface. For each action-profile, at least one event and one action must be specified. The actions are taken only when all of the events in the action profile are true. If more than one action is specified, all the actions are executed.

You might want to set a lower threshold for a specific action such as logging the error and set a higher threshold for another action such as sending a critical event TLV.

To specify the action, include the **action** statement at the **[edit protocols oam ethernet link-fault-management action-profile *profile-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management action-profile profile-name]
event {
 link-adjacency-loss;
 protocol-down;
}
action {
 syslog;
 link-down;
 send-critical-event;
}
```

To create a system log entry when the link-fault event occurs, include the **syslog** statement.

To administratively disable the link when the link-fault event occurs, include the **link-down** statement.

To send IEEE 802.3ah link event TLVs in the OAM PDU when a link-fault event occurs, include the **send-critical-event** statement.



**NOTE:** If multiple actions are specified in the action profile, all of the actions are executed in no particular order.

#### Related Documentation

- [action on page 4095](#)
- [syslog on page 4115](#)
- [link-down on page 4107](#)
- [send-critical-event on page 4114](#)
- [IEEE 802.3ah OAM Link-Fault Management Overview on page 4073](#)
- [Configuring IEEE 802.3ah OAM Link-Fault Management on page 4075](#)
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- [Configuring Link Discovery on page 4077](#)
- [Configuring the OAM PDU Interval on page 4078](#)
- [Configuring the OAM PDU Threshold on page 4079](#)
- [Configuring Threshold Values for Local Fault Events on an Interface on page 4079](#)
- [Disabling the Sending of Link Event TLVs on page 4080](#)
- [Detecting Remote Faults on page 4081](#)
- [Configuring an OAM Action Profile on page 4082](#)
- [Monitoring the Loss of Link Adjacency on page 4085](#)
- [Monitoring Protocol Status on page 4086](#)
- [Configuring Threshold Values for Fault Events in an Action Profile on page 4087](#)
- [Applying an Action Profile on page 4088](#)

- [Setting a Remote Interface into Loopback Mode on page 4088](#)
- [Enabling Remote Loopback Support on the Local Interface on page 4089](#)
- [Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 4090](#)
- *Junos® OS Ethernet Interfaces*

### Monitoring the Loss of Link Adjacency

You can specify actions be taken when link adjacency is lost. When link adjacency is lost, the system takes the action defined in the **action** statement of the action profile.

To configure the system to take action when link adjacency is lost, include the **link-adjacency-loss** statement at the **[edit protocols oam ethernet link-fault-management action-profile *profile-name* event]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management action-profile profile-name]
link-adjacency-loss;
```

#### Related Documentation

- [link-adjacency-loss on page 4107](#)
- [IEEE 802.3ah OAM Link-Fault Management Overview on page 4073](#)
- [Configuring IEEE 802.3ah OAM Link-Fault Management on page 4075](#)
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- [Configuring Link Discovery on page 4077](#)
- [Configuring the OAM PDU Interval on page 4078](#)
- [Configuring the OAM PDU Threshold on page 4079](#)
- [Configuring Threshold Values for Local Fault Events on an Interface on page 4079](#)
- [Disabling the Sending of Link Event TLVs on page 4080](#)
- [Detecting Remote Faults on page 4081](#)
- [Configuring an OAM Action Profile on page 4082](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083](#)
- [Monitoring Protocol Status on page 4086](#)
- [Configuring Threshold Values for Fault Events in an Action Profile on page 4087](#)
- [Applying an Action Profile on page 4088](#)
- [Setting a Remote Interface into Loopback Mode on page 4088](#)
- [Enabling Remote Loopback Support on the Local Interface on page 4089](#)
- [Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 4090](#)
- *Junos® OS Ethernet Interfaces*

## Monitoring Protocol Status

---

The CCC-DOWN flag is associated with a circuit cross-connect (CCC) connection, Layer 2 circuit, and Layer 2 VPN, which send the CCC-DOWN status to the kernel. The CCC-DOWN flag indicates that the CCC is down. The CCC-DOWN status is sent to the kernel when the CCC connection, Layer 2 circuit, or Layer 2 VPN is down. This in turn, brings down the CE-facing PE interface associated with the CCC connection, Layer 2 circuit, or Layer 2 VPN.

When the CCC-DOWN flag is signaled to the IEEE 802.3ah protocol, the system takes the action defined in the **action** statement of the action profile. For additional information about Layer 2 circuits, see the Junos OS Layer 2 Circuits Feature Guide, Junos OS VPNs Configuration Guide.

To monitor the IEEE 802.3ah protocol, on the CE-facing PE interface, include the **protocol-down** statement at the **[edit protocols oam ethernet link-fault-management action-profile *profile-name* event]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management action-profile profile-name]
protocol-down;
```

---



**NOTE:** If multiple events are specified in the action profile, all the events must occur before the specified action is taken.

---

### Related Documentation

- [protocol-down on page 4113](#)
- [IEEE 802.3ah OAM Link-Fault Management Overview on page 4073](#)
- [Configuring IEEE 802.3ah OAM Link-Fault Management on page 4075](#)
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- [Configuring Link Discovery on page 4077](#)
- [Configuring the OAM PDU Interval on page 4078](#)
- [Configuring the OAM PDU Threshold on page 4079](#)
- [Configuring Threshold Values for Local Fault Events on an Interface on page 4079](#)
- [Disabling the Sending of Link Event TLVs on page 4080](#)
- [Detecting Remote Faults on page 4081](#)
- [Configuring an OAM Action Profile on page 4082](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083](#)
- [Monitoring the Loss of Link Adjacency on page 4085](#)
- [Configuring Threshold Values for Fault Events in an Action Profile on page 4087](#)
- [Applying an Action Profile on page 4088](#)
- [Setting a Remote Interface into Loopback Mode on page 4088](#)
- [Enabling Remote Loopback Support on the Local Interface on page 4089](#)

- [Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 4090](#)
- [Junos® OS Ethernet Interfaces](#)

### Configuring Threshold Values for Fault Events in an Action Profile

You can configure link event thresholds for received error events that trigger the action specified in the **action** statement. You can then apply the action profile to one or more interfaces.

To configure link event thresholds, include the **link-event-rate** statement at the **[edit protocols oam ethernet link-fault-management action-profile *profile-name* event]** hierarchy level:

```
link-event-rate {
 frame-error count;
 frame-period count;
 frame-period-summary count;
 symbol-period count;
}
```

#### Related Documentation

- [link-event-rate on page 4108](#)
- [IEEE 802.3ah OAM Link-Fault Management Overview on page 4073](#)
- [Configuring IEEE 802.3ah OAM Link-Fault Management on page 4075](#)
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- [Configuring Link Discovery on page 4077](#)
- [Configuring the OAM PDU Interval on page 4078](#)
- [Configuring the OAM PDU Threshold on page 4079](#)
- [Configuring Threshold Values for Local Fault Events on an Interface on page 4079](#)
- [Disabling the Sending of Link Event TLVs on page 4080](#)
- [Detecting Remote Faults on page 4081](#)
- [Configuring an OAM Action Profile on page 4082](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083](#)
- [Monitoring the Loss of Link Adjacency on page 4085](#)
- [Monitoring Protocol Status on page 4086](#)
- [Applying an Action Profile on page 4088](#)
- [Setting a Remote Interface into Loopback Mode on page 4088](#)
- [Enabling Remote Loopback Support on the Local Interface on page 4089](#)
- [Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 4090](#)
- [Junos® OS Ethernet Interfaces](#)

## Applying an Action Profile

---

You can apply an action profile to one or more interfaces.

To apply an action profile to an interface, include the **apply-action-profile** statement at the **[edit protocols oam ethernet link-fault-management action-profile interface *interface-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
 apply-action-profile profile-name;
```

### Related Documentation

- [apply-action-profile on page 4097](#)
- [IEEE 802.3ah OAM Link-Fault Management Overview on page 4073](#)
- [Configuring IEEE 802.3ah OAM Link-Fault Management on page 4075](#)
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- [Configuring Link Discovery on page 4077](#)
- [Configuring the OAM PDU Interval on page 4078](#)
- [Configuring the OAM PDU Threshold on page 4079](#)
- [Configuring Threshold Values for Local Fault Events on an Interface on page 4079](#)
- [Disabling the Sending of Link Event TLVs on page 4080](#)
- [Detecting Remote Faults on page 4081](#)
- [Configuring an OAM Action Profile on page 4082](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083](#)
- [Monitoring the Loss of Link Adjacency on page 4085](#)
- [Monitoring Protocol Status on page 4086](#)
- [Configuring Threshold Values for Fault Events in an Action Profile on page 4087](#)
- [Setting a Remote Interface into Loopback Mode on page 4088](#)
- [Enabling Remote Loopback Support on the Local Interface on page 4089](#)
- [Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 4090](#)
- [Junos® OS Ethernet Interfaces](#)

## Setting a Remote Interface into Loopback Mode

---

You can configure the software to set the remote DTE into loopback mode on the following interfaces:

- IQ2 and IQ2-E Gigabit Ethernet interfaces
- Ethernet interfaces on the MX Series routers or EX Series switches

Junos OS can place a remote DTE into loopback mode (if remote-loopback mode is supported by the remote DTE). When you place a remote DTE into loopback mode, the

interface receives the remote-loopback request and puts the interface into remote-loopback mode. When the interface is in remote-loopback mode, all frames except OAM PDUs are looped back without any changes made to the frames. OAM PDUs continue to be sent to the management plane and processed.

To configure remote loopback, include the **remote-loopback** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
 remote-loopback;
```

To take the remote DTE out of loopback mode, remove the **remote-loopback** statement from the configuration.

#### Related Documentation

- [remote-loopback on page 4114](#)
- [IEEE 802.3ah OAM Link-Fault Management Overview on page 4073](#)
- [Configuring IEEE 802.3ah OAM Link-Fault Management on page 4075](#)
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- [Configuring Link Discovery on page 4077](#)
- [Configuring the OAM PDU Interval on page 4078](#)
- [Configuring the OAM PDU Threshold on page 4079](#)
- [Configuring Threshold Values for Local Fault Events on an Interface on page 4079](#)
- [Disabling the Sending of Link Event TLVs on page 4080](#)
- [Detecting Remote Faults on page 4081](#)
- [Configuring an OAM Action Profile on page 4082](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083](#)
- [Monitoring the Loss of Link Adjacency on page 4085](#)
- [Monitoring Protocol Status on page 4086](#)
- [Configuring Threshold Values for Fault Events in an Action Profile on page 4087](#)
- [Applying an Action Profile on page 4088](#)
- [Enabling Remote Loopback Support on the Local Interface on page 4089](#)
- [Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 4090](#)
- *Junos® OS Ethernet Interfaces*

#### Enabling Remote Loopback Support on the Local Interface

You can allow a remote DTE to set a local interface into remote loopback mode on IQ2 and IQ2-E Gigabit Ethernet interfaces and all Ethernet interfaces on the MX Series routers and EX Series switches. When a remote-loopback request is sent by a remote DTE, the Junos OS places the local interface into loopback mode. When an interface is in loopback mode, all frames except OAM PDUs are looped back without any changes to the frames.

OAM PDUs continue to be sent to the management plane and processed. By default, the remote loopback feature is not enabled.

To enable remote loopback, include the **allow-remote-loopback** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name* negotiation-options]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name
 negotiation-options]
allow-remote-loopback;
```



**NOTE:** Activation of OAM remote loopback may result in data frame loss.

#### Related Documentation

- [allow-remote-loopback on page 4097](#)
- [IEEE 802.3ah OAM Link-Fault Management Overview on page 4073](#)
- [Configuring IEEE 802.3ah OAM Link-Fault Management on page 4075](#)
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- [Configuring Link Discovery on page 4077](#)
- [Configuring the OAM PDU Interval on page 4078](#)
- [Configuring the OAM PDU Threshold on page 4079](#)
- [Configuring Threshold Values for Local Fault Events on an Interface on page 4079](#)
- [Disabling the Sending of Link Event TLVs on page 4080](#)
- [Detecting Remote Faults on page 4081](#)
- [Configuring an OAM Action Profile on page 4082](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083](#)
- [Monitoring the Loss of Link Adjacency on page 4085](#)
- [Monitoring Protocol Status on page 4086](#)
- [Configuring Threshold Values for Fault Events in an Action Profile on page 4087](#)
- [Applying an Action Profile on page 4088](#)
- [Setting a Remote Interface into Loopback Mode on page 4088](#)
- [Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 4090](#)
- [Junos® OS Ethernet Interfaces](#)

---

#### Example: Configuring IEEE 802.3ah OAM Support on an Interface

Configure 802.3ah OAM support on a 10-Gigabit Ethernet interface:

```
[edit]
protocols {
 oam {
```



```

ethernet {
 link-fault-management {
 interface xe-0/0/0 {
 link-discovery active;
 pdu-interval 800;
 pdu-threshold 4;
 remote-loopback;
 negotiation-options {
 allow-remote-loopback;
 }
 event-thresholds {
 frame-error 30;
 frame-period 50;
 frame-period summary 40;
 symbol-period 20;
 }
 }
 }
}

```

#### Related Documentation

- [link-fault-management on page 4093](#)
- [IEEE 802.3ah OAM Link-Fault Management Overview on page 4073](#)
- [Configuring IEEE 802.3ah OAM Link-Fault Management on page 4075](#)
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- [Configuring Link Discovery on page 4077](#)
- [Configuring the OAM PDU Interval on page 4078](#)
- [Configuring the OAM PDU Threshold on page 4079](#)
- [Configuring Threshold Values for Local Fault Events on an Interface on page 4079](#)
- [Disabling the Sending of Link Event TLVs on page 4080](#)
- [Detecting Remote Faults on page 4081](#)
- [Configuring an OAM Action Profile on page 4082](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083](#)
- [Monitoring the Loss of Link Adjacency on page 4085](#)
- [Monitoring Protocol Status on page 4086](#)
- [Configuring Threshold Values for Fault Events in an Action Profile on page 4087](#)
- [Applying an Action Profile on page 4088](#)
- [Setting a Remote Interface into Loopback Mode on page 4088](#)
- [Enabling Remote Loopback Support on the Local Interface on page 4089](#)
- *Junos® OS Ethernet Interfaces*

## Configuration Statements

- [\[edit services flow-monitoring\] Hierarchy Level on page 4094](#)

## link-fault-management

```
Syntax link-fault-management {
 action-profile profile-name {
 action {
 link-down;
 send-critical-event;
 syslog;
 }
 event {
 link-adjacency-loss;
 link-event-rate {
 frame-error count;
 frame-period count;
 frame-period-summary count;
 symbol-period count;
 }
 protocol-down;
 }
 }
 interface interface-name {
 apply-action-profile profile-name;
 link-discovery (active | passive);
 pdu-interval interval;
 pdu-threshold threshold-value;
 remote-loopback;
 event-thresholds {
 frame-error count;
 frame-period count;
 frame-period-summary count;
 symbol-period count;
 }
 negotiation-options {
 allow-remote-loopback;
 no-allow-link-events;
 }
 }
 }
```

**Hierarchy Level** [edit protocols oam [ethernet](#)]

**Release Information** Statement introduced in Junos OS Release 8.2.

**Description** For Ethernet interfaces on M320, M120, MX Series, and T Series routers and EX Series switches, specify fault signaling and detection for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

The remaining statements are explained separately.

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

- Related Documentation**
- [Enabling IEEE 802.3ah OAM Support on page 4076](#)

#### [\[edit services flow-monitoring\]](#) Hierarchy Level

---

```
services {
 flow-monitoring {
 version-ipfix {
 template template-name {
 flow-active-timeout seconds;
 flow-inactive-timeout seconds;
 ipv4-template;
 ipv6-template;
 option-refresh-rate packets packets seconds seconds;
 template-refresh-rate packets packets seconds seconds;
 }
 }
 }
 version9 {
 template template-name {
 flow-active-timeout seconds;
 flow-inactive-timeout seconds;
 ipv4-template {
 nexthop-options {
 mpls {
 label-position [positions];
 }
 }
 }
 ipv6-template;
 mpls-template {
 label-position [positions];
 }
 mpls-ipv4-template {
 label-position [positions];
 }
 option-refresh-rate {
 packets packets;
 seconds seconds;
 }
 peer-as-billing-template;
 template-refresh-rate {
 packets packets;
 seconds seconds;
 }
 peer-as-billing-template;
 option-refresh-rate packets;
 template-refresh-rate packets;
 }
 }
}
```

- Related Documentation**
- [Notational Conventions Used in Junos OS Configuration Hierarchies](#)
  - [\[edit services\] Hierarchy Level](#)

---

## action (OAM)

---

|                                 |                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>action {<br/>    link-down;<br/>    send-critical-event;<br/>    syslog;<br/>}</pre>                                                          |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">ethernet link-fault-management action-profile</a> ]                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                      |
| <b>Description</b>              | Define the action or actions to be taken when the OAM fault event occurs.                                                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083</a></li></ul> |

## action-profile (Defining for LFM)

---

**Syntax**    `action-profile profile-name {  
                  action {  
                    link-down;  
                    send-critical-event;  
                    syslog;  
                  }  
                  event {  
                    link-adjacency-loss;  
                    link-event-rate {  
                      frame-error count;  
                      frame-period count;  
                      frame-period-summary count;  
                      symbol-period count;  
                    }  
                    protocol-down;  
                  }  
                }`

**Hierarchy Level**    [edit protocols oam [ethernet link-fault-management](#)]

**Release Information**    Statement introduced in Junos OS Release 8.5.

**Description**    Configure a name, one or more actions, and the events that trigger the action for an action profile.

**Options**    *profile-name*—Name of the action profile.

The remaining statements are explained separately.

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

**Related Documentation**    • [Configuring an OAM Action Profile on page 4082](#)

## allow-remote-loopback

---

|                                 |                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | allow-remote-loopback;                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">link-fault-management interface</a> <i>interface-name</i> <a href="#">negotiation-options</a> ]                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.                                                                                                     |
| <b>Description</b>              | Enable the remote loopback on IQ2 and IQ2-E Gigabit Ethernet interfaces, and Ethernet interfaces on the MX Series routers and EX Series switches. |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling Remote Loopback Support on the Local Interface on page 4089</a></li> </ul>          |

## apply-action-profile

---

|                                 |                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | apply-action-profile <i>profile-name</i> ;                                                                              |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">ethernet link-fault-management interface</a> ]                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                           |
| <b>Description</b>              | Apply the specified action profile to the interface for link-fault management.                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Applying an Action Profile on page 4088</a></li> </ul>             |

## ethernet (Protocols OAM)

```

Syntax ethernet {
 connectivity-fault-management {
 action-profile profile-name {
 default-actions {
 interface-down;
 }
 }
 }
 performance-monitoring {
 delegate-server-processing;
 hardware-assisted-timestamping;
 sla-iterator-profiles {
 profile-name {
 disable;
 calculation-weight {
 delay delay-weight;
 delay-variation delay-variation-weight;
 }
 cycle-time milliseconds;
 iteration-period connections;
 measurement-type (loss | statistical-frame-loss | two-way-delay);
 }
 }
 }
 linktrace {
 age (30m | 10m | 1m | 30s | 10s);
 path-database-size path-database-size;
 }
 maintenance-domain domain-name {
 level number;
 name-format (character-string | none | dns | mac+2octet);
 maintenance-association ma-name {
 short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
 protect-maintenance-association protect-ma-name;
 remote-maintenance-association remote-ma-name;
 continuity-check {
 convey-loss-threshold;
 hold-interval minutes;
 interface-status-tlv;
 interval (10m | 10s | 1m | 1s | 100ms);
 loss-threshold number;
 port-status-tlv;
 }
 }
 mep mep-id {
 auto-discovery;
 direction (up | down);
 interface interface-name (protect | working);
 lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect |
 rem-err-xcon | xcon);
 priority number;
 remote-mep mep-id {
 action-profile profile-name;
 sla-iterator-profile profile-name {

```



```

 data-tlv-size size;
 iteration-count count-value;
 priority priority-value;
 }
}
}
}
}
}
}
evcs evc-id {
 evc-protocol cfm management-domain domain-id (management-association
 association-id | vpls (routing-instance instance-id);
 remote-uni-count count;
 multipoint-to-multipoint;
}
link-fault-management {
 action-profile profile-name {
 action {
 link-down;
 send-critical-event;
 syslog;
 }
 event {
 link-adjacency-loss;
 link-event-rate {
 frame-error count;
 frame-period count;
 frame-period-summary count;
 symbol-period count;
 }
 protocol-down;
 }
 }
}
interface interface-name {
 apply-action-profile;
 link-discovery (active | passive);
 pdu-interval interval;
 pdu-threshold threshold-value;
 remote-loopback;
 event-thresholds {
 frame-error count;
 frame-period count;
 frame-period-summary count;
 symbol-period count;
 }
 negotiation-options {
 allow-remote-loopback;
 no-allow-link-events;
 }
}
}
lmi {
 status-counter count;
 polling-verification-timer value;
 interface name {
 uni-id uni-name;

```

```
status-counter number;
polling-verification-timer value;
evc-map-type (all-to-one-bundling | bundling | service-multiplexing);
evc evc-name {
 default-evc;
 vlan-list vlan-id-list;
}
}
}
```

**Hierarchy Level** [edit protocols oam]

**Release Information** Statement introduced in Junos OS Release 8.2.

**Description** For Ethernet interfaces on EX Series switches, and M320, MX Series, and T Series routers, provide fault signaling and detection for 802.3ah Operation, Administration, and Management (OAM) support.

The remaining statements are explained separately.

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

**Related Documentation**

- [Enabling IEEE 802.3ah OAM Support on page 4076](#)
- *Example: Configuring Connectivity Fault Management for a PBB Network*

## event (LFM)


|                                 |                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> event {   link-adjacency-loss;   link-event-rate {     frame-error <i>count</i>;     frame-period <i>count</i>;     frame-period-summary <i>count</i>;     symbol-period <i>count</i>;   }   protocol-down; }</pre> |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">ethernet link-fault-management action-profile</a> ]                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                                             |
| <b>Description</b>              | <p>Configure threshold values for link events in an action profile.</p> <p>The remaining statements are explained separately.</p>                                                                                         |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Monitoring Protocol Status on page 4086</a></li> </ul>                                                                                                               |

## event-thresholds

|                                 |                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> event-thresholds {   frame-error <i>count</i>;   frame-period <i>count</i>;   frame-period-summary <i>count</i>;   symbol-period <i>count</i>; }</pre> |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">link-fault-management interface interface-name</a> ]                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.                                                                                                                |
| <b>Description</b>              | <p>Configure threshold limit values for link events in periodic OAM PDUs.</p> <p>The remaining statements are explained separately.</p>                      |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Threshold Values for Local Fault Events on an Interface on page 4079</a></li> </ul>         |

## fast-aps-switch

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | fast-aps-switch;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> sonet-options aps]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | (M320 routers with Channelized OC3/STM1 Circuit Emulation PIC with SFP only and EX Series switches) Reduce the Automatic Protection Switching (APS) switchover time in Layer 2 circuits.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|                                 | <div> <b>NOTE:</b><ul style="list-style-type: none"><li>• Configuring this statement reduces the APS switchover time only when the Layer 2 circuit encapsulation type for the interface receiving traffic from a Layer 2 circuit neighbor is SAToP.</li><li>• When the fast-aps-switch statement is configured in revertive APS mode, you must configure an appropriate value for revert time to achieve reduction in APS switchover time.</li><li>• To prevent the logical interfaces in the data path from being shut down, configure appropriate hold-time values on all the interfaces in the data path that support TDM.</li><li>• The fast-aps-switch statement cannot be configured when the APS annex-b option is configured.</li><li>• The interfaces that have the fast-aps-switch statement configured cannot be used in virtual private LAN service (VPLS) environments.</li></ul></div> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Reducing APS Switchover Time in Layer 2 Circuits</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## frame-error

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>frame-error count;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">ethernet link-fault-management action-profile event link-event-rate</a> ],<br>[edit protocols oam <a href="#">link-fault-management interface interface-name event-thresholds</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Threshold for sending frame error events or taking the action specified in the action profile.</p> <p>A frame error is any frame error on the underlying physical layer. The threshold is reached when the number of frame errors reaches the configured value within the window.</p> <p>The window or period during which frame errors are counted is 5 seconds or multiples of it (with a maximum value of 1 minute). This window denotes the duration as intervals of 100 milliseconds, encoded as a 16-bit unsigned integer. This window is not configurable in Junos OS. According to the IEEE 802.3ah standard, the default value of the frame-errors window is 1 second. This window has a lower bound of 1 second and an upper bound of 1 minute.</p> |
| <b>Options</b>                  | <p><b>count</b>—Threshold count for frame error events.</p> <p><b>Range:</b> 1 through 100</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Threshold Values for Local Fault Events on an Interface on page 4079</a></li> <li>• <a href="#">Configuring Threshold Values for Fault Events in an Action Profile on page 4087</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## frame-period

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>frame-period count;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">ethernet link-fault-management action-profile event link-event-rate</a> ],<br>[edit protocols oam <a href="#">link-fault-management interface interface-name event-thresholds</a> ]                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Threshold for sending frame period error events or taking the action specified in the action profile.</p> <p>A frame error is any frame error on the underlying physical layer. The frame period threshold is reached when the number of frame errors reaches the configured value within the period window. The default period window is the number of minimum-size frames that can be transmitted on the underlying physical layer in 1 second. The window is not configurable.</p> |
| <b>Options</b>                  | <p><b>count</b>—Threshold count for frame period error events.</p> <p><b>Range:</b> 1 through 100</p>                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Threshold Values for Local Fault Events on an Interface on page 4079</a></li><li>• <a href="#">Configuring Threshold Values for Fault Events in an Action Profile on page 4087</a></li></ul>                                                                                                                                                                                                                             |

## frame-period-summary

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>frame-period-summary count;</code>                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">ethernet link-fault-management action-profile event link-event-rate</a> ],<br>[edit protocols oam <a href="#">link-fault-management interface interface-name event-thresholds</a> ]                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | <p>Threshold for sending frame period summary error events or taking the action specified in the action profile.</p> <p>An errored frame second is any 1-second period that has at least one errored frame. This event is generated if the number of errored frame seconds is equal to or greater than the specified threshold for that period window. The default window is 60 seconds. The window is not configurable.</p> |
| <b>Options</b>                  | <p><b>count</b>—Threshold count for frame period summary error events.</p> <p><b>Range:</b> 1 through 100</p>                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Threshold Values for Local Fault Events on an Interface on page 4079</a></li> <li>• <a href="#">Configuring Threshold Values for Fault Events in an Action Profile on page 4087</a></li> </ul>                                                                                                                                                              |

## interface (OAM Link-Fault Management)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>interface <i>interface-name</i> {<br/>    <b>apply-action-profile</b> <i>profile-name</i>;<br/>    <b>link-discovery</b> (active   passive);<br/>    <b>pdu-interval</b> <i>interval</i>;<br/>    <b>pdu-threshold</b> <i>threshold-value</i>;<br/>    <b>remote-loopback</b>;<br/>    <b>event-thresholds</b> {<br/>        <b>frame-error</b> <i>count</i>;<br/>        <b>frame-period</b> <i>count</i>;<br/>        <b>frame-period-summary</b> <i>count</i>;<br/>        <b>symbol-period</b> <i>count</i>;<br/>    }<br/>    <b>negotiation-options</b> {<br/>        <b>allow-remote-loopback</b>;<br/>        <b>no-allow-link-events</b>;<br/>    }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit protocols oam <b>ethernet link-fault-management</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | For Ethernet interfaces on M320, MX Series, and T Series routers, configure IEEE 802.3ah Operation, Administration, and Management (OAM) support.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>interface</b> <i>interface-name</i>—Interface to be enabled for IEEE 802.3ah link fault management OAM support.</p> <p><b>Range:</b> 1 through 10 interfaces can be tracked.</p> <p>The remaining statements are described separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling IEEE 802.3ah OAM Support on page 4076</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |



---

## link-adjacency-loss

---

|                                 |                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-adjacency-loss;                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">ethernet link-fault-management action-profile event</a> ]                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                     |
| <b>Description</b>              | Loss of adjacency with IEEE 802.3ah link-fault management peer event. When included, the loss-of-adjacency event triggers the action specified under the <b>action</b> statement. |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Monitoring the Loss of Link Adjacency on page 4085</a></li></ul>                                                              |

---

## link-down

---

|                                 |                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-down;                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">ethernet link-fault-management</a> ]                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                      |
| <b>Description</b>              | Mark the interface down for transit traffic.                                                                                                       |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083</a></li></ul> |

## link-discovery

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-discovery (active   passive);                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">ethernet link-fault-management interface interface-name</a> ]                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | For Ethernet interfaces on EX Series switches, and M320, M120, MX Series, and T Series routers, specify the discovery mode used for IEEE 802.3ah Operation, Administration, and Management (OAM) support. The discovery process is triggered automatically when OAM 802.3ah functionality is enabled on a port. Link monitoring is done when the interface sends periodic OAM PDUs. |
| <b>Options</b>                  | (active   passive)—Passive or active mode. In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. Once the discovery process is initiated, both sides participate in discovery.                                                             |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Link Discovery on page 4077</a></li></ul>                                                                                                                                                                                                                                                                           |

## link-event-rate

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | link-event-rate {<br><a href="#">frame-error count</a> ;<br><a href="#">frame-period count</a> ;<br><a href="#">frame-period-summary count</a> ;<br><a href="#">symbol-period count</a> ;<br>} |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">ethernet link-fault-management action-profile event</a> ]                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                  |
| <b>Description</b>              | Configure the number of link-fault management events per second.                                                                                                                               |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Threshold Values for Fault Events in an Action Profile on page 4087</a></li></ul>                                              |

## negotiation-options

---

|                                 |                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | negotiation-options {<br><a href="#">allow-remote-loopback</a> ;<br><a href="#">no-allow-link-events</a> ;<br>}                                                   |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">link-fault-management interface</a> <i>interface-name</i> ]                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.                                                                                                                     |
| <b>Description</b>              | Enable and disable IEEE 802.3ah Operation, Administration, and Management (OAM) features for Ethernet interfaces.<br><br>The statements are explained separately. |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">IEEE 802.3ah OAM Link-Fault Management Overview on page 4073</a></li> </ul>                                  |

## no-allow-link-events

---

|                                 |                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-allow-link-events;                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">ethernet link-fault-management interface</a> <i>interface-name</i> <a href="#">negotiation-options</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.                                                                                            |
| <b>Description</b>              | Disable the sending of link event TLVs.                                                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Disabling the Sending of Link Event TLVs on page 4080</a></li> </ul>                |

## oam

---

```
Syntax oam {
 ethernet{
 connectivity-fault-management {
 action-profile profile-name {
 action {
 interface-down;
 }
 default-actions {
 interface-down;
 }
 event {
 adjacency-loss;
 }
 }
 }
 linktrace {
 age (30m | 10m | 1m | 30s | 10s);
 path-database-size path-database-size;
 }
 maintenance-domain domain-name {
 level number;
 mip-half-function (none | default | explicit);
 name-format (character-string | none | dns | mac+2oct);
 maintenance-association ma-name {
 continuity-check {
 hold-interval minutes;
 interface-status-tlv;
 interval (10m | 10s | 1m | 1s | 100ms);
 loss-threshold number;
 port-status-tlv;
 }
 mep mep-id {
 auto-discovery;
 direction down;
 interface interface-name;
 remote-mep mep-id {
 action-profile profile-name;
 }
 }
 }
 }
 }
 performance-monitoring {
 sla-iterator-profiles {
 profile-name {
 calculation-weight {
 delay delay-value;
 delay-variation delay-variation-value;
 }
 cycle-time cycle-time-value;
 iteration-period iteration-period-value;
 measurement-type two-way-delay;
 passive;
 }
 }
 }
 }
```

```

 }
 }
 traceoptions {
 file filename <files number> <match regex> <size size> <world-readable |
 no-world-readable>;
 flag flag ;
 no-remote-trace;
 }
}
link-fault-management {
 action-profile profile-name;
 action {
 syslog;
 link-down;
 send-critical-event
 }
 event {
 link-adjacency-loss;
 link-event-rate {
 frame-error count;
 frame-period count;
 frame-period-summary count;
 symbol-period count;
 }
 }
}
interface interface-name {
 link-discovery (active | passive);
 pdu-interval interval;
 pdu-threshold threshold-value;
 remote-loopback;
 event-thresholds {
 frame-error count;
 frame-period count;
 frame-period-summary count;
 symbol-period count;
 }
 negotiation-options {
 allow-remote-loopback;
 no-allow-link-events;
 }
}
traceoptions {
 file filename <files number> <match regex> <size size> <world-readable |
 no-world-readable>;
 flag flag ;
 no-remote-trace;
}
}
}

```

**Hierarchy Level** [edit protocols]

**Release Information** Statement introduced in Junos OS Release 9.4 for EX Series switches.  
**connectivity-fault-management** introduced in Junos OS Release 10.2 for EX Series switches.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b>              | Provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support for Ethernet interfaces on EX Series switches or configure connectivity fault management (CFM) for IEEE 802.1ag Operation, Administration, and Management (OAM) support on the switches.<br><br>The remaining statements are explained separately.                                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches</i></li><li>• <i>Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches</i></li><li>• <i>Configuring Ethernet OAM Link Fault Management (CLI Procedure)</i></li><li>• <i>Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)</i></li></ul> |

---

## pdu-interval

---

|                                 |                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>pdu-interval <i>interval</i>;</code>                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">ethernet link-fault-management interface</a> <i>interface-name</i> ]                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.                                                                                                                                                                                                    |
| <b>Description</b>              | For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, specify the periodic OAM PDU sending interval for fault detection. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support. |
| <b>Options</b>                  | <i>interval</i> —Periodic OAM PDU sending interval.<br><b>Range:</b> 100 through 1000 milliseconds<br><b>Default:</b> 1000 milliseconds                                                                                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the OAM PDU Interval on page 4078</a></li></ul>                                                                                                                                  |

## pdu-threshold

---

|                                 |                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>pdu-threshold <i>threshold-value</i>;</code>                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">ethernet link-fault-management interface <i>interface-name</i></a> ]                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.                                                                                                                                                                                                   |
| <b>Description</b>              | For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, specify the number of OAM PDUs to miss before an error is logged. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support. |
| <b>Options</b>                  | <p><b><i>threshold-value</i></b>—The number of PDUs missed before declaring the peer lost.</p> <p><b>Range:</b> 3 through 10 PDUs</p> <p><b>Default:</b> 3 PDUs</p>                                                                             |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the OAM PDU Threshold on page 4079</a></li> </ul>                                                                                                                              |

## protocol-down

---

|                                 |                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>protocol-down;</code>                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">ethernet link-fault-management action-profile event</a> ]                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                                                                |
| <b>Description</b>              | Upper layer indication of protocol down event. When the <b>protocol-down</b> statement is included, the protocol down event triggers the action specified under the <b>action</b> statement. |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring an OAM Action Profile on page 4082</a></li> </ul>                                                                           |

## remote-loopback

---

|                                 |                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | remote-loopback;                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">link-fault-management interface interface-name</a> ]                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, set the remote DTE into loopback mode. Remove the statement from the configuration to take the remote DTE out of loopback mode. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support. |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Setting a Remote Interface into Loopback Mode on page 4088</a></li></ul>                                                                                                                                                                                  |

## send-critical-event

---

|                                 |                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | send-critical-event;                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">ethernet link-fault-management action-profile action</a> ]                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                      |
| <b>Description</b>              | Send OAM PDUs with the critical event bit set.                                                                                                     |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083</a></li></ul> |



## symbol-period

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>symbol-period count;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">ethernet link-fault-management action-profile event</a> , <a href="#">link-event-rate</a> ],<br>[edit protocols oam <a href="#">link-fault-management interface interface-name event-thresholds</a> ]                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Configure the threshold for sending symbol period events or taking the action specified in the action profile.</p> <p>A symbol error is any symbol code error on the underlying physical layer. The symbol period threshold is reached when the number of symbol errors reaches the configured value within the period window. The default period window is the number of symbols that can be transmitted on the underlying physical layer in 1 second. The window is not configurable.</p> |
| <b>Options</b>                  | <p><b>count</b>—Threshold count for symbol period events.</p> <p><b>Range:</b> 1 through 100</p>                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Threshold Values for Local Fault Events on an Interface on page 4079</a></li> <li>• <a href="#">Configuring Threshold Values for Fault Events in an Action Profile on page 4087</a></li> </ul>                                                                                                                                                                                                                                |

## syslog (OAM Action)

|                                 |                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>syslog;</code>                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit protocols oam <a href="#">ethernet link-fault-management action-profile action</a> ]                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.5.                                                                                                        |
| <b>Description</b>              | Generate a syslog message for the Ethernet Operation, Administration, and Management (OAM) event.                                                    |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Specifying the Actions to Be Taken for Link-Fault Management Events on page 4083</a></li> </ul> |

## version-ipfix (Services)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>version-ipfix {<br/>  template <i>template-name</i> {<br/>    flow-active-timeout <i>seconds</i>;<br/>    flow-inactive-timeout <i>seconds</i>;<br/>    ipv4-template;<br/>    ipv6-template;<br/>    option-refresh-rate <i>packets packets seconds seconds</i>;<br/>    template-refresh-rate <i>packets packets seconds seconds</i>;<br/>  }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit services flow-monitoring]                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2.<br>Statement introduced in Junos OS Release 12.R3 for EX Series switches.                                                                                                                                                                                                                                         |
| <b>Description</b>              | Specify the output template properties to support inline flow monitoring. The remaining statements are explained separately.                                                                                                                                                                                                                                     |
| <b>Usage Guidelines</b>         | See <i>Configuring Inline Sampling</i> .                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                          |

## Administration

---

- [Routine Monitoring on page 4116](#)
- [Operational Commands: General on page 4119](#)
- [Operational Commands: Port Mirroring on page 4139](#)
- [Operational Commands: Ethernet OAM Link Fault Management on page 4141](#)

## Routine Monitoring

- [Monitoring Traffic Through the Router or Switch on page 4116](#)

### Monitoring Traffic Through the Router or Switch

---

To help with the diagnosis of a problem, display real-time statistics about the traffic passing through physical interfaces on the router or switch.

To display real-time statistics about physical interfaces, perform these tasks:

1. [Displaying Real-Time Statistics About All Interfaces on the Router or Switch on page 4116](#)
2. [Displaying Real-Time Statistics About an Interface on the Router or Switch on page 4117](#)

#### *Displaying Real-Time Statistics About All Interfaces on the Router or Switch*

|                |                                                                                                    |
|----------------|----------------------------------------------------------------------------------------------------|
| <b>Purpose</b> | Display real-time statistics about traffic passing through all interfaces on the router or switch. |
|----------------|----------------------------------------------------------------------------------------------------|

**Action** To display real-time statistics about traffic passing through all interfaces on the router or switch:

```
user@host> monitor interface traffic
```

### Sample Output

```
user@host> monitor interface traffic
host name Seconds:15 Time: 12:31:09
Interface Link Input packets (pps) Output packets (pps)
so-1/0/0 Down 0 (0) 0 (0)
so-1/1/0 Down 0 (0) 0 (0)
so-1/1/1 Down 0 (0) 0 (0)
so-1/1/2 Down 0 (0) 0 (0)
so-1/1/3 Down 0 (0) 0 (0)
t3-1/2/0 Down 0 (0) 0 (0)
t3-1/2/1 Down 0 (0) 0 (0)
t3-1/2/2 Down 0 (0) 0 (0)
t3-1/2/3 Down 0 (0) 0 (0)
so-2/0/0 Up 211035 (1) 36778 (0)
so-2/0/1 Up 192753 (1) 36782 (0)
so-2/0/2 Up 211020 (1) 36779 (0)
so-2/0/3 Up 211029 (1) 36776 (0)
so-2/1/0 Up 189378 (1) 36349 (0)
so-2/1/1 Down 0 (0) 18747 (0)
so-2/1/2 Down 0 (0) 16078 (0)
so-2/1/3 Up 0 (0) 80338 (0)
at-2/3/0 Up 0 (0) 0 (0)
at-2/3/1 Down 0 (0) 0 (0)
Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D
```

**Meaning** The sample output displays traffic data for active interfaces and the amount that each field has changed since the command started or since the counters were cleared by using the C key. In this example, the **monitor interface** command has been running for 15 seconds since the command was issued or since the counters last returned to zero.

### Displaying Real-Time Statistics About an Interface on the Router or Switch

**Purpose** Display real-time statistics about traffic passing through an interface on the router or switch.

**Action** To display traffic passing through an interface on the router or switch, use the following Junos OS CLI operational mode command:

```
user@host> monitor interface interface-name
```

### Sample Output

```
user@host> monitor interface so-0/0/1
Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
R1
Interface: so-0/0/1, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: OC3 Traffic statistics:
 Input bytes: 5856541 (88 bps)
 Output bytes: 6271468 (96 bps)
 Input packets: 157629 (0 pps)
 Output packets: 157024 (0 pps)
Encapsulation statistics:
```

```

Input keepalives: 42353
Output keepalives: 42320
LCP state: Opened
Error statistics:
Input errors: 0
Input drops: 0
Input framing errors: 0
Input runs: 0
Input giants: 0
Policed discards: 0
L3 incompletes: 0
L2 channel errors: 0
L2 mismatch timeouts: 0
Carrier transitions: 1
Output errors: 0
Output drops: 0
Aged packets: 0
Active alarms : None
Active defects: None
SONET error counts/seconds:
LOS count 1
LOF count 1
SEF count 1
ES-S 77
SES-S 77
SONET statistics:
BIP-B1 0
BIP-B2 0
REI-L 0
BIP-B3 0
REI-P 0
Received SONET overhead: F1 : 0x00 J0 : 0xZ

```

**Meaning** The sample output shows the input and output packets for a particular SONET interface (**so-0/0/1**). The information can include common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors. For more information, see *Checklist for Tracking Error Conditions*.

To control the output of the command while it is running, use the keys shown in [Table 303 on page 418](#).

**Table 303: Output Control Keys for the monitor interface Command**

| Action                                                                                                                                                                                                                       | Key      |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| Display information about the next interface. The <b>monitor interface</b> command scrolls through the physical or logical interfaces in the same order that they are displayed by the <b>show interfaces terse</b> command. | <b>N</b> |
| Display information about a different interface. The command prompts you for the name of a specific interface.                                                                                                               | <b>I</b> |
| Freeze the display, halting the display of updated statistics.                                                                                                                                                               | <b>F</b> |
| Thaw the display, resuming the display of updated statistics.                                                                                                                                                                | <b>T</b> |

Table 303: Output Control Keys for the monitor interface Command (*continued*)

| Action                                                                                                                          | Key |
|---------------------------------------------------------------------------------------------------------------------------------|-----|
| Clear (zero) the current delta counters since <b>monitor interface</b> was started. It does not clear the accumulative counter. | C   |
| Stop the <b>monitor interface</b> command.                                                                                      | Q   |

See the *Junos OS Operational Mode Commands* for details on using match conditions with the **monitor traffic** command.

## Operational Commands: General

## monitor traffic

---

**Syntax**    monitor traffic  
              <brief | detail | extensive>  
              <absolute-sequence>  
              <count *count*>  
              <interface *interface-name*>  
              <layer2-headers>  
              <matching *matching*>  
              <no-domain-names>  
              <no-promiscuous>  
              <no-resolve>  
              <no-timestamp>  
              <print-ascii>  
              <print-hex>  
              <resolve-timeout>  
              <size *size*>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                              Command introduced in Junos OS Release 9.0 for EX Series switches.  
                              Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Display packet headers or packets received and sent from the Routing Engine.



---

**NOTE:**

- Using the **monitor-traffic** command can degrade router or switch performance.
  - Delays from DNS resolution can be eliminated by using the **no-resolve** option.
- 



---

**NOTE:** This command is not supported on the QFabric system.

---

**Options**    **none**—(Optional) Display packet headers transmitted through **fxp0**. On a TX Matrix Plus router, display packet headers transmitted through **em0**.

**brief | detail | extensive**—(Optional) Display the specified level of output.

**absolute-sequence**—(Optional) Display absolute TCP sequence numbers.

**count *count***—(Optional) Specify the number of packet headers to display (0 through 1,000,000). The **monitor traffic** command quits automatically after displaying the number of packets specified.

**interface *interface-name***—(Optional) Specify the interface on which the **monitor traffic** command displays packet data. If no interface is specified, the **monitor traffic** command displays packet data arriving on the lowest-numbered interface.

**layer2-headers**—(Optional) Display the link-level header on each line.

**matching *matching***—(Optional) Display packet headers that match a regular expression. Use matching expressions to define the level of detail with which the **monitor traffic** command filters and displays packet data.

**no-domain-names**—(Optional) Suppress the display of the domain portion of hostnames. With the **no-domain-names** option enabled, the **monitor traffic** command displays only **team** for the hostname **team.company.net**.

**no-promiscuous**—(Optional) Do not put the interface into promiscuous mode.

**no-resolve**—(Optional) Suppress reverse lookup of the IP addresses.

**no-timestamp**—(Optional) Suppress timestamps on displayed packets.

**print-ascii**—(Optional) Display each packet in ASCII format.

**print-hex**—(Optional) Display each packet, except the link-level header, in hexadecimal format.

**resolve-timeout *timeout***—(Optional) Amount of time the router or switch waits for each reverse lookup before timing out. You can set the timeout for 1 through 4,294,967,295 seconds. The default is 4 seconds. To display each packet, use the **print-ascii**, **print-hex**, or **extensive** option.

**size *size***—(Optional) Read but do not display up to the specified number of bytes for each packet. When set to **brief** output, the default packet size is 96 bytes and is adequate for capturing IP, ICMP, UDP, and TCP packet data. When set to **detail** and **extensive** output, the default packet size is 1514. The **monitor traffic** command truncates displayed packets if the matched data exceeds the configured size.

**Additional Information** In the **monitor traffic** command, you can specify an expression to match by using the **matching** option and including the expression in quotation marks:

```
monitor traffic matching "expression"
```

Replace ***expression*** with one or more of the match conditions listed in [Table 304 on page 4122](#).

Table 304: Match Conditions for the monitor traffic Command

| Match Type    | Condition                                             | Description                                                                                                                                                                                                                                                                            |
|---------------|-------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Entity        | <b>host</b> [ <i>address</i>   <i>hostname</i> ]      | Matches packets that contain the specified address or hostname.<br><br>The protocol match conditions <b>arp</b> , <b>ip</b> , or <b>rarp</b> , or any of the directional match conditions can be prepended to the <b>host</b> match condition.                                         |
|               | <b>net</b> <i>address</i>                             | Matches packets with source or destination addresses containing the specified network address.                                                                                                                                                                                         |
|               | <b>net</b> <i>address mask mask</i>                   | Matches packets containing the specified network address and subnet mask.                                                                                                                                                                                                              |
|               | <b>port</b> ( <i>port-number</i>   <i>port-name</i> ) | Matches packets containing the specified source or destination TCP or UDP port number or port name.<br><br>In place of the numeric port address, you can specify a text synonym, such as <b>bgp</b> (179), <b>dhcp</b> (67), or <b>domain</b> (53) (the port numbers are also listed). |
| Directional   | <b>dst</b>                                            | Matches packets going to the specified destination. This match condition can be prepended to any of the entity type match conditions.                                                                                                                                                  |
|               | <b>src</b>                                            | Matches packets from a specified source. This match condition can be prepended to any of the entity type match conditions.                                                                                                                                                             |
|               | <b>src and dst</b>                                    | Matches packets that contain the specified source and destination addresses. This match condition can be prepended to any of the entity type match conditions.                                                                                                                         |
|               | <b>src or dst</b>                                     | Matches packets containing either of the specified addresses. This match condition can be prepended to any of the entity type match conditions.                                                                                                                                        |
| Packet Length | <b>less</b> <i>value</i>                              | Matches packets shorter than or equal to the specified value, in bytes.                                                                                                                                                                                                                |
|               | <b>greater</b> <i>value</i>                           | Matches packets longer than or equal to the specified value, in bytes.                                                                                                                                                                                                                 |



Table 304: Match Conditions for the monitor traffic Command (*continued*)

| Match Type | Condition                                                | Description                                                                                                                                                                                                                                                                                                                     |
|------------|----------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Protocol   | <b>amt</b>                                               | Matches all AMT packets. Use the extensive level of output to decode the inner IGMP packets in addition to the AMT outer packet.                                                                                                                                                                                                |
|            | <b>arp</b>                                               | Matches all ARP packets.                                                                                                                                                                                                                                                                                                        |
|            | <b>ether</b>                                             | Matches all Ethernet packets.                                                                                                                                                                                                                                                                                                   |
|            | <b>ether (broadcast   multicast)</b>                     | Matches broadcast or multicast Ethernet frames. This match condition can be prepended with <b>src</b> and <b>dst</b> .                                                                                                                                                                                                          |
|            | <b>ether protocol (address   (arp   ip   rarp))</b>      | Matches packets with the specified Ethernet address or Ethernet packets of the specified protocol type. The <b>ether protocol</b> arguments <b>arp</b> , <b>ip</b> , and <b>rarp</b> are also independent match conditions, so they must be preceded by a backslash (\) when used in the <b>ether protocol</b> match condition. |
|            | <b>icmp</b>                                              | Matches all ICMP packets.                                                                                                                                                                                                                                                                                                       |
|            | <b>ip</b>                                                | Matches all IP packets.                                                                                                                                                                                                                                                                                                         |
|            | <b>ip (broadcast   multicast)</b>                        | Matches broadcast or multicast IP packets.                                                                                                                                                                                                                                                                                      |
|            | <b>ip protocol (address   (icmp   igmp   tcp   udp))</b> | Matches packets with the specified address or protocol type. The <b>ip protocol</b> arguments <b>icmp</b> , <b>tcp</b> , and <b>udp</b> are also independent match conditions, so they must be preceded by a backslash (\) when used in the <b>ip protocol</b> match condition.                                                 |
|            | <b>isis</b>                                              | Matches all IS-IS routing messages.                                                                                                                                                                                                                                                                                             |
|            | <b>rarp</b>                                              | Matches all RARP packets.                                                                                                                                                                                                                                                                                                       |
|            | <b>tcp</b>                                               | Matches all TCP datagrams.                                                                                                                                                                                                                                                                                                      |
|            | <b>udp</b>                                               | Matches all UDP datagrams.                                                                                                                                                                                                                                                                                                      |

To combine expressions, use the logical operators listed in [Table 305 on page 4123](#).

Table 305: Logical Operators for the monitor traffic Command

| Logical Operator (Highest to Lowest Precedence) | Description                                                                          |
|-------------------------------------------------|--------------------------------------------------------------------------------------|
| <b>!</b>                                        | Logical NOT. If the first condition does not match, the next condition is evaluated. |

Table 305: Logical Operators for the monitor traffic Command (*continued*)

| Logical Operator (Highest to Lowest Precedence) | Description                                                                                                                                         |
|-------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| &&                                              | Logical AND. If the first condition matches, the next condition is evaluated. If the first condition does not match, the next condition is skipped. |
|                                                 | Logical OR. If the first condition matches, the next condition is skipped. If the first condition does not match, the next condition is evaluated.  |
| ()                                              | Group operators to override default precedence order. Parentheses are special characters, each of which must be preceded by a backslash (\).        |

You can use relational operators to compare arithmetic expressions composed of integer constants, binary operators, a length operator, and special packet data accessors. The arithmetic expression matching condition uses the following syntax:

```
monitor traffic matching "ether[0] & 1 != 0"arithmetic_expression relational_operator arithmetic_expression
```

The packet data accessor uses the following syntax:

```
protocol [byte-offset <size>]
```

The optional *size* field represents the number of bytes examined in the packet header. The available values are 1, 2, or 4 bytes. The following sample command captures all multicast traffic:

```
user@host> monitor traffic matching "ether[0] & 1 != 0"
```

To specify match conditions that have a numeric value, use the arithmetic and relational operators listed in [Table 306 on page 4125](#).



**NOTE:** Because the Packet Forwarding Engine removes Layer 2 header information before sending packets to the Routing Engine:

- The **monitor traffic** command cannot apply match conditions to inbound traffic.
- The **monitor traffic interface** command also cannot apply match conditions for Layer 3 and Layer 4 packet data, resulting in the match pipe option (`| match`) for this command for Layer 3 and Layer 4 packets not working either. Therefore, ensure that you specify match conditions as described in this command summary. For more information about match conditions, see [Table 304 on page 4122](#).
- The 802.1Q VLAN tag information included in the Layer 2 header is removed from all inbound traffic packets. Because the **monitor traffic interface ae[x]** command for aggregated Ethernet interfaces (such as ae0) only shows inbound traffic data, the command does not show VLAN tag information in the output.

**Table 306: Arithmetic and Relational Operators for the monitor traffic Command**

| Arithmetic or Relational Operator                         | Description                                                                         |
|-----------------------------------------------------------|-------------------------------------------------------------------------------------|
| <b>Arithmetic Operator</b>                                |                                                                                     |
| +                                                         | Addition operator.                                                                  |
| -                                                         | Subtraction operator.                                                               |
| /                                                         | Division operator.                                                                  |
| &                                                         | Bitwise AND.                                                                        |
| *                                                         | Bitwise exclusive OR.                                                               |
|                                                           | Bitwise inclusive OR.                                                               |
| <b>Relational Operator (Highest to Lowest Precedence)</b> |                                                                                     |
| <=                                                        | If the first expression is less than or equal to the second, the packet matches.    |
| >=                                                        | If the first expression is greater than or equal to the second, the packet matches. |
| <                                                         | If the first expression is less than the second, the packet matches.                |
| >                                                         | If the first expression is greater than the second, the packet matches.             |
| =                                                         | If the compared expressions are equal, the packet matches.                          |
| !=                                                        | If the compared expressions are unequal, the packet matches.                        |

**Required Privilege Level** trace  
maintenance

**List of Sample Output** [monitor traffic count on page 4126](#)  
[monitor traffic detail count on page 4126](#)  
[monitor traffic extensive \(Absolute Sequence\) on page 4126](#)  
[monitor traffic extensive \(Relative Sequence\) on page 4126](#)  
[monitor traffic extensive count on page 4126](#)  
[monitor traffic interface on page 4127](#)  
[monitor traffic matching on page 4127](#)  
[monitor traffic \(TX Matrix Plus Router\) on page 4127](#)  
[monitor traffic \(QFX3500 Switch\) on page 4128](#)

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

## Sample Output

### monitor traffic count

```
user@host> monitor traffic count 2
listening on fxp0
04:35:49.814125 In my-server.home.net.1295 > my-server.work.net.telnet: . ack
4122529478 win 16798 (DF)
04:35:49.814185
Out my-server.work.net.telnet > my-server.home.net.1295: P
1:38(37) ack 0 win 17680 (DF) [tos 0x10]
```

### monitor traffic detail count

```
user@host> monitor traffic detail count 2
listening on fxp0
04:38:16.265864 In my-server.home.net.1295 > my-server.work.net.telnet: . ack
4122529971 win 17678 (DF) (ttl 121, id 6812)
04:38:16.265926
Out my-server.work.net.telnet.telnet > my-server.home.net.1295: P 1:38(37) ack 0
win 17680 (DF) [tos 0x10] (ttl 6)
```

### monitor traffic extensive (Absolute Sequence)

```
user@host> monitor traffic extensive no-domain-names no-resolve no-timestamp count 20
matching "tcp" absolute-sequence
listening on fxp0
In 207.17.136.193.179 > 192.168.4.227.1024: . 4042780859:4042780859(0)
ack 1845421797 win 16384 <nop,nop,timestamp 4935628 965951> [tos 0xc0] (ttl)
In 207.17.136.193.179 > 192.168.4.227.1024: P 4042780859:4042780912(53)
ack 1845421797 win 16384
<nop,nop,timestamp 4935628 965951>:
BGP [|BGP UPDAT)
In 192.168.4.227.1024 > 207.17.136.193.179:
P 1845421797:1845421852(55) ack 4042780912 win 16384 <nop,nop,timestamp 965951
4935628>: BGP [|BGP UPDAT)
...
```

### monitor traffic extensive (Relative Sequence)

```
user@host> monitor traffic extensive no-domain-names no-resolve no-timestamp count 20
matching "tcp"
listening on fxp0
In 172.24.248.221.1680 > 192.168.4.210.23: . 396159737:396159737(0)
ack 1664980689 win 17574 (DF) (ttl 121, id 50003)
Out 192.168.4.210.23 > 172.24.248.221.1680: P 1:40(39)
ack 0 win 17680 (DF) [tos 0x10] (ttl 64, id 5394)
In 207.17.136.193.179 > 192.168.4.227.1024: P 4042775817:4042775874(57)
ack 1845416593 win 16384 <nop,nop,timestamp 4935379 965690>: BGP [|BGP UPDAT)
...
```

### monitor traffic extensive count

```
user@host> monitor traffic extensive count 5 no-domain-names no-resolve
listening on fxp013:18:17.406933
In 192.168.4.206.2723610880 > 172.17.28.8.2049:
40 null (ttl 64, id 38367)13:18:17.407577
In 172.17.28.8.2049 > 192.168.4.206.2723610880:
```

```

reply ok 28 null (ttl 61, id 35495)13:18:17.541140
In 0:e0:1e:42:9c:e0 0:e0:1e:42:9c:e0 9000 60:
0000 0100 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 000013:18:17.591513
In 172.24.248.156.4139 > 192.168.4.210.23:
3556964918:3556964918(0)
ack 295526518 win 17601 (DF)
(ttl 121, id 14)13:18:17.591568
Out 192.168.4.210.23 >
172.24.248.156.4139: P 1:40(39)
ack 0 win 17680 (DF) [tos 0x10]
(ttl 64, id 52376)

```

### monitor traffic interface

```

user@host> monitor traffic interface fxp0
listening on fxp0.0
18:17:28.800650 In server.home.net.723 > host1-0.lab.home.net.log
18:17:28.800733 Out host2-0.lab.home.net.login > server.home.net.7
18:17:28.817813 In host30.lab.home.net.syslog > host40.home0
18:17:28.817846 In host30.lab.home.net.syslog > host40.home0
...

```

### monitor traffic matching

```

user@host> monitor traffic matching "net 192.168.1.0/24"
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on fxp0, capture size 96 bytes

Reverse lookup for 192.168.1.255 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use no-resolve to avoid reverse lookups on IP addresses.

21:55:54.003511 In IP truncated-ip - 18 bytes missing!
192.168.1.17.netbios-ns > 192.168.1.255.netbios-ns: UDP, length 50
21:55:54.003585 Out IP truncated-ip - 18 bytes missing!
192.168.1.17.netbios-ns > 192.168.1.255.netbios-ns: UDP, length 50
21:55:54.003864 In arp who-has 192.168.1.17 tell 192.168.1.9
...

```

### monitor traffic (TX Matrix Plus Router)

```

user@host> monitor traffic
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on em0, capture size 96 bytes
04:11:59.862121 Out IP truncated-ip - 25 bytes missing!
summit-em0.englab.juniper.net.syslog > sv-log-01.englab.juniper.net.syslog:
SYSLOG kernel.info, length: 57
04:11:59.862303
Out IP truncated-ip - 25 bytes missing!
summit-em0.englab.juniper.net.syslog >
sv-log-02.englab.juniper.net.syslog: SYSLOG kernel.info, length: 57
04:11:59.923948
In IP aj-em0.englab.juniper.net.65235 >

```

```

summit-em0.englab.juniper.net.telnet: .
ack 1087492766 win 33304 <nop,nop,timestamp 42366734 993490>
04:11:59.923983 Out IP truncated-ip - 232 bytes missing!
summit-em0.englab.juniper.net.telnet > aj-em0.englab.juniper.net.65235: P
1:241(240) ack 0 win 33304
<nop,nop,timestamp 993590 42366734>
04:12:00.022900
In IP aj-em0.englab.juniper.net.65235 >
summit-em0.englab.juniper.net.telnet: . ack 241 win 33304 <nop,nop,timestamp
42366834 993590>
04:12:00.141204
In IP truncated-ip - 40 bytes missing!
ipg-lnx-shell11.juniper.net.46182 > summit-em0.englab.juniper.net.telnet: P
2950530356:2950530404(48) ack 485494987 win 63712
<nop,nop,timestamp 1308555294 987086>
04:12:00.141345
Out IP summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 1:6(5)
ack 48 win 33304
<nop,nop,timestamp 993809 1308555294>
04:12:00.141572
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: .
ack 6 win 63712
<nop,nop,timestamp 1308555294 993809>
04:12:00.141597
Out IP summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 6:10(4) ack 48 win 33304
<nop,nop,timestamp 993810 1308555294>
04:12:00.141821
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: .
ack 10 win 63712 <nop,nop,timestamp 1308555294 993810>
04:12:00.141837 Out IP truncated-ip - 2 bytes missing!
summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 10:20(10) ack 48 win 33304
<nop,nop,timestamp 993810 1308555294>
04:12:00.142072
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: . ack 20 win 63712
<nop,nop,timestamp 1308555294 993810>
04:12:00.142089 Out IP summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 20:28(8) ack 48 win 33304 <nop,nop,timestamp
993810 1308555294>
04:12:00.142321
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: .
ack 28 win 63712 <nop,nop,timestamp 1308555294 993810>
04:12:00.142337
Out IP truncated-ip - 1 bytes missing!
summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 28:37(9) ack 48 win 33304 <nop,nop,timestamp
993810 1308555294>
...

```

### monitor traffic (QFX3500 Switch)

```

user@switch> monitor traffic
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.

```

```
Listening on me4, capture size 96 bytes
Reverse lookup for 172.22.16.246 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use <no-resolve> to avoid reverse lookups on IP addresses.
16:35:32.240873 Out IP truncated-ip - 112 bytes missing!
labqfx-me0.lab4.juniper.net.ssh >
172.22.16.246.telefinder: P 4200727624:4200727756(132) ack 2889954831 win 65535
16:35:32.240900 Out IP truncated-ip - 176 bytes missing!
labqfx-me0.lab4.juniper.net.ssh >
172.22.16.246.telefinder: P 132:328(196) ack 1 win 65535
...
```

## ping

---

**Syntax**    `ping host`  
              `<bypass-routing>`  
              `<count requests>`  
              `<detail>`  
              `<do-not-fragment>`  
              `<inet | inet6>`  
              `<interface source-interface>`  
              `<interval seconds>`  
              `<logical-system logical-system-name>`  
              `<loose-source value>`  
              `<mac-address mac-address>`  
              `<no-resolve>`  
              `<pattern string>`  
              `<rapid>`  
              `<record-route>`  
              `<routing-instance routing-instance-name>`  
              `<size bytes>`  
              `<source source-address>`  
              `<strict >`  
              `<strict-source value.>`  
              `<tos type-of-service>`  
              `<ttl value>`  
              `<verbose>`  
              `<vpls instance-name>`  
              `<wait seconds>`

**Syntax (QFX Series)**    `ping host`  
                              `<bypass-routing>`  
                              `<count requests>`  
                              `<detail>`  
                              `<do-not-fragment>`  
                              `<inet>`  
                              `<interface source-interface>`  
                              `<interval seconds>`  
                              `<logical-system logical-system-name>`  
                              `<loose-source value>`  
                              `<mac-address mac-address>`  
                              `<no-resolve>`  
                              `<pattern string>`  
                              `<rapid>`  
                              `<record-route>`  
                              `<routing-instance routing-instance-name>`  
                              `<size bytes>`  
                              `<source source-address>`  
                              `<strict>`  
                              `< strict-source value>`  
                              `<tos type-of-service>`  
                              `<ttl value>`  
                              `<verbose>`  
                              `<wait seconds>`

**Release Information**    Command introduced before Junos OS Release 7.4.



Command introduced in Junos OS Release 9.0 for EX Series switches.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Check host reachability and network connectivity. The **ping** command sends Internet Control Message Protocol (ICMP) ECHO\_REQUEST messages to elicit ICMP ECHO\_RESPONSE messages from the specified host. Press Ctrl+c to interrupt a ping command.

**Options** **host**—IP address or hostname of the remote system to ping.

**bypass-routing**—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.

**count requests**—(Optional) Number of ping requests to send. The range of values is 1 through 2,000,000,000. The default value is an unlimited number of requests.

**detail**—(Optional) Include in the output the interface on which the ping reply was received.

**do-not-fragment**—(Optional) Set the do-not-fragment (DF) flag in the IP header of the ping packets. For IPv6 packets, this option disables fragmentation.



**NOTE:** In Junos OS Release 11.1 and later, when issuing the **ping** command for an IPv6 route with the **do-not-fragment** option, the maximum ping packet size is calculated by subtracting 48 bytes (40 bytes for the IPV6 header and 8 bytes for the ICMP header) from the MTU. Therefore, if the ping packet size (including the 48-byte header) is greater than the MTU, the ping operation might fail.

**inet**—(Optional) Ping Packet Forwarding Engine IPv4 routes.

**inet6**—(Optional) Ping Packet Forwarding Engine IPv6 routes.

**interface source-interface**—(Optional) Interface to use to send the ping requests.

**interval seconds**—(Optional) How often to send ping requests. The range of values, in seconds, is 1 through infinity. The default value is 1.

**logical-system logical-system-name**—(Optional) Name of logical system from which to send the ping requests.

Alternatively, enter the **set cli logical-system logical-system-name** command and then run the **ping** command. To return to the main router or switch, enter the **clear cli logical-system** command.

**loose-source value**—(Optional) Intermediate loose source route entry (IPv4). Open a set of values.

**mac-address *mac-address***—(Optional) Ping the physical or hardware address of the remote system you are trying to reach.

**no-resolve**—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

**pattern *string***—(Optional) Specify a hexadecimal fill pattern to include in the ping packet.

**rapid**—(Optional) Send ping requests rapidly. The results are reported in a single message, not in individual messages for each ping request. By default, five ping requests are sent before the results are reported. To change the number of requests, include the **count** option.

**record-route**—(Optional) Record and report the packet's path (IPv4).

**routing-instance *routing-instance-name***—(Optional) Name of the routing instance for the ping attempt.

**size *bytes***—(Optional) Size of ping request packets. The range of values, in bytes, is **0** through **65,468**. The default value is **56**, which is effectively 64 bytes because 8 bytes of ICMP header data are added to the packet.

**source *source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

**strict**—(Optional) Use the strict source route option (IPv4).

**strict-source *value***—(Optional) Intermediate strict source route entry (IPv4). Open a set of values.

**tos *type-of-service***—(Optional) Set the type-of-service (ToS) field in the IP header of the ping packets. The range of values is **0** through **255**.

**ttl *value***—(Optional) Time-to-live (TTL) value to include in the ping request (IPv6). The range of values is **0** through **255**.

**verbose**—(Optional) Display detailed output.

**vpls *instance-name***—(Optional) Ping the instance to which this VPLS belongs.

**wait *seconds***—(Optional) Maximum wait time, in seconds, after the final packet is sent. If this option is not specified, the default delay is **10** seconds. If this option is used without the count option, a default count of **5** packets is used.

**Required Privilege Level**

network

**Related Documentation**

- [Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages on page 1803](#)

**List of Sample Output**    [ping hostname on page 4133](#)

[ping hostname rapid on page 4133](#)

[ping hostname size count on page 4133](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. These packets are not counted in the received packets count. They are accounted for separately.

## Sample Output

### ping hostname

```
user@host> ping skye
PING skye.net (192.168.169.254): 56 data bytes
64 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.028 ms
64 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=1.053 ms
64 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.025 ms
64 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.098 ms
64 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=1.032 ms
64 bytes from 192.168.169.254: icmp_seq=5 ttl=253 time=1.044 ms
^C [abort]
```

### ping hostname rapid

```
user@host> ping skye rapid
PING skye.net (192.168.169.254): 56 data bytes
!!!!
--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.956/0.974/1.025/0.026 ms
```

### ping hostname size count

```
user@host> ping skye size 200 count 5
PING skye.net (192.168.169.254): 200 data bytes
208 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.759 ms
208 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=2.075 ms
208 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.843 ms
208 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.803 ms
208 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=17.898 ms

--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 1.759/5.075/17.898 ms
```

## traceroute

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceroute <i>host</i> &lt;as-number-lookup&gt; &lt;bypass-routing&gt; &lt;clns&gt; &lt;gateway <i>address</i>&gt; &lt;inet   inet6&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical system <i>logical-system-name</i>&gt; &lt;monitor <i>host</i>&gt; &lt;mpls (<i>ldp FEC address</i>   <i>rsvp label-switched-path-name</i>)&gt; &lt;no-resolve&gt; &lt;propagate-ttl&gt; &lt;routing-instance <i>routing-instance-name</i>&gt; &lt;source <i>source-address</i>&gt; &lt;tos <i>value</i>&gt; &lt;ttl <i>value</i>&gt; &lt;wait <i>seconds</i>&gt;</pre> |
| <b>Syntax (QFX Series)</b> | <pre>traceroute <i>host</i> &lt;as-number-lookup&gt; &lt;bypass-routing&gt; &lt;gateway <i>address</i>&gt; &lt;inet&gt; &lt;interface <i>interface-name</i>&gt; &lt;monitor <i>host</i>&gt; &lt;no-resolve&gt; &lt;routing-instance <i>routing-instance-name</i>&gt; &lt;source <i>source-address</i>&gt; &lt;tos <i>value</i>&gt; &lt;ttl <i>value</i>&gt; &lt;wait <i>seconds</i>&gt;</pre>                                                                                                                                                                           |
| <b>Release Information</b> | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>mpls</b> option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p><b>propagate-ttl</b> option introduced in Junos OS Release 12.1.</p>                                                                                                                                                                                                                           |
| <b>Description</b>         | Display the route that packets take to a specified network host. Use <b>traceroute</b> as a debugging tool to locate points of failure in a network.                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>             | <p><b>host</b>—IP address or name of remote host.</p> <p><b>as-number-lookup</b>—(Optional) Display the autonomous system (AS) number of each intermediate hop on the path from the host to the destination.</p> <p><b>bypass-routing</b>—(Optional) Bypass the normal routing tables and send requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to display a route to a local system through an interface that has no route through it.</p>                             |

**clns**—(Optional) Trace the route belonging to the Connectionless Network Service (CLNS).

**gateway address**—(Optional) Address of a router or switch through which the route transits.

**inet | inet6**—(Optional) Trace the route belonging to IPv4 or IPv6, respectively.

**interface *interface-name***—(Optional) Name of the interface over which to send packets.

**logical-system *logical-system-name***—(Optional) Perform this operation on all logical systems or on a particular logical system.

**monitor *host***—(Optional) Display real-time monitoring information for the specified host.

**mpls (*ldp FEC address* | *rsvp label-switched-path name*)**—(Optional) See *traceroute mpls ldp* and *traceroute mpls rsvp*.

**no-resolve**—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

**propagate-ttl**—(Optional) On the PE routing device, use this option to view locally generated Routing Engine transit traffic. This is applicable for MPLS L3VPN traffic only.

Use for troubleshooting, when you want to view hop-by-hop information from the local provider router to the remote provider router, when TTL decrementing is disabled on the core network using the **no-propagate-ttl** configuration statement.



**NOTE:** Using **propagate-ttl** with **traceroute** on the CE router does not show hop-by-hop information.

**routing-instance *routing-instance-name***—(Optional) Name of the routing instance for the traceroute attempt.

**source *source-address***—(Optional) Source address of the outgoing traceroute packets.

**tos *value***—(Optional) Value to include in the IP type-of-service (ToS) field. The range of values is 0 through 255.

**ttl *value***—(Optional) Maximum time-to-live value to include in the traceroute request. The range of values is 0 through 128.

**wait *seconds***—(Optional) Maximum time to wait for a response to the traceroute request.

**Required Privilege Level**

network

**Related Documentation**

• [traceroute monitor on page 4138](#)

**List of Sample Output** [traceroute on page 4136](#)

[traceroute as-number-lookup host on page 4136](#)

[traceroute no-resolve on page 4136](#)

[traceroute propagate-ttl on page 4137](#)

[traceroute \(Between CE Routers, Layer 3 VPN\) on page 4137](#)

[traceroute \(Through an MPLS LSP\) on page 4137](#)

**Output Fields** Table 307 on page 4136 describes the output fields for the **traceroute** command. Output fields are listed in the approximate order in which they appear.

**Table 307: traceroute Output Fields**

| Field Name             | Field Description                                             |
|------------------------|---------------------------------------------------------------|
| <b>traceroute to</b>   | IP address of the receiver.                                   |
| <b>hops max</b>        | Maximum number of hops allowed.                               |
| <b>byte packets</b>    | Size of packets being sent.                                   |
| <b>number-of-hops</b>  | Number of hops from the source to the named router or switch. |
| <b>router-name</b>     | Name of the router or switch for this hop.                    |
| <b>address</b>         | Address of the router or switch for this hop.                 |
| <b>Round trip time</b> | Average round-trip time, in milliseconds (ms).                |

## Sample Output

### traceroute

```
user@host> traceroute santacruz
traceroute to green.company.net (10.156.169.254), 30 hops max, 40 byte packets
 1 blue23 (10.168.1.254) 2.370 ms 2.853 ms 0.367 ms
 2 red14 (10.168.255.250) 0.778 ms 2.937 ms 0.446 ms
 3 yellow (10.156.169.254) 7.737 ms 89.905 ms 0.834 ms
```

### traceroute as-number-lookup host

```
user@host> traceroute as-number-lookup 10.100.1.1
traceroute to 10.100.1.1 (10.100.1.1), 30 hops max, 40 byte packets
 1 10.39.1.1 (10.39.1.1) 0.779 ms 0.728 ms 0.562 ms
 2 10.39.1.6 (10.39.1.6) [AS 32] 0.657 ms 0.611 ms 0.617 ms
 3 10.100.1.1 (10.100.1.1) [AS 10, 40, 50] 0.880 ms 0.808 ms 0.774 ms
```

### traceroute no-resolve

```
user@host> traceroute santacruz no-resolve
traceroute to green.company.net (10.156.169.254), 30 hops max, 40 byte packets
 1 10.168.1.254 0.458 ms 0.370 ms 0.365 ms
 2 10.168.255.250 0.474 ms 0.450 ms 0.444 ms
```

```
3 10.156.169.254 0.931 ms 0.876 ms 0.862 ms
```

### traceroute propagate-ttl

```
user@host> traceroute propagate-ttl 100.200.2.2 routing-instance VPN-A
traceroute to 100.200.2.2 (100.200.2.2) from 1.1.0.2, 30 hops max, 40 byte packets
```

```
1 1.2.0.2 (1.2.0.2) 2.456 ms 1.753 ms 1.672 ms
 MPLS Label=299776 CoS=0 TTL=1 S=0
 MPLS Label=299792 CoS=0 TTL=1 S=1
2 1.3.0.2 (1.3.0.2) 1.213 ms 1.225 ms 1.166 ms
 MPLS Label=299792 CoS=0 TTL=1 S=1
3 100.200.2.2 (100.200.2.2) 1.422 ms 1.521 ms 1.443 ms
```

### traceroute (Between CE Routers, Layer 3 VPN)

```
user@host> traceroute vpn09
traceroute to vpn09.skybank.net (10.255.14.179), 30 hops max, 40
byte packets
1 10.39.10.21 (10.39.10.21) 0.598 ms 0.500 ms 0.461 ms
2 10.39.1.13 (10.39.1.13) 0.796 ms 0.775 ms 0.806 ms
 MPLS Label=100006 CoS=0 TTL=1 S=1
3 vpn09.skybank.net (10.255.14.179) 0.783 ms 0.716 ms 0.686
```

### traceroute (Through an MPLS LSP)

```
user@host> traceroute mpls1
traceroute to 10.168.1.224 (10.168.1.224), 30 hops max, 40 byte packets
1 mpls1-sr0.company.net (10.168.200.101) 0.555 ms 0.393 ms 0.367 ms
 MPLS Label=1024 CoS=0 TTL=1
2 mpls5-lo0.company.net (10.168.1.224) 0.420 ms 0.394 ms 0.401 ms
```

## traceroute monitor

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>traceroute monitor <i>host</i></code><br><code>&lt;count <i>value</i>&gt;</code><br><code>&lt;inet   inet 6&gt;</code><br><code>&lt;interval <i>seconds</i>&gt;</code><br><code>&lt;no resolve&gt;</code><br><code>&lt;size <i>value</i>&gt;</code><br><code>&lt;source <i>source-address</i>&gt;</code><br><code>&lt;summary&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax (QFX Series)</b>      | <code>traceroute monitor <i>host</i></code><br><code>&lt;count <i>value</i>&gt;</code><br><code>&lt;inet&gt;</code><br><code>&lt;interval <i>seconds</i>&gt;</code><br><code>&lt;no resolve&gt;</code><br><code>&lt;size <i>value</i>&gt;</code><br><code>&lt;source <i>source-address</i>&gt;</code><br><code>&lt;summary&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Command introduced in Junos OS Release 8.0<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Display live monitoring of each hop in the route that packets take to a specified network host. Use as a debugging tool to locate points of failure in a network.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b><i>host</i></b>—IP address or name of remote host.</p> <p><b><i>count value</i></b>—Number of ping requests, in packets, to send in summary mode. The default value is <b>10</b>.</p> <p><b><i>inet   inet6</i></b>—(Optional) Trace the route belonging to IPv4 or IPv6, respectively.</p> <p><b><i>interval seconds</i></b>—(Optional) Number of seconds to wait before sending ping requests. The default value is <b>1</b>.</p> <p><b><i>no resolve</i></b>—(Optional) Do not attempt to display addresses symbolically.</p> <p><b><i>size value</i></b>—(Optional) Receive the specified number of bytes for each packet. The range is <b>0</b> through <b>65468</b> bytes. The default value is <b>64</b>.</p> <p><b><i>source source-address</i></b>—(Optional) Source address of the outgoing ping packets.</p> <p><b><i>summary</i></b>—(Optional) Generate and display a summary of live monitoring of each hop on the route that packets take to a specified network host.</p> |
| <b>Required Privilege Level</b> | network                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>    | <a href="#">traceroute monitor on page 4139</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>            | <a href="#">Table 308 on page 4139</a> describes the output fields for the <b>traceroute monitor</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |



Table 308: traceroute monitor Output Fields

| Field Name   | Field Description                                                                                                       |
|--------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>Host</b>  | Hostname or IP address of the router at each hop.                                                                       |
| <b>Loss%</b> | Percent of packet loss. The number of ping responses divided by the number of ping requests, specified as a percentage. |
| <b>Snt</b>   | Number of ping requests sent to the router at this hop.                                                                 |
| <b>Last</b>  | Most recent round-trip time, in milliseconds, to the router at this hop.                                                |
| <b>Avg</b>   | Average round-trip time, in milliseconds, to the router at this hop.                                                    |
| <b>Best</b>  | Shortest round-trip time, in milliseconds, to the router at this hop.                                                   |
| <b>Wrst</b>  | Longest round-trip time, in milliseconds, to the router at this hop.                                                    |
| <b>StDev</b> | Standard deviation of round-trip times, in milliseconds, to the router at this hop.                                     |

## Sample Output

### traceroute monitor

```
user@host> traceroute monitor 10.16.0.1
```

| Host              | Loss% | Snt | Last | Avg | Best | Wrst | StDev |
|-------------------|-------|-----|------|-----|------|------|-------|
| 1. 10.17.41.254   | 0.0%  | 17  | 0.7  | 1.0 | 0.6  | 5.4  | 1.2   |
| 2. secret.net     | 0.0%  | 17  | 0.6  | 1.0 | 0.6  | 6.6  | 1.4   |
| 3. top-secret.net | 0.0%  | 17  | 0.6  | 0.6 | 0.6  | 0.6  | 0.0   |

## Operational Commands: Port Mirroring

## show analyzer

|                                 |                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show analyzer <i>analyzer-name</i></b>                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.                                                                                                          |
| <b>Description</b>              | Display information about analyzers configured for port mirroring.                                                                                                          |
| <b>Options</b>                  | <b><i>analyzer-name</i></b> —(Optional) Displays the status of a specific analyzer on the switch.                                                                           |
| <b>Required Privilege Level</b> | view                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Understanding Port Mirroring on EX Series Switches</i></li> </ul>                                                                 |
| <b>List of Sample Output</b>    | <a href="#">show analyzer on page 4140</a>                                                                                                                                  |
| <b>Output Fields</b>            | <a href="#">Table 309 on page 4140</a> lists the output fields for the <b>command-name</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 309: show analyzer Output Fields**

| Field Name                          | Field Description                                                                                                                                                                                                                                                       |
|-------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Analyzer name</b>                | Displays the name of the analyzer.                                                                                                                                                                                                                                      |
| <b>Output interface</b>             | Specifies a local interface to which mirrored packets are sent. An analyzer can have output to either an interface or a VLAN, not both.                                                                                                                                 |
| <b>Output VLAN</b>                  | Specifies a VLAN to which mirrored packets are sent. An analyzer can have output to either an interface or a VLAN, not both.                                                                                                                                            |
| <b>Mirror ratio</b>                 | Displays the ratio of packets to be mirrored, between 1 and 2047 where 1 sends copies of all packets and 2047 sends copies of 1 out of every 2047 packets.                                                                                                              |
| <b>Loss priority</b>                | Displays the loss priority of mirrored packets. By default, loss priority is set to <b>low</b> , with mirrored traffic dropped in preference for regular traffic when capacity is exceeded. For analyzers with output to a VLAN, set the loss priority to <b>high</b> . |
| <b>Egress monitored interfaces</b>  | Displays interfaces for which traffic exiting the interfaces is mirrored.                                                                                                                                                                                               |
| <b>Ingress monitored interfaces</b> | Displays interfaces for which traffic entering the interfaces is mirrored.                                                                                                                                                                                              |
| <b>Ingress monitored VLANs</b>      | Displays VLANs for which traffic entering the VLAN is mirrored.                                                                                                                                                                                                         |

## Sample Output

### show analyzer

```
user@host> show analyzer
```

```
Analyzer name : employee-monitor
Output interface : ge-0/0/10.0
Output VLAN : remote-analyzer
Mirror ratio : 1
Loss priority : High
Egress monitored interfaces : ge-0/0/3.0
Ingress monitored interfaces : ge-0/0/0.0
Ingress monitored interfaces : ge-0/0/1.0
```

## Operational Commands: Ethernet OAM Link Fault Management

## show oam ethernet link-fault-management

|                                 |                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show oam ethernet link-fault-management<br><brief   detail><br><interface-name>                                                                                                                        |
| <b>Release Information</b>      | Command introduced in Junos OS Release 8.2.                                                                                                                                                            |
| <b>Description</b>              | On EX Series switches and M320, M120, MX Series, T320, and T640 routers, display Operation, Administration, and Management (OAM) link fault management information for Ethernet interfaces.            |
| <b>Options</b>                  | <b>brief   detail</b> —(Optional) Display the specified level of output.<br><br><b>interface-name</b> —(Optional) Display link fault management information for the specified Ethernet interface only. |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">show oam ethernet link-fault-management brief on page 4146</a><br><a href="#">show oam ethernet link-fault-management detail on page 4146</a>                                              |
| <b>Output Fields</b>            | <a href="#">Table 310 on page 4142</a> lists the output fields for the <b>show oam ethernet link-fault-management</b> command. Output fields are listed in the approximate order in which they appear. |

Table 310: show oam ethernet link-fault-management Output Fields

| Field Name             | Field Description                                                                                                                                                                                                                              | Level of Output |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Status</b>          | Indicates the status of the established link.<br><br><ul style="list-style-type: none"> <li>• <b>Fail</b>—A link fault condition exists.</li> <li>• <b>Running</b>—A link fault condition does not exist.</li> </ul>                           | All levels      |
| <b>Discovery state</b> | State of the discovery mechanism:<br><br><ul style="list-style-type: none"> <li>• <b>Passive Wait</b></li> <li>• <b>Send Any</b></li> <li>• <b>Send Local Remote</b></li> <li>• <b>Send Local Remote Ok</b></li> <li>• <b>Fault</b></li> </ul> | All levels      |
| <b>Peer address</b>    | Address of the OAM peer.                                                                                                                                                                                                                       | All levels      |

Table 310: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Level of Output |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Flags</b>                     | <p>Information about the interface. Possible values are described in the “Link Flags” section under “<a href="#">Common Output Fields Description</a>” on page 2376.</p> <ul style="list-style-type: none"> <li>• <b>Remote-Stable</b>—Indicates remote OAM client acknowledgment of and satisfaction with local OAM state information. <b>False</b> indicates that remote DTE either has not seen or is unsatisfied with local state information. <b>True</b> indicates that remote DTE has seen and is satisfied with local state information.</li> <li>• <b>Local-Stable</b>—Indicates local OAM client acknowledgment of and satisfaction with remote OAM state information. <b>False</b> indicates that local DTE either has not seen or is unsatisfied with remote state information. <b>True</b> indicates that local DTE has seen and is satisfied with remote state information.</li> <li>• <b>Remote-State-Valid</b>—Indicates the OAM client has received remote state information found within Local Information TLVs of received Information OAM PDUs. <b>False</b> indicates that OAM client has not seen remote state information. <b>True</b> indicates that the OAM client has seen remote state information.</li> </ul> | All levels      |
| <b>Remote loopback status</b>    | Indicates the remote loopback status. An OAM entity can put its remote peer into loopback mode using the Loopback control OAM PDU. In loopback mode, every frame received is transmitted back on the same port (except for OAM PDUs, which are needed to maintain the OAM session).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | All levels      |
| <b>Remote entity information</b> | <p>Remote entity information.</p> <ul style="list-style-type: none"> <li>• <b>Remote MUX action</b>—Indicates the state of the multiplexer functions of the OAM sublayer. Device is forwarding non-OAM PDUs to the lower sublayer or discarding non-OAM PDUs.</li> <li>• <b>Remote parser action</b>—Indicates the state of the parser function of the OAM sublayer. Device is forwarding non-OAM PDUs to higher sublayer, looping back non-OAM PDUs to the lower sublayer, or discarding non-OAM PDUs.</li> <li>• <b>Discovery mode</b>—Indicates whether discovery mode is active or inactive.</li> <li>• <b>Unidirectional mode</b>—Indicates the ability to operate a link in a unidirectional mode for diagnostic purposes.</li> <li>• <b>Remote loopback mode</b>—Indicates whether remote loopback is supported or unsupported.</li> <li>• <b>Link events</b>—Indicates whether interpreting link events is supported or unsupported on the remote peer.</li> <li>• <b>Variable requests</b>—Indicates whether variable requests are supported. The Variable Request OAM PDU, is used to request one or more MIB variables from the remote peer.</li> </ul>                                                                        | All levels      |
| <b>OAM Receive Statistics</b>    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                 |
| <b>Information</b>               | The total number of information PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail</b>   |
| <b>Event</b>                     | The total number of loopback control PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail</b>   |
| <b>Variable request</b>          | The total number of variable request PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail</b>   |
| <b>Variable response</b>         | The total number of variable response PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b>   |

Table 310: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name                                         | Field Description                                                                                                                                                                  | Level of Output |
|----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Loopback control</b>                            | The total number of loopback control PDUs received.                                                                                                                                | <b>detail</b>   |
| <b>Organization specific</b>                       | The total number of vendor organization specific PDUs received.                                                                                                                    | <b>detail</b>   |
| <b>OAM Transmit Statistics</b>                     |                                                                                                                                                                                    |                 |
| <b>Information</b>                                 | The total number of information PDUs transmitted.                                                                                                                                  | <b>detail</b>   |
| <b>Event</b>                                       | The total number of event notification PDUs transmitted.                                                                                                                           | <b>detail</b>   |
| <b>Variable request</b>                            | The total number of variable request PDUs transmitted.                                                                                                                             | <b>detail</b>   |
| <b>Variable response</b>                           | The total number of variable response PDUs transmitted.                                                                                                                            | <b>detail</b>   |
| <b>Loopback control</b>                            | The total number of loopback control PDUs transmitted.                                                                                                                             | <b>detail</b>   |
| <b>Organization specific</b>                       | The total number of vendor organization specific PDUs transmitted.                                                                                                                 | <b>detail</b>   |
| <b>OAM Received Symbol Error Event information</b> |                                                                                                                                                                                    |                 |
| <b>Events</b>                                      | The number of symbol error event TLVs that have been received since the OAM sublayer was reset.                                                                                    | <b>detail</b>   |
| <b>Window</b>                                      | The symbol error event window in the received PDU.<br><br>The protocol default value is the number of symbols that can be received in one second on the underlying physical layer. | <b>detail</b>   |
| <b>Threshold</b>                                   | The number of errored symbols in the period required for the event to be generated.                                                                                                | <b>detail</b>   |
| <b>Errors in period</b>                            | The number of symbol errors in the period reported in the received event PDU.                                                                                                      | <b>detail</b>   |
| <b>Total errors</b>                                | The number of errored symbols that have been reported in received event TLVs since the OAM sublayer was reset.<br><br>Symbol errors are coding symbol errors.                      | <b>detail</b>   |
| <b>OAM Received Frame Error Event Information</b>  |                                                                                                                                                                                    |                 |
| <b>Events</b>                                      | The number of errored frame event TLVs that have been received since the OAM sublayer was reset.                                                                                   | <b>detail</b>   |
| <b>Window</b>                                      | The duration of the window in terms of the number of 100 ms period intervals.                                                                                                      | <b>detail</b>   |
| <b>Threshold</b>                                   | The number of detected errored frames required for the event to be generated.                                                                                                      | <b>detail</b>   |
| <b>Errors in period</b>                            | The number of detected errored frames in the period.                                                                                                                               | <b>detail</b>   |

Table 310: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name                                               | Field Description                                                                                                                                                                       | Level of Output |
|----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Total errors</b>                                      | The number of errored frames that have been reported in received event TLVs since the OAM sublayer was reset.<br><br>A frame error is any frame error on the underlying physical layer. | <b>detail</b>   |
| <b>OAM Received Frame Period Error Event Information</b> |                                                                                                                                                                                         |                 |
| <b>Events</b>                                            | The number of frame seconds errors event TLVs that have been received since the OAM sublayer was reset.                                                                                 | <b>detail</b>   |
| <b>Window</b>                                            | The duration of the frame seconds window.                                                                                                                                               | <b>detail</b>   |
| <b>Threshold</b>                                         | The number of frame seconds errors in the period.                                                                                                                                       | <b>detail</b>   |
| <b>Errors in period</b>                                  | The number of frame seconds errors in the period.                                                                                                                                       | <b>detail</b>   |
| <b>Total errors</b>                                      | The number of frame seconds errors that have been reported in received event TLVs since the OAM sublayer was reset.                                                                     | <b>detail</b>   |
| <b>OAM Transmitted Symbol Error Event Information</b>    |                                                                                                                                                                                         |                 |
| <b>Events</b>                                            | The number of symbol error event TLVs that have been transmitted since the OAM sublayer was reset.                                                                                      | <b>detail</b>   |
| <b>Window</b>                                            | The symbol error event window in the transmitted PDU.                                                                                                                                   | <b>detail</b>   |
| <b>Threshold</b>                                         | The number of errored symbols in the period required for the event to be generated.                                                                                                     | <b>detail</b>   |
| <b>Errors in period</b>                                  | The number of symbol errors in the period reported in the transmitted event PDU.                                                                                                        | <b>detail</b>   |
| <b>Total errors</b>                                      | The number of errored symbols reported in event TLVs that have been transmitted since the OAM sublayer was reset.                                                                       | <b>detail</b>   |
| <b>OAM Current Symbol Error Event Information</b>        |                                                                                                                                                                                         |                 |
| <b>Events</b>                                            | The number of symbol error TLVs that have been generated regardless of whether the threshold for sending event TLVs has been crossed.                                                   | <b>detail</b>   |
| <b>Window</b>                                            | The symbol error event window in the transmitted PDU.                                                                                                                                   | <b>detail</b>   |
| <b>Threshold</b>                                         | The number of errored symbols in the period required for the event to be generated.                                                                                                     | <b>detail</b>   |
| <b>Errors in period</b>                                  | The total number of symbol errors in the period reported.                                                                                                                               | <b>detail</b>   |
| <b>Total errors</b>                                      | The number of errored symbols reported in event TLVs that have been generated regardless of whether the threshold for sending event TLVs has been crossed.                              | <b>detail</b>   |
| <b>OAM Transmitted Frame Error Event Information</b>     |                                                                                                                                                                                         |                 |

Table 310: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name                                       | Field Description                                                                                                                            | Level of Output |
|--------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Events</b>                                    | The number of errored frame event TLVs that have been transmitted since the OAM sublayer was reset.                                          | <b>detail</b>   |
| <b>Window</b>                                    | The duration of the window in terms of the number of 100 ms period intervals.                                                                | <b>detail</b>   |
| <b>Threshold</b>                                 | The number of detected errored frames required for the event to be generated.                                                                | <b>detail</b>   |
| <b>Errors in period</b>                          | The number of detected errored frames in the period.                                                                                         | <b>detail</b>   |
| <b>Total errors</b>                              | The number of errored frames that have been detected since the OAM sublayer was reset.                                                       | <b>detail</b>   |
| <b>OAM Current Frame Error Event Information</b> |                                                                                                                                              |                 |
| <b>Events</b>                                    | The number of errored frame event TLVs that have been generated regardless of whether the threshold for sending event TLVs has been crossed. | <b>detail</b>   |
| <b>Window</b>                                    | The duration of the window in terms of the number of 100 ms period intervals.                                                                | <b>detail</b>   |
| <b>Threshold</b>                                 | The number of detected errored frames required for the event to be generated.                                                                | <b>detail</b>   |
| <b>Errors in period</b>                          | The number of errored frames in the period.                                                                                                  | <b>detail</b>   |
| <b>Total errors</b>                              | The number of errored frames detected regardless of whether the threshold for transmitting event TLVs has been crossed.                      | <b>detail</b>   |

## Sample Output

### show oam ethernet link-fault-management brief

```

user@host> show oam ethernet link-fault-management brief
Interface: ge-3/1/3
Status: Running, Discovery state: Send Any
Peer address: 00:90:69:72:2c:83
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
Remote loopback status: Disabled on local port, Enabled on peer port
Remote entity information:
 Remote MUX action: discarding, Remote parser action: loopback
 Discovery mode: active, Unidirectional mode: unsupported
 Remote loopback mode: supported, Link events: supported
 Variable requests: unsupported

```

### show oam ethernet link-fault-management detail

```

user@host> show oam ethernet link-fault-management detail
Interface: ge-6/1/0
Status: Running, Discovery state: Send Any
Peer address: 00:90:69:0a:07:14
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
OAM receive statistics:
 Information: 186365, Event: 0, Variable request: 0, Variable response: 0
 Loopback control: 0, Organization specific: 0

```



OAM transmit statistics:  
Information: 186347, Event: 0, Variable request: 0, Variable response: 0  
Loopback control: 0, Organization specific: 0  
OAM received symbol error event information:  
Events: 0, Window: 0, Threshold: 0  
Errors in period: 0, Total errors: 0  
OAM received frame error event information:  
Events: 0, Window: 0, Threshold: 0  
Errors in period: 0, Total errors: 0  
OAM received frame period error event information:  
Events: 0, Window: 0, Threshold: 0  
Errors in period: 0, Total errors: 0  
OAM transmitted symbol error event information:  
Events: 0, Window: 0, Threshold: 1  
Errors in period: 0, Total errors: 0  
OAM current symbol error event information:  
Events: 0, Window: 0, Threshold: 1  
Errors in period: 0, Total errors: 0  
OAM transmitted frame error event information:  
Events: 0, Window: 0, Threshold: 1  
Errors in period: 0, Total errors: 0  
OAM current frame error event information:  
Events: 0, Window: 0, Threshold: 1  
Errors in period: 0, Total errors: 0  
Remote entity information:  
Remote MUX action: forwarding, Remote parser action: forwarding  
Discovery mode: active, Unidirectional mode: unsupported  
Remote loopback mode: supported, Link events: supported  
Variable requests: unsupported

## show interfaces (Fast Ethernet)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show interfaces <i>interface-type</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Display status information about the specified Fast Ethernet interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b><i>interface-type</i></b>—On M Series and T Series routers, the interface type is <b><i>fe-fpc/pic/port</i></b>. On the J Series routers, the interface type is <b><i>fe-pim/O/port</i></b>.</p> <p><b><i>brief   detail   extensive   terse</i></b>—(Optional) Display the specified level of output.</p> <p><b><i>descriptions</i></b>—(Optional) Display interface description strings.</p> <p><b><i>media</i></b>—(Optional) Display media-specific information about network interfaces.</p> <p><b><i>snmp-index snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b><i>statistics</i></b>—(Optional) Display static interface statistics.</p> |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>List of Sample Output</b>    | <p><a href="#">show interfaces (Fast Ethernet) on page 4162</a></p> <p><a href="#">show interfaces brief (Fast Ethernet) on page 4162</a></p> <p><a href="#">show interfaces detail (Fast Ethernet) on page 4162</a></p> <p><a href="#">show interfaces extensive (Fast Ethernet) on page 4163</a></p>                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Output Fields</b>            | <p><a href="#">Table 311 on page 4148</a> lists the output fields for the <b>show interfaces Fast Ethernet</b> command. Output fields are listed in the approximate order in which they appear.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

**Table 311: show interfaces Fast Ethernet Output Fields**

| Field Name                | Field Description                                                                                                                                             | Level of Output              |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Physical Interface</b> |                                                                                                                                                               |                              |
| <b>Physical interface</b> | Name of the physical interface.                                                                                                                               | All levels                   |
| <b>Enabled</b>            | State of the interface. Possible values are described in the "Enabled Field" section under " <a href="#">Common Output Fields Description</a> " on page 2376. | All levels                   |
| <b>Interface index</b>    | Index number of the physical interface, which reflects its initialization sequence.                                                                           | <b>detail extensive none</b> |
| <b>SNMP ifIndex</b>       | SNMP index number for the physical interface.                                                                                                                 | <b>detail extensive none</b> |

Table 311: show interfaces Fast Ethernet Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                   | Level of Output         |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Generation</b>       | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                   | <b>detail extensive</b> |
| <b>Link-level type</b>  | Encapsulation being used on the physical interface.                                                                                                                                                                                                 | All levels              |
| <b>MTU</b>              | Maximum transmission unit size on the physical interface.                                                                                                                                                                                           | All levels              |
| <b>Link-mode</b>        | Type of link connection configured for the physical interface: <b>Full-duplex</b> or <b>Half-duplex</b>                                                                                                                                             | <b>extensive</b>        |
| <b>Speed</b>            | Speed at which the interface is running.                                                                                                                                                                                                            | All levels              |
| <b>Loopback</b>         | Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .                                                                                                                      | All levels              |
| <b>Source filtering</b> | Source filtering status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                        | All levels              |
| <b>LAN-PHY mode</b>     | 10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.                                                       | All levels              |
| <b>WAN-PHY mode</b>     | 10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.                           | All levels              |
| <b>Unidirectional</b>   | Unidirectional link mode status for 10-Gigabit Ethernet interface: <b>Enabled</b> or <b>Disabled</b> for parent interface; <b>Rx-only</b> or <b>Tx-only</b> for child interfaces.                                                                   | All levels              |
| <b>Flow control</b>     | Flow control status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                            | All levels              |
| <b>Auto-negotiation</b> | (Gigabit Ethernet interfaces) Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                           | All levels              |
| <b>Remote-fault</b>     | (Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> <li>• <b>Online</b>—Autonegotiation is manually configured as online.</li> <li>• <b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul> | All levels              |
| <b>Device flags</b>     | Information about the physical device. Possible values are described in the "Device Flags" section under <a href="#">"Common Output Fields Description" on page 2376</a> .                                                                          | All levels              |
| <b>Interface flags</b>  | Information about the interface. Possible values are described in the "Interface Flags" section under <a href="#">"Common Output Fields Description" on page 2376</a> .                                                                             | All levels              |
| <b>Link flags</b>       | Information about the link. Possible values are described in the "Links Flags" section under <a href="#">"Common Output Fields Description" on page 2376</a> .                                                                                      | All levels              |
| <b>Wavelength</b>       | (10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).                                                                                                              | All levels              |

Table 311: show interfaces Fast Ethernet Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output              |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Frequency</b>               | (10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | All levels                   |
| <b>CoS queues</b>              | Number of CoS queues configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b> none |
| <b>Schedulers</b>              | (GigabitEthernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>extensive</b>             |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b>      |
| <b>Current address</b>         | Configured MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b> none |
| <b>Hardware address</b>        | Hardware MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b> none |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b> none |
| <b>Input Rate</b>              | Input rate in bits per second (bps) and packets per second (pps).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | None specified               |
| <b>Output Rate</b>             | Output rate in bps and pps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | None specified               |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b>      |
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the <a href="#">show interfaces (10-Gigabit Ethernet)</a> command.</p> | <b>detail extensive</b>      |

Table 311: show interfaces Fast Ethernet Output Fields (*continued*)

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

Table 311: show interfaces Fast Ethernet Output Fields (*continued*)

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

Table 311: show interfaces Fast Ethernet Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output              |
|-----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Active alarms and Active defects</b> | <p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the routing device configuration, an alarm can ring the red or yellow alarm bell on the routing device, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value <b>None</b> or <b>Link</b>.</p> <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul> | <b>detail extensive none</b> |
| <b>OTN FEC statistics</b>               | <p>The forward error correction (FEC) counters provide the following statistics:</p> <ul style="list-style-type: none"> <li>• <b>Corrected Errors</b>—The count of corrected errors in the last second.</li> <li>• <b>Corrected Error Ratio</b>—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.</li> </ul>                                                                                                                                                                                                                                                                                                                                                |                              |
| <b>PCS statistics</b>                   | <p>(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.</p> <ul style="list-style-type: none"> <li>• <b>Bit errors</b>—High bit error rate. Indicates the number of bit errors when the PCS receiver is operating in normal mode.</li> <li>• <b>Errored blocks</b>—Loss of block lock. The number of errored blocks when PCS receiver is operating in normal mode.</li> </ul>                                                                                                                                                                                                                                                    | <b>detail extensive</b>      |

Table 311: show interfaces Fast Ethernet Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| MAC statistics                 | <p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> <li>• <b>Total octets</b> and <b>total packets</b>—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the <a href="#">show interfaces (10-Gigabit Ethernet)</a> command.</li> <li>• <b>Unicast packets</b>, <b>Broadcast packets</b>, and <b>Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>• <b>FIFO error</b>—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—Number of frames that exceed 1518 octets.</li> <li>• <b>Jabber frames</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> <li>• <b>Fragment frames</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>VLAN tagged frames</b>—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not.</li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."</li> </ul> | extensive       |
| OTN Received Overhead Bytes    | APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | extensive       |
| OTN Transmitted Overhead Bytes | APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | extensive       |



Table 311: show interfaces Fast Ethernet Output Fields (*continued*)

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

Table 311: show interfaces Fast Ethernet Output Fields (*continued*)

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

Table 311: show interfaces Fast Ethernet Output Fields (*continued*)

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

Table 311: show interfaces Fast Ethernet Output Fields (*continued*)

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

Table 311: show interfaces Fast Ethernet Output Fields (*continued*)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output              |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>CoS information</b>   | Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> <li>• <b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li>• <b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li>• <b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li>• <b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li>• <b>Buffer usec</b>—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li>• <b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li>• <b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul> | <b>extensive</b>             |
| <b>Logical Interface</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                              |
| <b>Logical interface</b> | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels                   |
| <b>Index</b>             | Index number of the logical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive none</b> |
| <b>SNMP ifIndex</b>      | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive none</b> |
| <b>Generation</b>        | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b>      |
| <b>Flags</b>             | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under “ <a href="#">Common Output Fields Description</a> ” on page 2376.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | All levels                   |

Table 311: show interfaces Fast Ethernet Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output                    |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------|
| <b>VLAN-Tag</b>                | <p>Rewrite profile applied to incoming or outgoing frames on the outer (<b>Out</b>) VLAN tag or for both the outer and inner (<b>In</b>) VLAN tags.</p> <ul style="list-style-type: none"> <li>• <b>push</b>—An outer VLAN tag is pushed in front of the existing VLAN tag.</li> <li>• <b>pop</b>—The outer VLAN tag of the incoming frame is removed.</li> <li>• <b>swap</b>—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information.</li> <li>• <b>push</b>—An outer VLAN tag is pushed in front of the existing VLAN tag.</li> <li>• <b>push-push</b>—Two VLAN tags are pushed in from the incoming frame.</li> <li>• <b>swap-push</b>—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.</li> <li>• <b>swap-swap</b>—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value.</li> <li>• <b>pop-swap</b>—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame.</li> <li>• <b>pop-pop</b>—Both the outer and inner VLAN tags of the incoming frame are removed.</li> </ul> | <b>brief detail extensive none</b> |
| <b>Demux:</b>                  | <p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> <li>• Source Family Inet</li> <li>• Destination Family Inet</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail extensive none</b>       |
| <b>Encapsulation</b>           | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels                         |
| <b>Protocol</b>                | Protocol family. Possible values are described in the “Protocol Field” section under “ <a href="#">Common Output Fields Description</a> ” on page 2376.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive none</b>       |
| <b>MTU</b>                     | Maximum transmission unit size on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive none</b>       |
| <b>Maximum labels</b>          | Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail extensive none</b>       |
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the interface set</li> <li>• <b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface set.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b>            |
| <b>IPv6 transit statistics</b> | Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>extensive</b>                   |
| <b>Local statistics</b>        | Number and rate of bytes and packets destined to the routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>extensive</b>                   |

Table 311: show interfaces Fast Ethernet Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Transit statistics</b>       | Number and rate of bytes and packets transiting the switch.<br><br><b>NOTE:</b> For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the <b>Output bytes</b> and <b>Output packets</b> interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler. | <b>extensive</b>             |
| <b>Generation</b>               | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b>      |
| <b>Route Table</b>              | Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive none</b> |
| <b>Flags</b>                    | Information about protocol family flags. Possible values are described in the “Family Flags” section under “ <a href="#">Common Output Fields Description</a> ” on page 2376.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b>      |
| <b>Donor interface</b>          | (Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive none</b> |
| <b>Preferred source address</b> | (Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive none</b> |
| <b>Input Filters</b>            | Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b>      |
| <b>Output Filters</b>           | Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b>      |
| <b>Mac-Validate Failures</b>    | Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive none</b> |
| <b>Addresses, Flags</b>         | Information about the address flags. Possible values are described in the “Addresses Flags” section under “ <a href="#">Common Output Fields Description</a> ” on page 2376.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive none</b> |
| <b><i>protocol-family</i></b>   | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>brief</b>                 |
| <b>Flags</b>                    | Information about address flag (possible values are described in the “Addresses Flags” section under “ <a href="#">Common Output Fields Description</a> ” on page 2376.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail extensive none</b> |
| <b>Destination</b>              | IP address of the remote side of the connection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive none</b> |
| <b>Local</b>                    | IP address of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive none</b> |
| <b>Broadcast</b>                | Broadcast address of the logical interlace.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive none</b> |

Table 311: show interfaces Fast Ethernet Output Fields (*continued*)

| Field Name | Field Description                                                 | Level of Output  |
|------------|-------------------------------------------------------------------|------------------|
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |

## Sample Output

### show interfaces (Fast Ethernet)

```

user@host> show interfaces fe-0/0/0
Physical interface: fe-0/0/0, Enabled, Physical link is Up
 Interface index: 128, SNMP ifIndex: 22
 Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 CoS queues : 4 supported, 4 maximum usable queues
 Current address: 00:05:85:02:38:00, Hardware address: 00:05:85:02:38:00
 Last flapped : 2006-01-20 14:50:58 PST (2w4d 00:44 ago)
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Active alarms : None
 Active defects : None
 Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198)
 Flags: SNMP-Traps Encapsulation: ENET2
 Protocol inet, MTU: 1500
 Flags: None
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255

```

### show interfaces brief (Fast Ethernet)

```

user@host> show interfaces fe-0/0/0 brief
Physical interface: fe-0/0/0, Enabled, Physical link is Up
 Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 Logical interface fe-0/0/0.0
 Flags: SNMP-Traps Encapsulation: ENET2
 inet 10.10.10.1/24

```

### show interfaces detail (Fast Ethernet)

```

user@host> show interfaces fe-0/0/0 detail
Physical interface: fe-0/0/0, Enabled, Physical link is Up
 Interface index: 128, SNMP ifIndex: 22, Generation: 5391
 Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 CoS queues : 4 supported, 4 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:05:85:02:38:00, Hardware address: 00:05:85:02:38:00
 Last flapped : 2006-01-20 14:50:58 PST (2w4d 00:45 ago)
 Statistics last cleared: Never
 Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 42 0 bps

```



```

Input packets: 0 0 pps
Output packets: 1 0 pps
Active alarms : None
Active defects : None
Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
 Flags: SNMP-Traps Encapsulation: ENET2
 Protocol inet, MTU: 1500, Generation: 105, Route table: 0
 Flags: Is-Primary, Mac-Validate-Strict
 Mac-Validate Failures: Packets: 0, Bytes: 0
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255,
 Generation: 136

```

### show interfaces extensive (Fast Ethernet)

```

user@host> show interfaces fe-0/0/0 extensive
Physical interface: fe-0/0/0, Enabled, Physical link is Up
Interface index: 128, SNMP ifIndex: 22, Generation: 5391
Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed:
100mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
CoS queues : 4 supported, 4 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:05:85:02:38:00, Hardware address: 00:05:85:02:38:00
Last flapped : 2006-01-20 14:50:58 PST (2w4d 00:46 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 42 0 bps
Input packets: 0 0 pps
Output packets: 1 0 pps
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
 L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
 FIFO errors: 0, Resource errors: 0
Output errors:
 Carrier transitions: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

 FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Active alarms : None
Active defects : None
MAC statistics:

```

|                    | Receive | Transmit |
|--------------------|---------|----------|
| Total octets       | 0       | 64       |
| Total packets      | 0       | 1        |
| Unicast packets    | 0       | 0        |
| Broadcast packets  | 0       | 1        |
| Multicast packets  | 0       | 0        |
| CRC/Align errors   | 0       | 0        |
| FIFO errors        | 0       | 0        |
| MAC control frames | 0       | 0        |
| MAC pause frames   | 0       | 0        |
| Oversized frames   | 0       |          |
| Jabber frames      | 0       |          |
| Fragment frames    | 0       |          |
| VLAN tagged frames | 0       |          |
| Code violations    | 0       |          |

```

Filter statistics:
Input packet count 0
Input packet rejects 0

```

```
Input DA rejects 0
Input SA rejects 0
Output packet count 1
Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link partner:
 Link partner: Full-duplex, Flow control: None, Remote fault: Ok
Local resolution:
Packet Forwarding Engine configuration:
 Destination slot: 0
CoS information:
 Bandwidth Buffer Priority Limit
 % bps % usec
0 best-effort 95 950000000 95 0 low none
3 network-control 5 500000000 5 0 low none
Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
Flags: SNMP-Traps Encapsulation: ENET2
Protocol inet, MTU: 1500, Generation: 105, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
 Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255,
 Generation: 136
```

## show interfaces (10-Gigabit Ethernet)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>show interfaces <i>xe-fpc/pic/port</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced in Junos OS Release 8.0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | (M320, M120, MX Series, and T Series routers and EX Series switches only) Display status information about the specified 10-Gigabit Ethernet interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><i>xe-fpc/pic/port</i>—Display standard information about the specified 10-Gigabit Ethernet interface.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p>                                                                                                                  |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>List of Sample Output</b>    | <p><a href="#">show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2) on page 4180</a></p> <p><a href="#">show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode) on page 4183</a></p> <p><a href="#">show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC) on page 4185</a></p> <p><a href="#">show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode) on page 4187</a></p> <p><a href="#">show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only) on page 4187</a></p> <p><a href="#">show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only) on page 4188</a></p> |
| <b>Output Fields</b>            | See <a href="#">Table 135 on page 1957</a> for the output fields for the <b>show interfaces</b> (10-Gigabit Ethernet) command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

Table 312: show interfaces Gigabit Ethernet Output Fields

| Field Name                | Field Description                                                                                                                                                                                                                               | Level of Output              |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Physical Interface</b> |                                                                                                                                                                                                                                                 |                              |
| <b>Physical interface</b> | Name of the physical interface.                                                                                                                                                                                                                 | All levels                   |
| <b>Enabled</b>            | State of the interface. Possible values are described in the “Enabled Field” section under <a href="#">“Common Output Fields Description” on page 2376</a> .                                                                                    | All levels                   |
| <b>Interface index</b>    | Index number of the physical interface, which reflects its initialization sequence.                                                                                                                                                             | <b>detail extensive none</b> |
| <b>SNMP ifIndex</b>       | SNMP index number for the physical interface.                                                                                                                                                                                                   | <b>detail extensive none</b> |
| <b>Generation</b>         | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                               | <b>detail extensive</b>      |
| <b>Link-level type</b>    | Encapsulation being used on the physical interface.                                                                                                                                                                                             | All levels                   |
| <b>MTU</b>                | Maximum transmission unit size on the physical interface.                                                                                                                                                                                       | All levels                   |
| <b>Speed</b>              | Speed at which the interface is running.                                                                                                                                                                                                        | All levels                   |
| <b>Loopback</b>           | Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .                                                                                                                  | All levels                   |
| <b>Source filtering</b>   | Source filtering status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                    | All levels                   |
| <b>LAN-PHY mode</b>       | 10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.                                                   | All levels                   |
| <b>WAN-PHY mode</b>       | 10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.                       | All levels                   |
| <b>Unidirectional</b>     | Unidirectional link mode status for 10-Gigabit Ethernet interface: <b>Enabled</b> or <b>Disabled</b> for parent interface; <b>Rx-only</b> or <b>Tx-only</b> for child interfaces.                                                               | All levels                   |
| <b>Flow control</b>       | Flow control status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                        | All levels                   |
| <b>Auto-negotiation</b>   | (Gigabit Ethernet interfaces) Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                       | All levels                   |
| <b>Remote-fault</b>       | (Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> <li><b>Online</b>—Autonegotiation is manually configured as online.</li> <li><b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul> | All levels                   |
| <b>Device flags</b>       | Information about the physical device. Possible values are described in the “Device Flags” section under <a href="#">“Common Output Fields Description” on page 2376</a> .                                                                      | All levels                   |
| <b>Interface flags</b>    | Information about the interface. Possible values are described in the “Interface Flags” section under <a href="#">“Common Output Fields Description” on page 2376</a> .                                                                         | All levels                   |

Table 312: show interfaces Gigabit Ethernet Output Fields (*continued*)

| Field Name                         | Field Description                                                                                                                                                                                                                                                           | Level of Output       |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| <b>Link flags</b>                  | Information about the link. Possible values are described in the “Links Flags” section under “ <a href="#">Common Output Fields Description</a> ” on page 2376.                                                                                                             | All levels            |
| <b>Wavelength</b>                  | (10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).                                                                                                                                      | All levels            |
| <b>Frequency</b>                   | (10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).                                                                                                                                            | All levels            |
| <b>CoS queues</b>                  | Number of CoS queues configured.                                                                                                                                                                                                                                            | detail extensive none |
| <b>Schedulers</b>                  | (Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.                                                                                                                                                                         | extensive             |
| <b>Hold-times</b>                  | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                                         | detail extensive      |
| <b>Current address</b>             | Configured MAC address.                                                                                                                                                                                                                                                     | detail extensive none |
| <b>Hardware address</b>            | Hardware MAC address.                                                                                                                                                                                                                                                       | detail extensive none |
| <b>Last flapped</b>                | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .                        | detail extensive none |
| <b>Input Rate</b>                  | Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. | None specified        |
| <b>Output Rate</b>                 | Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.                                        | None specified        |
| <b>Statistics last cleared</b>     | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                           | detail extensive      |
| <b>Egress accounting overhead</b>  | Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.                                                                                                                                                                                 | detail extensive      |
| <b>Ingress accounting overhead</b> | Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.                                                                                                                                                                                | detail extensive      |

Table 312: show interfaces Gigabit Ethernet Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output         |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Traffic statistics</b> | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see <a href="#">Table 135 on page 1957</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail extensive</b> |
| <b>Input errors</b>       | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the <b>ignore-l3-incompletes</b> statement.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b>        |

Table 312: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 312: show interfaces Gigabit Ethernet Output Fields (*continued*)

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



Table 312: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 312: show interfaces Gigabit Ethernet Output Fields (*continued*)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Level of Output  |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>MAC statistics</b>                 | <p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> <li>• <b>Total octets</b> and <b>total packets</b>—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see <a href="#">Table 136 on page 1971</a></li> <li>• <b>Unicast packets</b>, <b>Broadcast packets</b>, and <b>Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>• <b>FIFO error</b>—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—Number of frames that exceed 1518 octets.</li> <li>• <b>Jabber frames</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> <li>• <b>Fragment frames</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>VLAN tagged frames</b>—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not.</li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."</li> </ul> | <b>extensive</b> |
| <b>OTN Received Overhead Bytes</b>    | APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>extensive</b> |
| <b>OTN Transmitted Overhead Bytes</b> | APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>extensive</b> |

Table 312: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 312: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 312: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 312: show interfaces Gigabit Ethernet Output Fields (*continued*)

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

Table 312: show interfaces Gigabit Ethernet Output Fields (*continued*)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output              |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>CoS information</b>   | Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> <li>• <b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li>• <b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li>• <b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li>• <b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li>• <b>Buffer usec</b>—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li>• <b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li>• <b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul> | <b>extensive</b>             |
| <b>Logical Interface</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                              |
| <b>Logical interface</b> | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels                   |
| <b>Index</b>             | Index number of the logical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive</b> none |
| <b>SNMP ifIndex</b>      | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b> none |
| <b>Generation</b>        | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b>      |
| <b>Flags</b>             | Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under " <a href="#">Common Output Fields Description</a> " on page 2376.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | All levels                   |

Table 312: show interfaces Gigabit Ethernet Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output                    |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------|
| <b>VLAN-Tag</b>                | <p>Rewrite profile applied to incoming or outgoing frames on the outer (<b>Out</b>) VLAN tag or for both the outer and inner (<b>In</b>) VLAN tags.</p> <ul style="list-style-type: none"> <li><b>push</b>—An outer VLAN tag is pushed in front of the existing VLAN tag.</li> <li><b>pop</b>—The outer VLAN tag of the incoming frame is removed.</li> <li><b>swap</b>—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information.</li> <li><b>push</b>—An outer VLAN tag is pushed in front of the existing VLAN tag.</li> <li><b>push-push</b>—Two VLAN tags are pushed in from the incoming frame.</li> <li><b>swap-push</b>—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.</li> <li><b>swap-swap</b>—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value.</li> <li><b>pop-swap</b>—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame.</li> <li><b>pop-pop</b>—Both the outer and inner VLAN tags of the incoming frame are removed.</li> </ul> | <b>brief detail extensive none</b> |
| <b>Demux:</b>                  | <p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> <li>Source Family Inet</li> <li>Destination Family Inet</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive none</b>       |
| <b>Encapsulation</b>           | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | All levels                         |
| <b>Protocol</b>                | Protocol family. Possible values are described in the “Protocol Field” section under “ <a href="#">Common Output Fields Description</a> ” on page 2376.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive none</b>       |
| <b>MTU</b>                     | Maximum transmission unit size on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b>       |
| <b>Maximum labels</b>          | Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive none</b>       |
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> <li><b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.</li> <li><b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface set.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b>            |
| <b>IPv6 transit statistics</b> | Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>extensive</b>                   |
| <b>Local statistics</b>        | Number and rate of bytes and packets destined to the routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>extensive</b>                   |



Table 312: show interfaces Gigabit Ethernet Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Transit statistics</b>       | Number and rate of bytes and packets transiting the switch.<br><br><b>NOTE:</b> For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the <b>Output bytes</b> and <b>Output packets</b> interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler. | <b>extensive</b>             |
| <b>Generation</b>               | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b>      |
| <b>Route Table</b>              | Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive none</b> |
| <b>Flags</b>                    | Information about protocol family flags. Possible values are described in the “Family Flags” section under “ <a href="#">Common Output Fields Description</a> ” on page 2376.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b>      |
| <b>Donor interface</b>          | (Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive none</b> |
| <b>Preferred source address</b> | (Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive none</b> |
| <b>Input Filters</b>            | Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b>      |
| <b>Output Filters</b>           | Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b>      |
| <b>Mac-Validate Failures</b>    | Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive none</b> |
| <b>Addresses, Flags</b>         | Information about the address flags. Possible values are described in the “Addresses Flags” section under “ <a href="#">Common Output Fields Description</a> ” on page 2376.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive none</b> |
| <b><i>protocol-family</i></b>   | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>brief</b>                 |
| <b>Flags</b>                    | Information about address flag (possible values are described in the “Addresses Flags” section under “ <a href="#">Common Output Fields Description</a> ” on page 2376.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail extensive none</b> |
| <b>Destination</b>              | IP address of the remote side of the connection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive none</b> |
| <b>Local</b>                    | IP address of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive none</b> |
| <b>Broadcast</b>                | Broadcast address of the logical interlace.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive none</b> |

Table 312: show interfaces Gigabit Ethernet Output Fields (*continued*)

| Field Name        | Field Description                                                 | Level of Output         |
|-------------------|-------------------------------------------------------------------|-------------------------|
| <b>Generation</b> | Unique number for use by Juniper Networks technical support only. | <b>detail extensive</b> |

For Gigabit Ethernet IQ PICs, traffic and MAC statistics output varies. [Table 136 on page 1971](#) describes the traffic and MAC statistics for two sample interfaces, each of which is sending traffic in packets of 500 bytes (including 478 bytes for the Layer 3 packet, 18 bytes for the Layer 2 VLAN traffic header, and 4 bytes for cyclic redundancy check [CRC] information). In [Table 136 on page 1971](#), the **ge-0/3/0** interface is the inbound physical interface, and the **ge-0/0/0** interface is the outbound physical interface. On both interfaces, traffic is carried on logical unit .50 (VLAN 50).

Table 313: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

| Interface Type              | Sample Command                               | Byte and Octet Counts Include                                                                                                                                                                                 | Comments                                                                                                                                      |
|-----------------------------|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| Inbound physical interface  | <b>show interfaces ge-0/3/0 extensive</b>    | Traffic statistics:<br><br>Input bytes: 496 bytes per packet, representing the Layer 2 packet<br><br>MAC statistics:<br><br>Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes  | The additional 4 bytes are for the CRC.                                                                                                       |
| Inbound logical interface   | <b>show interfaces ge-0/3/0.50 extensive</b> | Traffic statistics:<br><br>Input bytes: 478 bytes per packet, representing the Layer 3 packet                                                                                                                 |                                                                                                                                               |
| Outbound physical interface | <b>show interfaces ge-0/0/0 extensive</b>    | Traffic statistics:<br><br>Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes<br><br>MAC statistics:<br><br>Received octets: 478 bytes per packet, representing the Layer 3 packet | For input bytes, the additional 12 bytes includes 6 bytes for the destination MAC address + 4 bytes for VLAN + 2 bytes for the Ethernet type. |
| Outbound logical interface  | <b>show interfaces ge-0/0/0.50 extensive</b> | Traffic statistics:<br><br>Input bytes: 478 bytes per packet, representing the Layer 3 packet                                                                                                                 |                                                                                                                                               |

## Sample Output

### show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2)

```

user@host> show interfaces xe-5/0/0 extensive
Physical interface: xe-5/0/0, Enabled, Physical link is Up
 Interface index: 177, SNMP ifIndex: 99, Generation: 178
 Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Speed: 10Gbps, Loopback:

```

```

None, Source filtering: Enabled,
Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags : None
CoS queues : 8 supported, 4 maximum usable queues
Schedulers : 1024
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:14:f6:b9:f1:f6, Hardware address: 00:14:f6:b9:f1:f6
Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes : 6970332384 0 bps
Output bytes : 0 0 bps
Input packets: 81050506 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Ingress traffic statistics at Packet Forwarding Engine:
Input bytes : 6970299398 0 bps
Input packets: 81049992 0 pps
Drop bytes : 0 0 bps
Drop packets: 0 0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0,
MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets

0 best-effort 81049992 81049992 0
1 expedited-fo 0 0 0
2 assured-forw 0 0 0
3 network-cont 0 0 0

Egress queues: 4 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets

0 best-effort 0 0 0
1 expedited-fo 0 0 0
2 assured-forw 0 0 0
3 network-cont 0 0 0

Active alarms : None
Active defects : None
PCS statistics Seconds
Bit errors 0
Errored blocks 0

```

```

MAC statistics:
Total octets 6970332384
Total packets 81050506
Unicast packets 81050000
Broadcast packets 506
Multicast packets 0
CRC/Align errors 0
FIFO errors 0
MAC control frames 0
MAC pause frames 0
Oversized frames 0
Jabber frames 0
Fragment frames 0
VLAN tagged frames 0
Code violations 0

Filter statistics:
Input packet count 81050506
Input packet rejects 506
Input DA rejects 0
Input SA rejects 0
Output packet count 0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0

Packet Forwarding Engine configuration:
Destination slot: 5

CoS information:
Direction : Output
CoS transmit queue Bandwidth Buffer Priority Limit
 % bps % usec
0 best-effort 95 950000000 95 0 low none
3 network-control 5 50000000 5 0 low none

Direction : Input
CoS transmit queue Bandwidth Buffer Priority Limit
 % bps % usec
0 best-effort 95 950000000 95 0 low none
3 network-control 5 50000000 5 0 low none

Logical interface xe-5/0/0.0 (Index 71) (SNMP ifIndex 95) (Generation 195)
Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.100] Encapsulation: ENET2
Egress accounting overhead: 100
Ingress accounting overhead: 90

Traffic statistics:
Input bytes : 0
Output bytes : 46
Input packets: 0
Output packets: 1

IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Local statistics:
Input bytes : 0
Output bytes : 46
Input packets: 0
Output packets: 1

Transit statistics:
Input bytes : 0
Output bytes : 0

```

```

Input packets: 0 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Protocol inet, MTU: 1500, Generation: 253, Route table: 0
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 192.1.1/24, Local: 192.1.1.1, Broadcast: 192.1.1.255,
Generation: 265
Protocol multiservice, MTU: Unlimited, Generation: 254, Route table: 0
 Flags: None
 Policer: Input: __default_arp_policer__

```

### show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode)

```

user@host> show interfaces xe-1/0/0 extensive
Physical interface: xe-1/0/0, Enabled, Physical link is Up
Interface index: 141, SNMP ifIndex: 34, Generation: 47
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Loopback: Disabled
WAN-PHY mode
Source filtering: Disabled, Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps 16384
Link flags : None
CoS queues : 4 supported
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:05:85:a2:10:9d, Hardware address: 00:05:85:a2:10:9d
Last flapped : 2005-07-07 11:22:34 PDT (3d 12:28 ago)
Statistics last cleared: Never
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
 L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
 HS Link CRC errors: 0, HS Link FIFO overflows: 0,
 Resource errors: 0
Output errors:
 Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0,
 Aged packets: 0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
 Resource errors: 0
Queue counters:
 Queued packets Transmitted packets Dropped packets
0 best-effort 0 0 0
1 expedited-fo 0 0 0
2 assured-forw 0 0 0
3 network-cont 0 0 0
Active alarms : LOL, LOS, LBL
Active defects: LOL, LOS, LBL, SEF, AIS-L, AIS-P
PCS statistics
 Seconds Count
Bit errors 0 0
Errored blocks 0 0
MAC statistics:
 Receive Transmit
Total octets 0 0
Total packets 0 0
Unicast packets 0 0
Broadcast packets 0 0
Multicast packets 0 0

```

```

CRC/Align errors 0 0
FIFO errors 0 0
MAC control frames 0 0
MAC pause frames 0 0
Oversized frames 0
Jabber frames 0
Fragment frames 0
VLAN tagged frames 0
Code violations 0
Filter statistics:
 Input packet count 0
 Input packet rejects 0
 Input DA rejects 0
 Input SA rejects 0
 Output packet count 0
 Output packet pad count 0
 Output packet error count 0
CAM destination filters: 0, CAM source filters: 0
PMA PHY:
 Seconds Count State
 PLL lock 0 0 OK
 PHY light 63159 1 Light Missing
WIS section:
 BIP-B1 0 0
 SEF 434430 434438 Defect Active
 LOS 434430 1 Defect Active
 LOF 434430 1 Defect Active
 ES-S 434430
 SES-S 434430
 SEFS-S 434430
WIS line:
 BIP-B2 0 0
 REI-L 0 0
 RDI-L 0 0 OK
 AIS-L 434430 1 Defect Active
 BERR-SF 0 0 OK
 BERR-SD 0 0 OK
 ES-L 434430
 SES-L 434430
 UAS-L 434420
 ES-LFE 0
 SES-LFE 0
 UAS-LFE 0
WIS path:
 BIP-B3 0 0
 REI-P 0 0
 LOP-P 0 0 OK
 AIS-P 434430 1 Defect Active
 RDI-P 0 0 OK
 UNEQ-P 0 0 OK
 PLM-P 0 0 OK
 ES-P 434430
 SES-P 434430
 UAS-P 434420
 ES-PFE 0
 SES-PFE 0
 UAS-PFE 0
Received path trace:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Transmitted path trace: orissa so-1/0/0
6f 72 69 73 73 61 20 73 6f 2d 31 2f 30 2f 30 00 orissa so-1/0/0.
Packet Forwarding Engine configuration:

```

```

Destination slot: 1
CoS information:
 CoS transmit queue Bandwidth Buffer Priority Limit
 % bps % bytes
0 best-effort 95 950000000 95 0 low none
3 network-control 5 50000000 5 0 low none

```

### show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC)

```

user@host> show interfaces ge-7/0/0 extensive
Physical interface: ge-7/0/0, Enabled, Physical link is Down
Interface index: 143, SNMP ifIndex: 508, Generation: 208
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Enabled
Device flags : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
Link flags : None
Wavelength : 1550.12 nm, Frequency: 193.40 THz
CoS queues : 8 supported, 8 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:05:85:70:2b:72, Hardware address: 00:05:85:70:2b:72
Last flapped : 2011-04-20 15:48:54 PDT (18:39:49 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets : 0 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets : 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 2, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets

0 best-effort 0 0 0

1 expedited-fo 0 0 0

2 assured-forw 0 0 0

3 network-cont
Queue number: Mapped forwarding classes
0 best-effort
1 expedited-forwarding
2 assured-forwarding
3 network-control
Active alarms : LINK
Active defects : LINK
MAC statistics:
Total octets Receive Transmit
Total packets 0 0

```

```

Unicast packets 0 0
Broadcast packets 0 0
Multicast packets 0 0
CRC/Align errors 0 0
FIFO errors 0 0
MAC control frames 0 0
MAC pause frames 0 0
Oversized frames 0
Jabber frames 0
Fragment frames 0
VLAN tagged frames 0
Code violations 0
Total octets 0 0
Total packets 0 0
Unicast packets 0 0
Broadcast packets 0 0
Multicast packets 0 0
CRC/Align errors 0 0
FIFO errors 0 0
MAC control frames 0 0
MAC pause frames 0 0
Oversized frames 0
Jabber frames 0
Fragment frames 0
VLAN tagged frames 0
Code violations 0
OTN alarms : None
OTN defects : None
OTN FEC Mode : GFEC
OTN Rate : Fixed Stuff Bytes 11.0957Gbps
OTN Line Loopback : Enabled
OTN FEC statistics :
 Corrected Errors 0
 Corrected Error Ratio (0 sec average) 0e-0
OTN FEC alarms: Seconds Count State
 FEC Degrade 0 0 OK
 FEC Excessive 0 0 OK
OTN OC: Seconds Count State
 LOS 2 1 OK
 LOF 67164 2 Defect Active
 LOM 67164 71 Defect Active
 Wavelength Lock 0 0 OK
OTN OTU:
 AIS 0 0 OK
 BDI 65919 4814 Defect Active
 IAE 67158 1 Defect Active
 TTIM 7 1 OK
 SF 67164 2 Defect Active
 SD 67164 3 Defect Active
 TCA-ES 0 0 OK
 TCA-SES 0 0 OK
 TCA-UAS 80 40 OK
 TCA-BBE 0 0 OK
 BIP 0 0 OK
 BBE 0 0 OK
 ES 0 0 OK
 SES 0 0 OK
 UAS 587 0 OK
Received DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00
Received SAPI:

```



```

00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Transmitted DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Transmitted SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
OTN Received Overhead Bytes:
 APS/PCC0: 0x02, APS/PCC1: 0x42, APS/PCC2: 0xa2, APS/PCC3: 0x48
 Payload Type: 0x03
OTN Transmitted Overhead Bytes:
 APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
 Payload Type: 0x03
Filter statistics:
 Input packet count 0
 Input packet rejects 0
 Input DA rejects 0
 Input SA rejects 0
 Output packet count 0
 Output packet pad count 0
 Output packet error count 0
 CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
 Destination slot: 7
CoS information:
 Direction : Output
 CoS transmit queue Bandwidth Buffer Priority
Limit % bps % usec low
 0 best-effort 95 9500000000 95 0 low
none
 3 network-control 5 500000000 5 0 low
none
...

```

#### show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode)

```

user@host> show interfaces xe-7/0/0 extensive
Physical interface: xe-7/0/0, Enabled, Physical link is Up
 Interface index: 173, SNMP ifIndex: 212, Generation: 174
 Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
 Unidirectional: Enabled,
 Loopback: None, Source filtering: Disabled, Flow control: Enabled
 Device flags : Present Running
...

```

#### show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only)

```

user@host> show interfaces xe-7/0/0-tx extensive
Physical interface: xe-7/0/0-tx, Enabled, Physical link is Up
 Interface index: 176, SNMP ifIndex: 137, Generation: 177
 Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
 Unidirectional: Tx-Only
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 Link flags : None
 CoS queues : 8 supported, 8 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:05:85:73:e4:83, Hardware address: 00:05:85:73:e4:83
 Last flapped : 2007-06-01 09:08:19 PDT (3d 02:31 ago)
 Statistics last cleared: Never
Traffic statistics:
 Input bytes : 0 0 bps

```

```

Output bytes : 322891152287160 9627472888 bps
Input packets: 0 0 pps
Output packets: 328809727380 1225492 pps

...

Filter statistics:
 Output packet count 328810554250
 Output packet pad count 0
 Output packet error count 0
...

Logical interface xe-7/0/0-tx.0 (Index 73) (SNMP ifIndex 138) (Generation 139)

Flags: SNMP-Traps Encapsulation: ENET2
Egress accounting overhead: 100
Ingress accounting overhead: 90
Traffic statistics:
 Input bytes : 0
 Output bytes : 322891152287160
 Input packets: 0
 Output packets: 328809727380
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 322891152287160 9627472888 bps
 Input packets: 0 0 pps
 Output packets: 328809727380 1225492 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Protocol inet, MTU: 1500, Generation: 147, Route table: 0
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 10.11.12/24, Local: 10.11.12.13, Broadcast: 10.11.12.255,
 Generation: 141
Protocol multiservice, MTU: Unlimited, Generation: 148, Route table: 0
 Flags: None
 Policer: Input: __default_arp_policer__

```

#### show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only)

```

user@host> show interfaces xe-7/0/0-rx extensive
Physical interface: xe-7/0/0-rx, Enabled, Physical link is Up
 Interface index: 174, SNMP ifIndex: 118, Generation: 175
 Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
 Unidirectional: Rx-Only
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 Link flags : None
 CoS queues : 8 supported, 8 maximum usable queues

```

```

Hold-times : Up 0 ms, Down 0 ms
Current address: 00:05:85:73:e4:83, Hardware address: 00:05:85:73:e4:83
Last flapped : 2007-06-01 09:08:22 PDT (3d 02:31 ago)
Statistics last cleared: Never
Traffic statistics:
 Input bytes : 322857456303482 9627496104 bps
 Output bytes : 0 0 bps
 Input packets: 328775413751 1225495 pps
 Output packets: 0 0 pps

...

Filter statistics:
 Input packet count 328775015056
 Input packet rejects 1
 Input DA rejects 0

...

Logical interface xe-7/0/0-rx.0 (Index 72) (SNMP ifIndex 120) (Generation 138)

Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
 Input bytes : 322857456303482
 Output bytes : 0
 Input packets: 328775413751
 Output packets: 0
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Transit statistics:
 Input bytes : 322857456303482 9627496104 bps
 Output bytes : 0 0 bps
 Input packets: 328775413751 1225495 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Protocol inet, MTU: 1500, Generation: 145, Route table: 0
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 192.1.1/24, Local: 192.1.1.1, Broadcast: 192.1.1.255,
Generation: 139
 Protocol multiservice, MTU: Unlimited, Generation: 146, Route table: 0
 Flags: None
 Policer: Input: __default_arp_policer__

```



## CHAPTER 17

# Port Mirroring

- [Overview on page 4191](#)
- [Configuration on page 4207](#)
- [Administration on page 4263](#)

## Overview

---

- [Port Mirroring Overview on page 4191](#)
- [Configuration Guidelines for Layer 2 Port Mirroring on page 4201](#)
- [Behavior of Layer 2 Port Mirroring at Physical Interfaces on page 4204](#)
- [Behavior of Layer 2 Port Mirroring on PE Routers and PE Switches on page 4205](#)

## Port Mirroring Overview

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Layer 2 Port Mirroring Properties on page 4192](#)
- [Layer 2 Port Mirroring Global Instance on page 4194](#)
- [Layer 2 Port Mirroring Named Instances on page 4194](#)
- [Layer 2 Port Mirroring of PE Router Logical Interfaces on page 4196](#)
- [Layer 2 Port Mirroring Firewall Filters on page 4198](#)
- [Layer 2 Port Mirroring to Multiple Destinations Using Next-Hop Groups on page 4200](#)

## Layer 2 Port Mirroring Overview

---

On routing platforms and switches that contain an Internet Processor II ASIC, you can send a copy of any incoming packet from the routing platform or switch to an external host address or a packet analyzer for analysis. This is known as *port mirroring*. In Junos OS Release 9.3 and later, Juniper Networks MX Series 3D Universal Edge Routers in a Layer 2 environment support port mirroring for Layer 2 bridging traffic and virtual private LAN service (VPLS) traffic. In Junos OS Release 9.4 and later, MX Series routers in a Layer 2 environment also support port mirroring for Layer 2 VPN traffic over a circuit cross-connect (CCC) that transparently connects logical interfaces of the same type. In Junos OS Release 12.3R2, Juniper Networks EX Series switches support port mirroring for Layer 2 bridging traffic.

Layer 2 port mirroring enables you to specify the manner in which incoming and outgoing packets at specified ports are monitored and the manner in which copies of selected packets are forwarded to another destination, where the packets can be analyzed. MX Series routers and EX Series switches support Layer 2 port mirroring by performing flow monitoring functions using a class-of-service (CoS) architecture that is in concept similar to, but in particulars different from, other routing platforms and switches.

Like the M120 Multiservice Edge Router and M320 Multiservice Edge Routers, MX Series routers and EX Series switches support port mirroring of IPv4, IPv6, and VPLS packets simultaneously. However, the *Junos OS Layer 2 Configuration Guide* describes port mirroring only for Layer 2 bridging traffic (**family ethernet-switching**), Layer 2 VPLS traffic (**family vpls**) through an MX Series router, and Layer 2 VPN traffic that passes through a CCC (**family ccc**).

For general information about packet flow within MX Series routers and other routers, see the *Junos OS Class of Service Configuration Guide*.

In a Layer 3 environment, MX Series routers and EX Series switches support port mirroring of IPv4 (**family inet**) and IPv6 (**family inet6**) traffic. For information about Layer 3 port mirroring, see the *Routing Policy Configuration Guide*.

### Related Documentation

- [Layer 2 Port Mirroring Properties on page 4192](#)
- [Restrictions on Layer 2 Port Mirroring](#)
- [Application of Layer 2 Port Mirroring Types](#)

## Layer 2 Port Mirroring Properties

---

Port mirroring specifies the following types of properties:

- [Packet-Selection Properties on page 4193](#)
- [Packet Address Family on page 4193](#)
- [Mirror Destination Properties on page 4193](#)
- [Mirror-Once Option on page 4194](#)

### ***Packet-Selection Properties***

The packet-selection properties of Layer 2 port-mirroring specify how the sampled packets are to be selected for mirroring:

- The number of packets in each sample.
- The number of packets to mirror from each sample.
- The length to which mirrored packets are to be truncated.

### ***Packet Address Family***

The packet address family type specifies the type of traffic to be mirrored. In a Layer 2 environment, MX Series routers and EX Series switches support port mirroring for the following packet address families:

- Family type **ethernet-switching**—For mirroring VPLS traffic when the physical interface is configured with encapsulation type **ethernet-bridge**.
- Family type **ccc**—For mirroring Layer 2 VPN traffic.
- Family type **vpls**—For mirroring VPLS traffic.



**NOTE:** In typical applications, you send mirrored packets directly to an analyzer or a workstation for analysis, not to another router or switch. If you must send mirrored packets over a network, you should use tunnels. For Layer 2 VPN implementations, you can use the Layer 2 VPN routing instance type **l2vpn** to tunnel the packets to a remote destination.

For information about configuring a routing instance for Layer 2 VPN, see the *Junos OS VPNs Configuration Guide*. For a detailed Layer 2 VPN example configuration, see the *Junos OS Feature Guides*. For information about tunnel interfaces, see the *Junos® OS Network Interfaces*.

### ***Mirror Destination Properties***

For a given packet address family, the mirror destination properties of a Layer 2 port-mirroring instance specify how the selected packets are to be sent on a particular physical interface:

- The physical interface on which to send the selected packets.
- Whether filter checking is to be disabled for the mirror destination interface. By default, filter checking is enabled on all



**NOTE:** If you apply a filter to an interface that is also a Layer 2 port-mirroring destination, a commit failure occurs unless you have disabled filter checking for that mirror destination interface.

### **Mirror-Once Option**

If port mirroring is enabled at both ingress and egress interfaces, you can prevent the MX Series router and an EX Series switch from sending duplicate packets to the same destination (which would complicate the analysis of the mirrored traffic).



**NOTE:** The mirror-once port-mirroring option is a global setting. The option is independent of the packet selection properties and the packet family type-specific mirror destination properties.

#### **Related Documentation**

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Restrictions on Layer 2 Port Mirroring](#)
- [Application of Layer 2 Port Mirroring Types](#)

---

### **Layer 2 Port Mirroring Global Instance**

On an MX Series router and on an EX Series switch, you can configure a set of port-mirroring properties that implicitly apply to packets received on all ports in the router (or switch) chassis. This set of port-mirroring properties is the *global instance* of Layer 2 port mirroring for the router or switch.

Within the global instance configuration, you can specify a set of mirror destination properties for each packet address family supported by Layer 2 port mirroring.

For a general description of Layer 2 port-mirroring properties, see “[Layer 2 Port Mirroring Properties](#)” on page 4192. For a comparison of the types of Layer 2 port mirroring available on an MX Series router and on an EX Series switch, see [Application of Layer 2 Port Mirroring Types](#).

#### **Related Documentation**

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Configuring the Global Instance of Layer 2 Port Mirroring](#)
- [Examples: Layer 2 Port-Mirroring at Multiple Levels of the Chassis on page 4218](#)
- [Example: Layer 2 Port Mirroring with Multiple Instances](#)
- [Example: Layer 2 Port Mirroring to Multiple Destinations](#)

---

### **Layer 2 Port Mirroring Named Instances**

This topic describes the following information:

- [Layer 2 Port Mirroring Named Instances Overview on page 4195](#)
- [Mirroring at Ports Grouped at the FPC Level on page 4195](#)
- [Mirroring at Ports Grouped at the PIC Level on page 4196](#)
- [Mirroring at a Group of Ports Bound to Multiple Named Instances on page 4196](#)



### Layer 2 Port Mirroring Named Instances Overview

On an MX Series router and on an EX Series switch, you can define a set of port-mirroring properties that you can explicitly bind to physical ports on the router or switch. This set of port mirroring properties is known as a *named instance* of Layer 2 port mirroring.

You can bind a named instance of Layer 2 port mirroring to physical ports associated with an MX Series router's or an EX Series switch's Packet Forwarding Engine components at different levels of the router (or switch) chassis:

- At the FPC level—You can bind a named instance to the physical ports associated with a specific Dense Port Concentrator (DPC) or to the physical ports associated with a specific Flexible Port Concentrator (FPC).
- At the PIC level—You can bind a named instance of port mirroring to a specific Packet Forwarding Engine (on a specific DPC) or to a specific PIC.



**NOTE:** MX Series routers support DPCs as well as FPCs and PICs. Unlike FPCs, DPCs do not support PICs. In the Junos OS CLI, however, you use FPC and PIC syntax to configure or display information about DPCs and the Packet Forwarding Engines on the DPCs.

The following points summarize the behavior of Layer 2 port mirroring based on named instances:

- The scope of packet selection is determined by the target of the binding—At the ports (or port) bound to a named instance of Layer 2 port mirroring, the router or switch selects input packets according to the packet-selection properties in the named instance.
- The destination of a selected packet is determined by the packet address family—Of the packets selected, the router or switch mirrors only the packets belonging to an address family for which the named instance of Layer 2 port mirroring specifies a set of mirror destination properties. In a Layer 2 environment, MX Series routers and EX Series switches support port mirroring of VPLS (**family ethernet-switching** or **family vpls**) traffic and Layer 2 VPN traffic with **family ccc**.

For a general description of Layer 2 port-mirroring properties, see "[Layer 2 Port Mirroring Properties](#)" on page 4192. For a comparison of the types of Layer 2 port mirroring available on an MX Series router and on an EX Series switch, see *Application of Layer 2 Port Mirroring Types*.

### Mirroring at Ports Grouped at the FPC Level

On an MX Series router and on an EX Series switch, you can bind a named instance of Layer 2 port mirroring to a specific DPC or FPC installed in the router (or switch) chassis. The port mirroring properties in the instance are applied to all Packet Forwarding Engines (and their associated ports) on the specified DPC or to all PICs (and their associated ports) installed in the specified FPC. Port mirroring properties that are bound to a DPC

or FPC override any port-mirroring properties bound at the global level or the MX Series router (or switch) chassis.

#### ***Mirroring at Ports Grouped at the PIC Level***

On an MX Series router and on an EX Series switch, you can bind a named instance of Layer 2 port mirroring to a specific Packet Forwarding Engine or PIC. The port-mirroring properties in that instance are applied to all ports associated with the specified Packet Forwarding Engine or PIC. Port-mirroring properties that are bound to a Packet Forwarding Engine or PIC override any port-mirroring properties bound at the DPC or FPC that contains them.



**NOTE:** For MX960 routers, there is a one-to-one mapping of Packet Forwarding Engines to Ethernet ports. Therefore, on MX960 routers only, you can configure port-specific bindings of port-mirroring instances.

#### ***Mirroring at a Group of Ports Bound to Multiple Named Instances***

On an MX Series router and on an EX Series switch, you can apply up to two named instances of Layer 2 port mirroring to the same group of ports within the router (or switch) chassis. By applying two different port-mirroring instances to the same DPC, FPC, Packet Forwarding Engine, or PIC, you can bind two distinct Layer 2 port mirroring specifications to a single group of ports.



**NOTE:** You can configure only one global instance of Layer 2 port mirroring on an MX Series router and on an EX Series switch.



**NOTE:** You can configure more than two port mirroring instances for each FPC by configuring inline port mirroring. For information on inline port mirroring, see [“Configuring Inline Port Mirroring” on page 4252](#).

#### **Related Documentation**

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Defining a Named Instance of Layer 2 Port Mirroring](#)
- [Binding Layer 2 Port Mirroring to Ports Grouped at the FPC Level on page 4214](#)
- [Binding Layer 2 Port Mirroring to Ports Grouped at the PIC Level on page 4216](#)
- [Examples: Layer 2 Port-Mirroring at Multiple Levels of the Chassis on page 4218](#)
- [Example: Layer 2 Port Mirroring with Multiple Instances](#)

#### **Layer 2 Port Mirroring of PE Router Logical Interfaces**

For an MX Series router or an EX Series switch configured as a provider edge (PE) router or PE switch on the customer-facing edge of a service provider network, you can apply a Layer 2 port-mirroring firewall filter at the following ingress and egress points to mirror

the traffic between the MX Series router (or an EX Series switch) and customer edge (CE) devices, such as routers and Ethernet switches.

[Table 314 on page 4197](#) describes the ways in which you can apply Layer 2 port-mirroring firewall filters to an MX Series router or an EX Series switch configured as a PE router or PE switch.

**Table 314: Application of Layer 2 Port Mirroring Firewall Filters on PE Routers and PE Switches**

| Point of Application                                             | Scope of Mirroring                                                                                                                                                                                                                                                                                                                                                                                             | Notes                                                                                                                                                                                                                                                                                                                                                                                                                  | Configuration Details                                                                                                                                                                                                                                                |
|------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Ingress Customer-Facing Logical Interface                        | Packets originating within a service provider customer's network, sent first to a CE device, and sent next to an MX Series router or an EX Series switch acting as a PE router or PE switch.                                                                                                                                                                                                                   | You can also configure aggregated Ethernet interfaces between CE devices and PE routers or PE switches for VPLS routing instances. Traffic is load-balanced across all of the links in the aggregated interface.                                                                                                                                                                                                       | See <i>Applying Layer 2 Port Mirroring to a Logical Interface</i> .<br><br>For more information about VPLS routing instances, see <i>Configuring a VPLS Routing Instance</i> and <i>Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances</i> . |
| Egress Customer-Facing Logical Interface                         | Unicast packets being forwarded by the MX Series router or the EX Series switch to another PE router or PE switch.<br><br><b>NOTE:</b> If you apply a port-mirroring filter to the output for a logical interface, only unicast packets are mirrored. To mirror multicast, unknown unicast, and broadcast packets, apply a filter to the input to the flood table of a bridge domain or VPLS routing instance. | Traffic received on an aggregated Ethernet interface is forwarded over a different interface based on a lookup of the destination MAC (DMAC) address:<br><br><ul style="list-style-type: none"> <li>• Packets destined for a local site are sent out of the load-balanced child interface.</li> <li>• Packets destined for the remote site are encapsulated and forwarded over a label-switched path (LSP).</li> </ul> | See <i>Applying Layer 2 Port Mirroring to a Logical Interface</i> .                                                                                                                                                                                                  |
| Input to a Bridge Domain Forwarding Table or Flood Table         | Forwarding traffic or flood traffic sent to the bridge domain from a CE device.                                                                                                                                                                                                                                                                                                                                | Forwarding and flood traffic typically consists of broadcast packets, multicast packets, unicast packets with an unknown destination MAC address, or packets with a MAC entry in the DMAC routing table.                                                                                                                                                                                                               | See <i>Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a Bridge Domain</i> . For information about flooding behavior in VPLS, see the <i>Junos OS VPNs Configuration Guide</i> .                                                                  |
| Input to a VPLS Routing Instance Forwarding Table or Flood Table | Forwarding traffic or flood traffic sent to the VPLS routing instance from a CE device.                                                                                                                                                                                                                                                                                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                        | See <i>Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a VPLS Routing Instance</i> . For information about flooding behavior in VPLS, see the <i>Junos OS VPNs Configuration Guide</i> .                                                          |

**Related Documentation** • [Layer 2 Port Mirroring Overview on page 4192](#)

- *Layer 2 Port Mirroring Firewall Filters*
- *Defining a Layer 2 Port-Mirroring Firewall Filter*
- *Example: Layer 2 Port Mirroring at a Logical Interface*

## **Layer 2 Port Mirroring Firewall Filters**

---

This topic describes the following information:

- [Layer 2 Port Mirroring Firewall Filters Overview on page 4198](#)
- [Mirroring of Packets Received or Sent on a Logical Interface on page 4199](#)
- [Mirroring of Packets Forwarded or Flooded to a VLAN on page 4199](#)
- [Mirroring of Packets Forwarded or Flooded to a VPLS Routing Instance on page 4200](#)

### **Layer 2 Port Mirroring Firewall Filters Overview**

On an MX Series router and on an EX Series switch, you can configure a firewall filter *term* to specify that Layer 2 port mirroring is to be applied to all packets at the interface to which the firewall filter is applied.

You can apply a Layer 2 port-mirroring firewall filter to the input or output logical interfaces (including aggregated Ethernet logical interfaces), to traffic forwarded or flooded to a VLAN, or traffic forwarded or flooded to a VPLS routing instance.

MX Series routers and EX Series switches support Layer 2 port mirroring of VPLS (**family ethernet-switching** or **family vpls**) traffic and Layer 2 VPN traffic with **family ccc** in a Layer 2 environment.

Within a firewall filter **term**, you can specify the Layer 2 port-mirroring properties under the **then** statement in either of the following ways:

- Implicitly reference the Layer 2 port mirroring properties in effect on the port.
- Explicitly reference a particular named instance of Layer 2 port mirroring.



**NOTE:** When configuring a Layer 2 port-mirroring firewall filter, do not include the optional **from** statement that specifies match conditions based on the route source address. Omit this statement so that all packets are considered to match and all *actions* and *action-modifiers* specified in the **then** statement are taken.

If you want to mirror all incoming packets, then you must not use the **from** statement; /\* comment: one configure filter terms with **from** if they are interested in mirroring only a subset of packet.

---

For a general description of Layer 2 port-mirroring properties, see “[Layer 2 Port Mirroring Properties](#)” on page 4192. For a comparison of the types of Layer 2 port mirroring available on an MX Series router and on an EX Series switch, see *Application of Layer 2 Port Mirroring Types*.



**NOTE:** If you associate integrated routing and bridging (IRB) with the VLAN (or VPLS routing instance), and also configure within the VLAN (or VPLS routing instance) a forwarding table filter with the `port-mirror` or `port-mirror-instance` action, then the IRB packet is mirrored as a Layer 2 packet. You can disable this behavior by configuring the `no-irb-layer-2-copy` statement in the VLAN (or VPLS routing instance).

For a detailed description of how to configure a Layer 2 port-mirroring firewall filter, see *Defining a Layer 2 Port-Mirroring Firewall Filter*.

For detailed information about how you can use Layer 2 port-mirroring firewall filters with MX Routers and EX Series switches configured as provider edge (PE) routers or PE switches, see “[Layer 2 Port Mirroring of PE Router Logical Interfaces](#)” on page 4196. For detailed information about configuring firewall filters in general (including in a Layer 3 environment), see the *Routing Policy Configuration Guide*.

#### ***Mirroring of Packets Received or Sent on a Logical Interface***

To mirror Layer 2 traffic received or sent on a logical interface, apply a port-mirroring firewall filter to the input or output of the interface.

A port-mirroring firewall filter can also be applied to an aggregated-Ethernet logical interface. For details, see *Layer 2 Port Mirroring of PE Router Aggregated Ethernet Interfaces*.



**NOTE:** If port-mirroring firewall filters are applied at both the input and output of a logical interface, two copies of each packet are mirrored. To prevent the router or switch from forwarding duplicate packets to the same destination, you can enable the “mirror-once” option for Layer 2 port mirroring in the global instance for the Layer 2 packet address family.

#### ***Mirroring of Packets Forwarded or Flooded to a VLAN***

To mirror Layer 2 traffic forwarded to or flooded to a VLAN, apply a port-mirroring firewall filter to the input to the forwarding table or flood table. Any packet received for the VLAN forwarding or flood table and that matches the filter conditions is mirrored.

For more information about VLANs, see *Layer 2 Bridge Domains Overview*. For information about flooding behavior in a VLAN, see *Layer 2 Learning and Forwarding for Bridge Domains Overview*.



**NOTE:** When you configure port mirroring on any interface under one VLAN, the mirrored packet can move to an external analyzer located under different VLANs.

### ***Mirroring of Packets Forwarded or Flooded to a VPLS Routing Instance***

To mirror Layer 2 traffic forwarded to or flooded to a VPLS routing instance, apply a port-mirroring firewall filter to the input to the forwarding table or flood table. Any packet received for the VPLS routing instance forwarding or flood table and that matches the filter condition is mirrored.

For more information about VPLS routing instances, see *Configuring a VPLS Routing Instance* and *Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances*. For information about flooding behavior in VPLS, see the *Junos OS VPNs Configuration Guide*.

#### **Related Documentation**

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Defining a Layer 2 Port-Mirroring Firewall Filter](#)
- [Example: Layer 2 Port Mirroring at a Logical Interface](#)
- [Example: Layer 2 Port Mirroring for a Layer 2 VPN](#)
- [Example: Layer 2 Port Mirroring for a Layer 2 VPN with LAG Links](#)
- [Example: Layer 2 Port Mirroring to Multiple Destinations](#)

### **Layer 2 Port Mirroring to Multiple Destinations Using Next-Hop Groups**

On an MX Series router and on an EX Series switch, you can mirror traffic to multiple destinations by configuring next-hop groups in Layer 2 port-mirroring firewall filters applied to tunnel interfaces. The mirroring of packets to multiple destinations is also known as *multipacket port mirroring*,



**NOTE:** Junos OS Release 9.5 introduced support for Layer 2 port mirroring using next-hop groups on MX Series routers, but required installation of a Tunnel PIC. Beginning in Junos OS Release 9.6, Layer 2 port mirroring using next-hop groups on MX Series routers does not require Tunnel PICs.

On MX Series routers and on EX Series switches, you can define a firewall filter for mirroring packets to a next-hop group. The next-hop group can contain Layer 2 members, Layer 3 members, and subgroups that are either unit list (mirroring packets to each interface) or load-balanced (mirroring packets to one of several interfaces). The MX Series router and the EX Series switch supports up to 30 next-hop groups. Each next-hop group supports up to 16 next-hop addresses. Each next-hop group must specify at least two addresses.

To enable port mirroring to the members of a next-hop group, you specify the next-hop group as the filter action of a firewall filter, and then you apply the firewall filter to logical tunnel interfaces (**lt-**) or virtual tunnel interfaces (**vt-**) on the MX Series router or on the EX Series switch.



NOTE: The use of subgroups for load-balancing mirrored traffic is not supported.

- Related Documentation
- [Layer 2 Port Mirroring Overview on page 4192](#)[Layer 2 Port Mirroring Overview on page 4192](#)
  - [Defining a Layer 2 Port-Mirroring Firewall Filter](#)
  - [Defining a Next-Hop Group for Layer 2 Port Mirroring on page 4245](#)
  - [Example: Layer 2 Port Mirroring to Multiple Destinations](#)

Configuration Guidelines for Layer 2 Port Mirroring

- [Application of Layer 2 Port Mirroring Types on page 4201](#)
- [Restrictions on Layer 2 Port Mirroring on page 4203](#)

Application of Layer 2 Port Mirroring Types

You can apply different sets of Layer 2 port-mirroring properties to the VPLS packets at different ingress or egress points of an MX Series or of an EX Series route.

[Table 315 on page 4201](#) describes the three types of Layer 2 port mirroring you can configure on an MX Series router and EX Series switch: the global instance, named instances, and firewall filters.

Table 315: Application of Layer 2 Port Mirroring Types

| Type of Layer2PortMirroring Definition | Point of Application                                  | Scope of Mirroring                                                             | Description                                                                                                                                       | Configuration Details                                                |
|----------------------------------------|-------------------------------------------------------|--------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------|
| Global Instance of Layer2PortMirroring | All ports in the MX Series router (or switch) chassis | VPLS packets received on all ports in the MX Series router (or switch) chassis | If configured, the global port-mirroring properties implicitly apply to all VPLS packets received on all ports in the router (or switch) chassis. | See <i>Configuring the Global Instance of Layer 2 Port Mirroring</i> |

Table 315: Application of Layer 2 Port Mirroring Types (*continued*)

| Type of Layer2PortMirroring Definition | Point of Application                                                                                                                                                 | Scope of Mirroring                                                                                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Configuration Details                                                                                                                                                                                                                                                                                                                         |
|----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Named Instance of Layer2PortMirroring  | Ports grouped at the FPC level<br><br>See <a href="#">“Binding Layer 2 Port Mirroring to Ports Grouped at the FPC Level”</a> on page 4214.                           | VPLS packets received on ports associated with a specific DPC or FPC and its Packet Forwarding Engines. | Overrides any port-mirroring properties configured by the global port-mirroring instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | See <i>Defining a Named Instance of Layer 2 Port Mirroring</i> .<br><br><b>NOTE:</b> The number of port-mirroring destinations supported for an MX Series router and for an EX Series switch is limited to the number of Packet Forwarding Engines contained on the DPCs or FPCs installed in the router or switch chassis.                   |
|                                        | Ports grouped at the PIC level<br><br>See <a href="#">“Binding Layer 2 Port Mirroring to Ports Grouped at the PIC Level”</a> on page 4216.                           | VPLS packets received on ports associated with a specific Packet Forwarding Engine.                     | Overrides any port-mirroring properties configured at the FPC level or in the global port-mirroring instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                                                                                                                                                                                                                                                                                                                                               |
| Layer2Port-Mirroring Firewall Filter   | Logical interface (including an aggregated Ethernet interface)<br><br>See <i>Applying Layer 2 Port Mirroring to a Logical Interface</i> .                            | VPLS packets received or sent on a logical interface.                                                   | In the firewall filter configuration, include <i>action</i> and <i>action-modifier</i> terms to apply to the packets selected for mirroring:<br><br>1. The <b>accept</b> action is recommended.<br><br>2. Specify port mirroring by including one of the following modifiers:<br><br>• The <b>port-mirror</b> modifier implicitly references the port-mirroring properties currently bound to the underlying physical interfaces.<br><br>• The <b>port-mirror-instance pm-instance-name</b> modifier explicitly references a named instance of port mirroring.<br><br>3. (Optional) For tunnel interface input packets only, to mirror the packets to additional destinations, include the <b>next-hop-group next-hop-group-name</b> modifier. This modifier references a next-hop-group that specifies the next-hop addresses (for sending additional copies of packets to an analyzer). | See <i>Defining a Layer 2 Port-Mirroring Firewall Filter</i> .<br><br><b>NOTE:</b> Layer 2 port-mirroring firewall filters are not supported for logical systems.<br><br>For mirroring tunnel interface input packets to multiple destinations, also see <a href="#">“Defining a Next-Hop Group for Layer 2 Port Mirroring”</a> on page 4245. |
|                                        | VLAN forwarding table or flood table<br><br>See <i>Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a Bridge Domain</i> .                          | Layer 2 traffic forwarded or flooded to a VLAN                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                               |
|                                        | VPLS routing instance forwarding table or flood table<br><br>See <i>Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a VPLS Routing Instance</i> . | Layer 2 traffic forwarded or flooded to a VPLS routing instance                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                               |



## Related Documentation

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Restrictions on Layer 2 Port Mirroring](#)
- [Precedence of Multiple Levels of Layer 2 Port Mirroring on a Physical Interface on page 4204](#)
- [Layer 2 Port Mirroring of PE Router Logical Interfaces on page 4196](#)
- [Layer 2 Port Mirroring of PE Router Aggregated Ethernet Interfaces](#)

## Restrictions on Layer 2 Port Mirroring

The following restrictions apply to Layer 2 port mirroring:

- Only Layer 2 transit data (packets that contain chunks of data transiting the routing platform or switch as they are forwarded from a source to a destination) can be mirrored. Layer 2 local data (packets that contain chunks of data that are destined for or sent by the Routing Engine, such as Layer 2 control packets) are not mirrored.
- If you apply a port-mirroring filter to the output of a logical interface, only unicast packets are mirrored. To mirror broadcast packets, multicast packets, unicast packets with an unknown destination media access control (MAC) address, or packets with MAC entry in the destination MAC (DMAC) routing table, apply a filter to the input to the flood table of a VLAN or virtual private LAN service (VPLS) routing instance.
- The mirror destination device should be on a dedicated VLAN and should not participate in any bridging activity: The mirror destination device should not have a bridge to the ultimate traffic destination, and the mirror destination device should not send the mirrored packets back to the source address.
- For either the global port-mirroring instance or a named port-mirroring instance, you can configure only one mirror output interface per port-mirroring instance and packet address family. If you include more than one **interface** statement under the **family (ethernet-switching | ccc | vpls) output** statement, the previous **interface** statement is overridden.
- Layer 2 port-mirroring firewall filtering is not supported for logical systems.

In a Layer 2 port-mirroring firewall filter definition, the filter **action-modifier (port-mirror or port-mirror-instance *pm-instance-name*)** relies on port-mirroring properties defined in the global instance or named instances of Layer 2 port mirroring, which are configured under the **[edit forwarding-options port-mirroring]** hierarchy. Therefore, the filter **term** cannot support Layer 2 port mirroring for logical systems.

- For a Layer 2 port mirroring firewall filter in which you implicitly reference Layer 2 port mirroring properties by including the **port-mirror** statement, if multiple named instances of Layer 2 port mirroring are bound to the underlying physical interface, then only the first binding in the stanza (or the only binding) is used at the logical interface. This is done mainly for backward compatibility.
- Layer 2 port-mirroring firewall filters do not support the use of next-hop subgroups for load-balancing mirrored traffic.

**Related Documentation**

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Application of Layer 2 Port Mirroring Types](#)
- [Precedence of Multiple Levels of Layer 2 Port Mirroring on a Physical Interface on page 4204](#)
- [Layer 2 Port Mirroring of PE Router Logical Interfaces on page 4196](#)
- [Layer 2 Port Mirroring of PE Router Aggregated Ethernet Interfaces](#)

## Behavior of Layer 2 Port Mirroring at Physical Interfaces

- [Precedence of Multiple Levels of Layer 2 Port Mirroring on a Physical Interface on page 4204](#)

### Precedence of Multiple Levels of Layer 2 Port Mirroring on a Physical Interface

You can bind different sets of Layer 2 port mirroring properties (the global instance and one or more named instances) at various levels of an MX Series router or of an EX Series switch chassis (at the chassis level, at the FPC level, or at the PIC level). Therefore, it is possible for a single group of physical interfaces to be bound to multiple Layer 2 port mirroring definitions.

If a group of ports (or, in the case of a PIC-level binding in an MX960 router, a single port) is bound to multiple Layer 2 port mirroring definitions, the router (or switch) applies the Layer 2 port-mirroring properties to those ports as follows:

1. **Chassis-level port-mirroring properties implicitly apply to all ports in the chassis.** If an MX Series router or an EX Series switch is configured with the global port-mirroring instance, those port mirroring properties apply to all ports. See *Configuring the Global Instance of Layer 2 Port Mirroring*.
2. **FPC-level port-mirroring properties override chassis-level properties.** If a DPC or FPC is bound to a named instance of port mirroring, those port mirroring properties apply to all ports associated with that DPC or FPC, overriding any port mirroring properties bound at the chassis level. See [“Binding Layer 2 Port Mirroring to Ports Grouped at the FPC Level” on page 4214](#).
3. **PIC-level port-mirroring properties override FPC-level properties.** If a Packet Forwarding Engine or PIC is bound to a named instance of port mirroring, those port mirroring properties apply to all ports associated with the Packet Forwarding Engine or PIC, overriding any port mirroring properties bound to those ports at the FPC level. See [“Binding Layer 2 Port Mirroring to Ports Grouped at the PIC Level” on page 4216](#).

**Related Documentation**

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Restrictions on Layer 2 Port Mirroring](#)
- [Application of Layer 2 Port Mirroring Types](#)
- [Layer 2 Port Mirroring of PE Router Logical Interfaces on page 4196](#)
- [Layer 2 Port Mirroring of PE Router Aggregated Ethernet Interfaces](#)

## Behavior of Layer 2 Port Mirroring on PE Routers and PE Switches

- [Layer 2 Port Mirroring of PE Router or PE Switch Logical Interfaces on page 4205](#)
- [Layer 2 Port Mirroring of PE Router or PE Switch Aggregated Ethernet Interfaces on page 4206](#)

### Layer 2 Port Mirroring of PE Router or PE Switch Logical Interfaces

For a router or switch configured as a provider edge (PE) device on the customer-facing edge of a service provider network, you can apply a Layer 2 port-mirroring firewall filter at the following ingress and egress points to mirror the traffic between the router or switch and customer edge (CE) devices, which are typically also routers and Ethernet switches.

[Table 314 on page 4197](#) describes the ways in which you can apply Layer 2 port-mirroring firewall filters to a router or switch configured as a PE device.

**Table 316: Application of Layer 2 Port Mirroring Firewall Filters on PE Devices**

| Point of Application                      | Scope of Mirroring                                                                                                                                                                                                                                                                                                                                               | Notes                                                                                                                                                                                                                                                                                                                                                                 | Configuration Details                                                                                                                                                                                                                                                     |
|-------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Ingress Customer-Facing Logical Interface | Packets originating within a service provider customer's network, sent first to a CE device, and sent next to the PE device.                                                                                                                                                                                                                                     | <p>You can also configure aggregated Ethernet interfaces between CE devices and PE devices for VPLS routing instances. Traffic is load-balanced across all of the links in the aggregated interface.</p> <p>Traffic received on an aggregated Ethernet interface is forwarded over a different interface based on a lookup of the destination MAC (DMAC) address:</p> | <p>See <i>Applying Layer 2 Port Mirroring to a Logical Interface</i>.</p> <p>For more information about VPLS routing instances, see <i>Configuring a VPLS Routing Instance</i> and <i>Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances</i>.</p> |
| Egress Customer-Facing Logical Interface  | <p>Unicast packets being forwarded by the PE device to another PE device.</p> <p><b>NOTE:</b> If you apply a port-mirroring filter to the output for a logical interface, only unicast packets are mirrored. To mirror multicast, unknown unicast, and broadcast packets, apply a filter to the input to the flood table of a VLAN or VPLS routing instance.</p> | <ul style="list-style-type: none"> <li>• Packets destined for a local site are sent out of the load-balanced child interface.</li> <li>• Packets destined for the remote site are encapsulated and forwarded over a label-switched path (LSP).</li> </ul>                                                                                                             | See <i>Applying Layer 2 Port Mirroring to a Logical Interface</i> .                                                                                                                                                                                                       |

Table 316: Application of Layer 2 Port Mirroring Firewall Filters on PE Devices (*continued*)

| Point of Application                                             | Scope of Mirroring                                                                      | Notes                                                                                                                                                                                                    | Configuration Details                                                                                                                                                                                       |
|------------------------------------------------------------------|-----------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Input to a VLAN Forwarding Table or Flood Table                  | Forwarding traffic or flood traffic sent to the VLAN from a CE device.                  | Forwarding and flood traffic typically consists of broadcast packets, multicast packets, unicast packets with an unknown destination MAC address, or packets with a MAC entry in the DMAC routing table. | See <i>Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a Bridge Domain</i> . For information about flooding behavior in VPLS, see the <i>Junos OS VPNs Configuration Guide</i> .         |
| Input to a VPLS Routing Instance Forwarding Table or Flood Table | Forwarding traffic or flood traffic sent to the VPLS routing instance from a CE device. |                                                                                                                                                                                                          | See <i>Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a VPLS Routing Instance</i> . For information about flooding behavior in VPLS, see the <i>Junos OS VPNs Configuration Guide</i> . |

**Related Documentation**

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Layer 2 Port Mirroring Firewall Filters](#)
- [Defining a Layer 2 Port-Mirroring Firewall Filter](#)
- [Example: Layer 2 Port Mirroring at a Logical Interface](#)

**Layer 2 Port Mirroring of PE Router or PE Switch Aggregated Ethernet Interfaces**

An aggregated Ethernet interface is a virtual aggregated link that consists of a set of physical interfaces of the same speed and operating in full-duplex link connection mode. You can configure aggregated Ethernet interfaces between CE devices and PE devices for VPLS routing instances. Traffic is load-balanced across all of the links in the aggregated interface. If one or more links in the aggregated interface fails, the traffic is switched to the remaining links.

You can apply a Layer 2 port-mirroring firewall filter to an aggregated Ethernet interface to configure port mirroring at the parent interface. However, if any child interfaces are bound to different Layer 2 port-mirroring instances, packets received at the child interfaces will be mirrored to the destinations specified by their respective port-mirroring instances. Thus, multiple child interfaces can mirror packets to multiple destinations.

For example, suppose the parent aggregated Ethernet interface instance **ae0** has two child interfaces:

- **xe-2/0/0**
- **xe-3/1/2**

Suppose that these child interfaces on **ae0** are bound to two different Layer 2 port-mirroring instances:

- **pm\_instance\_A**—A named instance of Layer 2 port mirroring, bound to child interface **xe-2/0/0**.

- **pm\_instance\_B**—A named instance of Layer 2 port mirroring, bound to child interface **xe-3/1/2**.

Now suppose you apply a Layer 2 port-mirroring firewall filter to the Layer 2 traffic sent on **ae0.0** (logical unit **0** on the aggregated Ethernet interface instance **0**). This enables port mirroring on **ae0.0**, which has the following effect on the processing of traffic received on the child interfaces for which Layer 2 port-mirroring properties are specified:

- The packets received on **xe-2/0/0.0** are mirrored to the output interfaces configured in port-mirroring instance **pm\_instance\_A**.
- The packets received on **xe-3/1/2.0** are mirrored to the output interfaces configured in port-mirroring instance **pm\_instance\_B**.

Because **pm\_instance\_A** and **pm\_instance\_B** can specify different packet-selection properties or mirror destination properties, the packets received on **xe-2/0/0.0** and **xe-3/1/2.0** can mirror different packets to different destinations.

#### Related Documentation

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Layer 2 Port Mirroring Firewall Filters](#)
- [Defining a Layer 2 Port-Mirroring Firewall Filter](#)
- [Example: Layer 2 Port Mirroring for a Layer 2 VPN](#)

## Configuration

- [Configuration Tasks for Layer 2 Port Mirroring at Physical Interfaces on page 4207](#)
- [Examples for Layer 2 Port Mirroring at Physical Interfaces on page 4218](#)
- [Configuration Tasks for Layer 2 Port Mirroring at Logical Interfaces on page 4224](#)
- [Examples for Layer 2 Port Mirroring at Logical Interfaces on page 4234](#)
- [Configuration Tasks for Layer 2 Port Mirroring at Multiple Destinations on page 4242](#)
- [Examples of Layer 2 Port Mirroring at Multiple Destinations on page 4249](#)
- [Configuring Inline Port Mirroring on page 4252](#)
- [Configuration Statements on page 4253](#)

### Configuration Tasks for Layer 2 Port Mirroring at Physical Interfaces

- [Configuring the Global Instance of Layer 2 Port Mirroring on page 4208](#)
- [Defining a Named Instance of Layer 2 Port Mirroring on page 4210](#)
- [Binding Layer 2 Port Mirroring to Ports Grouped at the FPC Level on page 4214](#)
- [Binding Layer 2 Port Mirroring to Ports Grouped at the PIC Level on page 4216](#)
- [Disabling Layer 2 Port Mirroring Instances on page 4217](#)

### Configuring the Global Instance of Layer 2 Port Mirroring

---

On an MX Series router and on an EX Series switch, you can configure a set of Layer 2 port-mirroring properties that implicitly apply to packets received on all ports in the router (or switch) chassis.

To configure the global instance of Layer 2 port mirroring on an MX Series router and on an EX Series switch:

1. Enable configuration of the Layer 2 port mirroring:

```
[edit]
user@host# edit forwarding-options port-mirroring
```

2. Enable configuration of the packet-selection properties:

```
[edit forwarding-options port-mirroring]
user@host# edit input
```

3. Specify global-level packet-selection properties.

- a. Specify the number of packets to select:

```
[edit forwarding-options port-mirroring input]
user@host# set rate number
```

The valid range is 1 through 65535.

- b. Specify the number of packets to mirror from each selection:

```
[edit forwarding-options port-mirroring input]
user@host# set run-length number
```

The valid range is 0 through 20. The default value is 0.

- c. Specify the length to which mirrored packets are to be truncated:

```
[edit forwarding-options port-mirroring input]
user@host# set maximum-packet-length number
```

The valid range is 0 through 9216. The default value is 0, which means the mirrored packets are not truncated.

4. Specify the global-level Layer 2 address-type family from which traffic is to be selected for mirroring:

```
[edit forwarding-options port-mirroring input]
user@host# up
[edit forwarding-options port-mirroring]
user@host# edit family family
```

The value of the *family* option can be **ethernet-switching**, **ccc**, or **vpls**.



**NOTE:** Under the [edit forwarding-options port-mirroring] hierarchy level, the protocol family statement family ethernet-switching is an alias for family vpls. The command-line interface (CLI) displays Layer 2 port-mirroring configurations as family vpls, even for Layer 2 port-mirroring configured as family ethernet-switching. Use family ethernet-switching when the physical interface is configured with encapsulation ethernet-bridge.

5. Enable configuration of global-level mirror destination properties for this address family:

```
[edit forwarding-options port-mirroring family family]
user@host# edit output
```

6. Specify global-level mirror destination properties for this address family.

- a. Specify the physical interface on which to send the mirrored packets:

```
[edit forwarding-options port-mirroring family family output]
user@host# set interface interface-name
```

You can also specify an integrated routing and bridging (IRB) interface as the output interface.

- b. (Optional) Allow configuration of filters on the destination interface for the named port-mirroring instance:

```
[edit forwarding-options port-mirroring family family output]
user@host# set no-filter-check
```

7. (Optional) Specify that any packets selected for mirroring are to be mirrored only once to any mirroring destination:

```
[edit forwarding-options port-mirroring family family output]
user@host# up 2
[edit forwarding-options port-mirroring]
user@host# set mirror-once
```



**TIP:** Enable the mirror-once option when an MX Series router or an EX Series switch is configured to perform Layer 2 port mirroring at both ingress and egress interfaces, which could result in sending duplicate packets to the same destination (which would complicate the analysis of the mirrored traffic).

8. Verify the minimum configuration of the global instance of Layer 2 port mirroring:

```
[edit forwarding-options ...]
user@host# top
[edit]
user@host# show forwarding-options

forwarding-options {
```

```
port-mirroring {
 input { # Global packet-selection properties.
 maximum-packet-length number; # Default is 0.
 rate number;
 run-length number;
 }
 family (ccc | vpls) { # Address- type 'ethernet-switching' displays as 'vpls'.
 output { # Global mirror destination properties.
 interface interface-name;
 no-filter-check; # Optional. Allow filters on interface.
 }
 }
 mirror-once; # Optional. Mirror destinations do not receive duplicate packets.
}
```

**Related  
Documentation**

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Layer 2 Port Mirroring Global Instance on page 4194](#)
- [Examples: Layer 2 Port-Mirroring at Multiple Levels of the Chassis on page 4218](#)
- [Example: Layer 2 Port Mirroring with Multiple Instances](#)

---

**Defining a Named Instance of Layer 2 Port Mirroring**

---

On an MX Series router and on an EX Series switch, you can define a set of Layer 2 port-mirroring properties that you can bind to a particular Packet Forwarding Engine (at the PIC level of the router or switch chassis) or to a group of Packet Forwarding Engines (at the DPC or FPC level of the chassis).

To define a named instance of Layer 2 port mirroring on an MX Series router or on an EX Series switch:

1. Enable configuration of a named instance of Layer 2 port mirroring :

```
[edit]
user@host# edit forwarding-options port-mirroring instance pm-instance-name
```

2. Enable configuration of the packet-sampling properties:

```
[edit forwarding-options port-mirroring instance pm-instance-name]
user@host# edit input
```



3. Specify packet-selection properties:

a. Specify the number of packets to select:

```
[edit forwarding-options port-mirroring instance pm-instance-name input]
user@host# set rate number
```

The valid range is 1 through 65535.

b. Specify the number of packets to mirror from each selection:

```
[edit forwarding-options port-mirroring instance pm-named-instance input]
user@host# set run-length number
```

The valid range is 0 through 20. The default value is 0.



**NOTE:** The `run-length` statement is not supported on MX80 routers.

c. Specify the length to which mirrored packets are to be truncated:

```
[edit forwarding-options port-mirroring instance pm-instance-name input]
user@host# set maximum-packet-length number
```

The valid range is 0 through 9216. The default value is 0, which means the mirrored packets are not truncated.



**NOTE:** The `maximum-packet-length` statement is not supported on MX80 routers.

4. Enable configuration of the mirror destination properties for Layer 2 packets that are part of bridging domain, Layer 2 switching cross-connects, or virtual private LAN service (VPLS):

a. Specify the Layer 2 address family type of traffic to be mirrored:

```
[edit forwarding-options port-mirroring instance pm-instance-name input]
user@host# up
[edit forwarding-options port-mirroring instance pm-instance-name]
user@host# edit family family
```

The value of the `family` option can be `ethernet-switching`, `ccc`, or `vpls`.



**NOTE:** Under the `[edit forwarding-options port-mirroring]` hierarchy level, the protocol family statement `family ethernet-switching` is an alias for `family vpls`. The command-line interface (CLI) displays Layer 2 port-mirroring configurations as `family vpls`, even for Layer 2 port-mirroring configured as `family ethernet-switching`. Use `family ethernet-switching` when the physical interface is configured with `encapsulation ethernet-bridge`.

Enable configuration of the mirror destination properties:

- b. [edit forwarding-options port-mirroring instance *pm-instance-name* family *family*]  
user@host# edit output

5. Specify mirror destination properties.

- a. Specify the physical interface on which to send the mirrored packets:

```
[edit forwarding-options port-mirroring instance pm-instance-name family family
output]
user@host# set interface interface-name
```

- b. (Optional) Allow configuration of filters on the destination interface for the global port-mirroring instance:

```
[edit forwarding-options port-mirroring instance pm-instance-name family family
output]
user@host# set no-filter-check
```



**NOTE:** You cannot configure port mirroring instances on MX80 routers. You can only configure port mirroring at the global level on MX80 routers.

- 6. (Optional) Specify that any packets selected for mirroring are to be mirrored only once to any mirroring destination:

```
[edit forwarding-options port-mirroring instance pm-instance-name family family
output]
user@host# up 3
[edit forwarding-options port-mirroring]
user@host# set mirror-once
```



**TIP:** Enable the global mirror-once option when an MX Series router or an EX Series switch is configured to perform Layer 2 port mirroring at both ingress and egress interfaces, which could result in sending duplicate packets to the same destination (which in turn would complicate the analysis of the mirrored traffic).

- 7. To configure a mirroring destination for a different packet family type, repeat steps 4 through 6.

- 8. Verify the minimum configuration of the named instances of Layer 2 port mirroring:

```
[edit forwarding-options ...]
user@host# top
[edit]
user@host# show forwarding-options

forwarding-options {
 port-mirroring {
 ... optional-global-port-mirroring-configuration ...
 instance {
 pm-instance-name (# A named instance of port mirroring
 input { # Packet-selection properties
```

```

 maximum-packet-length number; # Default is 0.
 rate number;
 run-length number;
 }
 family (ccc | vpls) { # Address- type 'ethernet-switching' displays as 'vpls'.
 output { # Mirror destination properties
 interface interface-name;
 no-filter-check; # Optional. Allow filters on interface.
 }
 }
}
mirror-once; # Optional. Mirror destinations do not receive duplicate packets.
}

```

#### Related Documentation

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Layer 2 Port Mirroring Named Instances](#)
- [Binding Layer 2 Port Mirroring to Ports Grouped at the FPC Level on page 4214](#)
- [Binding Layer 2 Port Mirroring to Ports Grouped at the PIC Level on page 4216](#)
- [Examples: Layer 2 Port-Mirroring at Multiple Levels of the Chassis on page 4218](#)
- [Example: Layer 2 Port Mirroring with Multiple Instances](#)

### Binding Layer 2 Port Mirroring to Ports Grouped at the FPC Level

---

On an MX Series router and on an EX Series switch, you can bind a named instance of Layer 2 port mirroring to a specific DPC or to a specific FPC in the router (or switch) chassis. This is known as binding a named instance of Layer 2 port mirroring *at the FPC level* of the router (or switch) chassis. The port mirroring properties specified in the named instance are applied to all physical ports associated with all Packet Forwarding Engines on the specified DPC or FPC.



**NOTE:** You can also bind a named instance of Layer 2 port mirroring to a specific Packet Forwarding Engine on a DPC or FPC in the router (or switch) chassis.

For any packet-type family supported by Layer 2 port mirroring

- Port mirroring properties bound to a specific DPC or FPC override any port-mirroring properties configured at the global level.
- Port mirroring properties bound to a specific Packet Forwarding Engine override any port-mirroring properties configured at the DPC or FPC level.

You can apply up to two named instances of Layer 2 port mirroring to the same group of ports within the router (or switch) chassis. By applying two different port-mirroring instances to the same DPC or FPC, you can bind two distinct Layer 2 port mirroring specifications to a single group of ports.

Before you begin, complete the following tasks:

- Define a named instance of Layer 2 port mirroring. See *Defining a Named Instance of Layer 2 Port Mirroring*.
- Display information about the number and types of DPCs or FPCs in the MX Series router and in the EX Series switch, the number of Packet Forwarding Engines on each, and the number and types of ports per Packet Forwarding Engine. See *Displaying Information About DPCs or FPCs in an MX Series Router*.

To bind a named instance of Layer 2 port-mirroring to a DPC or FPC and its Packet Forwarding Engines:

1. Enable configuration of the router (or switch) chassis properties:

```
[edit]
user@host# edit chassis
```

2. Enable configuration of a DPC (and its corresponding Packet Forwarding Engines) or an FPC (and its installed PICs):

```
[edit chassis]
user@host# edit fpc slot-number
```

3. Bind a named instance of Layer 2 port mirroring (*pm-instance-name*) to the DPC or FPC:

```
[edit chassis fpc slot-number]
user@host# set port-mirror-instance pm-instance-name
```

4. (Optional) To bind a second named instance of Layer 2 port mirroring to the same DPC or FPC, repeat step 3 and specify a different named instance of Layer 2 port mirroring.
5. Verify the minimum configuration of the binding:

```
[edit chassis fpc slot-number port-mirror-instance pm-instance-name]
user@host# top
[edit]
user@host# show chassis

chassis {
 fpc slot-number { # Bind two port mirroring named instances at the FPC level.
 port-mirror-instance pm-instance-name-1;
 port-mirror-instance pm-instance-name-2;
 }
}
```

#### Related Documentation

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Layer 2 Port Mirroring Named Instances](#)
- [Defining a Named Instance of Layer 2 Port Mirroring](#)
- [Binding Layer 2 Port Mirroring to Ports Grouped at the PIC Level on page 4216](#)
- [Examples: Layer 2 Port-Mirroring at Multiple Levels of the Chassis on page 4218](#)
- [Example: Layer 2 Port Mirroring with Multiple Instances](#)

### Binding Layer 2 Port Mirroring to Ports Grouped at the PIC Level

---

On an MX Series router and on an EX Series switch, you can bind a named instance of Layer 2 port mirroring to the ports associated with a specific Packet Forwarding Engine (on a DPC) or to the ports associated with a specific PIC (installed in an FPC). This is known as binding a named instance of Layer 2 port mirroring *at the PIC level* of the router (or switch) chassis. The port-mirroring properties specified in the named instance are applied to all physical ports associated with the specified Packet Forwarding Engine.



**NOTE:** You can also bind a named instance of Layer 2 port mirroring to a specific DPC or FPC in the router (or switch) chassis.

For any packet-type family supported by Layer 2 port mirroring:

- Port mirroring properties bound to a specific Packet Forwarding Engine override any port-mirroring properties configured at the DPC or FPC level.
- Port mirroring properties bound to a specific DPC or FPC override any port-mirroring properties configured at the global level.

You can apply up to two named instances of Layer 2 port mirroring to the same group of ports within the router (or switch) chassis. By applying two different port-mirroring instances to the same Packet Forwarding Engine or PIC, you can bind two distinct Layer 2 port mirroring specifications to a single group of ports.

For MX960 routers, there is a one-to-one mapping of Packet Forwarding Engines to Ethernet ports. Therefore, on MX960 routers only, you can bind a named instance of Layer 2 port mirroring to a *specific port* by binding the instance to the Packet Forwarding Engine associated with the port.

Before you begin, complete the following tasks:

- Define a named instance of Layer 2 port mirroring. See *Defining a Named Instance of Layer 2 Port Mirroring*.
- Display information about the number and types of DPCs in the MX Series router or in the EX Series switch, the number of Packet Forwarding Engines on each DPC, and the number and types of ports per Packet Forwarding Engine. See *Displaying Information About DPCs or FPCs in an MX Series Router*.

To bind a named instance of Layer 2 port-mirroring to a Packet Forwarding Engine:

1. Enable configuration of the router (or switch) chassis properties:

```
[edit]
user@host# edit chassis
```

2. Enable configuration of a Packet Forwarding Engine or PIC:

```
[edit chassis]
user@host# edit fpc slot-number
user@host# edit pic slot-number
```

3. Bind a named instance of Layer 2 port mirroring (*pm-instance-name*) to the Packet Forwarding Engine or PIC:

```
[edit chassis fpc slot-number pic slot-number]
user@host# set port-mirror-instance pm-instance-name
```

4. (Optional) To bind a second named instance of Layer 2 port mirroring to the same Packet Forwarding Engine or PIC, repeat step 3 and specify a different named instance of Layer 2 port mirroring.
5. Verify the minimum configuration of the binding:

```
[edit forwarding-options ...]
user@host# top
[edit]
user@host# show chassis
chassis {
 fpc slot-number {
 ... optional-binding-of-a-port-mirroring-instance-at-the-dpc-level ...
 pic slot-number { # Bind two port-mirroring named instances at the PIC level.
 port-mirror-instance pm-instance-name-1;
 port-mirror-instance pm-instance-name-2;
 }
 }
}
```

#### Related Documentation

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Layer 2 Port Mirroring Named Instances](#)
- [Defining a Named Instance of Layer 2 Port Mirroring](#)
- [Binding Layer 2 Port Mirroring to Ports Grouped at the FPC Level on page 4214](#)
- [Examples: Layer 2 Port-Mirroring at Multiple Levels of the Chassis on page 4218](#)
- [Example: Layer 2 Port Mirroring with Multiple Instances](#)

### Disabling Layer 2 Port Mirroring Instances

You can disable the global instance of Layer 2 port mirroring, a particular named instance, or all instances of port mirroring:

- To disable the global instance of Layer 2 port mirroring, include the **disable** statement at the **[edit forwarding-options port-mirroring]** hierarchy level:

```
[edit]
forwarding-options {
 port-mirroring {
 disable; Disables the global instance of Layer 2 port mirroring.
 ...global-instance-of-layer-2-port-mirroring-configuration...
 }
}
```

- To disable the definition of a particular named instance of Layer 2 port mirroring, include the **disable** statement at the **[edit forwarding-options port-mirroring instance instance-name]** hierarchy level:

```
[edit]
forwarding-options {
 port-mirroring {
 ...optional-configuration-of-the-global-instance-of-layer-2-port-mirroring...
 instance {
 port-mirroring-instance-name {
 disable; Disables this named instance of Layer 2 port mirroring.
 ...definition-of-a-named-instance-of-layer-2-port-mirroring...
 }
 }
 }
}
```

- To disable the global instance and all named instances of Layer 2 port mirroring, include the **disable-all-instances** statement at the **[edit forwarding-options port-mirroring]** hierarchy level:

```
[edit]
forwarding-options {
 port-mirroring {
 disable-all-instances; Disables all instances of Layer 2 port mirroring.
 ...optional-configuration-of-the-global-instance-of-layer-2-port-mirroring...
 instance {
 port-mirroring-instance-name {
 ...definition-of-a-named-instance-of-layer-2-port-mirroring...
 }
 }
 }
}
```

#### Related Documentation

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Layer 2 Port Mirroring Global Instance on page 4194](#)
- [Layer 2 Port Mirroring Named Instances](#)
- [Displaying Layer 2 Port-Mirroring Instance Settings and Status on page 4263](#)

## Examples for Layer 2 Port Mirroring at Physical Interfaces

- [Examples: Layer 2 Port-Mirroring at Multiple Levels of the Chassis on page 4218](#)
- [Example: Layer 2 Port Mirroring with Multiple Instances on page 4220](#)

### Examples: Layer 2 Port-Mirroring at Multiple Levels of the Chassis

On an MX Series router or on an EX Series switch, you can apply named instances of Layer 2 port mirroring at the FPC or DPC level of the chassis or at the PIC level of the



chassis. However, you can configure (and implicitly apply) only one global instance of Layer 2 port mirroring to the entire chassis.

- [Layer 2 Port Mirroring at the FPC Level on page 4219](#)
- [Layer 2 Port Mirroring at the PIC Level on page 4219](#)
- [Layer 2 Port Mirroring at the FPC and PIC Levels on page 4219](#)

### ***Layer 2 Port Mirroring at the FPC Level***

In this example configuration of an MX Series router or of an EX Series switch chassis, a named instance of Layer 2 port mirroring (**pm1**) is bound to physical ports grouped at the FPC level:

```
[edit]
chassis {
 fpc 2 {
 port-mirror-instance pm1;
 }
}
```

This is not a complete configuration. The physical interfaces associated with the FPC or DPC in slot 2 must be configured at the **[edit interfaces]** hierarchy level. The Layer 2 port mirroring named instance **pm1** must be configured at the **[edit forwarding-options port-mirroring instance]** hierarchy level.

### ***Layer 2 Port Mirroring at the PIC Level***

In this example configuration of an MX Series router or of an EX Series switch chassis, a named instance of Layer 2 port mirroring (**pm2**) is bound to the physical ports grouped at the PIC level:

```
[edit]
chassis {
 fpc 2 {
 pic 0 {
 port-mirror-instance pm2;
 }
 }
}
```

This is not a complete configuration. The physical interfaces associated with the FPC or DPC in slot 2 must be configured at the **[edit interfaces]** hierarchy level. The Layer 2 port mirroring named instance **pm2** must be configured at the **[edit forwarding-options port-mirroring instance]** hierarchy level.

### ***Layer 2 Port Mirroring at the FPC and PIC Levels***

In this example configuration of an MX Series router chassis or an EX Series switch, one named instance of Layer 2 port mirroring (**pm1**) is applied at the FPC level of the router (or switch) chassis. A second named instance (**pm2**) is applied at the PIC level:

```
[edit]
chassis {
 fpc 2 {
 port-mirror-instance pm1;
 }
 pic 0 {
 port-mirror-instance pm2;
 }
}
```

```

 }
 }
}

```

This is not a complete configuration. Physical interfaces associated with the FPC or DPC in slot 2, including physical interfaces associated with **pic 0**, must be configured at the **[edit interfaces]** hierarchy level. The Layer 2 port mirroring named instances **pm1** and **pm2** must be configured at the **[edit forwarding-options port-mirroring instance]** hierarchy level.

#### Related Documentation

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Layer 2 Port Mirroring Global Instance on page 4194](#)
- [Layer 2 Port Mirroring Named Instances](#)
- [Configuring the Global Instance of Layer 2 Port Mirroring](#)
- [Defining a Named Instance of Layer 2 Port Mirroring](#)

#### Example: Layer 2 Port Mirroring with Multiple Instances

Because you can configure more than one port-mirroring instance, care is required when specifying which instance is meant. This topic contains the following information:

- [Example: Configuring Multiple Instances of Layer 2 Port Mirroring on page 4220](#)
- [Explicit Reference of a Port Mirroring Instance on page 4223](#)
- [Implicit Reference of Port Mirroring on the Underlying Physical Interface on page 4223](#)

#### Example: Configuring Multiple Instances of Layer 2 Port Mirroring

This configuration example illustrates the configuration of Layer 2 port mirroring at the physical interfaces associated with FPC 2, PIC 0 and at two logical interfaces on one of those ports.

At the physical interface levels of the router (or switch) chassis, two named instances of port mirroring are configured and then bound to the group of physical ports associated with FPC 2, PIC 0.

At two of the logical interfaces on physical interface **ge-2/0/1**, two Layer 2 port-mirroring firewall filters are applied to the input traffic. One filter *explicitly* references the port mirroring properties specified in one of the named instances of port mirroring. The other filter *implicitly* references the port mirroring properties in effect on the underlying physical interface **ge-2/0/1**.

The resulting configuration is an example of the relationships that can exist between multiple instances of Layer 2 port mirroring applied to an MX Series router or an EX Series switch.

1. Configure two named instances of Layer 2 port mirroring (**pm\_instance\_1** and **pm\_instance\_2**), and include mirror destination properties for VLAN traffic (**family ethernet-switching**):

```

[edit]
forwarding-options {

```

```

port-mirroring {
 instance {
 pm_instance_1 {
 input {
 ... packet-selection-properties-configuration ...
 }
 family ethernet-switching {
 ... mirror-destination-properties-configuration ...
 }
 }
 pm_instance_2 {
 input {
 ... packet-selection-properties-configuration ...
 }
 family ethernet-switching {
 ... mirror-destination-properties-configuration ...
 }
 }
 }
}

```



**NOTE:** In this example, no global port-mirroring properties are configured on the router (or switch).

2. Apply the Layer 2 port mirroring instances to the same group of ports in the router (or switch) chassis. In this example, the named instances of Layer 2 port mirroring are applied to the same group of physical interfaces specified at the PIC level of the chassis:

```

[edit]
chassis {
 fpc 2 {
 pic 0 {
 port-mirror-instance pm_instance_1;
 port-mirror-instance pm_instance_2;
 }
 }
}

```

Note that, in this example, two named instances of Layer 2 port mirroring are bound to the PIC level of the chassis at the same group of ports.

3. Configure two Layer 2 port-mirroring firewall filters, both for VLAN traffic and with one of the filters explicitly referencing one of the named instances of Layer 2 port mirroring:
  - Configure the filter **pm\_filter\_1** to use the Layer 2 port-mirroring properties configured in the named port-mirroring instance **pm\_instance\_1**. To refer to the Layer 2 port mirroring properties configured in a particular named instance of port mirroring, use the **port-mirror-instance** *port-mirroring-instance-name* statement.
  - Configure the filter **pm\_filter\_2** to use the Layer 2 port mirroring properties in effect on the underlying physical interface of the logical interface to which the filter is applied. To refer to the Layer 2 port mirroring properties in effect on the underlying physical interface, use the **port-mirror** statement. If two instances of port mirroring are bound to that port, then the firewall filter uses the first instance bound within the **[edit chassis fpc slot-number]** or **[edit chassis fpc slot-number pic slot-number]** hierarchy level.

```
[edit]
firewall {
 family ethernet-switching {
 filter pm_filter_1 {
 term pm {
 then port-mirror-instance pm_instance_1;
 }
 }
 filter pm_filter_2 {
 term pm {
 then port-mirror;
 }
 }
 }
}
```



**NOTE:** Because the **port-mirror** filter action modifier relies on the port-mirroring properties defined at the **[edit forwarding-options port-mirroring]** hierarchy level, the **port-mirror** filter action is not supported for logical systems.

4. Apply the two Layer 2 port-mirroring firewall filters to logical interfaces on interface **ge-2/0/1**:

```
[edit]
interfaces {
 ge-2/0/1 {
 flexible-vlan-tagging;
 encapsulation ethernet-bridge;
 unit 0 {
 vlan-id 201;
 family ethernet-switching {
 filter { # Explicitly references a named instance of port mirroring.
 input pm_filter_1;
 }
 }
 }
 unit 1 {
```

```

vlan-id 202;
family ethernet-switching {
 filter { # Implicitly references the underlying port mirroring.
 input pm_filter_2;
 }
}
}
}
}

```

### ***Explicit Reference of a Port Mirroring Instance***

On logical interface **ge-2/0/1.0**, the **port-mirror-instance** statement explicitly references the Layer 2 port mirroring properties in the named instance **pm\_instance\_1**. In this example, the port mirroring properties specified in **pm\_instance\_1** remain in effect at logical interface **ge-2/0/1.0** under the following conditions:

- The firewall filter **pm\_filter\_1** remains configured (as shown in step 3).
- The named instance **pm\_instance\_1** remains configured (as shown in step 1).

Even if the named instance **pm\_instance\_1** is no longer configured or no longer bound to the router (or switch) chassis at FPC 2, PIC 0, the port mirroring properties specified in **pm\_instance\_1** remain in effect at logical interface **fe-2/0/1.0** through firewall filter **pm\_filter\_1**.

### ***Implicit Reference of Port Mirroring on the Underlying Physical Interface***

On logical interface **ge-2/0/1.1**, the **port-mirror** statement implicitly references the Layer 2 port mirroring properties in effect at the underlying physical interface **ge-2/0/1**. In this example, the port mirroring properties specified in **pm\_instance\_2** remain in effect at the ports associated with FPC 2, PIC 0 under the following conditions:

- The firewall filter **pm\_filter\_2** remains configured (as shown in step 3).
- The named instance **pm\_instance\_2** remains configured (as shown in step 1).
- The named instance **pm\_instance\_2** remains bound to the router (or switch) chassis at FPC 2, PIC 0 (as shown in step 2).

If you disable the named instance **pm\_instance\_2** or delete its binding to the physical ports associated with FPC 2, PIC 0, then—if global Layer 2 port mirroring properties had been configured—the global port mirroring properties would be used at logical interface **ge-2/0/1.1** through firewall filter **pm\_filter\_2**.



**NOTE:** There is a limitation to a Layer 2 port mirroring firewall filter in which you implicitly reference Layer 2 port mirroring properties by including the **port-mirror** statement. If multiple named instances of Layer 2 port mirroring are bound to the underlying physical interface, then only the first binding in the stanza (or the only binding) is used at the logical interface. This is done mainly for backward compatibility.

In the example above, filter **pmff\_bd\_filter\_2** uses the **port-mirror** statement, and so the filter action uses the mirroring properties of the first port mirroring instance applied to the router (or switch) chassis at the **[edit chassis fpc 2 pic 0]** hierarchy level, which is the instance **pm\_instance\_1**.

**Related  
Documentation**

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Layer 2 Port Mirroring Named Instances](#)
- [Layer 2 Port Mirroring Firewall Filters](#)
- [Defining a Named Instance of Layer 2 Port Mirroring](#)
- [Defining a Layer 2 Port-Mirroring Firewall Filter](#)

## Configuration Tasks for Layer 2 Port Mirroring at Logical Interfaces

- [Defining a Layer 2 Port-Mirroring Firewall Filter on page 4224](#)
- [Applying Layer 2 Port Mirroring to a Logical Interface on page 4228](#)
- [Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a VLAN on page 4231](#)
- [Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a VPLS Routing Instance on page 4233](#)

### Defining a Layer 2 Port-Mirroring Firewall Filter

---

For virtual private LAN service (VPLS) traffic (**family ethernet-switching** or **family vpls**) and for Layer 2 VPNs with **family ccc** on MX Series routers and on EX Series switches only, you can define a firewall filter that specifies Layer 2 port mirroring as the action to be performed if a packet matches the conditions configured in the firewall filter term.

You can use a Layer 2 port-mirroring firewall filter in the following ways:

- To mirror packets received or sent on a logical interface.
- To mirror packets forwarded or flooded to a VLAN.
- To mirror packets forwarded or flooded to a VPLS routing instance.
- To mirror tunnel interface input packets only to multiple destinations.

For a summary of the three types of Layer 2 port-mirroring you can configure on an MX Series router and on an EX Series switch, see *Application of Layer 2 Port Mirroring Types*.

For information about configuring firewall filters in general (including in a Layer 3 environment), see “[Stateless Firewall Filter Overview](#)” on page 4463 and “[How Standard Firewall Filters Evaluate Packets](#)” on page 4475 in the *Junos OS Firewall Filters and Traffic Policers Configuration Guide*.

To define a firewall filter with a Layer 2 port-mirroring action:

1. Enable configuration of firewall filters for Layer 2 packets that are part of a VLAN, a Layer 2 switching cross-connect, or a virtual private LAN service (VPLS):

```
[edit]
user@host# edit firewall family family
```

The value of the **family** option can be **ethernet-switching**, **ccc**, or **vpls**.

2. Enable configuration of a firewall filter **pm-filter-name**:

```
[edit firewall family family]
user@host# edit filter pm-filter-name
```

3. Enable configuration of a firewall filter term **pm-filter-term-name**:

```
[edit firewall family family filter pm-filter-name]
user@host# edit term pm-filter-term-name
```

For more information about firewall filter terms in general (including in a Layer 3 environment), see [“Guidelines for Configuring Standard Firewall Filters” on page 4478](#) in the *Junos OS Firewall Filters and Traffic Policers Configuration Guide*.

4. (Optional) Specify the firewall filter match conditions based on the route source address *only if* you want to mirror a subset of the sampled packets.

For information about configuring firewall filter match conditions in general (including in a Layer 3 environment), see [“Firewall Filter Match Conditions Based on Numbers or Text Aliases” on page 4488](#), [“Firewall Filter Match Conditions Based on Bit-Field Values” on page 4489](#), [“Firewall Filter Match Conditions Based on Address Fields” on page 4493](#), and [“Firewall Filter Match Conditions Based on Address Classes” on page 4502](#), in the *Junos OS Firewall Filters and Traffic Policers Configuration Guide*.

- For detailed information about Layer 2 bridging firewall filter match conditions (which are supported on MX Series routers and EX Series switches only), see [“Standard Firewall Filter Match Conditions for Layer 2 Bridging Traffic” on page 4737](#).
- For detailed information about VPLS firewall filter match conditions, see [“Standard Firewall Filter Match Conditions for VPLS Traffic” on page 4727](#).
- For detailed information about Layer 2 circuit cross-connect (CCC) firewall filter match conditions, see [“Standard Firewall Filter Match Conditions for Layer 2 CCC Traffic” on page 4734](#).



**NOTE:** If you want all sampled packets to be considered to match (and be subjected to the actions specified in the **then** statement), then omit the **from** statement altogether.

5. Enable configuration of the **action** and **action-modifier** to apply to matching packets:

```
[edit firewall family family filter pm-filter-name term pm-filter-term-name]
user@host# edit then
```

6. Specify the actions to be taken on matching packets:

```
[edit firewall family family filter pm-filter-name term pm-filter-term-name then]
user@host# set action
```

The recommended value for the **action** is **accept**. If you do not specify an action, or if you omit the **then** statement entirely, all packets that match the conditions in the **from** statement are accepted.

7. Specify Layer 2 port mirroring or a next-hop group as the **action-modifier**:

- To reference the Layer 2 port mirroring properties currently in effect for the Packet Forwarding Engine or PIC associated with the underlying physical interface, use the **port-mirror** statement:

```
[edit firewall family family filter pm-filter-name term pm-filter-term-name then]
user@host# set port-mirror
```

- To reference the Layer 2 port mirroring properties configured in a specific named instance, use the **port-mirror-instance *pm-instance-name*** action modifier:

```
[edit firewall family family filter pm-filter-name term pm-filter-term-name then]
user@host# set port-mirror-instance pm-instance-name
```

If the underlying physical interface is not bound to a named instance of Layer 2 port mirroring but instead is implicitly bound to the global instance of Layer 2 port mirroring, then traffic at the logical interface is mirrored according to the properties specified in the named instance referenced by the **port-mirror-instance** action modifier.

- To reference a next-hop group that specifies the next-hop addresses (for sending additional copies of packets to an analyzer), use the **next-hop-group *pm-next-hop-group-name*** action modifier:

```
[edit firewall family family filter pm-filter-name term pm-filter-term-name then]
user@host# set next-hop-group pm-next-hop-group-name
```

For configuration information about next-hop groups, see [“Defining a Next-Hop Group for Layer 2 Port Mirroring” on page 4245](#). If you specify a next-hop group for Layer 2 port mirroring, the firewall filter term applies to the tunnel interface input only.

## 8. Verify the minimum configuration of the Layer 2 port-mirroring firewall filter:

```
[edit firewall ...]
user@host# top
[edit]
user@host# show firewall
```

```
family (ethernet-switching | ccc | vpls) { # Type of packets to mirror
 filter pm-filter-name { # Firewall filter name
 term pm-filter-term-name {
 from { # Do not specify match conditions based on route source address
 }
 then {
 action; # Recommended action is 'accept'
 action-modifier; # Three options for Layer 2 port mirroring
 }
 }
 }
}
```

In the firewall filter term **then** statement, the **action-modifier** can be **port-mirror**, **port-mirror-instance *pm-instance-name***, or **next-hop-group *pm-next-hop-group-name***.



**Related  
Documentation**

- [Layer 2 Port Mirroring Overview on page 4192](#)
- *Layer 2 Port Mirroring Firewall Filters*
- [Layer 2 Port Mirroring to Multiple Destinations Using Next-Hop Groups on page 4200](#)
- *Example: Layer 2 Port Mirroring at a Logical Interface*
- *Example: Layer 2 Port Mirroring for a Layer 2 VPN*
- *Example: Layer 2 Port Mirroring for a Layer 2 VPN with LAG Links*
- *Example: Layer 2 Port Mirroring to Multiple Destinations*

## Applying Layer 2 Port Mirroring to a Logical Interface

---

You can apply a Layer 2 port-mirroring firewall filter to the input or to the output of a logical interface, including an aggregated Ethernet logical interface. Only packets of the address-type family specified by the filter action are mirrored.

Before you begin, complete the following task:

- Define a Layer 2 port-mirroring firewall filter to be applied to the input to a logical interface or output to a logical interface. For details, see *Defining a Layer 2 Port-Mirroring Firewall Filter*.



**NOTE:** This configuration task shows two Layer 2 port-mirroring firewall filters: one filter applied to the logical interface ingress traffic, and one filter applied to the logical interface egress traffic.

To apply a Layer 2 port-mirroring firewall filter to an input or output logical interface:

1. Configure the underlying physical interface for the logical interface.

- a. Enable configuration of the underlying physical interface:

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



**NOTE:** A port-mirroring firewall filter can also be applied to an aggregated-Ethernet logical interface.

- b. For Fast Ethernet and Gigabit Ethernet interfaces and aggregated Ethernet interfaces configured for VPLS, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface:

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

- c. For Ethernet interfaces that have IEEE 802.1Q VLAN tagging and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID, set the logical link-layer encapsulation type:

```
[edit interfaces interface-name]
user@host# set encapsulation extended-vlan-ethernet-switching
```

2. Configure the logical interface to which you want to apply a Layer 2 port-mirroring firewall filter.

- a. Specify the logical unit number:

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

- b.

For a Fast Ethernet, Gigabit Ethernet, or Aggregated Ethernet interface, bind an 802.1Q VLAN tag ID to the logical interface:

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

3. Enable specification of an input or output filter to be applied to Layer 2 packets that are part of bridging domain, Layer 2 switching cross-connects, or virtual private LAN service (VPLS).

- If the filter is to be evaluated when packets are received on the interface:

```
[edit interfaces interface-name unit logical-unit-number]
user@host# set family family filter input pm-filter-name-a
```

- If the filter is to be evaluated when packets are sent on the interface:

```
[edit interfaces interface-name unit logical-unit-number]
user@host# set family family filter output pm-filter-name-b
```

The value of the *family* option can be **ethernet-switching**, **ccc**, or **vpls**.



**NOTE:** If port-mirroring firewall filters are applied at both the input and output of a logical interface, two copies of each packet are mirrored. To prevent the router or switch from forwarding duplicate packets to the same destination, include the optional **mirror-once** statement at the [edit forwarding-options] hierarchy level.

4. Verify the minimum configuration for applying a named Layer 2 port mirroring firewall filter to a logical interface:

```
[edit interfaces interface-name unit logical-unit-number family family filter ...]
user@host# top
[edit]
user@host# show interfaces
```

```
interfaces {
 interface-name {
 vlan-tagging;
 encapsulation extended-vlan-ethernet-switching;
 unit number { # Apply a filter to the input of this interface
 vlan-id number;
 family (ethernet-switching | ccc | vpls) {
 filter {
 input pm-filter-for-logical-interface-input;
 }
 }
 }
 unit number { # Apply a filter to the output of this interface
 vlan-id number;
 family (ethernet-switching | ccc | vpls) {
 filter {
 output pm-filter-for-logical-interface-output;
 }
 }
 }
 }
}
```

```
 }
 }
}
```

**Related  
Documentation**

- [Layer 2 Port Mirroring Overview on page 4192](#)
- *Layer 2 Port Mirroring Firewall Filters*
- *Defining a Layer 2 Port-Mirroring Firewall Filter*
- *Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a Bridge Domain*
- *Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a VPLS Routing Instance*
- *Example: Layer 2 Port Mirroring at a Logical Interface*
- *Example: Layer 2 Port Mirroring for a Layer 2 VPN*
- *Example: Layer 2 Port Mirroring for a Layer 2 VPN with LAG Links*

### Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a VLAN

You can apply a Layer 2 port-mirroring firewall filter to traffic being forwarded or flooded to a VLAN. Only packets of the specified family type and forwarded or flooded to that VLAN are mirrored.

Before you begin, complete the following task:

- Define a Layer 2 port-mirroring firewall filter to be applied to the traffic being forwarded to a VLAN or flooded to a VLAN. For details, see *Defining a Layer 2 Port-Mirroring Firewall Filter*.



**NOTE:** This configuration task shows two Layer 2 port-mirroring firewall filters: one filter applied to the VLAN forwarding table ingress traffic, and one filter applied to the VLAN flood table ingress traffic.

To apply a Layer 2 port-mirroring firewall filter to the forwarding table or flood table of a VLAN:

1. Enable configuration of the VLAN **bridge-domain-name** to which you want to apply a Layer 2 port-mirroring firewall filter for forwarded or flooded traffic:

- For a VLAN:

```
[edit]
user@host# edit bridge-domains bridge-domain-name
```

- For a VLAN under a routing instance:

```
[edit]
user@host# edit routing-instances routing-instance-name bridge-domains
bridge-domain-name
user@host# set instance-type virtual-switch
```

For more detailed configuration information, see *Configuring a VPLS Routing Instance*.

2. Configure the VLAN:

```
[edit]
user@host# set domain-type bridge
user@host# set interface interface-name
user@host# set routing-interface routing-interface-name
```

For more detailed configuration information, see *Configuring a Bridge Domain* and *Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances*.

3. Enable configuration of traffic forwarding on the VLAN:

```
[edit ... bridge-domains bridge-domain-name]
user@host# edit forwarding-options
```

4. Apply a Layer 2 port-mirroring firewall filter to the VLAN forwarding table or flood table.

- To mirror packets being forwarded to the VLAN:

```
[edit ... bridge-domains bridge-domain-name forwarding-options]
user@host# set filter input pm-filter-for-bd-ingress-forwarded
```

- To mirror packets being flooded to the VLAN:

```
[edit ... bridge-domains bridge-domain-name forwarding-options]
user@host# set flood input pm-filter-for-bd-ingress-flooded
```

5. Verify the minimum configuration for applying a Layer 2 port-mirroring firewall filter to the forwarding table or flood table of the VLAN.

- a. Navigate to the hierarchy level at which the VLAN is configured:

- **[edit]**
- **[edit routing-instances *routing-instance-name*]**

- b. Display the VLAN configurations:

```
user@host# show vlans
```

```
vlans {
 vlan-name {
 instance-type virtual-switch; # For a bridge domain under a routing instance.
 domain-type bridge;
 interface interface-name;
 forwarding-options {
 filter { # Mirror ingress forwarded traffic
 input pm-filter-for-bd-ingress-forwarded;
 }
 flood { # Mirror ingress flooded traffic
 input pm-filter-for-bd-ingress-flooded;
 }
 }
 }
}
```

#### Related Documentation

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Layer 2 Port Mirroring Firewall Filters](#)
- [Defining a Layer 2 Port-Mirroring Firewall Filter](#)
- [Applying Layer 2 Port Mirroring to a Logical Interface](#)
- [Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a VPLS Routing Instance](#)
- [Example: Layer 2 Port Mirroring at a Logical Interface](#)
- [Example: Layer 2 Port Mirroring for a Layer 2 VPN](#)
- [Example: Layer 2 Port Mirroring for a Layer 2 VPN with LAG Links](#)

## Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a VPLS Routing Instance

You can apply a Layer 2 port-mirroring firewall filter to traffic being forwarded or flooded to a VPLS routing instance. Only packets of the specified family type and forwarded or flooded to that VPLS routing instance are mirrored.

Before you begin, complete the following task:

- Define a Layer 2 port-mirroring firewall filter to be applied to the traffic being forwarded to a VPLS routing instance or flooded to a VLAN. For details, see *Defining a Layer 2 Port-Mirroring Firewall Filter*.



**NOTE:** This configuration task shows two Layer\_2 port-mirroring firewall filters: one filter applied to the VPLS routing instance forwarding table ingress traffic, and one filter applied to the VPLS routing instance flood table ingress traffic.

To apply a Layer 2 port-mirroring firewall filter to the forwarding table or flood table of a VPLS routing instance:

1. Enable configuration of the VPLS routing instance to which you want to apply a Layer 2 port-mirroring firewall filter for forwarded or flooded traffic:

```
[edit]
user@host# edit routing-instances routing-instance-name
user@host# set instance-type vpls
user@host# set interface interface-name
user@host# set route-distinguisher (as-number:number | ip-address:number)
user@host# set vrf-import [policy-names]
user@host# set vrf-export [policy-names]
user@host# edit protocols vpls
user@host@ ... vpls-configuration ...
```

For more detailed configuration information, see *Configuring a VPLS Routing Instance*.

2. Enable configuration of traffic forwarding on the VPLS routing instance:

```
[edit routing-instances routing-instance-name protocols vpls]
user@host# up 2
[edit routing-instances routing-instance-name]
user@host# edit forwarding-options
```

3. Apply a Layer 2 port-mirroring firewall filter to the VPLS routing instance forwarding table or flood table.

- To mirror packets being forwarded to the VPLS routing instance:

```
[edit routing-instances routing-instance-name forwarding-options]
user@host# set filter input pm-filter-for-vpls-ri-forwarded
```

- To mirror packets being flooded to the VPLS routing instance:

```
[edit routing-instances routing-instance-name forwarding-options]
```

```
user@host# set flood input pm-filter-for-vpls-ri-flooded
```

4. Verify the minimum configuration for applying a Layer 2 port-mirroring firewall filter to the forwarding table or flood table of the VPLS routing instance:

```
[edit routing-instances routing-instance-name forwarding-options]
user@host# top
[edit]
user@host# show routing-instances
```

```
routing-instances {
 routing-instance-name {
 instance-type vpls;
 interface interface-name;
 route-distinguisher (as-number:number | ip-address:number);
 vrf-import [policy-names];
 vrf-export [policy-names];
 protocols {
 vpls {
 ...vpls-configuration...
 }
 }
 forwarding-options {
 family vpls {
 filter { # Mirror ingress forwarded traffic
 input pm-filter-for-vpls-ri-forwarded;
 }
 flood { # Mirror ingress flooded traffic
 input pm-filter-for-vpls-ri-flooded;
 }
 }
 }
 }
}
```

#### Related Documentation

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Layer 2 Port Mirroring Firewall Filters](#)
- [Defining a Layer 2 Port-Mirroring Firewall Filter](#)
- [Applying Layer 2 Port Mirroring to a Logical Interface](#)
- [Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a Bridge Domain](#)
- [Example: Layer 2 Port Mirroring at a Logical Interface](#)
- [Example: Layer 2 Port Mirroring for a Layer 2 VPN](#)
- [Example: Layer 2 Port Mirroring for a Layer 2 VPN with LAG Links](#)

## Examples for Layer 2 Port Mirroring at Logical Interfaces

- [Example: Layer 2 Port Mirroring at a Logical Interface on page 4235](#)
- [Example: Layer 2 Port Mirroring for a Layer 2 VPN on page 4237](#)
- [Example: Layer 2 Port Mirroring for a Layer 2 VPN with LAG Links on page 4239](#)



### Example: Layer 2 Port Mirroring at a Logical Interface

The following steps describe an example in which the global port-mirroring instance and a port-mirroring firewall filter are used to configure Layer 2 port mirroring for the input to a logical interface.

1. Configure the VLAN **example-bd-with-analyzer**, which contains the external packet analyzer, and the VLAN **example-bd-with-traffic**, which contains the source and destination of the Layer 2 traffic being mirrored:

```
[edit]
bridge-domains {
 example-bd-with-analyzer { # Contains an external traffic analyzer
 vlan-id 1000;
 interface ge-2/0/0.0; # External analyzer
 }
 example-bd-with-traffic { # Contains traffic input and output interfaces
 vlan-id 1000;
 interface ge-2/0/6.0; # Traffic input port
 interface ge-3/0/1.2; # Traffic output port
 }
}
```

Assume that logical interface **ge-2/0/0.0** is associated with an external traffic analyzer that is to receive port-mirrored packets. Assume that logical interfaces **ge-2/0/6.0** and **ge-3/0/1.2** will be traffic input and output ports, respectively.

2. Configure Layer 2 port-mirroring for the global instance, with the port-mirroring destination being the VLAN interface associated with the external analyzer (logical interface **ge-2/0/0.0** on VLAN **example-bd-with-analyzer**). Be sure to enable the option that allows filters to be applied to this port-mirroring destination:

```
[edit]
forwarding-options {
 port-mirroring {
 input {
 rate 10;
 run-length 5;
 }
 family ethernet-switching {
 output {
 interface ge-2/0/0.0; # Mirror packets to the external analyzer
 no-filter-check; # Allow filters on the mirror destination interface
 }
 }
 }
}
```

The **input** statement at the **[edit forwarding-options port-mirroring]** hierarchy level specifies that sampling begins every tenth packet and that each of the first five packets selected are to be mirrored.

The **output** statement at the **[edit forwarding-options port-mirroring family ethernet-switching]** hierarchy level specifies the output mirror interface for Layer 2 packets in a bridging environment:

- Logical interface **ge-2/0/0.0**, which is associated with the external packet analyzer, is configured as the port-mirroring destination.
- The optional **no-filter-check** statement allows filters to be configured on this destination interface.

3. Configure the Layer 2 port-mirroring firewall filter **example-bridge-pm-filter**:

```
[edit]
firewall {
 family ethernet-switching {
 filter example-bridge-pm-filter {
 term example-filter-terms {
 then {
 accept;
 port-mirror;
 }
 }
 }
 }
}
```

When this firewall filter is applied to the input or output of a logical interface for traffic in a bridging environment, Layer 2 port mirroring is performed according to the input packet-sampling properties and mirror destination properties configured for the Layer 2 port mirroring global instance. Because this firewall filter is configured with the single, default filter action **accept**, all packets selected by the **input** properties (**rate = 10** and **run-length = 5**) match this filter.

4. Configure the logical interfaces:

```
[edit]
interfaces {
 ge-2/0/0 { # Define the interface to the external analyzer
 encapsulation ethernet-bridge;
 unit 0 {
 family ethernet-switching;
 }
 }
 ge-2/0/6 { # Define the traffic input port
 flexible-vlan-tagging;
 encapsulation extended-vlan-bridge;
 unit 0 {
 vlan-id 100;
 family ethernet-switching {
 filter {
 input example-bridge-pm-filter; # Apply the port-mirroring firewall filter
 }
 }
 }
 }
 ge-3/0/1 { # Define the traffic output port
```

```

flexible-vlan-tagging;
encapsulation extended-vlan-bridge;
unit 2 {
 vlan-tags outer 10 inner 20;
 family ethernet-switching;
}
}
}

```

Packets received at logical interface **ge-2/0/6.0** on VLAN **example-bd-with-traffic** are evaluated by the port-mirroring firewall filter **example-bridge-pm-filter**. The firewall filter acts on the input traffic according to the filter actions configured in the firewall filter itself plus the input packet-sampling properties and mirror destination properties configured in the global port-mirroring instance:

- All packets received at **ge-2/0/6.0** are forwarded to their (assumed) normal destination at logical interface **ge-3/0/1.2**.
- For every ten input packets, copies of the first five packets in that selection are forwarded to the external analyzer at logical interface **ge-0/0/0.0** in the other VLAN, **example-bd-with-analyzer**.

If you configure the port-mirroring firewall filter **example-bridge-pm-filter** to take the **discard** action instead of the **accept** action, all original packets are discarded while copies of the packets selected using the global port-mirroring **input** properties are sent to the external analyzer.

#### Related Documentation

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Layer 2 Port Mirroring Firewall Filters](#)
- [Defining a Layer 2 Port-Mirroring Firewall Filter](#)

#### Example: Layer 2 Port Mirroring for a Layer 2 VPN

The following example is not a complete configuration, but shows all the steps needed to configure port mirroring on an L2VPN using **family ccc**.

1. Configure the VLAN **port-mirror-bd**, which contains the external packet analyzer:

```

[edit]
vllns {
 port-mirror-vlln { # Contains an external traffic analyzer
 interface ge-2/2/9.0; # External analyzer
 }
}

```

2. Configure the Layer 2 VPN CCC to connect logical interface **ge-2/0/1.0** and logical interface **ge-2/0/1.1**:

```

[edit]
protocols {
 mpls {
 interface all;
 }
}

```

```
connections {
 interface-switch if_switch {
 interface ge-2/0/1.0;
 interface ge-2/0/1.1;
 }
}
```

3. Configure Layer 2 port mirroring for the global instance, with the port-mirroring destination being the VLAN interface associated with the external analyzer (logical interface **ge-2/2/9.0** on VLAN **example-bd-with-analyzer**):

```
[edit]
forwarding-options {
 port-mirroring {
 input {
 rate 1;
 maximum-packet-length 200;
 }
 family ccc {
 output {
 interface ge-2/2/9.0; # Mirror packets to the external analyzer
 }
 }
 instance {
 inst1 {
 input {
 rate 1;
 maximum-packet-length 300;
 }
 family ccc {
 output {
 interface ge-2/2/9.0;
 }
 }
 }
 }
 }
}
```

4. Define the Layer 2 port-mirroring firewall filter **pm\_filter\_ccc** for **family ccc**:

```
[edit]
firewall {
 family ccc {
 filter pm_filter_ccc {
 term pm {
 then port-mirror;
 }
 }
 }
}
```

5. Apply the port mirror instance to the chassis:

```
[edit]
chassis {
```

```
fpc 2 {
 port-mirror-instance inst1;
}
```

6. Configure interface **ge-2/2/9** for the VLANs, and configure interface **ge-2/0/1** for port mirroring with the **pm\_filter\_ccc** firewall filter:

```
[edit]
interfaces {
 ge-2/2/9 {
 encapsulation ethernet-bridge;
 unit 0 {
 family ethernet-switching;
 }
 }
 ge-2/0/1 {
 vlan-tagging;
 encapsulation extended-vlan-ccc;
 unit 0 {
 vlan-id 10;
 family ccc {
 filter {
 input pm_filter_ccc;
 }
 }
 }
 unit 1 {
 vlan-id 20;
 family ccc {
 filter {
 output pm_filter_ccc;
 }
 }
 }
 }
}
```

#### Related Documentation

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Layer 2 Port Mirroring Firewall Filters](#)
- [Defining a Layer 2 Port-Mirroring Firewall Filter](#)

#### Example: Layer 2 Port Mirroring for a Layer 2 VPN with LAG Links

The following example is not a complete configuration, but shows all the steps needed to configure port mirroring on an L2VPN using **family ccc** and aggregated Ethernet links.

1. Configure the VLAN **port\_mirror\_bd**, which contains the external packet analyzer:

```
[edit]
vlands {
 port_mirror_vlan { # Contains an external traffic analyzer
 interface ge-2/2/8.0; # External analyzer
 }
}
```

```
}
```

2. Configure the Layer 2 VPN CCC to connect interface **ae0.0** and interface **ae0.1**:

```
[edit]
protocols {
 mpls {
 interface all;
 }
 connections {
 interface-switch if_switch {
 interface ae0.0;
 interface ae0.1;
 }
 }
}
```

3. Configure Layer 2 port mirroring for the global instance, with the port-mirroring destination being the VLAN interface associated with the external analyzer (logical interface **ge-2/2/9.0** on VLAN **example\_bd\_with\_analyzer**):

```
[edit]
forwarding-options {
 port-mirroring {
 input {
 rate 1;
 maximum-packet-length 200;
 }
 family ccc {
 output {
 interface ge-2/2/8.0; # Mirror packets to the external analyzer
 }
 }
 instance {
 pm_instance_1 {
 input {
 rate 1;
 maximum-packet-length 300;
 }
 family ccc {
 output {
 interface ge-2/2/8.0;
 }
 }
 }
 }
 }
}
```

4. Configure the firewall filter **pm\_ccc** for **family ccc**:

```
[edit]
firewall {
 family ccc {
 filter pm_ccc {
 term pm {
 then port-mirror;
 }
 }
 }
}
```

```

 }
 }
}

```

5. Apply the aggregated Ethernet interfaces and port mirror instance to the chassis:

```

[edit]
chassis {
 aggregated-devices {
 ethernet {
 device-count 10;
 }
 }
 fpc 2 {
 port-mirror-instance pm_instance_1;
 }
}

```

6. Configure interfaces **ae0** and **ge-2/0/2** (for aggregated Ethernet) and **ge-2/2/8** (for port mirroring) with the **pm\_ccc** filter:

```

[edit]
interfaces {
 ae0 {
 vlan-tagging;
 encapsulation extended-vlan-ccc;
 unit 0 {
 vlan-id 10;
 family ccc {
 filter {
 input pm_ccc;
 }
 }
 }
 }
 unit 1 {
 vlan-id 20;
 family ccc {
 filter {
 output pm_ccc;
 }
 }
 }
}
ge-2/0/2 {
 gigether-options {
 802.3ad ae0;
 }
}
ge-2/2/8 {
 encapsulation ethernet-bridge;
 unit 0 {
 family ethernet-switching;
 }
}
}

```

- Related Documentation**
- [Layer 2 Port Mirroring Overview on page 4192](#)
  - [Layer 2 Port Mirroring Firewall Filters](#)
  - [Defining a Layer 2 Port-Mirroring Firewall Filter](#)

## Configuration Tasks for Layer 2 Port Mirroring at Multiple Destinations

- [Defining a Layer 2 Port-Mirroring Firewall Filter on page 4242](#)
- [Defining a Next-Hop Group for Layer 2 Port Mirroring on page 4245](#)
- [Applying Layer 2 Port Mirroring to a Logical Interface on page 4247](#)

### Defining a Layer 2 Port-Mirroring Firewall Filter

---

For virtual private LAN service (VPLS) traffic (**family ethernet-switching** or **family vpls**) and for Layer 2 VPNs with **family ccc** on MX Series routers and on EX Series switches only, you can define a firewall filter that specifies Layer 2 port mirroring as the action to be performed if a packet matches the conditions configured in the firewall filter term.

You can use a Layer 2 port-mirroring firewall filter in the following ways:

- To mirror packets received or sent on a logical interface.
- To mirror packets forwarded or flooded to a VLAN.
- To mirror packets forwarded or flooded to a VPLS routing instance.
- To mirror tunnel interface input packets only to multiple destinations.

For a summary of the three types of Layer 2 port-mirroring you can configure on an MX Series router and on an EX Series switch, see *Application of Layer 2 Port Mirroring Types*.

For information about configuring firewall filters in general (including in a Layer 3 environment), see “[Stateless Firewall Filter Overview](#)” on page 4463 and “[How Standard Firewall Filters Evaluate Packets](#)” on page 4475 in the *Junos OS Firewall Filters and Traffic Policers Configuration Guide*.

To define a firewall filter with a Layer 2 port-mirroring action:

1. Enable configuration of firewall filters for Layer 2 packets that are part of a VLAN, a Layer 2 switching cross-connect, or a virtual private LAN service (VPLS):

```
[edit]
user@host# edit firewall family family
```

The value of the **family** option can be **ethernet-switching**, **ccc**, or **vpls**.

2. Enable configuration of a firewall filter **pm-filter-name**:

```
[edit firewall family family]
user@host# edit filter pm-filter-name
```

3. Enable configuration of a firewall filter term **pm-filter-term-name**:

```
[edit firewall family family filter pm-filter-name]
user@host# edit term pm-filter-term-name
```



For more information about firewall filter terms in general (including in a Layer 3 environment), see [“Guidelines for Configuring Standard Firewall Filters” on page 4478](#) in the *Junos OS Firewall Filters and Traffic Policers Configuration Guide*.

4. (Optional) Specify the firewall filter match conditions based on the route source address *only if* you want to mirror a subset of the sampled packets.

For information about configuring firewall filter match conditions in general (including in a Layer 3 environment), see [“Firewall Filter Match Conditions Based on Numbers or Text Aliases” on page 4488](#), [“Firewall Filter Match Conditions Based on Bit-Field Values” on page 4489](#), [“Firewall Filter Match Conditions Based on Address Fields” on page 4493](#), and [“Firewall Filter Match Conditions Based on Address Classes” on page 4502](#), in the *Junos OS Firewall Filters and Traffic Policers Configuration Guide*.

- For detailed information about Layer 2 bridging firewall filter match conditions (which are supported on MX Series routers and EX Series switches only), see [“Standard Firewall Filter Match Conditions for Layer 2 Bridging Traffic” on page 4737](#).
- For detailed information about VPLS firewall filter match conditions, see [“Standard Firewall Filter Match Conditions for VPLS Traffic” on page 4727](#).
- For detailed information about Layer 2 circuit cross-connect (CCC) firewall filter match conditions, see [“Standard Firewall Filter Match Conditions for Layer 2 CCC Traffic” on page 4734](#).



**NOTE:** If you want all sampled packets to be considered to match (and be subjected to the actions specified in the **then** statement), then omit the **from** statement altogether.

5. Enable configuration of the **action** and **action-modifier** to apply to matching packets:

```
[edit firewall family family filter pm-filter-name term pm-filter-term-name]
user@host# edit then
```

6. Specify the actions to be taken on matching packets:

```
[edit firewall family family filter pm-filter-name term pm-filter-term-name then]
user@host# set action
```

The recommended value for the **action** is **accept**. If you do not specify an action, or if you omit the **then** statement entirely, all packets that match the conditions in the **from** statement are accepted.

7. Specify Layer 2 port mirroring or a next-hop group as the **action-modifier**:

- To reference the Layer 2 port mirroring properties currently in effect for the Packet Forwarding Engine or PIC associated with the underlying physical interface, use the **port-mirror** statement:

```
[edit firewall family family filter pm-filter-name term pm-filter-term-name then]
user@host# set port-mirror
```

- To reference the Layer 2 port mirroring properties configured in a specific named instance, use the **port-mirror-instance** *pm-instance-name* action modifier:

```
[edit firewall family family filter pm-filter-name term pm-filter-term-name then]
user@host# set port-mirror-instance pm-instance-name
```

If the underlying physical interface is not bound to a named instance of Layer 2 port mirroring but instead is implicitly bound to the global instance of Layer 2 port mirroring, then traffic at the logical interface is mirrored according to the properties specified in the named instance referenced by the **port-mirror-instance** action modifier.

- To reference a next-hop group that specifies the next-hop addresses (for sending additional copies of packets to an analyzer), use the **next-hop-group** *pm-next-hop-group-name* action modifier:

```
[edit firewall family family filter pm-filter-name term pm-filter-term-name then]
user@host# set next-hop-group pm-next-hop-group-name
```

For configuration information about next-hop groups, see [“Defining a Next-Hop Group for Layer 2 Port Mirroring” on page 4245](#). If you specify a next-hop group for Layer 2 port mirroring, the firewall filter term applies to the tunnel interface input only.

- Verify the minimum configuration of the Layer 2 port-mirroring firewall filter:

```
[edit firewall ...]
user@host# top
[edit]
user@host# show firewall

family (ethernet-switching | ccc | vpls) { # Type of packets to mirror
 filter pm-filter-name { # Firewall filter name
 term pm-filter-term-name {
 from { # Do not specify match conditions based on route source address
 }
 then {
 action; # Recommended action is 'accept'
 action-modifier; # Three options for Layer 2 port mirroring
 }
 }
 }
}
```

In the firewall filter term **then** statement, the *action-modifier* can be **port-mirror**, **port-mirror-instance** *pm-instance-name*, or **next-hop-group** *pm-next-hop-group-name*.

#### Related Documentation

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Layer 2 Port Mirroring Firewall Filters](#)
- [Layer 2 Port Mirroring to Multiple Destinations Using Next-Hop Groups on page 4200](#)
- [Example: Layer 2 Port Mirroring at a Logical Interface](#)
- [Example: Layer 2 Port Mirroring for a Layer 2 VPN](#)
- [Example: Layer 2 Port Mirroring for a Layer 2 VPN with LAG Links](#)

- *Example: Layer 2 Port Mirroring to Multiple Destinations*

### Defining a Next-Hop Group for Layer 2 Port Mirroring

On MX Series routers and EX Series switches, you can mirror tunnel interface input traffic to multiple destinations. To this form of *multipacket port mirroring*, you specify two or more additional destinations in a next-hop group, define a firewall filter that references the next-hop group as the filter action, and then apply the filter to a logical tunnel interface (**lt-**) or virtual tunnel interface (**vt-**) on the MX Series router and on an EX Series switch.



**NOTE:** This topic describes how to define a next-hop group for Layer 2 port mirroring to multiple destinations. For detailed information about defining a firewall filter for Layer 2 port mirroring to multiple destinations, see *Defining a Layer 2 Port-Mirroring Firewall Filter*.

To define a next-hop group for a Layer 2 port-mirroring firewall filter action:

1. Enable configuration of Layer 2 forwarding options.

- To enable Layer 2 forwarding options at the top level:

```
[edit]
user@host edit forwarding-options port-mirroring family (ccc | vpls) output
```

- To enable Layer 2 forwarding options for a routing instance:

```
[edit]
user@host edit forwarding-options port-mirroring instance instance-name
family (ccc | vpls) output
```

2. Enable configuration of a next-hop-group for Layer 2 port mirroring:

```
[edit forwarding-options port-mirroring ... family (ccc | vpls) output]
user@host# edit next-hop-group pm-next-hop-group-name
```

3. Specify the type of addresses to be used in the next-hop group configuration. By default, the next-hop group is specified using Layer 3 addresses (**group-type inet**). To specify the next-hop group using Layer 2 addresses instead, you must include the **group-type layer-2** statement:

```
[edit forwarding-options port-mirroring ... family (ccc | vpls) output next-hop-group
pm-next-hop-group-name]
user@host# set group-type layer-2
```

4. Specify the logical interfaces of the next-hop route (or switch) r:

```
[edit forwarding-options port-mirroring ... family (ccc | vpls) output next-hop-group
pm-next-hop-group-name]
user@host# set interface logical-interface-name-1
user@host# set interface logical-interface-name-2
```

The MX Series router and the EX Series switch supports up to 30 next-hop groups. Each next-hop group supports up to 16 next-hop addresses. Each next-hop group must specify at least two addresses.

## 5. Verify the configuration of the next-hop group:

```
[edit forwarding-options port-mirroring ... family (ccc | vpls) output next-hop-group
 pm-next-hop-group-name]
user@host# top
[edit]
user@host# show forwarding-options

...
next-hop-group pm-next-hop-group-name { # Next-hop group on a bridge domain.
 group-type layer-2;
 interface logical-interface-name-1;
 interface logical-interface-name-2;
}
...
```

**Related  
Documentation**

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Layer 2 Port Mirroring to Multiple Destinations Using Next-Hop Groups on page 4200](#)
- *Defining a Layer 2 Port-Mirroring Firewall Filter*
- [Displaying Next-Hop Group Settings and Status on page 4264](#)
- *Example: Layer 2 Port Mirroring to Multiple Destinations*

## Applying Layer 2 Port Mirroring to a Logical Interface

You can apply a Layer 2 port-mirroring firewall filter to the input or to the output of a logical interface, including an aggregated Ethernet logical interface. Only packets of the address-type family specified by the filter action are mirrored.

Before you begin, complete the following task:

- Define a Layer 2 port-mirroring firewall filter to be applied to the input to a logical interface or output to a logical interface. For details, see *Defining a Layer 2 Port-Mirroring Firewall Filter*.



**NOTE:** This configuration task shows two Layer 2 port-mirroring firewall filters: one filter applied to the logical interface ingress traffic, and one filter applied to the logical interface egress traffic.

To apply a Layer 2 port-mirroring firewall filter to an input or output logical interface:

1. Configure the underlying physical interface for the logical interface.

- a. Enable configuration of the underlying physical interface:

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



**NOTE:** A port-mirroring firewall filter can also be applied to an aggregated-Ethernet logical interface.

- b. For Fast Ethernet and Gigabit Ethernet interfaces and aggregated Ethernet interfaces configured for VPLS, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface:

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

- c. For Ethernet interfaces that have IEEE 802.1Q VLAN tagging and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID, set the logical link-layer encapsulation type:

```
[edit interfaces interface-name]
user@host# set encapsulation extended-vlan-ethernet-switching
```

2. Configure the logical interface to which you want to apply a Layer 2 port-mirroring firewall filter.

- a. Specify the logical unit number:

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

- b.

For a Fast Ethernet, Gigabit Ethernet, or Aggregated Ethernet interface, bind an 802.1Q VLAN tag ID to the logical interface:

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

3. Enable specification of an input or output filter to be applied to Layer 2 packets that are part of bridging domain, Layer 2 switching cross-connects, or virtual private LAN service (VPLS).

- If the filter is to be evaluated when packets are received on the interface:

```
[edit interfaces interface-name unit logical-unit-number]
user@host# set family family filter input pm-filter-name-a
```

- If the filter is to be evaluated when packets are sent on the interface:

```
[edit interfaces interface-name unit logical-unit-number]
user@host# set family family filter output pm-filter-name-b
```

The value of the *family* option can be **ethernet-switching**, **ccc**, or **vpls**.



**NOTE:** If port-mirroring firewall filters are applied at both the input and output of a logical interface, two copies of each packet are mirrored. To prevent the router or switch from forwarding duplicate packets to the same destination, include the optional **mirror-once** statement at the [edit forwarding-options] hierarchy level.

4. Verify the minimum configuration for applying a named Layer 2 port mirroring firewall filter to a logical interface:

```
[edit interfaces interface-name unit logical-unit-number family family filter ...]
user@host# top
[edit]
user@host# show interfaces
```

```
interfaces {
 interface-name {
 vlan-tagging;
 encapsulation extended-vlan-ethernet-switching;
 unit number { # Apply a filter to the input of this interface
 vlan-id number;
 family (ethernet-switching | ccc | vpls) {
 filter {
 input pm-filter-for-logical-interface-input;
 }
 }
 }
 unit number { # Apply a filter to the output of this interface
 vlan-id number;
 family (ethernet-switching | ccc | vpls) {
 filter {
 output pm-filter-for-logical-interface-output;
 }
 }
 }
 }
}
```

```

 }
 }
}

```

#### Related Documentation

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Layer 2 Port Mirroring Firewall Filters](#)
- [Defining a Layer 2 Port-Mirroring Firewall Filter](#)
- [Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a Bridge Domain](#)
- [Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a VPLS Routing Instance](#)
- [Example: Layer 2 Port Mirroring at a Logical Interface](#)
- [Example: Layer 2 Port Mirroring for a Layer 2 VPN](#)
- [Example: Layer 2 Port Mirroring for a Layer 2 VPN with LAG Links](#)

## Examples of Layer 2 Port Mirroring at Multiple Destinations

- [Example: Layer 2 Port Mirroring to Multiple Destinations on page 4249](#)

### Example: Layer 2 Port Mirroring to Multiple Destinations

On MX Series routers and EX Series switches, you can mirror traffic to multiple destinations by configuring next-hop groups in Layer 2 port-mirroring firewall filters applied to tunnel interfaces.

1. Configure the chassis to support tunnel services at PIC 0 on FPC 2. This configuration includes two logical tunnel interfaces on FPC 2, PIC 0, port 10.

```

[edit]
chassis {
 fpc 2 {
 pic 0 {
 tunnel-services {
 bandwidth 1g;
 }
 }
 }
}

```

2. Configure the physical and logical interfaces for three VLANs and one Layer 2 VPN CCC:
  - VLAN **bd** will span logical interfaces **ge-2/0/1.0** and **ge-2/0/1.1**.
  - VLAN **bd\_next\_hop\_group** will span logical interfaces **ge-2/2/9.0** and **ge-2/0/2.0**.
  - VLAN **bd\_port\_mirror** will use the logical tunnel interface **lt-2/0/10.2**.
  - Layer 2 VPN CCC **if\_switch** will connect logical interfaces **ge-2/0/1.2** and **lt-2/0/10.1**.

```

[edit]
interfaces {
 ge-2/0/1 {
 flexible-vlan-tagging;
 }
}

```

```
encapsulation flexible-ethernet-services;
unit 0 { # An interface on bridge domain 'bd'.
 encapsulation vlan-bridge;
 vlan-id 200;
 family ethernet-switching {
 filter {
 input pm_bridge;
 }
 }
}
unit 1 { # An interface on bridge domain 'bd'.
 encapsulation vlan-bridge;
 vlan-id 201;
 family ethernet-switching {
 filter {
 input pm_bridge;
 }
 }
}
unit 2 {
 encapsulation vlan-ccc;
 vlan-id 1000;
}
}
ge-2/0/2 { # For 'bd_next_hop_group'
 encapsulation ethernet-bridge;
 unit 0 {
 family ethernet-switching;
 }
}
lt-2/0/10 {
 unit 1 {
 encapsulation ethernet-ccc;
 peer-unit 2;
 }
 unit 2 {
 encapsulation ethernet-bridge;
 peer-unit 1;
 family ethernet-switching {
 filter {
 output redirect_to_nhg;
 }
 }
 }
}
}
ge-2/2/9 {
 encapsulation ethernet-bridge;
 unit 0 { # For 'bd_next_hop_group'
 family ethernet-switching;
 }
}
}
```



## 3. Configure the three VLANs and the Layer 2 VPN switching CCC:

- VLAN **bd** spans logical interfaces **ge-2/0/1.0** and **ge-2/0/1.1**.
- VLAN **bd\_next\_hop\_group** spans logical interfaces **ge-2/2/9.0** and **ge-2/0/2.0**.
- VLAN **bd\_port\_mirror** uses the logical tunnel interface **lt-2/0/10.2**.
- Layer 2 VPN CCC **if\_switch** connects interfaces **ge-2/0/1.2** and **lt-2/0/10.1**.

```
[edit]
vpls {
 vlans {
 vlans {
 interface ge-2/0/1.0;
 interface ge-2/0/1.1;
 }
 bd_next_hop_group {
 interface ge-2/2/9.0;
 interface ge-2/0/2.0;
 }
 bd_port_mirror {
 interface lt-2/0/10.2;
 }
 }
}
protocols {
 mpls {
 interface all;
 }
}
connections {
 interface-switch if_switch {
 interface ge-2/0/1.2;
 interface lt-2/0/10.1;
 }
}
```

For detailed information about configuring the CCC connection for Layer 2 switching cross-connects, see the *Junos OS MPLS Applications Configuration Guide*.

## 4. Configure forwarding options:

- Configure global port mirroring properties to mirror **family vpls** traffic to an interface on the bridge domain **bd\_port\_mirror**.
- Configure the next-hop group **nhg\_mirror\_to\_bd** to forward Layer 2 traffic to the VLAN **bd\_next\_hop\_group**.

Both of these forwarding options will be referenced by the port-mirroring firewall filter:

```
[edit]
forwarding-options {
 port-mirroring { # Global port mirroring properties.
 input {
 rate 1;
 }
 family vpls {
 output {
 interface lt-2/0/10.2; # Interface on 'bd_port_mirror' bridge domain.
 no-filter-check;
 }
 }
 }
}
```

```

 }
 }
}
next-hop-group nhg_mirror_to_bd { # Configure a next-hop group.
 group-type layer-2; # Specify 'layer-2' for Layer 2; default 'inet' is for Layer 3.
 interface ge-2/0/2.0; # Interface on 'bd_next_hop_group' bridge domain.
 interface ge-2/2/9.0; # Interface on 'bd_next_hop_group' bridge domain.
}
}

```

5. Configure two Layer 2 port-mirroring firewall filters for **family bridge** traffic:

- **filter\_pm\_bridge**—Sends all **family bridge** traffic to the global port mirroring destination.
- **filter\_redirect\_to\_nhg**—Sends all **family bridge** traffic to the final next-hop group **nhg\_mirror\_to\_bd**.

Layer 2 port-mirroring firewall filters for **family bridge** traffic applies to traffic on a physical interface configured with encapsulation **ethernet-bridge**.

```

[edit]
firewall {
 family bridge {
 filter filter_pm_bridge {
 term term_port_mirror {
 then port-mirror;
 }
 }
 filter filter_redirect_to_nhg {
 term term_nhg {
 then next-hop-group nhg_mirror_to_bd;
 }
 }
 }
}

```

#### Related Documentation

- [Layer 2 Port Mirroring Overview on page 4192](#)
- [Layer 2 Port Mirroring to Multiple Destinations Using Next-Hop Groups on page 4200](#)
- [Defining a Layer 2 Port-Mirroring Firewall Filter](#)
- [Defining a Next-Hop Group for Layer 2 Port Mirroring on page 4245](#)
- [Displaying Next-Hop Group Settings and Status on page 4264](#)

## Configuring Inline Port Mirroring

- [Configuring Inline Port Mirroring on page 4252](#)

### Configuring Inline Port Mirroring

Inline port mirroring provides you with the ability to specify instances that are not bound to the flexible PIC concentrator (FPC) in the firewall filter's **then port-mirror-instance** action. This way, you are not limited to only two port-mirror instances per FPC. Inline port mirroring decouples the port-mirror destination from the input parameters like **rate**. While

the input parameters are programmed in the switch interface board, the next-hop destination of the mirrored packet is available in the packet itself. Inline port mirroring is supported only on Trio-based modular port concentrators (MPCs).

Using inline port mirroring, a port-mirror instance will have an option to inherit input parameters from another instance that specifies it, as shown in the following CLI configuration example:

```
instance pm2 {
 + input-parameters-instance pm1;
 family inet {
 output {
 interface ge-1/2/3.0 {
 next-hop 50.0.0.3;
 }
 }
 }
}
```

Multiple levels of inheritance are not allowed. One instance can be referred by multiple instances. An instance can refer to another instance that is defined before it. Forward references are not allowed and an instance cannot refer to itself, doing so will cause an error during configuration parsing.

The user can specify an instance that is not bound to the FPC in the firewall filter. The specified filter should inherit one of the two instances that have been bound to the FPC. If it does not, the packet is not marked for port-mirroring. If it does, then the packet will be sampled using the input parameters specified by the referred instance but the copy will be sent to the its own destination.

#### Related Documentation

- [Layer 2 Port Mirroring Overview on page 4192](#)
- *Defining a Layer 2 Port-Mirroring Firewall Filter*
- *Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a Bridge Domain*
- *Applying Layer 2 Port Mirroring to Traffic Forwarded or Flooded to a VPLS Routing Instance*
- *Example: Layer 2 Port Mirroring at a Logical Interface*
- *Example: Layer 2 Port Mirroring for a Layer 2 VPN*
- *Example: Layer 2 Port Mirroring for a Layer 2 VPN with LAG Links*

## Configuration Statements

- [\[edit forwarding-options port-mirroring\] Hierarchy Level on page 4253](#)

### [\[edit forwarding-options port-mirroring\] Hierarchy Level](#)

```
forwarding-options {
 port-mirroring {
 disable;
 disable-all-instances;
 family (ccc | ethernet-switching | inet | inet6 | vpls) {
 output {
```

```
(interface interface-name | next-hop-group group-name);
no-filter-check;
routing-instance instance-name;
vlan (vlan-name | vlan-id) <no-tag>;
}
}
family ccc {
 output {
 interface interface-name
 next-hop-group group-name;
 no-filter-check;
 }
}
family ethernet-switching {
 output {
 interface interface-name
 next-hop-group group-name;
 no-filter-check;
 }
}
family inet {
 output {
 interface interface-name {
 next-hop ipv4-address;
 }
 next-hop-group group-name;
 no-filter-check;
 }
}
family inet6 {
 output {
 interface interface-name {
 next-hop ipv6-address;
 }
 no-filter-check;
 }
}
family vpls {
 output {
 interface interface-name
 next-hop-group group-name;
 no-filter-check;
 }
}
input {
 maximum-packet-length bytes;
 rate rate;
}
instance instance-name {
 disable;
 family family-name {
 ... same statements as at the [edit forwarding-options port-mirroring family (ccc |
 inet | inet6 | vpls)] hierarchy levels ...
 }
 input {
```

```

 ...same statements as at the [edit forwarding-options port-mirroring input] hierarchy
 level ...
 }
}
mirror-once;
traceoptions {
 file <filename> <files number> <match regular-expression> <size maximum-file-size>
 <world-readable | no-world-readable>;
 no-remote-trace;
}
}
}

```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
  - *[edit forwarding-options] Hierarchy Level*

## disable (Forwarding Options)

|                                 |                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable;                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit forwarding-options port-mirror],<br>[edit forwarding-options port-mirror instance <i>instance-name</i> ],<br>[edit forwarding-options sampling],<br>[edit forwarding-options sampling instance <i>instance-name</i> ],<br>[edit forwarding-options sampling family (inet  inet6  mpls) ],<br>[edit forwarding-options sampling family (inet  inet6  mpls) output file] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement added to <b>port-mirror</b> hierarchy in Junos OS Release 9.6.                                                                                                                                                                                                                                                |
| <b>Description</b>              | Disable traffic accounting, port mirroring, or sampling.                                                                                                                                                                                                                                                                                                                     |
| <b>Usage Guidelines</b>         | See <i>Configuring Traffic Sampling</i> or <i>Configuring Port Mirroring</i> .                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                      |

## disable-all-instances

|                                 |                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable-all-instances;                                                                                                  |
| <b>Hierarchy Level</b>          | [edit forwarding-options port-mirror]                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.                                                                           |
| <b>Description</b>              | Disable all port mirroring instances globally.                                                                          |
| <b>Usage Guidelines</b>         | See <i>Configuring Port Mirroring</i> .                                                                                 |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration. |

## forwarding-options

---

|                                 |                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | forwarding-options { ... }                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit]                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 11.3 for QFX Series switches.                                       |
| <b>Description</b>              | Configure traffic forwarding.<br><br>The statements are explained separately.                                                                                     |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Traffic Forwarding and Monitoring</i></li><li>• <i>[edit forwarding-options] Hierarchy Level</i></li></ul> |

## family (Port Mirroring)

---

|                                 |                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>family (inet   inet6) {<br/>    output {<br/>        interface <i>interface-name</i> {<br/>            next-hop <i>address</i>;<br/>        }<br/>        no-filter-check;<br/>    }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit forwarding-options port-mirroring]                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                     |
| <b>Description</b>              | Configure the protocol family to be sampled. Only IPv4 ( <b>inet</b> ) and IPv6 ( <b>inet6</b> ) are supported.<br><br>The statements are explained separately.                                       |
| <b>Usage Guidelines</b>         | See <i>Configuring Port Mirroring</i> .                                                                                                                                                               |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                               |

---

## input (Port Mirroring)

---

|                                 |                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>input {<br/>    rate <i>number</i>;<br/>    run-length <i>number</i>;<br/>}</pre>                                                                                                 |
| <b>Hierarchy Level</b>          | [edit forwarding-options port-mirroring],<br>[edit forwarding-options port-mirroring instance <i>instance-name</i> ]<br>[edit forwarding-options port-mirroring family (inet   inet6)] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                      |
| <b>Description</b>              | Configure port mirroring on a logical interface.<br><br>The statements are explained separately.                                                                                       |
| <b>Usage Guidelines</b>         | See <i>Configuring Port Mirroring</i> .                                                                                                                                                |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                |

## instance

---

**Syntax**

```
instance {
 instance-name {
 input {
 maximum-packet-length bytes;
 rate number;
 run-length number;
 }
 family (ccc| inet | inet6 | vpls) {
 output {
 interface interface-name {
 next-hop address;
 }
 no-filter-check;
 }
 }
 }
}
```

**Hierarchy Level** [edit [forwarding-options](#) port-mirroring]

**Release Information** Statement introduced in Junos OS Release 9.3 (MX Series routers only). Support extended to M120 and M320 routers in Junos OS Release 9.5. **maximum-packet-length** and **ccc** options introduced in Junos OS Release 9.6 for M120 and M320 routers only.

**Description** Configure a port-mirroring instance.

**Options** *port-mirroring-instance-name*—Name of the port-mirroring instance.

The remaining statements are explained separately.

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

**Related Documentation**

- *Configuring Port Mirroring*



## interface (Port Mirroring)

---

|                                 |                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>interface <i>interface-name</i> {<br/>    next-hop <i>address</i>;<br/>}</code>                                              |
| <b>Hierarchy Level</b>          | [edit forwarding-options port-mirroring family (inet   inet6) output]                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                  |
| <b>Description</b>              | Specify the output interface for sending copies of packets elsewhere to be analyzed.                                               |
| <b>Options</b>                  | <p><i>interface-name</i>—Name of the interface.</p> <p>The remaining statements are explained separately.</p>                      |
| <b>Usage Guidelines</b>         | See <i>Configuring Port Mirroring</i> .                                                                                            |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |


## interface (Next-Hop Group)

---

|                                 |                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>interface <i>interface-name</i> {<br/>    next-hop <i>address</i>;<br/>}</code>                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit forwarding-options next-hop-group <i>group-name</i> ]                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Specify the output interface for sending copies of packets elsewhere to be analyzed.</p> <p>The commit operation fails when a next-hop group has only one interface configured. It is implicitly assumed that a subgroup is up only if more than one interface in the subgroup is up.</p> |
| <b>Options</b>                  | <p><i>interface-name</i>—Name of the interface.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Next-Hop Groups</i></li> </ul>                                                                                                                                                                                                       |

## maximum-packet-length

---

|                                                                                                                                                                                 |                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                   | maximum-packet-length <i>bytes</i> ;                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>                                                                                                                                                          | [edit <a href="#">forwarding-options</a> port-mirroring input],<br>[edit <a href="#">forwarding-options</a> port-mirroring <a href="#">instance</a> <i>instance-name</i> input],<br>[edit <a href="#">forwarding-options</a> sampling input],<br>[edit forwarding-options sampling instance <i>instance-name</i> input] |
| <b>Release Information</b>                                                                                                                                                      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches.                                                                                                                                                                                |
| <b>Description</b>                                                                                                                                                              | Set the maximum length of the packet used for port mirroring or traffic sampling. Packets with lengths greater than the specified maximum are truncated.                                                                                                                                                                |
| <div> <b>NOTE:</b> The maximum-packet-length statement is not supported on MX80 routers.</div> |                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                                                                                                                                                  | <i>bytes</i> —Maximum length (in bytes) of the mirrored packet or the sampled packet.<br><b>Range:</b> 0 through 9192                                                                                                                                                                                                   |
| <b>Required Privilege Level</b>                                                                                                                                                 | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                 |
| <b>Related Documentation</b>                                                                                                                                                    | <ul style="list-style-type: none"><li>• <i>Configuring Port Mirroring</i></li><li>• <i>Configuring Traffic Sampling</i></li></ul>                                                                                                                                                                                       |

## mirror-once

---

|                                 |                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | mirror-once;                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit <a href="#">forwarding-options</a> port-mirroring]                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3 (MX Series routers only).<br>Support extended to M120 routers in Junos OS Release 9.5.<br>Statement introduced in Junos OS Release 12.1 for PTX Packet Transport Switches. |
| <b>Description</b>              | Configure the router to mirror packets only once. This feature is useful if you configure port mirroring on both ingress and egress interfaces, which could result in the same packet being mirrored twice.             |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Port Mirroring</i></li></ul>                                                                                                                                     |

## next-hop-group (Port Mirroring)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>next-hop-group <i>group-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit forwarding-options port-mirroring family (inet   vpls) output],<br>[edit forwarding-options port-mirroring instance <i>instance-name</i> family (inet   vpls) output]                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Specify the next-hop address for sending copies of packets to an analyzer. This configuration enables multipacket port mirroring on MX Series routers and EX Series switches without the use of a Tunnel PIC.<br><br>The commit operation fails when a next-hop group has only one interface configured. It is implicitly assumed that a subgroup is up only if more than one interface in the subgroup is up. |
| <b>Options</b>                  | <i>group-name</i> —Name of next-hop group.                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Usage Guidelines</b>         | See <i>Port Mirroring with Next-Hop Groups</i> .                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                        |

## no-filter-check

|                                 |                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>no-filter-check;</code>                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit forwarding-options port-mirroring family (inet   inet6) output]                                                                                                           |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                               |
| <b>Description</b>              | Disable filter checking on the port-mirroring interface.<br><br>This statement is required when you send port-mirrored traffic to a Tunnel PIC that has a filter applied to it. |
| <b>Usage Guidelines</b>         | See <i>Configuring Port Mirroring</i> .                                                                                                                                         |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                         |

## output (Port Mirroring)

---

|                                 |                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>output {<br/>    interface <i>interface-name</i> {<br/>        next-hop <i>address</i>;<br/>    }<br/>    no-filter-check;<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit forwarding-options port-mirroring family (inet   inet6)]                                                                              |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                           |
| <b>Description</b>              | Configure output interfaces and flow properties.<br><br>The statements are explained separately.                                            |
| <b>Usage Guidelines</b>         | See <i>Configuring Port Mirroring</i> .                                                                                                     |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                     |

## rate

---

|                                 |                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>rate <i>number</i>;</pre>                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit forwarding-options port-mirroring <b>input</b> ],<br>[edit forwarding-options sampling <b>input</b> ],<br>[edit forwarding-options sampling instance <i>instance-name</i> <b>input</b> ],<br>[edit forwarding-options port-mirroring family (inet inet6) <b>input</b> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.                                                                                                                                                                                                                             |
| <b>Description</b>              | Set the ratio of the number of packets to be sampled. For example, if you specify a rate of 10, every tenth packet (1 packet out of 10) is sampled.                                                                                                                           |
| <b>Options</b>                  | <b>number</b> —Denominator of the ratio.<br><b>Range:</b> 1 through 65,535                                                                                                                                                                                                    |
| <b>Usage Guidelines</b>         | See <i>Configuring Port Mirroring</i> or <i>Configuring Traffic Sampling</i> .                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                       |

## run-length

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>run-length <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit forwarding-options port-mirroring <a href="#">input</a> ],<br>[edit <a href="#">forwarding-options</a> port-mirroring <a href="#">instance</a> <i>port-mirroring-instance-name</i> input],<br>[edit forwarding-options port-mirroring family (inet inet6) <a href="#">input</a> ],<br>[edit forwarding-options sampling <a href="#">input</a> ],<br>[edit forwarding-options sampling instance <i>instance-name</i> <a href="#">input</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches.                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Set the number of samples following the initial trigger event. This allows you to sample packets following those already being sampled.                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <i>number</i> —Number of samples.<br><b>Range:</b> 0 through 20<br><b>Default:</b> 0                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Applying Filters to Forwarding Tables</a></li> <li>• <a href="#">Configuring Port Mirroring</a></li> <li>• <a href="#">Configuring Traffic Sampling</a></li> </ul>                                                                                                                                                                                                                           |

## Administration

- [Displaying Information on page 4263](#)
- [Operational Mode Commands for Packet Forwarding Engine Components on page 4264](#)
- [Operational Mode Commands for Layer 2 Port-Mirroring Instances on page 4445](#)
- [Operational Mode Commands for Firewall Filter Statistics and Logs on page 4447](#)
- [Operational Mode Commands for Next-Hop Groups for Layer 2 Port Mirroring on page 4459](#)

## Displaying Information

- [Displaying Layer 2 Port-Mirroring Instance Settings and Status on page 4263](#)
- [Displaying Next-Hop Group Settings and Status on page 4264](#)

### Displaying Layer 2 Port-Mirroring Instance Settings and Status

To display the current state of port-mirroring instances, use the `show forwarding-options port-mirroring <terse | detail> <instance-name>` operational command.

For more information about displaying port mirroring instance settings and status, see the *Junos OS System Basics Configuration Guide*.

- Related Documentation**
- [Layer 2 Port Mirroring Overview on page 4192](#)
  - [Layer 2 Port Mirroring Global Instance on page 4194](#)
  - [Layer 2 Port Mirroring Named Instances](#)
  - [Configuring the Global Instance of Layer 2 Port Mirroring](#)
  - [Defining a Named Instance of Layer 2 Port Mirroring](#)
  - [Disabling Layer 2 Port Mirroring Instances on page 4217](#)
  - [Examples: Layer 2 Port-Mirroring at Multiple Levels of the Chassis on page 4218](#)
  - [Example: Layer 2 Port Mirroring with Multiple Instances](#)

---

### Displaying Next-Hop Group Settings and Status

To display the current state of next-hop groups, use the **show forwarding-options next-hop-group** *<terse | brief | detail>* *<group-name>* operational command.

For more information, see the *Junos OS Operational Mode Commands*.

- Related Documentation**
- [Layer 2 Port Mirroring Overview on page 4192](#)
  - [Layer 2 Port Mirroring to Multiple Destinations Using Next-Hop Groups on page 4200](#)
  - [Defining a Layer 2 Port-Mirroring Firewall Filter](#)
  - [Defining a Next-Hop Group for Layer 2 Port Mirroring on page 4245](#)
  - [Example: Layer 2 Port Mirroring to Multiple Destinations](#)

## Operational Mode Commands for Packet Forwarding Engine Components

## show chassis fabric fpcs

|                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                        | show chassis fabric fpcs<br><fcc <i>number</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax (MX Series Routers)</b>                    | show chassis fabric fpcs                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax (EX Series Switches)</b>                   | show chassis fabric fpcs<br><all-members><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax (T4000 Core Router)</b>                    | show chassis fabric fpcs                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Syntax (PTX Series Packet Transport Switches)</b> | show chassis fabric fpcs <slot <i>fpc-slot</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>                           | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.4 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches.<br>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.<br>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                                   | (M320, MX Series, and T Series routers, EX8200 switches, and PTX Series Packet Transport Switches only) Display the state of the electrical switch fabric links between the Flexible PIC Concentrators (FPCs) and the Switch Interface Boards (SIBs).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                       | <p><b>none</b>—Display the switch fabric link state. On a TX Matrix router, display the switching fabric link states for the FPCs in all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display the switching fabric link states for the FPCs in all T1600 routers connected to the TX Matrix Plus router.</p> <p><b>all-members</b>—(EX Series switches only) (Optional) Display the switching fabric link states for the FPCs in all members of the Virtual Chassis configuration.</p> <p><b>fcc <i>number</i></b>—(TX Matrix and TX Matrix Plus router only) (Optional) On a TX Matrix router, display the switch fabric link state for the FPCs in the specified T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display the switch fabric link state for the FPCs in the specified T1600 router (or line-card chassis) that is connected to the TX Matrix Plus router. Replace <b><i>number</i></b> with a value from 0 through 3.</p> <p><b>local</b>—(EX Series switches only) (Optional) Display the switching fabric link states for the FPCs in the local Virtual Chassis member.</p> <p><b>member <i>member-id</i></b>—(EX Series switches only) (Optional) Display the switching fabric link states for the FPCs in the specified member of the Virtual Chassis configuration. Replace <b><i>member-id</i></b> with a value of 0 or 1.</p> |

**slot *fpc-slot***—(PTX Series Packet Transport Switches only) (Optional) Display the fabric state of the specified FPC slot. If no value is provided, display the status of all FPCs.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>    | <a href="#">show chassis fabric fpcs (M320 Router) on page 4267</a><br><a href="#">show chassis fabric fpcs (MX240 Router) on page 4268</a><br><a href="#">show chassis fabric fpcs (MX480 Router) on page 4268</a><br><a href="#">show chassis fabric fpcs (MX960 Router) on page 4269</a><br><a href="#">show chassis fabric fpcs (MX240 with AS MLC Modular Carrier Card) on page 4270</a><br><a href="#">show chassis fabric fpcs (MX480 with AS MLC Modular Carrier Card) on page 4271</a><br><a href="#">show chassis fabric fpcs (MX480 Router with MPC4E) on page 4272</a><br><a href="#">show chassis fabric fpcs (MX960 Router with AS MLC Modular Carrier Card on page 4273</a><br><a href="#">show chassis fabric fpcs (MX2010 Router) on page 4275</a><br><a href="#">show chassis fabric fpcs (MX2020 Router) on page 4278</a><br><a href="#">show chassis fabric fpcs (MX2020 Router with MPC4E) on page 4281</a><br><a href="#">show chassis fabric fpcs (T320 Router) on page 4282</a><br><a href="#">show chassis fabric fpcs (T640 Router) on page 4283</a><br><a href="#">show chassis fabric fpcs (TX Matrix Router) on page 4283</a><br><a href="#">show chassis fabric fpcs (T1600 Router) on page 4285</a><br><a href="#">show chassis fabric fpcs (T4000 Core Router) on page 4286</a><br><a href="#">show chassis fabric fpcs (TX Matrix Plus Router) on page 4287</a><br><a href="#">show chassis fabric fpcs lcc (TX Matrix Plus Router) on page 4295</a><br><a href="#">show chassis fabric fpcs (EX8200 Switch) on page 4296</a><br><a href="#">show chassis fabric fpcs (PTX Series Packet Transport Switches) on page 4297</a> |
| <b>Output Fields</b>            | Table 317 on page 4267 lists the output fields for the <b>show chassis fabric fpcs</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |



Table 317: show chassis fabric fpcs Output Fields

| Field Name                         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Fabric management FPC state</b> | <p>Switching fabric link (link from SIB to FPC) state for each FPC:</p> <ul style="list-style-type: none"> <li>• <b>Unused</b>—FPC is not present.<br/>(On MX240 and MX480 routers with AS- MLC modular carrier card or MPC4E only) the fabric plane from the pair that share physical links (1 and 5, and 3 and 7) is inactive.</li> <li>• <b>Destination error on PFEs <i>list of PFE numbers</i></b>—Destination errors to the listed Packet Forwarding Engines. Indicates that the link is not carrying traffic to the listed Packet Forwarding Engines.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 9.6 and later, the list of Packet Forwarding Engines with destination errors is displayed in the output.</p> <p>In Junos OS Releases before 9.6, the output only indicates that there are destination errors. However, the list of Packet Forwarding Engines with destination errors is not displayed.</p> <ul style="list-style-type: none"> <li>• <b>Links ok</b>—Link between the spare SIB and FPC is eligible to carry traffic.</li> <li>• <b>Link error</b>—Link between the SIB and FPC has CRC errors. However, the link is still eligible to carry traffic.</li> <li>• <b>Plane disabled</b>—Fabric plane has been disabled for the following reasons: <ul style="list-style-type: none"> <li>• Destination errors have exceeded the thresholds.</li> <li>• Run-time link errors have exceeded the thresholds.</li> <li>• Initialization time link errors detected, and link training was unsuccessful.</li> <li>• <b>Plane Disabled, Links Error</b> (PTX Series Packet Transport Switches only)—The plane is disabled because of link errors detected at the FPC RX.</li> </ul> </li> <li>• <b>Plane Disabled, Links Down</b> (PTX Series Packet Transport Switches only)—The plane is disabled because of link errors detected at the SIB RX.</li> <li>• <b>Plane enabled</b>—Link between the active SIB and FPC is eligible to carry traffic.</li> </ul> <p><b>NOTE:</b> On the Enhanced MX SCB with MPC, a maximum of 4 planes are operational and running. On all the other SCBs with MPC, all the planes are operational and running.</p> <ul style="list-style-type: none"> <li>• <b>Plane Enabled, Links OK</b> (PTX Series Packet Transport Switches only)—The FPC CCL RX link is eligible to carry traffic.</li> </ul> |

## Sample Output

### show chassis fabric fpcs (M320 Router)

```

user@host> show chassis fabric fpcs
Fabric management FPC state:
FPC #2
 PFE #1
 SIB #0
 Plane enabled

```

```
SIB #1
 Plane enabled
SIB #2
 Plane enabled
SIB #3
 Plane enabled
```

#### show chassis fabric fpcs (MX240 Router)

```
user@host> show chassis fabric fpcs
Fabric management FPC state:
FPC 2
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Links ok
 Plane 5: Links ok
 Plane 6: Links ok
 Plane 7: Links ok
 PFE #1
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Links ok
 Plane 5: Links ok
 Plane 6: Links ok
 Plane 7: Links ok
 PFE #2
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Links ok
 Plane 5: Links ok
 Plane 6: Links ok
 Plane 7: Links ok
 PFE #3
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Links ok
 Plane 5: Links ok
 Plane 6: Links ok
 Plane 7: Links ok
```

#### show chassis fabric fpcs (MX480 Router)

```
user@host> show chassis fabric fpcs
FPC 0
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Links ok
 Plane 5: Links ok
```

```

 Plane 6: Links ok
 Plane 7: Links ok
PFE #1
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Links ok
 Plane 5: Links ok
 Plane 6: Links ok
 Plane 7: Links ok
PFE #2
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Links ok
 Plane 5: Links ok
 Plane 6: Links ok
 Plane 7: Links ok
PFE #3
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Links ok
 Plane 5: Links ok
 Plane 6: Links ok
 Plane 7: Links ok
FPC 1
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled
 PFE #1
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled

```

#### show chassis fabric fpcs (MX960 Router)

```

user@host> show chassis fabric fpcs
FPC 0
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Links ok
 Plane 5: Links ok
 PFE #1

```

```
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
FPC 1
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
FPC 2
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
...
```

#### show chassis fabric fpcs (MX240 with AS MLC Modular Carrier Card)

In the following output, FPC 1 is the AS MLC modular carrier card (AS MCC).

```

user@host>show chassis fabric fpcs
FPC 1
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Plane enabled
 Plane 5: Unused
 Plane 6: Plane enabled
 Plane 7: Unused
FPC 2
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled

```

#### show chassis fabric fpcs (MX480 with AS MLC Modular Carrier Card)

In the following output, FPC 5 is the AS MLC modular carrier card (AS MCC).

```

user@host>show chassis fabric fpcs
FPC 2
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled
FPC 4
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Links ok
 Plane 5: Links ok
 Plane 6: Links ok
 Plane 7: Links ok
 PFE #2
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Links ok
 Plane 5: Links ok
 Plane 6: Links ok
 Plane 7: Links ok
FPC 5
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled

```

```
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Unused
Plane 6: Plane enabled
Plane 7: Unused
```

### show chassis fabric fpcs (MX480 Router with MPC4E)

In the following output, FPC 4 is the MPC4E (MPC4E-3D-32XGE-SFPP) card.

```
user@host>show chassis fabric fpcs
```

```
Fabric management FPC state:
```

```
FPC 0
```

```
PFE #0
```

```
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Links ok
Plane 6: Plane enabled
Plane 7: Links ok
```

```
PFE #1
```

```
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Links ok
Plane 6: Plane enabled
Plane 7: Links ok
```

```
FPC 1
```

```
PFE #0
```

```
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Links ok
Plane 6: Plane enabled
Plane 7: Links ok
```

```
PFE #1
```

```
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Links ok
Plane 6: Plane enabled
Plane 7: Links ok
```

```
PFE #2
```

```
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Links ok
Plane 6: Plane enabled
Plane 7: Links ok
```

```
PFE #3
```

```
Plane 0: Links ok
```

```

Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Links ok
Plane 6: Plane enabled
FPC 3
PFE #0
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Links ok
Plane 6: Plane enabled
Plane 7: Links ok
FPC 4
PFE #0
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Unused
Plane 6: Plane enabled
Plane 7: Unused
PFE #1
Plane 0: Links ok
Plane 1: Links ok
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Unused
Plane 6: Plane enabled
Plane 7: Unused

```

### show chassis fabric fpcs (MX960 Router with AS MLC Modular Carrier Card)

In the following output, FPC 5 is the AS MLC modular carrier card (AS MCC).

```

user@host>show chassis fabric fpcs
Fabric management FPC state:
FPC 0
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
FPC 1
PFE #0
Plane 0: Plane enabled

```

```
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
FPC 4
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
FPC 5
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
FPC 8
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
```



```

Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Links ok
Plane 5: Links ok

```

### show chassis fabric fpcs (MX2010 Router)

```

user@host> show chassis fabric fpcs
Fabric management FPC state:
FPC 0
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane disabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled
 PFE #1
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane disabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled
FPC 1
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane disabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled
FPC 2
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane disabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled
 PFE #1
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane disabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled

```

```

 Plane 6: Plane enabled
 Plane 7: Plane enabled
FPC 3
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane disabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled
 PFE #1
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane disabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled
 PFE #2
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane disabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled
 PFE #3
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane disabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled
FPC 4
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane disabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled
FPC 5
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane disabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled
 PFE #1
 Plane 0: Plane enabled

```

```
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 6
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 7
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
```

```
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 8
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 9
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane disabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
```

#### show chassis fabric fpcs (MX2020 Router)

```
user@host> show chassis fabric fpcs
Fabric management FPC state:
FPC 0
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
```

```
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 1
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 2
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
```

```
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 3
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
```

```

 Plane 7: Plane enabled
FPC 4
...

```

### show chassis fabric fpcs (MX2020 Router with MPC4E)

In the following output, FPC 14 is the MPC4E (MPC4E-3D-2CGE-8XGE) card.

```

user@host>show chassis fabric fpcs
Fabric management FPC state:
FPC 0
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled
 PFE #1
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled
FPC 9
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled
 PFE #1
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled
FPC 10
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Plane enabled
 Plane 5: Plane enabled
 Plane 6: Plane enabled
 Plane 7: Plane enabled
FPC 14
 PFE #0
 Plane 0: Plane enabled

```

```
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
FPC 19
PFE #0
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #1
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #2
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
PFE #3
Plane 0: Plane enabled
Plane 1: Plane enabled
Plane 2: Plane enabled
Plane 3: Plane enabled
Plane 4: Plane enabled
Plane 5: Plane enabled
Plane 6: Plane enabled
Plane 7: Plane enabled
```

#### show chassis fabric fpcs (T320 Router)

```
user@host> show chassis fabric fpcs
FPC #3
PFE #1
SIB #0
Links ok
```



```

SIB #1
 Plane enabled
SIB #2
 Plane enabled
FPC #5
 PFE #1
 SIB #0
 Links ok
 SIB #1
 Plane enabled
 SIB #2
 Plane enabled
FPC #7
 PFE #1
 SIB #0
 Links ok
 SIB #1
 Plane enabled
 SIB #2
 Plane enabled

```

#### show chassis fabric fpcs (T640 Router)

```

user@host> show chassis fabric fpcs
Fabric management FPC state:

```

```

FPC #2
 PFE #1
 SIB #0
 Links ok
 SIB #1
 Plane enabled
 SIB #2
 Plane enabled
 SIB #3
 Plane enabled
 SIB #4
 Plane enabled
FPC #3
 PFE #1
 SIB #2
 Plane enabled
 SIB #3
 Link error
 Destination error on PFEs
 8 9 10 11 12 13 14 15 16 17 18 19 20 21
 SIB #4
 Destination error on PFEs
 8 9 10 11 12 13 14 15 16 17 18 19 20 21
...

```

#### show chassis fabric fpcs (TX Matrix Router)

```

user@host> show chassis fabric fpcs
lcc0-re0:

Fabric management FPC state:
FPC #0
 PFE #1
 SIB #0
 Links ok

```

```

SIB #2
 Links ok
SIB #3
 Links ok
SIB #4
 Links ok
FPC #2
PFE #1
 SIB #0
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
FPC #3
PFE #1
 SIB #2
 Plane enabled
 SIB #3
 Link error
 Destination error on PFes
 0 1 2 3 4 5 6 7
 8 9 10 11 12 13 14 15 16 17 18 19 20 21
 SIB #4
 Destination error on PFes
 0 1 2 3 4 5 6 7
 8 9 10 11 12 13 14 15 16 17 18 19 20 21
...
FPC #4
PFE #0
 SIB #4 Links ok
PFE #1
 SIB #4 Links ok
FPC #5
PFE #1
 SIB #4 Links ok
FPC #6
PFE #1
 SIB #4 Links ok

lcc2-re0:

Fabric management FPC state:
FPC #0
PFE #1
 SIB #4 Links ok
FPC #1
PFE #1
 SIB #4 Links ok
FPC #2
PFE #0
 SIB #4 Links ok
PFE #1
 SIB #4 Links ok
FPC #4
PFE #0
 SIB #4 Links ok
PFE #1
 SIB #4 Links ok
FPC #5

```

```
PFE #1
SIB #4 Links ok
```

### show chassis fabric fpcs (T1600 Router)

```
user@host> show chassis fabric fpcs
Fabric management FPC state:
FPC #0
 PFE #0
 SIB #0 Links ok
 SIB #1 Plane enabled
 SIB #2 Plane enabled
 SIB #3 Plane enabled
 SIB #4 Plane enabled
 PFE #1
 SIB #0 Links ok
 SIB #1 Plane enabled
 SIB #2 Plane enabled
 SIB #3 Plane enabled
 SIB #4 Plane enabled
FPC #1
 PFE #0
 SIB #0 Links ok
 SIB #1 Plane enabled
 SIB #2 Plane enabled
 SIB #3 Plane enabled
 SIB #4 Plane enabled
 PFE #1
 SIB #0 Links ok
 SIB #1 Plane enabled
 SIB #2 Plane enabled
 SIB #3 Plane enabled
 SIB #4 Plane enabled
FPC #2
 PFE #0
 SIB #0 Links ok
 SIB #1 Plane enabled
 SIB #2 Plane enabled
```

```

SIB #3
 Plane enabled
SIB #4
 Plane enabled
FPC #4
PFE #0
 SIB #0
 Links ok
 SIB #1
 Plane enabled
 SIB #2
 Plane enabled
 SIB #3
 Plane enabled
 SIB #4
 Plane enabled
PFE #1
 SIB #0
 Links ok
 SIB #1
 Plane enabled
 SIB #2
 Plane enabled
 SIB #3
 Plane enabled
 SIB #4
 Plane enabled
FPC #3
PFE #1
 SIB #2
 Plane enabled
 SIB #3
 Link error
 Destination error on PFes
 8 9 10 11 12 13 14 15 16 17 18 19 20 21
 0 1 2 3 4 5 6 7
 SIB #4
 Destination error on PFes
 8 9 10 11 12 13 14 15 16 17 18 19 20 21
 0 1 2 3 4 5 6 7

```

#### show chassis fabric fpcs (T4000 Core Router)

```

Fabric management FPC state:
FPC #2
PFE #0
 SIB #0
 Links ok
 SIB #1
 Plane enabled
 SIB #2
 Plane enabled
 SIB #3
 Plane enabled
 SIB #4
 Plane enabled
FPC #3
PFE #0
 SIB #0
 Links ok
 SIB #1
 Plane enabled
 SIB #2

```

```

 Plane enabled
 SIB #3 Plane enabled
 SIB #4 Plane enabled
FPC #5
 PFE #0
 SIB #0 Links ok
 SIB #1 Plane enabled
 SIB #2 Plane enabled
 SIB #3 Plane enabled
 SIB #4 Plane enabled
 PFE #1
 SIB #0 Links ok
 SIB #1 Plane enabled
 SIB #2 Plane enabled
 SIB #3 Plane enabled
 SIB #4 Plane enabled
FPC #6
 PFE #0
 SIB #0 Links ok
 SIB #1 Plane enabled
 SIB #2 Plane enabled
 SIB #3 Plane enabled
 SIB #4 Plane enabled
 PFE #1
 SIB #0 Links ok
 SIB #1 Plane enabled
 SIB #2 Plane enabled
 SIB #3 Plane enabled
 SIB #4 Plane enabled

```

#### show chassis fabric fpcs (TX Matrix Plus Router)

```

user@host> show chassis fabric fpcs
lcc0-re0:

Fabric management FPC state:
FPC #0
 PFE #1
 SIB #0

```

```

 Unused
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
FPC #2
PFE #0
 SIB #0
 Unused
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
PFE #1
 SIB #0
 Unused
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
FPC #3
PFE #1
 SIB #2
 Plane enabled
 SIB #3
 Link error
 Destination error on PFes 0 1 2 3 4 5 6 7
 8 9 10 11 12 13 14 15 16 17 18 19 20 21
 SIB #4
 Destination error on PFes 0 1 2 3 4 5 6 7
 8 9 10 11 12 13 14 15 16 17 18 19 20 21
FPC #4
PFE #0
 SIB #0
 Unused
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
PFE #1
 SIB #0
 Unused
 SIB #1
 Links ok
 SIB #2
```

```

 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
FPC #6
 PFE #0
 SIB #0
 Unused
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
 PFE #1
 SIB #0
 Unused
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
FPC #7
 PFE #0
 SIB #0
 Unused
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok

```

```
lcc1-re0:
```

```

Fabric management FPC state:
```

```

FPC #2
 PFE #0
 SIB #0
 Links ok
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
 PFE #1
 SIB #0
 Links ok
 SIB #1
 Links ok

```

```
SIB #2
 Links ok
SIB #3
 Links ok
SIB #4
 Links ok
FPC #4
 PFE #0
 SIB #0
 Links ok
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
 PFE #1
 SIB #0
 Links ok
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Destination error on PFES 1 8 9 29 40 65 72 73
 93 104
 SIB #4
 Links ok
FPC #6
 PFE #0
 SIB #0
 Links ok
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
 PFE #1
 SIB #0
 Links ok
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
FPC #7
 PFE #0
 SIB #0
 Links ok
 SIB #1
 Links ok
 SIB #2
```



```

SIB #3 Links ok
SIB #4 Links ok
SIB #4 Links ok
```

```
lcc2-re0:
```

```

Fabric management FPC state:
```

```
FPC #0
```

```
PFE #0
```

```

SIB #0 Links ok
SIB #1 Links ok
SIB #2 Links ok
SIB #3 Links ok
SIB #4 Links ok
```

```
PFE #1
```

```

SIB #0 Links ok
SIB #1 Links ok
SIB #2 Links ok
SIB #3 Links ok
SIB #4 Links ok
```

```
FPC #2
```

```
PFE #0
```

```

SIB #0 Links ok
SIB #1 Links ok
SIB #2 Links ok
SIB #3 Links ok
SIB #4 Links ok
```

```
PFE #1
```

```

SIB #0 Links ok
SIB #1 Links ok
SIB #2 Links ok
SIB #3 Links ok
SIB #4 Links ok
```

```
FPC #4
```

```
PFE #0
```

```

SIB #0 Links ok
SIB #1 Links ok
```

```
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
FPC #5
PFE #0
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
PFE #1
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
FPC #6
PFE #0
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
PFE #1
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
Links ok
SIB #4
Links ok
FPC #7
PFE #0
SIB #0
Links ok
SIB #1
Links ok
SIB #2
Links ok
SIB #3
```

```

 Links ok
 SIB #4
 Links ok

lcc3-re0:

Fabric management FPC state:
FPC #0
 PFE #0
 SIB #0
 Links ok
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
 PFE #1
 SIB #0
 Links ok
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
FPC #2
 PFE #0
 SIB #0
 Links ok
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
 PFE #1
 SIB #0
 Links ok
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
FPC #4
 PFE #0
 SIB #0
 Links ok
 SIB #1
 Links ok
 SIB #2
 Links ok
```

```
SIB #3
 Links ok
SIB #4
 Links ok
PFE #1
 SIB #0
 Links ok
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
FPC #5
 PFE #0
 SIB #0
 Links ok
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
 PFE #1
 SIB #0
 Links ok
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
FPC #6
 PFE #0
 SIB #0
 Links ok
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
 PFE #1
 SIB #0
 Links ok
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok
```

```

FPC #7
 PFE #0
 SIB #0
 Links ok
 SIB #1
 Links ok
 SIB #2
 Links ok
 SIB #3
 Links ok
 SIB #4
 Links ok

```

### show chassis fabric fpcs lcc (TX Matrix Plus Router)

```

user@host> show chassis fabric fpcs lcc 0
lcc0-re1:

```

```

Fabric management FPC state:

```

```

FPC #3
 PFE #1
 SIB #2
 Plane enabled
 SIB #3
 Link error
 Destination error on PFES
 8 9 10 11 12 13 14 15 16 17 18 19 20 21
 SIB #4
 Destination error on PFES
 8 9 10 11 12 13 14 15 16 17 18 19 20 21
FPC #4
 PFE #0
 SIB #0 Links ok
 SIB #1 Links ok
 SIB #2 Links ok
 SIB #3 Links ok
 SIB #4 Links ok
 PFE #1
 SIB #0 Links ok
 SIB #1 Links ok
 SIB #2 Links ok
 SIB #3 Links ok
 SIB #4 Links ok
FPC #6
 PFE #0
 SIB #0 Links ok
 SIB #1 Links ok
 SIB #2 Links ok
 SIB #3 Links ok
 SIB #4 Links ok
 PFE #1
 SIB #0 Links ok
 SIB #1 Links ok
 SIB #2 Links ok
 SIB #3 Links ok
 SIB #4 Links ok
FPC #7
 PFE #0
 SIB #0 Links ok
 SIB #1 Links ok
 SIB #2 Links ok

```

```
SIB #3 Links ok
SIB #4 Links ok
```

### show chassis fabric fpcs (EX8200 Switch)

```
user@host> show chassis fabric fpcs
Fabric management FPC state
FPC 6
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Links ok
 Plane 5: Links ok
 Plane 6: Links ok
 Plane 7: Links ok
 Plane 8: Plane enabled
 Plane 9: Plane enabled
 Plane 10: Plane enabled
 Plane 11: Plane enabled
 PFE #1
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Links ok
 Plane 5: Links ok
 Plane 6: Links ok
 Plane 7: Links ok
 Plane 8: Plane enabled
 Plane 9: Plane enabled
 Plane 10: Plane enabled
 Plane 11: Plane enabled
FPC 7
 PFE #0
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Links ok
 Plane 5: Links ok
 Plane 6: Links ok
 Plane 7: Links ok
 Plane 8: Plane enabled
 Plane 9: Plane enabled
 Plane 10: Plane enabled
 Plane 11: Plane enabled
 PFE #1
 Plane 0: Plane enabled
 Plane 1: Plane enabled
 Plane 2: Plane enabled
 Plane 3: Plane enabled
 Plane 4: Links ok
 Plane 5: Links ok
 Plane 6: Links ok
 Plane 7: Links ok
 Plane 8: Plane enabled
 Plane 9: Plane enabled
```

Plane 10: Plane enabled  
Plane 11: Plane enabled

### show chassis fabric fpcs (PTX Series Packet Transport Switches)

```
user@host> show chassis fabric fpcs slot 0
Fabric management FPC state:
FPC #0
 PFE #0
 SIB0_Fcore0 (plane 0) Plane Enabled, Links OK
 SIB0_Fcore1 (plane 1) Plane Enabled, Links OK
 SIB1_Fcore0 (plane 2) Plane Disabled, Links Down
 SIB1_Fcore1 (plane 3) Plane Enabled, Links OK
 SIB2_Fcore0 (plane 4) Plane Enabled, Links OK
 SIB2_Fcore1 (plane 5) Plane Enabled, Links OK
 SIB3_Fcore0 (plane 6) Plane Enabled, Links OK
 SIB3_Fcore1 (plane 7) Plane Enabled, Links OK
 SIB5_Fcore0 (plane 10) Plane Enabled, Links OK
 SIB5_Fcore1 (plane 11) Plane Enabled, Links OK
 SIB6_Fcore0 (plane 12) Plane Enabled, Links OK
 SIB6_Fcore1 (plane 13) Plane Enabled, Links OK
 SIB7_Fcore0 (plane 14) Plane Enabled, Links OK
 SIB7_Fcore1 (plane 15) Plane Enabled, Links OK
 SIB8_Fcore0 (plane 16) Plane Enabled, Links OK
 SIB8_Fcore1 (plane 17) Plane Enabled, Links OK
```

## show chassis fpc

---

|                                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                               | show chassis fpc<br><detail <slot>>   <pic-status <slot>>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax (EX Series Switches)</b>                          | show chassis fpc<br><detail <fpc-slot>>   <pic-status <fpc-slot>><br><fpc-slot>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax (T4000 Routers)</b>                               | show chassis fpc<br><detail <fpc-slot>><br><pic-status <fpc-slot>>                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax (TX Matrix and TX Matrix Plus Routers)</b>        | show chassis fpc<br><detail <fpc-slot>>   <pic-status <fpc-slot>><br><slot>                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Syntax (MX Series Routers and EX Series switches)</b>    | show chassis fpc<br><detail <slot>>   <pic-status <slot>><br><all-members><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax (MX2010 and MX2020 3D Universal Edge Routers)</b> | show chassis fpc<br><slot> detail   <detail <slot>>   <pic-status <slot>><br><fpc-slot>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax (QFX Series)</b>                                  | show chassis fpc<br><detail><br><interconnect-device <i>name</i> <fpc-slot fpc-slot>><br><node-device <i>name</i> >                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Syntax (PTX Series Packet Transport Switches)</b>        | show chassis fpc<br><detail <fpc-slot>>   <pic-status <fpc-slot>><br><fpc-slot>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax (ACX Series Universal Access Routers)</b>         | show chassis fpc<br><detail <fpc-slot>>   <pic-status <fpc-slot>><br><fpc-slot>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>                                  | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for QFX Series.<br>Command introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches.<br>Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.<br>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.<br>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers. |
| <b>Description</b>                                          | Display status information about the installed Flexible PIC Concentrators (FPCs) and PICs.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |



**Options** **none**—Display status information for all FPCs. On a TX Matrix router, display status information for all FPCs on the attached T640 routers in the routing matrix. On a TX Matrix Plus router, display status information for all FPCs on the attached T1600 routers in the routing matrix.



**NOTE:** In EX8200 switches, line cards initialize Packet Forwarding Engine during startup. If an error occurs during hardware initialization, the FPCs with bad hardware parts power down after transferring the debug information to the Routing Engine. The Routing Engine marks the FPC offline, logs the error in system log messages (/var/log/messages), and generates an alarm to inform the user.

See the following sample output:

```
user@host> show chassis fpc
```

| Utilization (%) |         | Temp                 | CPU Utilization (%) |           | Memory    |      |
|-----------------|---------|----------------------|---------------------|-----------|-----------|------|
| Slot            | State   | (C)                  | Total               | Interrupt | DRAM (MB) | Heap |
| Buffer          |         |                      |                     |           |           |      |
| 0               | Empty   |                      |                     |           |           |      |
| 1               | Empty   |                      |                     |           |           |      |
| 2               | Empty   |                      |                     |           |           |      |
| 3               | Empty   |                      |                     |           |           |      |
| 4               | Empty   |                      |                     |           |           |      |
| 5               | Offline | ---Hard FPC error--- |                     |           |           |      |
| 6               | Empty   |                      |                     |           |           |      |
| 7               | Online  | 26                   | 4                   | 0         | 1024      | 0    |
| 32              |         |                      |                     |           |           |      |

The following sample output shows the alarm raised for the failed FPCs.

```
user@host > show chassis alarms
4 alarms currently active
```

| Alarm time              | Class | Description                          |
|-------------------------|-------|--------------------------------------|
| 2011-03-24 00:52:51 UTC | Major | FPC 5 Hard errors                    |
| 2011-03-24 00:52:31 UTC | Major | Fan Tray Failure                     |
| 2011-03-24 00:52:31 UTC | Major | Fan Tray Failure                     |
| 2011-03-24 00:51:26 UTC | Minor | Loss of communication with Backup RE |



**NOTE:** On T4000 routers, when you include the **enhanced-mode** statement at the **[edit chassis network-services]** hierarchy level and reboot the system, only the T4000 Type 5 FPCs present on the router become online while the remaining FPCs are offline, and FPC misconfiguration alarms are generated. The **show chassis alarm** command output displays FPC misconfiguration (**FPC *fpc-slot* misconfig**) as the reason for the generation the alarms.

The following sample output shows the FPC status after the **enhanced-mode** statement is configured on the T4000 router. The T4000 Type 5 FPC present in slot 5 becomes online while the remaining FPCs are offline.

```
user@host> show chassis fpc
```

|                 | Temp | CPU Utilization (%)        |           | Memory    |      |
|-----------------|------|----------------------------|-----------|-----------|------|
| Utilization (%) |      |                            |           |           |      |
| Slot State      | (C)  | Total                      | Interrupt | DRAM (MB) | Heap |
| Buffer          |      |                            |           |           |      |
| 0 offline       |      | ---FPC misconfiguration--- |           |           |      |
| 1 offline       |      | ---FPC misconfiguration--- |           |           |      |
| 2 offline       |      | ---FPC misconfiguration--- |           |           |      |
| 3 Empty         |      |                            |           |           |      |
| 4 Empty         |      |                            |           |           |      |
| 5 Online        | 66   | 50                         | 0         | 2816      | 29   |
| 27              |      |                            |           |           |      |

The following sample output shows FPC misconfiguration alarms.

```
user@host > show chassis alarms
```

3 alarms currently active

| Alarm time              | Class | Description     |
|-------------------------|-------|-----------------|
| 2011-03-24 00:52:51 PST | Major | FPC 1 misconfig |
| 2011-03-24 00:52:31 PST | Major | FPC 2 misconfig |
| 2011-03-24 00:52:31 PST | Major | FPC 3 misconfig |

**detail**—(Optional) Display detailed status information for all FPCs or for the FPC in the specified slot (see ***fpc-slot*** or ***slot***).

**all-members**—(MX Series routers and EX Series switches only) (Optional) Display status information for all FPCs on all members of the Virtual Chassis configuration.

**interconnect-device *name***—(QFabric systems only) (Optional) Display status information for all FPCs on the Interconnect device.

***fpc-slot***—(Optional) FPC slot number:

- (TX Matrix and TX Matrix Plus router only)—On a TX Matrix router, if you specify the number of the T640 router (or line-card chassis) by using the ***lcc number*** option (the recommended method), replace ***fpc-slot*** with a value from 0 through 7. Otherwise, replace ***fpc-slot*** with a value from 0 through 31. Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 router (or line-card chassis)

by using the **lcc number** option (the recommended method), replace **fpc-slot** with a value from 0 through 7. Otherwise, replace **fpc-slot** with a value from 0 through 31. For example, the following commands have the same result:

```
user@host> show chassis fpc detail 1 lcc 1
user@host> show chassis fpc detail 9
```

- M120 router—Replace **fpc-slot** with a value from 0 through 5.
- MX80 router—Replace **fpc-slot** with a value from 0 through 1.
- MX240 router—Replace **fpc-slot** with a value from 0 through 2.
- MX480 router—Replace **fpc-slot** with a value from 0 through 5.
- MX-960 router—Replace **fpc-slot** with a value from 0 through 11.
- MX2010 router—Replace **fpc-slot-number** with a value from 0 through 9.
- MX2020 router—Replace **fpc-slot-number** with a value from 0 through 19.
- Other routers—Replace **fpc-slot** with a value from 0 through 7.
- EX Series switches:
  - EX3200 switches and EX4200 standalone switches—Replace **fpc-slot** with 0.
  - EX4200 switches in a Virtual Chassis configuration—Replace **fpc-slot** with a value from 0 through 9.
  - EX6210 switches—Replace **fpc-slot** with a value from 0 through 9.
  - EX8208 switches—Replace **fpc-slot** with a value from 0 through 7.
  - EX8216 switches—Replace **fpc-slot** with a value from 0 through 15.
- QFX Series:
  - QFX3500 switches—Replace **fpc-slot** with 0.
  - QFabric systems—Replace **fpc-slot** with 0 through 31 on the Interconnect device.
- PTX Series Packet Transport Switches:
  - PTX5000 Packet Transport Switch—Replace **fpc-slot** with a value from 0 through 7.
- ACX Series Universal Access Routers:
  - ACX1000 and ACX2000 Universal Access Routers—Replace **fpc-slot** with 0.

**local**—(MX Series routers and EX Series switches only) (Optional) Display status information for all FPCs on the local Virtual Chassis member.

**member member-id**—(MX Series routers and EX Series switches only) (Optional) Display status information for all FPCs on the specified member of the Virtual Chassis configuration. Replace **member-id** with a value of 0 or 1.

**node-device name**—(QFabric systems only) (Optional) Display status information for each Node device. Each Node device is equivalent to an FPC.

**pic-status**—(Optional) Display status information for all PICs or for the PIC in the specified slot (see *fpc-slot*).



**NOTE:** On T1600 routers, Type 4 FPCs with ASICs based on the SL2.0 chipset do not support the 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (10x10GE [LAN/WAN] SFPP). If you issue the `show chassis fpc` command with the `pic-status` option, the CLI displays the string “Not Supported” for 10x10GE(LAN/WAN) SFPP PICs installed on such FPCs. The following is a sample output:

```
user@host> show chassis fpc pic-status
Slot 0 Online E2-FPC Type 1
 PIC 0 Online 1x G/E SFP, 1000 BASE
 PIC 1 Online Adaptive Services-II
 PIC 2 Online 1x G/E IQ, 1000 BASE
 PIC 3 Online 1x G/E IQ, 1000 BASE
Slot 1 Online FPC Type 3-ES
 PIC 0 Present UNUSED- Not Supported
Slot 2 Online FPC Type 4-ES
 PIC 0 Offline 4x OC-192 SONET XFP
 PIC 1 Present 10x10GE(LAN/WAN) SFPP- Not Supported
<<<<<<
Slot 4 Offline FPC Type 1-ES
Slot 5 Offline FPC Type 2-ES
Slot 6 Online E2-FPC Type 3
 PIC 0 Online 1x OC-192 SONET XFP
 PIC 1 Online 4x OC-48 SONET
 PIC 2 Online 4x OC-48 SONET
 PIC 3 Online MultiServices 500
Slot 7 Online FPC Type 4-ES
 PIC 0 Online 4x 10GE (LAN/WAN) XFP
 PIC 1 Online 4x 10GE (LAN/WAN) XFP
```

In addition, an entry is logged in the system log messages (`/var/log/messages`) that the PIC is not supported. The following is a sample message logged in the system log:

```
Apr 5 08:47:36 router1 chassisd[2770]: CHASSISD_UNSUPPORTED_PIC:
PIC 1 in FPC 2 (type 763, version 257) is not supported
```

**lcc number**—(TX Matrix and TX Matrix Plus router only) (Optional) On a TX Matrix router, display status information for a T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display status information for a T1600 router (or line-card chassis) that is connected to the TX Matrix Plus router. Replace *number* with a value from 0 through 3.

**Required Privilege Level** view

**Related Documentation**

- *request chassis fpc*
- *show chassis fpc-feb-connectivity*
- [show chassis fabric fpcs on page 4265](#)

- *Configuring the Junos OS to Resynchronize FPC Sequence Numbers with Active FPCs when an FPC Comes Online*
- *MX960 Flexible PIC Concentrator Description*
- *ACX2000 and ACX2100 Routers Hardware and CLI Terminology Mapping*
- *enhanced-mode*

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Sample Output</b> | <a href="#">show chassis fpc (EX6210 Switch) on page 4306</a><br><a href="#">show chassis fpc (M10 Router) on page 4306</a><br><a href="#">show chassis fpc (M20 Router) on page 4306</a><br><a href="#">show chassis fpc detail (M Series Routers) on page 4306</a><br><a href="#">show chassis fpc detail (MX80 Router) on page 4307</a><br><a href="#">show chassis fpc (MX240 Router) on page 4307</a><br><a href="#">show chassis fpc (EX Series Switch) on page 4307</a><br><a href="#">show chassis fpc (MX480 Router) on page 4307</a><br><a href="#">show chassis fpc (MX480 Router with 100-Gigabit Ethernet CFP) on page 4307</a><br><a href="#">show chassis fpc pic-status (MX480 Router with 100-Gigabit Ethernet CFP) on page 4308</a><br><a href="#">show chassis fpc pic-status (EX Series Switch) on page 4308</a><br><a href="#">show chassis fpc (MX480 Router with MPC4E) on page 4308</a><br><a href="#">show chassis fpc detail (MX480 Router with MPC4E) on page 4308</a><br><a href="#">show chassis fpc pic-status (MX480 Router with MPC4E) on page 4309</a><br><a href="#">show chassis fpc (MX960 Router) on page 4309</a><br><a href="#">show chassis fpc (MX240, MX480, MX960 Routers with Application Services Modular Line Card) on page 4309</a><br><a href="#">show chassis fpc (MX240, MX480, MX960 with Application Services Modular Line Card) on page 4310</a><br><a href="#">show chassis fpc (MX2010 Router) on page 4310</a><br><a href="#">show chassis fpc (MX2020 Router) on page 4310</a><br><a href="#">show chassis fpc (MX2020 Router with MPC4E) on page 4310</a><br><a href="#">show chassis fpc detail (MX2020 Router with MPC4E) on page 4311</a><br><a href="#">show chassis fpc pic-status (MX2020 Router with MPC4E) on page 4312</a><br><a href="#">show chassis fpc detail (MX Series Routers) on page 4312</a><br><a href="#">show chassis fpc detail (EX Series Switches) on page 4312</a><br><a href="#">show chassis fpc (Hardware Not Supported) on page 4313</a><br><a href="#">show chassis fpc detail (Hardware Not Supported) on page 4313</a><br><a href="#">show chassis fpc pic-status on page 4313</a><br><a href="#">show chassis fpc pic-status (M Series Routers) on page 4313</a><br><a href="#">show chassis fpc pic-status (M120 Router) on page 4314</a><br><a href="#">show chassis fpc pic-status (MX240, MX480, and MX960 Routers with Application Services Modular Line Card) on page 4314</a><br><a href="#">show chassis fpc lcc (TX Matrix Router) on page 4314</a><br><a href="#">show chassis fpc pic-status (TX Matrix Router) on page 4315</a><br><a href="#">show chassis fpc pic-status lcc (TX Matrix Router) on page 4315</a><br><a href="#">show chassis fpc (TX Matrix Plus Router) on page 4315</a><br><a href="#">show chassis fpc lcc (TX Matrix Plus Router) on page 4316</a><br><a href="#">show chassis fpc detail (TX Matrix Plus Router) on page 4316</a><br><a href="#">show chassis fpc pic-status (TX Matrix Plus Router) on page 4319</a><br><a href="#">show chassis fpc (T1600 Router) on page 4319</a> |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

[show chassis fpc detail \(T1600 Router\) on page 4320](#)  
[show chassis fpc <fpc-slot> \(EX Series Switch\) on page 4320](#)  
[show chassis fpc slot \(T1600 Router\) on page 4320](#)  
[show chassis fpc pic-status \(T1600 Router\) on page 4321](#)  
[show chassis fpc \(T4000 Router\) on page 4321](#)  
[show chassis fpc detail \(T4000 Router\) on page 4321](#)  
[show chassis fpc pic-status \(T4000 Router\) on page 4322](#)  
[show chassis fpc \(QFX Series\) on page 4322](#)  
[show chassis fpc detail \(QFX3500 Switches\) on page 4322](#)  
[show chassis fpc pic-status \(QFX3500 Switches\) on page 4322](#)  
[show chassis fpc interconnect-device \(QFabric System\) on page 4322](#)  
[show chassis fpc interconnect-device \(QFabric System\) on page 4323](#)  
[show chassis fpc interconnect-device detail \(QFabric System\) on page 4323](#)  
[show chassis fpc pic-status interconnect-device \(QFabric System\) on page 4323](#)  
[show chassis fpc pic-status node-device \(QFabric System\) on page 4324](#)  
[show chassis fpc \(PTX5000 Packet Transport Switch\) on page 4324](#)  
[show chassis fpc detail \(PTX5000 Packet Transport Switch\) on page 4324](#)  
[show chassis fpc pic-status \(PTX5000 Packet Transport Switch\) on page 4325](#)  
[show chassis fpc \(ACX2000 Universal Access Router\) on page 4325](#)  
[show chassis fpc 0 \(ACX2000 Universal Access Router\) on page 4325](#)  
[show chassis fpc detail \(ACX2000 Universal Access Router\) on page 4326](#)  
[show chassis fpc pic-status \(ACX2000 Universal Access Router\) on page 4326](#)  
[show chassis fpc 1 \(MX Routers with Media Services Blade \[MSB\]\) on page 4326](#)  
[show chassis fpc 1 detail \(MX Routers with Media Services Blade \[MSB\]\) on page 4326](#)

**Output Fields** [Table 318 on page 4304](#) lists the output fields for the **show chassis fpc** command. Output fields are listed in the approximate order in which they appear.

**Table 318: show chassis fpc Output Fields**

| Field Name                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output |
|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Slot</b> or <b>Slot State</b> | Slot number and state. The state can be one of the following conditions: <ul style="list-style-type: none"> <li>• <b>Dead</b>—Held in reset because of errors.</li> <li>• <b>Diag</b>—Slot is being ignored while the FPC is running diagnostics.</li> <li>• <b>Dormant</b>—Held in reset.</li> <li>• <b>Empty</b>—No FPC is present.</li> <li>• <b>Offline</b>—(PTX Series Packet Transport Switches only) One of the following two states is displayed:               <ul style="list-style-type: none"> <li>• <b>FPC offlined due to unreachable destinations</b></li> <li>• <b>FPC Offlined due to degraded FPC action</b></li> </ul> </li> <li>• <b>Online</b>—FPC is online and running.</li> <li>• <b>Present</b>—FPC is detected by the chassis daemon but either is not supported by the current version of Junos OS or is inserted in the wrong slot. The output also states either <b>Hardware Not Supported</b> or <b>Hardware Not In Right Slot</b>. The FPC is coming up but not yet online.</li> <li>• <b>Probed</b>—Probe is complete; awaiting restart of the Packet Forwarding Engine.</li> <li>• <b>Probe-wait</b>—Waiting to be probed.</li> </ul> | all levels      |
| <b>Logical slot</b>              | Slot number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | all levels      |

Table 318: show chassis fpc Output Fields (*continued*)

| Field Name                           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output          |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|
| <b>Temp (C) or Temperature</b>       | Temperature of the air passing by the FPC, in degrees Celsius or in both Celsius and Fahrenheit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | all levels<br>all levels |
| <b>Temperature (PTX Series)</b>      | On PTX Series Packet Transport Switches, temperature details are provided in degrees Celsius and Fahrenheit. Output includes: <ul style="list-style-type: none"> <li>• Temperature (PMB)—Temperature of the air passing by the Processor Mezzanine Board (PMB) at the bottom of the FPC.</li> <li>• Temperature (Intake)—Temperature of the air flowing into the chassis.</li> <li>• Temperature (Exhaust)—Exhaust temperatures for multiple zones (Exhaust A and Exhaust B).</li> <li>• Temperature (TLn)—Temperature of the specified Lookup ASIC (TL) of the packet forwarding engine on the FPC.</li> <li>• Temperature (TQn)—Temperature of the specified Queuing and Memory Interface ASIC (TQ) of the packet forwarding engine on the FPC.</li> </ul> | <b>detail</b>            |
| <b>Total CPU Utilization (%)</b>     | Total percentage of CPU being used by the FPC's processor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | all levels               |
| <b>Interrupt CPU Utilization (%)</b> | Of the total CPU being used by the FPC's processor, the percentage being used for interrupts.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | none specified           |
| <b>Memory DRAM (MB)</b>              | Total DRAM, in megabytes, available to the FPC's processor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | none specified           |
| <b>Heap Utilization (%)</b>          | Percentage of heap space (dynamic memory) being used by the FPC's processor. If this number exceeds 80 percent, there may be a software problem (memory leak).<br><br><b>NOTE:</b> On MX Series routers and EX Series switches in a broadband edge environment, heap utilization levels higher than 70 percent can affect unified ISSU, router stability, or scaling capability.                                                                                                                                                                                                                                                                                                                                                                             | none specified           |
| <b>Buffer Utilization (%)</b>        | Percentage of buffer space being used by the FPC's processor for buffering internal messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | none specified           |
| <b>Total CPU DRAM</b>                | Amount of DRAM available to the FPC's CPU.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail</b>            |
| <b>Total RLDRAM</b>                  | Amount of reduced latency dynamic random access memory (RLDRAM) available to the FPC CPU.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail</b>            |
| <b>Total DDR DRAM</b>                | Amount of double data rate dynamic random access memory (DDR DRAM) available to the FPC CPU.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail</b>            |
| <b>Total SRAM</b>                    | Amount of static RAM (SRAM) used by the FPC's CPU.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail</b>            |
| <b>Total SDRAM</b>                   | Total amount of memory used for storing packets and notifications.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail</b>            |
| <b>I/O Manager ASICs information</b> | I/O Manager version number, manufacturer, and part number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail</b>            |

Table 318: show chassis fpc Output Fields (*continued*)

| Field Name        | Field Description                                                                                                   | Level of Output |
|-------------------|---------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Start time</b> | Time when the Routing Engine detected that the FPC was running.                                                     | <b>detail</b>   |
| <b>Uptime</b>     | How long the Routing Engine has been connected to the FPC and, therefore, how long the FPC has been up and running. | <b>detail</b>   |
| <b>PIC type</b>   | (pic-status output only) Type of PIC.                                                                               | none specified  |

## Sample Output

### show chassis fpc (EX6210 Switch)

```

user@switch> show chassis fpc

```

| Slot | State  | Temp<br>(C) | CPU Utilization (%)<br>Total Interrupt | Memory<br>DRAM (MB) | Utilization (%)<br>Heap | Buffer |
|------|--------|-------------|----------------------------------------|---------------------|-------------------------|--------|
| 0    | Empty  |             |                                        |                     |                         |        |
| 1    | Online | 7           | 5 0                                    | 1024                | 0                       | 32     |
| 2    | Empty  |             |                                        |                     |                         |        |
| 3    | Empty  |             |                                        |                     |                         |        |
| 4    | Online | 25          | 17 2                                   | 2048                | 0                       | 30     |
| 5    | Online | 25          | 3 0                                    | 2048                | 0                       | 24     |
| 6    | Online | 6           | 5 0                                    | 1024                | 0                       | 32     |
| 7    | Empty  |             |                                        |                     |                         |        |
| 8    | Empty  |             |                                        |                     |                         |        |
| 9    | Online | 8           | 7 0                                    | 1024                | 0                       | 32     |

### show chassis fpc (M10 Router)

```

user@host> show chassis fpc
FPC status:

```

| Slot | State  | Temp<br>(C) |
|------|--------|-------------|
| 0    | Online | 27          |
| 1    | Online | 28          |

### show chassis fpc (M20 Router)

```

user@host> show chassis fpc
FPC status:

```

| Slot | State  | Temp<br>(C) | CPU Utilization (%)<br>Total Interrupt | Memory<br>DRAM (MB) | Utilization (%)<br>Heap | Buffer |
|------|--------|-------------|----------------------------------------|---------------------|-------------------------|--------|
| 0    | Empty  | 0           | 0 0                                    | 0                   | 0                       | 0      |
| 1    | Online | 38          | 0 0                                    | 8                   | 0                       | 4      |
| 2    | Online | 35          | 0 0                                    | 8                   | 0                       | 3      |
| 3    | Empty  | 0           | 0 0                                    | 0                   | 0                       | 0      |

### show chassis fpc detail (M Series Routers)

```

user@host> show chassis fpc detail 1
Slot 1 information:
State Online
Temperature 48 degrees C
Total CPU DRAM 32 MB
Total SRAM 4 MB
Total SDRAM 256 MB
I/O Manager ASICs information Version 2.0, Foundry IBM, Part number 0

```



```

I/O Manager ASICs information Version 2.0, Foundry IBM, Part number 0
Start time 2000-02-08 02:18:49 UTC
Uptime 14 hours, 41 minutes, 41 seconds

```

#### show chassis fpc detail (MX80 Router)

```

user@host> show chassis fpc detail
Slot 0 information:
 State Online
 Temperature 47 degrees C / 116 degrees F
 Total CPU DRAM 1024 MB
 Total SRAM 331 MB
 Total SDRAM 1280 MB
 Start time 2010-02-08 12:25:33 PST
 Uptime 2 hours, 13 minutes, 19 seconds
Slot 1 information:
 State Online
 Temperature 47 degrees C / 116 degrees F
 Total CPU DRAM 1024 MB
 Total SRAM 331 MB
 Total SDRAM 1280 MB
 Start time 2010-02-08 12:25:33 PST
 Uptime 2 hours, 13 minutes, 19 seconds

```

#### show chassis fpc (MX240 Router)

```

user@host> show chassis fpc

```

| Slot | State  | Temp (C) | CPU Utilization (%) | Memory Utilization (%) |
|------|--------|----------|---------------------|------------------------|
|      |        |          | Total               | DRAM (MB) Heap Buffer  |
| 0    | Empty  |          |                     |                        |
| 1    | Online | 34       | 6                   | 1024 18 30             |
| 2    | Online | 33       | 9                   | 1024 24 30             |

#### show chassis fpc (EX Series Switch)

```

user@host> show chassis fpc

```

| Slot | State  | Temp (C) | CPU Utilization (%) | Memory Utilization (%) |
|------|--------|----------|---------------------|------------------------|
|      |        |          | Total               | DRAM (MB) Heap Buffer  |
| 0    | Empty  |          |                     |                        |
| 1    | Online | 41       | 13                  | 2048 19 14             |
| 2    | Online | 42       | 12                  | 2048 19 14             |

#### show chassis fpc (MX480 Router)

```

user@host> show chassis fpc

```

| Slot | State  | Temp (C) | CPU Utilization (%) | Memory Utilization (%) |
|------|--------|----------|---------------------|------------------------|
|      |        |          | Total               | DRAM (MB) Heap Buffer  |
| 0    | Empty  |          |                     |                        |
| 1    | Online | 36       | 9                   | 1024 17 57             |
| 2    | Empty  |          |                     |                        |
| 3    | Empty  |          |                     |                        |
| 4    | Empty  |          |                     |                        |
| 5    | Empty  |          |                     |                        |

#### show chassis fpc (MX480 Router with 100-Gigabit Ethernet CFP)

```

user@host> show chassis fpc

```

| Slot | State  | Temp (C) | CPU Utilization (%) | Memory Utilization (%) |
|------|--------|----------|---------------------|------------------------|
|      |        |          | Total               | DRAM (MB) Heap Buffer  |
| 0    | Online | 33       | 4                   | 2048 10 13             |
| 1    | Online | 36       | 7                   | 2048 16 13             |
| 2    | Online | 29       | 6                   | 1024 27 29             |

|   |        |    |    |    |      |    |    |
|---|--------|----|----|----|------|----|----|
| 3 | Online | 33 | 0  | 0  | 0    | 0  | 0  |
| 4 | Online | 36 | 7  | 0  | 2048 | 19 | 13 |
| 5 | Online | 34 | 31 | 11 | 2048 | 14 | 13 |

**show chassis fpc pic-status (MX480 Router with 100-Gigabit Ethernet CFP)**

```

user@host> show chassis fpc pic-status
Slot 1 Online MPC Type 3
PIC 2 Online 1X100GE CFP
Slot 2 Online DPCE 40x 1GE R EQ
PIC 0 Online 10x 1GE(LAN) EQ
PIC 1 Online 10x 1GE(LAN) EQ
PIC 2 Online 10x 1GE(LAN) EQ
PIC 3 Online 10x 1GE(LAN) EQ
Slot 3 Online MPC Type 3
PIC 0 Online 1X100GE CFP
PIC 2 Online 1X100GE CFP
Slot 4 Online MPC Type 3
PIC 0 Online 1X100GE CFP
PIC 2 Online 1X100GE CFP
Slot 5 Online MPC Type 2 3D EQ
PIC 0 Online 2x 10GE XFP
PIC 1 Online 2x 10GE XFP
PIC 2 Online 10x 1GE(LAN) SFP
PIC 3 Online 10x 1GE(LAN) SFP

```

**show chassis fpc pic-status (EX Series Switch)**

```

user@host> show chassis fpc pic-status
Slot 1 Online EX9200 32x10G SFP
PIC 0 Online 8X10GE SFPP
PIC 1 Online 8X10GE SFPP
PIC 2 Online 8X10GE SFPP
PIC 3 Online 8X10GE SFPP
Slot 2 Online EX9200 32x10G SFP
PIC 0 Online 8X10GE SFPP
PIC 1 Online 8X10GE SFPP
PIC 2 Online 8X10GE SFPP
PIC 3 Online 8X10GE SFPP

```

**show chassis fpc (MX480 Router with MPC4E)**

```

user@host> show chassis fpc

```

| Slot  | Temp   | CPU | Utilization (%) | Memory    | Utilization (%) |      |        |
|-------|--------|-----|-----------------|-----------|-----------------|------|--------|
| State |        |     | (C) Total       | Interrupt | DRAM (MB)       | Heap | Buffer |
| 0     | Empty  |     |                 |           |                 |      |        |
| 1     | Empty  |     |                 |           |                 |      |        |
| 2     | Online |     | 38 7            | 0         | 2048            | 19   | 14     |
| 3     | Online |     | 39 8            | 0         | 2048            | 18   | 14     |
| 4     | Online |     | 39 7            | 0         | 2048            | 17   | 14     |
| 5     | Empty  |     |                 |           |                 |      |        |

**show chassis fpc detail (MX480 Router with MPC4E)**

```

user@host> show chassis fpc detail
Slot 2 information:
State Online
Temperature 38
Total CPU DRAM 2048 MB
Total RLDRAM 1036 MB
Total DDR DRAM 11264 MB
Start time: 2013-02-18 05:06:57 PST

```

```

Uptime: 17 hours, 41 minutes, 9 seconds
Max Power Consumption 610 Watts
Slot 3 information:
State Online
Temperature 38
Total CPU DRAM 2048 MB
Total RLDRAM 1036 MB
Total DDR DRAM 11264 MB
Start time: 2013-02-18 05:07:00 PST
Uptime: 17 hours, 41 minutes, 6 seconds
Max Power Consumption 610 Watts
Slot 4 information:
State Diagnostics
Temperature 37
Total CPU DRAM 0 MB
Total RLDRAM 0 MB
Total DDR DRAM 0 MB
Max Power Consumption 520 Watts

```

#### show chassis fpc pic-status (MX480 Router with MPC4E)

```

user@host> show chassis fpc pic-status
Slot 2 Online MPC4E 3D 32XGE
PIC 0 Online 8X10GE SFPP
PIC 1 Online 8X10GE SFPP
PIC 2 Online 8X10GE SFPP
PIC 3 Online 8X10GE SFPP
Slot 3 Online MPC4E 3D 2CGE+8XGE
PIC 0 Online 4x10GE SFPP
PIC 1 Online 1X100GE CFP
PIC 2 Online 4x10GE SFPP
PIC 3 Online 1X100GE CFP
Slot 4 Diagnostics MPCE Type 3 3D

```

#### show chassis fpc (MX960 Router)

```

user@host> show chassis fpc

```

| Slot | State  | Temp (C) | CPU Total | Utilization (%) Interrupt | Memory DRAM (MB) | Utilization (%) Heap | Utilization (%) Buffer |
|------|--------|----------|-----------|---------------------------|------------------|----------------------|------------------------|
| 0    | Empty  |          |           |                           |                  |                      |                        |
| 1    | Empty  |          |           |                           |                  |                      |                        |
| 2    | Empty  |          |           |                           |                  |                      |                        |
| 3    | Online | 25       | 19        | 0                         | 1024             | 15                   | 57                     |
| 4    | Empty  |          |           |                           |                  |                      |                        |
| 5    | Online | 26       | 27        | 0                         | 1024             | 15                   | 57                     |
| 6    | Empty  |          |           |                           |                  |                      |                        |
| 7    | Empty  |          |           |                           |                  |                      |                        |
| 8    | Empty  |          |           |                           |                  |                      |                        |
| 9    | Empty  |          |           |                           |                  |                      |                        |
| 10   | Empty  |          |           |                           |                  |                      |                        |
| 11   | Empty  |          |           |                           |                  |                      |                        |

#### show chassis fpc (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```

user@host> show chassis fpc 1

```

| Slot | State  | Temp (C) | CPU Total | Utilization (%) Interrupt | Memory DRAM (MB) | Utilization (%) Heap | Utilization (%) Buffer |
|------|--------|----------|-----------|---------------------------|------------------|----------------------|------------------------|
| 1    | Online | 34       | 5         | 0                         | 3072             | 5                    | 13                     |

**show chassis fpc (MX240, MX480, MX960 with Application Services Modular Line Card)**

```

user@host>show chassis fpc 1 detail
Slot 1 information:
 State Online
 Temperature 34
 Total CPU DRAM 3072 MB
 Total RLDRAM 259 MB
 Total DDR DRAM 4864 MB
 Start time: 2012-06-19 10:51:43 PDT
 Uptime: 16 minutes, 48 seconds
 Max Power Consumption 550 Watts

```

**show chassis fpc (MX2010 Router)**

```

user@host show chassis fpc
 Temp CPU Utilization (%) Memory Utilization (%)
Slot State (C) Total Interrupt DRAM (MB) Heap Buffer
0 Online 34 9 0 2048 18 13
1 Online 32 9 0 2048 15 13
2 Empty
3 Empty
4 Empty
5 Empty
6 Empty
7 Empty
8 Online 31 13 0 2048 11 13
9 Online 33 10 0 2048 18 13

```

**show chassis fpc (MX2020 Router)**

```

user@host show chassis fpc
 Temp CPU Utilization (%) Memory Utilization (%)
Slot State (C) Total Interrupt DRAM (MB) Heap Buffer
0 Online 10 12 0 2048 18 13
1 Online 8 9 0 2048 18 13
2 Online 7 9 0 2048 18 13
3 Online 8 10 0 2048 18 13
4 Online 9 10 0 2048 18 13
5 Online 8 9 0 2048 18 13
6 Online 8 10 0 2048 18 13
7 Online 9 9 0 2048 18 13
8 Online 9 10 0 2048 18 13
9 Online 10 9 0 2048 18 13
10 Online 16 8 0 2048 18 13
11 Online 11 10 0 2048 18 13
12 Online 10 10 0 2048 18 13
13 Online 11 9 0 2048 18 13
14 Online 12 10 0 2048 18 13
15 Online 13 9 0 2048 18 13
16 Online 13 9 0 2048 18 13
17 Online 12 9 0 2048 18 13
18 Online 12 8 0 2048 18 13
19 Online 14 10 0 2048 18 13

```

**show chassis fpc (MX2020 Router with MPC4E)**

```

user@host> show chassis fpc
 Temp CPU Utilization (%) Memory Utilization (%)
Slot State (C) Total Interrupt DRAM (MB) Heap Buffer
0 Online 33 12 2 2048 11 13

```

|    |        |    |    |   |      |    |    |
|----|--------|----|----|---|------|----|----|
| 1  | Empty  |    |    |   |      |    |    |
| 2  | Empty  |    |    |   |      |    |    |
| 3  | Empty  |    |    |   |      |    |    |
| 4  | Empty  |    |    |   |      |    |    |
| 5  | Empty  |    |    |   |      |    |    |
| 6  | Empty  |    |    |   |      |    |    |
| 7  | Empty  |    |    |   |      |    |    |
| 8  | Empty  |    |    |   |      |    |    |
| 9  | Online | 31 | 10 | 0 | 2048 | 11 | 13 |
| 10 | Online | 32 | 7  | 0 | 2048 | 14 | 13 |
| 11 | Empty  |    |    |   |      |    |    |
| 12 | Empty  |    |    |   |      |    |    |
| 13 | Empty  |    |    |   |      |    |    |
| 14 | Online | 28 | 12 | 0 | 2048 | 15 | 14 |
| 15 | Empty  |    |    |   |      |    |    |
| 16 | Empty  |    |    |   |      |    |    |
| 17 | Empty  |    |    |   |      |    |    |
| 18 | Empty  |    |    |   |      |    |    |
| 19 | Online | 38 | 8  | 0 | 2048 | 18 | 13 |

#### show chassis fpc detail (MX2020 Router with MPC4E)

```

user@host> show chassis fpc detail
Slot 0 information:
 State Online
 Temperature 34
 Total CPU DRAM 2048 MB
 Total RLDRAM 806 MB
 Total DDR DRAM 2632 MB
 Start time: 2013-02-17 08:17:35 PST
 Uptime: 1 day, 14 hours, 50 minutes, 39 seconds
 Max Power Consumption 368 Watts
Slot 9 information:
 State Online
 Temperature 32
 Total CPU DRAM 2048 MB
 Total RLDRAM 806 MB
 Total DDR DRAM 2632 MB
 Start time: 2013-02-17 08:17:43 PST
 Uptime: 1 day, 14 hours, 50 minutes, 31 seconds
 Max Power Consumption 368 Watts
Slot 10 information:
 State Online
 Temperature 37
 Total CPU DRAM 2048 MB
 Total RLDRAM 1036 MB
 Total DDR DRAM 6656 MB
 Start time: 2013-02-17 08:17:54 PST
 Uptime: 1 day, 14 hours, 50 minutes, 20 seconds
 Max Power Consumption 520 Watts
Slot 14 information:
 State Online
 Temperature 32
 Total CPU DRAM 2048 MB
 Total RLDRAM 1036 MB
 Total DDR DRAM 11264 MB
 Start time: 2013-02-17 08:18:01 PST
 Uptime: 1 day, 14 hours, 50 minutes, 13 seconds
 Max Power Consumption 610 Watts
Slot 19 information:
 State Online

```

```
Temperature 38
Total CPU DRAM 2048 MB
Total RLD RAM 1324 MB
Total DDR DRAM 5120 MB
Start time: 2013-02-17 08:18:08 PST
Uptime: 1 day, 14 hours, 50 minutes, 6 seconds
Max Power Consumption 440 Watts
```

#### show chassis fpc pic-status (MX2020 Router with MPC4E)

```
user@host> show chassis fpc pic-status
Slot 0 Online MPC Type 2 3D EQ
 PIC 0 Online MIC-3D-10C192-XFP
 PIC 2 Online MIC-3D-80C30C12-40C48
Slot 9 Online MPC Type 2 3D EQ
 PIC 0 Online 2x 10GE XFP
 PIC 1 Online 2x 10GE XFP
 PIC 2 Online MIC-3D-10C192-XFP
Slot 10 Online MPCE Type 3 3D
 PIC 0 Online 1x 10GE XFP
 PIC 1 Online 1x 10GE XFP
Slot 14 Online MPC4E 3D 2CGE+8XGE
 PIC 0 Online 4x10GE SFPP
 PIC 1 Online 1X100GE CFP
 PIC 2 Online 4x10GE SFPP
 PIC 3 Online 1X100GE CFP
Slot 19 Online MPC 3D 16x 10GE
 PIC 0 Online 4x 10GE(LAN) SFP+
 PIC 1 Online 4x 10GE(LAN) SFP+
 PIC 2 Online 4x 10GE(LAN) SFP+
 PIC 3 Online 4x 10GE(LAN) SFP+
```

#### show chassis fpc detail (MX Series Routers)

```
user@host> show chassis fpc detail 2
Slot 0 information:
State Online
Temperature 36 degrees C / 96 degrees F
Total CPU DRAM 1024 MB
Total RLD RAM 256 MB
Total DDR DRAM 4096 MB
Start time: 2009-08-11 21:20:30 PDT
Uptime: 2 hours, 8 minutes, 50 seconds
Max Power Consumption 335 Watts
```

#### show chassis fpc detail (EX Series Switches)

```
user@host> show chassis fpc detail 2
Slot 1 information:
State Online
Temperature 41
Total CPU DRAM 2048 MB
Total RLD RAM 1036 MB
Total DDR DRAM 11264 MB
Start time: 2013-04-02 00:04:52 PDT
Uptime: 7 days, 9 hours, 47 minutes, 46 seconds
Max Power Consumption 610 Watts
Slot 2 information:
State Online
Temperature 41
Total CPU DRAM 2048 MB
Total RLD RAM 1036 MB
```

```

Total DDR DRAM 11264 MB
Start time: 2013-04-02 00:04:56 PDT
Uptime: 7 days, 9 hours, 47 minutes, 42 seconds
Max Power Consumption 610 Watts

```

### show chassis fpc (Hardware Not Supported)

```

user@host> show chassis fpc
show chassis fpc

```

| Slot | State   | Temp (C) | CPU Utilization (%) Total | Interrupt | Memory Utilization (%) CPU less FPC | DRAM (MB) | Heap | Buffer |
|------|---------|----------|---------------------------|-----------|-------------------------------------|-----------|------|--------|
| 0    | Online  |          |                           |           |                                     |           |      |        |
| 1    | Present |          |                           |           |                                     |           |      |        |
| 2    | Online  |          | 0                         | 0         | 0                                   | 0         | 0    | 0      |
| 3    | Present |          |                           |           |                                     |           |      |        |
| 4    | Empty   |          |                           |           |                                     |           |      |        |
| 5    | Empty   |          |                           |           |                                     |           |      |        |
| 6    | Online  |          | 0                         | 0         | 0                                   | 0         | 0    | 0      |

### show chassis fpc detail (Hardware Not Supported)

```

user@host> show chassis fpc detail
Slot 0 information:
 State Online
 Total CPU DRAM ----- CPU less FPC -----
 Start time 2006-07-07 03:21:00 UTC
 Uptime 27 minutes, 51 seconds
Slot 1 information:
 State Present
 Reason --- Hardware Not In Right Slot ---
Slot 2 information:
 State Online
 Total CPU DRAM 32 MB
 Start time 2006-07-07 03:20:59 UTC
 Uptime 27 minutes, 52 seconds
Slot 3 information:
 State Present
 Reason --- Hardware Not Supported ---
 Total CPU DRAM 0 MB
Slot 6 information:
 State Online
 Total CPU DRAM 32 MB
 Start time 2006-07-07 03:21:01 UTC
 Uptime 27 minutes, 50 seconds

```

### show chassis fpc pic-status

```

user@host> show chassis fpc pic-status
Slot 0 Online
 PIC 1 1x OC-12 ATM, MM
 PIC 2 1x OC-12 ATM, MM
 PIC 3 1x OC-12 ATM, MM
Slot 1 Online
 PIC 0 1x OC-48 SONET, SMIR
Slot 2 Online
 PIC 0 1x OC-192 SONET, SMSR

```

### show chassis fpc pic-status (M Series Routers)

```

user@host> show chassis fpc pic-status

```

```

Slot 1 Online FPC Type 1
 PIC 0 Present 2x OC-3 ATM, MM- Hardware Error
 PIC 1 Online 4x OC-3 SONET, SMIR
Slot 2 Online E-FPC Type 2
 PIC 0 Online 4x G/E, 1000 BASE-SX
 PIC 1 Online 2x G/E SFP, 1000 BASE
 PIC 3 Online 1x Tunnel
Slot 3 Online E-FPC Type 1
 PIC 0 Online 1x G/E IQ, 1000 BASE
 PIC 2 Online 1x G/E SFP, 1000 BASE
Slot 4 Online E-FPC Type 2
 PIC 0 Online 4x G/E SFP, 1000 BASE
 PIC 1 Online 4x G/E SFP, 1000 BASE
 PIC 2 Online 4x G/E SFP, 1000 BASE
 PIC 3 Online 4x G/E SFP, 1000 BASE
Slot 5 Online FPC Type 2
...

```

#### show chassis fpc pic-status (M120 Router)

```

user@host> show chassis fpc pic-status
Slot 1 Online M120 CFPC 10GE
 PIC 0 Online 1x 10GE(LAN/WAN) XFP
Slot 3 Online M120 FPC Type 2 (proto)
 PIC 0 Online 2x G/E IQ, 1000 BASE
 PIC 1 Online 4x OC-3 SONET, SMIR
 PIC 2 Online 2x G/E IQ, 1000 BASE
 PIC 3 Online 8x 1GE(LAN), IQ2
Slot 4 Online M120 FPC Type 3 (proto)
 PIC 0 Online 10x 1GE(LAN), 1000 BASE
Slot 5 Online M120 FPC Type 1 (proto)
 PIC 0 Present 1x G/E, 1000 BASE-LX- Not Supported
 PIC 1 Online 1x CHOC3 IQ SONET, SMLR
 PIC 2 Online 4x CHDS3 IQ
 PIC 3 Online 1x G/E SFP, 1000 BASE

```

#### show chassis fpc pic-status (MX240, MX480, and MX960 Routers with Application Services Modular Line Card)

In the following output **Slot 1** and **Slot 5** are the Application Services Modular Carrier Cards (AS MCC), **PIC 0** is the Application Services Modular Storage Card (AS MSC), and **PIC 2** is the Application Services Modular Processing Card (AS MXC).

```

user@host> show chassis fpc pic-status
Slot 2 Online MPC Type 1 3D Q
Slot 1 Online AS-MCC
 PIC 0 Online AS-MSC
 PIC 2 Online AS-MXC
Slot 4 Offline MPC 3D 16x 10GE
Slot 5 Offline AS-MCC

```

#### show chassis fpc lcc (TX Matrix Router)

```

user@host> show chassis fpc lcc 0
lcc0-re0:

```

| Slot | State  | Temp (C) | CPU Total | Utilization (%) Interrupt | Utilization (%) DRAM (MB) | Memory Heap | Utilization (%) Buffer |
|------|--------|----------|-----------|---------------------------|---------------------------|-------------|------------------------|
| 0    | Empty  |          |           |                           |                           |             |                        |
| 1    | Online | 27       | 2         | 0                         | 256                       | 8           | 44                     |
| 2    | Online | 27       | 3         | 0                         | 256                       | 15          | 44                     |
| 3    | Empty  |          |           |                           |                           |             |                        |



```

4 Empty
5 Empty
6 Empty
7 Empty

```

### show chassis fpc pic-status (TX Matrix Router)

```

user@host> show chassis fpc pic-status
lcc0-re0:

```

```

Slot 0 Online FPC Type 3
 PIC 0 Online 1x OC-192 SM SR1
 PIC 1 Online 1x OC-192 SM SR2
 PIC 2 Online 1x OC-192 SM SR1
 PIC 3 Online 1x Tunnel
Slot 1 Online FPC Type 2
 PIC 0 Online 1x OC-48 SONET, SMSR
 PIC 1 Online 1x OC-48 SONET, SMSR

```

```

lcc1-re0:

```

```

lcc2-re0:

```

```

Slot 1 Online FPC Type 3
 PIC 0 Online 1x OC-192 SM SR1
Slot 5 Online FPC Type 2
 PIC 0 Online 1x OC-48 SONET, SMSR
 PIC 1 Online 2x G/E, 1000 BASE-LX
 PIC 2 Online 2x G/E, 1000 BASE-LX
 PIC 3 Online 1x OC-48 SONET, SMSR

```

```

lcc3-re0:

```

### show chassis fpc pic-status lcc (TX Matrix Router)

```

user@host> show chassis fpc pic-status lcc 0
lcc0-re0:

```

```

Slot 0 Online FPC Type 3
 PIC 0 Online 1x OC-192 SM SR2
Slot 1 Online FPC Type 2
 PIC 0 Online 2x OC-12 ATM2 IQ, MM
 PIC 1 Online 1x OC-48 SONET, SMSR
 PIC 2 Online 1x OC-48 SONET, SMSR
 PIC 3 Online 4x G/E, 1000 BASE-SX

```

### show chassis fpc (TX Matrix Plus Router)

```

user@host> show chassis fpc
lcc0-re0:

```

```

Slot State Temp CPU Utilization (%) Memory Utilization (%)
 (C) Total Interrupt DRAM (MB) Heap Buffer
0 Empty
1 Online 38 4 0 2048 3 24
2 Online 43 8 0 2048 6 24
3 Empty
4 Online 43 6 0 2048 6 24
5 Empty
6 Online 42 13 0 2048 6 24

```

```

7 Online 45 7 0 2048 3 24

```

```
lcc2-re0:
```

```

Slot State Temp CPU Utilization (%) Memory Utilization (%)
 (C) Total Interrupt DRAM (MB) Heap Buffer
0 Online 42 10 0 2048 6 24
1 Empty
2 Online 42 11 0 2048 6 24
3 Online 40 5 0 2048 3 24
4 Online 33 26 0 1024 8 49
5 Empty
6 Online 43 8 0 2048 6 24
7 Online 46 6 0 2048 3 24

```

```
lcc3-re0:
```

```

Slot State Temp CPU Utilization (%) Memory Utilization (%)
 (C) Total Interrupt DRAM (MB) Heap Buffer
0 Empty
1 Empty
2 Online 39 30 0 2048 7 24
3 Empty
4 Online 41 8 0 2048 6 24
5 Online 41 12 0 2048 6 24
6 Online 40 8 0 2048 6 24
7 Online 42 4 0 2048 3 24

```

#### show chassis fpc lcc (TX Matrix Plus Router)

```
user@host> show chassis fpc lcc 0
```

```
lcc0-re0:
```

```

Slot State Temp CPU Utilization (%) Memory Utilization (%)
 (C) Total Interrupt DRAM (MB) Heap Buffer
0 Empty
1 Online 38 4 0 2048 3 24
2 Online 43 8 0 2048 6 24
3 Empty
4 Online 43 6 0 2048 6 24
5 Empty
6 Online 42 14 0 2048 6 24
7 Online 45 6 0 2048 3 24

```

#### show chassis fpc detail (TX Matrix Plus Router)

```
user@host> show chassis fpc details
```

```
lcc0-re0:
```

```
Slot 1 information:
```

```

State Online
Temperature 38 degrees C / 100 degrees F
Total CPU DRAM 2048 MB
Total SRAM 64 MB
Total SDRAM 1280 MB
Start time 2010-10-04 20:06:22 PDT
Uptime 1 hour, 32 minutes, 51 seconds

```

```
Slot 2 information:
```

```

State Online
Temperature 43 degrees C / 109 degrees F

```

```

Total CPU DRAM 2048 MB
Total SRAM 128 MB
Total SDRAM 2560 MB
Start time 2010-10-04 20:06:37 PDT
Uptime 1 hour, 32 minutes, 36 seconds
Slot 4 information:
State Online
Temperature 43 degrees C / 109 degrees F
Total CPU DRAM 2048 MB
Total SRAM 128 MB
Total SDRAM 2560 MB
Start time 2010-10-04 20:06:40 PDT
Uptime 1 hour, 32 minutes, 33 seconds
Slot 6 information:
State Online
Temperature 42 degrees C / 107 degrees F
Total CPU DRAM 2048 MB
Total SRAM 128 MB
Total SDRAM 2560 MB
Start time 2010-10-04 20:06:42 PDT
Uptime 1 hour, 32 minutes, 31 seconds
Slot 7 information:
State Online
Temperature 45 degrees C / 113 degrees F
Total CPU DRAM 2048 MB
Total SRAM 64 MB
Total SDRAM 1280 MB
Start time 2010-10-04 20:06:43 PDT
Uptime 1 hour, 32 minutes, 30 seconds

```

lcc2-re0:

```

Slot 0 information:
State Online
Temperature 42 degrees C / 107 degrees F
Total CPU DRAM 2048 MB
Total SRAM 128 MB
Total SDRAM 2560 MB
Start time 2010-10-04 20:06:35 PDT
Uptime 1 hour, 32 minutes, 38 seconds
Slot 2 information:
State Online
Temperature 42 degrees C / 107 degrees F
Total CPU DRAM 2048 MB
Total SRAM 128 MB
Total SDRAM 2560 MB
Start time 2010-10-04 20:06:37 PDT
Uptime 1 hour, 32 minutes, 36 seconds
Slot 3 information:
State Online
Temperature 40 degrees C / 104 degrees F
Total CPU DRAM 2048 MB
Total SRAM 64 MB
Total SDRAM 1280 MB
Start time 2010-10-04 20:06:28 PDT
Uptime 1 hour, 32 minutes, 45 seconds
Slot 4 information:
State Online
Temperature 33 degrees C / 91 degrees F
Total CPU DRAM 1024 MB
Total SRAM 64 MB

```

```
Total SDRAM 1280 MB
Start time 2010-10-04 20:08:03 PDT
Uptime 1 hour, 31 minutes, 10 seconds
Slot 6 information:
State Online
Temperature 43 degrees C / 109 degrees F
Total CPU DRAM 2048 MB
Total SRAM 128 MB
Total SDRAM 2560 MB
Start time 2010-10-04 20:06:44 PDT
Uptime 1 hour, 32 minutes, 29 seconds
Slot 7 information:
State Online
Temperature 46 degrees C / 114 degrees F
Total CPU DRAM 2048 MB
Total SRAM 64 MB
Total SDRAM 1280 MB
Start time 2010-10-04 20:06:46 PDT
Uptime 1 hour, 32 minutes, 27 seconds
```

lcc3-re0:

```

Slot 2 information:
State Online
Temperature 38 degrees C / 100 degrees F
Total CPU DRAM 2048 MB
Total SRAM 128 MB
Total SDRAM 2560 MB
Start time 2010-10-04 20:17:31 PDT
Uptime 1 hour, 21 minutes, 42 seconds
Slot 4 information:
State Online
Temperature 41 degrees C / 105 degrees F
Total CPU DRAM 2048 MB
Total SRAM 128 MB
Total SDRAM 2560 MB
Start time 2010-10-04 20:17:34 PDT
Uptime 1 hour, 21 minutes, 39 seconds
Slot 5 information:
State Online
Temperature 41 degrees C / 105 degrees F
Total CPU DRAM 2048 MB
Total SRAM 128 MB
Total SDRAM 2560 MB
Start time 2010-10-04 20:17:36 PDT
Uptime 1 hour, 21 minutes, 37 seconds
Slot 6 information:
State Online
Temperature 40 degrees C / 104 degrees F
Total CPU DRAM 2048 MB
Total SRAM 128 MB
Total SDRAM 2560 MB
Start time 2010-10-04 20:17:39 PDT
Uptime 1 hour, 21 minutes, 34 seconds
Slot 7 information:
State Online
Temperature 42 degrees C / 107 degrees F
Total CPU DRAM 2048 MB
Total SRAM 64 MB
Total SDRAM 1280 MB
```

```

Start time 2010-10-04 20:17:41 PDT
Uptime 1 hour, 21 minutes, 32 seconds

```

### show chassis fpc pic-status (TX Matrix Plus Router)

```
user@host> show chassis fpc pic-status
```

```
lcc0-re0:
```

```

Slot 1 Online FPC Type 2-ES
 PIC 0 Online 8x 1GE(LAN), IQ2
Slot 2 Online FPC Type 4-ES
 PIC 0 Online 4x 10GE (LAN/WAN) XFP
Slot 4 Online FPC Type 4-ES
 PIC 0 Online 4x 10GE (LAN/WAN) XFP
Slot 6 Online FPC Type 4-ES
 PIC 0 Online 4x 10GE (LAN/WAN) XFP
 PIC 1 Online 4x 10GE (LAN/WAN) XFP
Slot 7 Online FPC Type 3-ES
 PIC 0 Online 10x 1GE(LAN), 1000 BASE
 PIC 2 Online 1x OC-192 SM SR2
 PIC 3 Online 10x 1GE(LAN), 1000 BASE

```

```
lcc2-re0:
```

```

Slot 0 Online FPC Type 4-ES
 PIC 0 Online 4x 10GE (LAN/WAN) XFP
Slot 2 Online FPC Type 4-ES
 PIC 0 Online 4x 10GE (LAN/WAN) XFP
 PIC 1 Online 4x 10GE (LAN/WAN) XFP
Slot 3 Online FPC Type 2-ES
 PIC 0 Online 8x 1GE(LAN), IQ2
Slot 4 Online FPC Type 4
 PIC 0 Online 10x10GE(LAN/WAN) SFPP
Slot 6 Online FPC Type 4-ES
 PIC 0 Online 4x OC-192 SONET XFP
Slot 7 Online FPC Type 3-ES
 PIC 0 Online 10x 1GE(LAN), 1000 BASE
 PIC 1 Offline 1x 10GE(LAN/WAN) IQ2E
 PIC 2 Online 1x OC-192 SM SR2
 PIC 3 Online 1x Tunnel

```

```
lcc3-re0:
```

```

Slot 2 Online FPC Type 4-ES
 PIC 0 Online 10x10GE(LAN/WAN) SFPP
Slot 4 Online FPC Type 4-ES
 PIC 0 Online 4x OC-192 SONET XFP
Slot 5 Online FPC Type 4-ES
 PIC 0 Online 4x OC-192 SONET XFP
 PIC 1 Online 4x 10GE (LAN/WAN) XFP
Slot 6 Online FPC Type 4-ES
 PIC 1 Online 4x 10GE (LAN/WAN) XFP
Slot 7 Online FPC Type 3-ES
 PIC 0 Online 10x 1GE(LAN), 1000 BASE
 PIC 1 Online 8x 1GE(TYPE3), IQ2E
 PIC 2 Online 4x OC-48 SONET

```

### show chassis fpc (TI600 Router)

```
user@host> show chassis fpc
```

| Slot | State  | Temp<br>(C) | CPU Utilization (%)<br>Total Interrupt | Memory<br>DRAM (MB) | Utilization (%)<br>Heap Buffer |
|------|--------|-------------|----------------------------------------|---------------------|--------------------------------|
| 0    | Empty  |             |                                        |                     |                                |
| 1    | Empty  |             |                                        |                     |                                |
| 2    | Online | 49          | 3 0                                    | 2048                | 3 24                           |
| 3    | Online | 46          | 6 0                                    | 2048                | 6 24                           |
| 4    | Empty  |             |                                        |                     |                                |
| 5    | Online | 46          | 5 0                                    | 2048                | 3 24                           |
| 6    | Empty  |             |                                        |                     |                                |
| 7    | Online | 44          | 8 0                                    | 1024                | 7 49                           |

### show chassis fpc detail (TI600 Router)

```
user@host> show chassis fpc detail
```

```
show chassis fpc detail
```

```
Slot 2 information:
```

```
State Online
Temperature 49 degrees C / 120 degrees F
Total CPU DRAM 2048 MB
Total SRAM 64 MB
Total SDRAM 1280 MB
Start time 2010-10-04 21:12:52 PDT
Uptime 32 minutes, 9 seconds
```

```
Slot 3 information:
```

```
State Online
Temperature 47 degrees C / 116 degrees F
Total CPU DRAM 2048 MB
Total SRAM 128 MB
Total SDRAM 2560 MB
Start time 2010-10-04 21:13:06 PDT
Uptime 31 minutes, 55 seconds
```

```
Slot 5 information:
```

```
State Online
Temperature 46 degrees C / 114 degrees F
Total CPU DRAM 2048 MB
Total SRAM 64 MB
Total SDRAM 1280 MB
Start time 2010-10-04 21:12:56 PDT
Uptime 32 minutes, 5 seconds
```

```
Slot 7 information:
```

```
State Online
Temperature 44 degrees C / 111 degrees F
Total CPU DRAM 1024 MB
Total SRAM 64 MB
Total SDRAM 1280 MB
Start time 2010-10-04 21:14:34 PDT
Uptime 30 minutes, 27 seconds
```

### show chassis fpc <fpc-slot> (EX Series Switch)

```
user@host> show chassis fpc 2
```

| Slot | State  | Temp<br>(C) | CPU Utilization (%)<br>Total Interrupt | Memory<br>DRAM (MB) | Utilization (%)<br>Heap Buffer |
|------|--------|-------------|----------------------------------------|---------------------|--------------------------------|
| 2    | Online | 40          | 12 0                                   | 2048                | 19 14                          |

### show chassis fpc slot (TI600 Router)

```
user@host> show chassis fpc slot 2
```

```
Temp CPU Utilization (%) Memory Utilization (%)
```

| Slot | State  | (C) | Total | Interrupt | DRAM (MB) | Heap | Buffer |
|------|--------|-----|-------|-----------|-----------|------|--------|
| 2    | Online | 49  | 3     | 0         | 2048      | 3    | 24     |

### show chassis fpc pic-status (T1600 Router)

```
user@host> show chassis fpc pic-status
```

```
Slot 2 Online FPC Type 1-ES
PIC 0 Online Load Type 1
PIC 1 Online 4x 1GE(LAN), IQ2E
PIC 3 Online 1x OC-12-3 SFP
Slot 3 Online FPC Type 4-ES
PIC 0 Online 4x 10GE (LAN/WAN) XFP
PIC 1 Online 4x OC-192 SONET XFP
Slot 5 Online FPC Type 2-ES
PIC 0 Online Load Type 2
PIC 1 Online 8x 1GE(LAN), IQ2E
PIC 2 Online 8x 1GE(LAN), IQ2E
PIC 3 Online 1x OC-48-12-3 SFP
Slot 7 Online FPC Type 4
PIC 0 Online 4x 10GE (LAN/WAN) XFP
```

### show chassis fpc (T4000 Router)

```
user@host> show chassis fpc
```

```
regress@stymphalian# run show chassis fpc
```

| Slot | State  | Temp<br>(C) | CPU<br>Total | Utilization (%)<br>Interrupt | Memory<br>DRAM (MB) | Utilization (%)<br>Heap | Buffer |
|------|--------|-------------|--------------|------------------------------|---------------------|-------------------------|--------|
| 0    | Online | 48          | 15           | 0                            | 2816                | 21                      | 27     |
| 1    | Empty  |             |              |                              |                     |                         |        |
| 2    | Empty  |             |              |                              |                     |                         |        |
| 3    | Online | 51          | 15           | 0                            | 2816                | 21                      | 27     |
| 4    | Empty  |             |              |                              |                     |                         |        |
| 5    | Online | 39          | 8            | 0                            | 2048                | 6                       | 23     |
| 6    | Online | 49          | 15           | 0                            | 2816                | 21                      | 27     |
| 7    | Empty  |             |              |                              |                     |                         |        |

### show chassis fpc detail (T4000 Router)

```
user@host> show chassis fpc detail
```

Slot 0 information:

```
State Online
Temperature 48 degrees C / 118 degrees F
Total CPU DRAM 2816 MB
Total SRAM 1554 MB
Total SDRAM 10752 MB
Start time 2012-02-09 22:56:25 PST
Uptime 2 hours, 40 minutes, 52 seconds
```

Slot 3 information:

```
State Online
Temperature 51 degrees C / 123 degrees F
Total CPU DRAM 2816 MB
Total SRAM 1554 MB
Total SDRAM 10752 MB
Start time 2012-02-09 22:56:22 PST
Uptime 2 hours, 40 minutes, 55 seconds
```

Slot 5 information:

```
State Online
Temperature 39 degrees C / 102 degrees F
Total CPU DRAM 2048 MB
Total SRAM 128 MB
```

```

Total SDRAM 2560 MB
Start time 2012-02-09 22:51:27 PST
Uptime 2 hours, 45 minutes, 50 seconds
Slot 6 information:
State Online
Temperature 49 degrees C / 120 degrees F
Total CPU DRAM 2816 MB
Total SRAM 1554 MB
Total SDRAM 10752 MB
Start time 2012-02-09 22:56:29 PST
Uptime 2 hours, 40 minutes, 48 seconds

```

#### show chassis fpc pic-status (T4000 Router)

```

user@host> show chassis fpc pic-status
Slot 0 Online FPC Type 5-3D
PIC 0 Online 12x10GE (LAN/WAN) SFPP
PIC 1 Online 12x10GE (LAN/WAN) SFPP
Slot 3 Online FPC Type 5-3D
PIC 0 Online 1x100GE
PIC 1 Online 12x10GE (LAN/WAN) SFPP
Slot 5 Online FPC Type 4-ES
PIC 0 Online 100GE
PIC 1 Online 100GE CFP
Slot 6 Online FPC Type 5-3D
PIC 0 Online 12x10GE (LAN/WAN) SFPP
PIC 1 Online 12x10GE (LAN/WAN) SFPP

```

#### show chassis fpc (QFX Series)

```

user@switch> show chassis fpc
Temp CPU Utilization (%) Memory Utilization (%)
Slot State (C) Total Interrupt DRAM (MB) Heap Buffer
0 Online 26 2 0 2820 0 49

```

#### show chassis fpc detail (QFX3500 Switches)

```

user@switch> show chassis fpc detail
Slot 0 information:
State Online
Temperature 28 degrees C / 82 degrees F
Total CPU DRAM 2820 MB
Total SRAM 0 MB
Total SDRAM 0 MB
Start time 2010-09-20 01:34:13 PDT
Uptime 3 days, 3 hours, 31 minutes, 48 seconds

```

#### show chassis fpc pic-status (QFX3500 Switches)

```

user@switch> show chassis fpc pic-status
Slot 0 Online QFX 48x10G 4x40G Switch
PIC 0 Online 48x 10G-SFP+
PIC 1 Online 15x 10G-SFP+

```

#### show chassis fpc interconnect-device (QFabric System)

```

user@switch> show chassis fpc interconnect-device interconnect1
FPC status:
Slot State Temp
 (C)
0 Online 0
1 Online 0

```



|    |        |   |
|----|--------|---|
| 2  | Online | 0 |
| 3  | Online | 0 |
| 4  | Online | 0 |
| 5  | Online | 0 |
| 6  | Online | 0 |
| 7  | Online | 0 |
| 8  | Online | 0 |
| 9  | Online | 0 |
| 10 | Online | 0 |
| 11 | Online | 0 |
| 12 | Online | 0 |
| 13 | Online | 0 |
| 14 | Online | 0 |
| 15 | Online | 0 |

### show chassis fpc interconnect-device (QFabric System)

```
user@switch> show chassis fpc interconnect-device interconnect1 3
FPC status:
```

| Slot | State  | Temp<br>(C) |
|------|--------|-------------|
| 3    | Online | 0           |

### show chassis fpc interconnect-device detail (QFabric System)

```
user@switch> show chassis fpc interconnect-device interconnect1 3 detail
Slot 3 information:
```

|             |                            |
|-------------|----------------------------|
| State       | Online                     |
| Temperature | 0 degrees C / 32 degrees F |
| Start time  | 2011-08-18 10:45:04 PDT    |
| Uptime      | 1 minute, 49 seconds       |

### show chassis fpc pic-status interconnect-device (QFabric System)

```
user@switch> show chassis fpc pic-status interconnect-device interconnect1
```

|        |        |                              |
|--------|--------|------------------------------|
| Slot 0 | Online | QFX 16-port QSFP+ Front Card |
| PIC 0  | Online | 16x 40G-QSFP+                |
| PIC 1  | Online | 16x 40G-GE                   |
| Slot 1 | Online | QFX 16-port QSFP+ Front Card |
| PIC 0  | Online | 16x 40G-QSFP+                |
| PIC 1  | Online | 16x 40G-GE                   |
| Slot 2 | Online | QFX 16-port QSFP+ Front Card |
| PIC 0  | Online | 16x 40G-QSFP+                |
| PIC 1  | Online | 16x 40G-GE                   |
| Slot 3 | Online | QFX 16-port QSFP+ Front Card |
| PIC 0  | Online | 16x 40G-QSFP+                |
| PIC 1  | Online | 16x 40G-GE                   |
| Slot 4 | Online | QFX 16-port QSFP+ Front Card |
| PIC 0  | Online | 16x 40G-QSFP+                |
| PIC 1  | Online | 16x 40G-GE                   |
| Slot 5 | Online | QFX 16-port QSFP+ Front Card |
| PIC 0  | Online | 16x 40G-QSFP+                |
| PIC 1  | Online | 16x 40G-GE                   |
| Slot 6 | Online | QFX 16-port QSFP+ Front Card |
| PIC 0  | Online | 16x 40G-QSFP+                |
| PIC 1  | Online | 16x 40G-GE                   |
| Slot 7 | Online | QFX 16-port QSFP+ Front Card |
| PIC 0  | Online | 16x 40G-QSFP+                |
| PIC 1  | Online | 16x 40G-GE                   |
| Slot 8 | Online | QFX Fabric Rear Card         |
| PIC 0  | Online | 16x 40G-GE                   |
| Slot 9 | Online | QFX Fabric Rear Card         |

```

PIC 0 Online 16x 40G-GE
Slot 10 Online QFX Fabric Rear Card
PIC 0 Online 16x 40G-GE
Slot 11 Online QFX Fabric Rear Card
PIC 0 Online 16x 40G-GE
Slot 12 Online QFX Fabric Rear Card
PIC 0 Online 16x 40G-GE
Slot 13 Online QFX Fabric Rear Card
PIC 0 Online 16x 40G-GE
Slot 14 Online QFX Fabric Rear Card
PIC 0 Online 16x 40G-GE
Slot 15 Online QFX Fabric Rear Card
PIC 0 Online 16x 40G-GE

```

#### show chassis fpc pic-status node-device (QFabric System)

```

user@switch> show chassis fpc pic-status node-device node1
Slot node1 Online QFX 48x10G 4x40G Switch
PIC 0 Online 48x 10G-SFP+
PIC 1 Online 4x 40G-QSFP+

```

#### show chassis fpc (PTX5000 Packet Transport Switch)

```

user@host> show chassis fpc

```

| Slot | State  | Temp (C) | CPU Utilization (%) | Memory DRAM (MB) | Utilization (%) |
|------|--------|----------|---------------------|------------------|-----------------|
|      |        |          | Total Interrupt     | Heap             | Buffer          |
| 0    | Empty  |          |                     |                  |                 |
| 1    | Empty  |          |                     |                  |                 |
| 2    | Online | 50       | 6 0                 | 2816             | 5 27            |
| 3    | Empty  |          |                     |                  |                 |
| 4    | Empty  |          |                     |                  |                 |
| 5    | Online | 48       | 9 0                 | 2816             | 5 27            |
| 6    | Empty  |          |                     |                  |                 |
| 7    | Online | 49       | 8 0                 | 2816             | 5 27            |

#### show chassis fpc detail (PTX5000 Packet Transport Switch)

```

user@host> show chassis fpc detail
Slot 2 information:
State Online
Temperature 35 degrees C / 95 degrees F (PMB)
Temperature 35 degrees C / 95 degrees F (Intake)
Temperature 50 degrees C / 122 degrees F (Exhaust A)
Temperature 54 degrees C / 129 degrees F (Exhaust B)
Temperature 54 degrees C / 129 degrees F (TL0)
Temperature 52 degrees C / 125 degrees F (TQ0)
Temperature 61 degrees C / 141 degrees F (TL1)
Temperature 58 degrees C / 136 degrees F (TQ1)
Temperature 57 degrees C / 134 degrees F (TL2)
Temperature 58 degrees C / 136 degrees F (TQ2)
Temperature 62 degrees C / 143 degrees F (TL3)
Temperature 61 degrees C / 141 degrees F (TQ3)
Total CPU DRAM 2816 MB
Total SRAM 0 MB
Total SDRAM 0 MB
Start time 2012-01-12 12:05:42 PST
Uptime 3 hours, 14 minutes, 7 seconds
Slot 5 information:
State Online
Temperature 35 degrees C / 95 degrees F (PMB)
Temperature 34 degrees C / 93 degrees F (Intake)

```

```

Temperature 48 degrees C / 118 degrees F (Exhaust A)
Temperature 53 degrees C / 127 degrees F (Exhaust B)
Temperature 54 degrees C / 129 degrees F (TL0)
Temperature 52 degrees C / 125 degrees F (TQ0)
Temperature 69 degrees C / 156 degrees F (TL1)
Temperature 56 degrees C / 132 degrees F (TQ1)
Temperature 54 degrees C / 129 degrees F (TL2)
Temperature 56 degrees C / 132 degrees F (TQ2)
Temperature 59 degrees C / 138 degrees F (TL3)
Temperature 60 degrees C / 140 degrees F (TQ3)
Total CPU DRAM 2816 MB
Total SRAM 0 MB
Total SDRAM 0 MB
Start time 2012-01-12 12:05:43 PST
Uptime 3 hours, 14 minutes, 6 seconds
Slot 7 information:
State Online
Temperature 35 degrees C / 95 degrees F (PMB)
Temperature 33 degrees C / 91 degrees F (Intake)
Temperature 50 degrees C / 122 degrees F (Exhaust A)
Temperature 55 degrees C / 131 degrees F (Exhaust B)
Temperature 56 degrees C / 132 degrees F (TL0)
Temperature 56 degrees C / 132 degrees F (TQ0)
Temperature 61 degrees C / 141 degrees F (TL1)
Temperature 57 degrees C / 134 degrees F (TQ1)
Temperature 55 degrees C / 131 degrees F (TL2)
Temperature 59 degrees C / 138 degrees F (TQ2)
Temperature 62 degrees C / 143 degrees F (TL3)
Temperature 62 degrees C / 143 degrees F (TQ3)
Total CPU DRAM 2816 MB
Total SRAM 0 MB
Total SDRAM 0 MB
Start time 2012-01-12 12:05:44 PST
Uptime 3 hours, 14 minutes, 5 seconds

```

### show chassis fpc pic-status (PTX5000 Packet Transport Switch)

```

user@host> show chassis fpc pic-status
Slot 2 Online FPC
PIC 0 Online 24x 10GE(LAN) SFP+
PIC 1 Online 24x 10GE(LAN) SFP+
Slot 5 Online FPC
PIC 0 Online 24x 10GE(LAN) SFP+
PIC 1 Online 2x 40GE CFP
Slot 7 Online FPC
PIC 0 Online 24x 10GE(LAN) SFP+
PIC 1 Online 2x 40GE CFP

```

### show chassis fpc (ACX2000 Universal Access Router)

```

user@host> show chassis fpc

```

| Slot | State  | Temp<br>(C) | CPU Utilization (%)<br>Total Interrupt | Memory<br>DRAM (MB) | Utilization (%)<br>Heap Buffer |
|------|--------|-------------|----------------------------------------|---------------------|--------------------------------|
| 0    | Online | 61          | 17 6                                   | 512                 | 21 37                          |

### show chassis fpc 0 (ACX2000 Universal Access Router)

```

user@host> show chassis fpc 0

```

| Slot | State  | Temp<br>(C) | CPU Utilization (%)<br>Total Interrupt | Memory<br>DRAM (MB) | Utilization (%)<br>Heap Buffer |
|------|--------|-------------|----------------------------------------|---------------------|--------------------------------|
| 0    | Online | 61          | 17 6                                   | 512                 | 21 37                          |

**show chassis fpc detail (ACX2000 Universal Access Router)**

```
user@host> show chassis fpc detail
Slot 0 information:
 State Online
 Temperature 61 degrees C / 141 degrees F
 Total CPU DRAM 512 MB
 Start time 2012-05-29 02:52:06 PDT
 Uptime 27 minutes, 17 seconds
```

**show chassis fpc pic-status (ACX2000 Universal Access Router)**

```
user@host> show chassis fpc pic-status
Slot 0 Online
 PIC 0 Online 16x CHE1T1, RJ48
 PIC 1 Online 8x 1GE(LAN) RJ45
 PIC 2 Online 2x 1GE(LAN) SFP
 PIC 3 Online 2x 10GE(LAN) SFP+
```

**show chassis FPC 1 (MX Routers with Media Services Blade [MSB])**

```
user@switch> show chassis fpc 1
 Temp CPU Utilization (%) Memory Utilization (%)
Slot State (C) Total Interrupt DRAM (MB) Heap Buffer
 1 Online 34 5 0 3072 5 13
```

**show chassis FPC 1 detail (MX Routers with Media Services Blade [MSB])**

```
user@switch> show chassis fpc 1 detail
Slot 1 information:
 State Online
 Temperature 34
 Total CPU DRAM 3072 MB
 Total RLDRAM 259 MB
 Total DDR DRAM 4864 MB
 Start time: 2012-06-19 10:51:43 PDT
 Uptime: 16 minutes, 48 seconds
 Max Power Consumption 550 Watts
```

## show chassis hardware

---

|                                                             |                                                                                                                                                |
|-------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                               | show chassis hardware<br><detail   extensive><br><clei-models><br><models>                                                                     |
| <b>Syntax (EX Series)</b>                                   | show chassis hardware<br><clei-models><br><detail   extensive><br><models>                                                                     |
| <b>Syntax (T4000 Router)</b>                                | show chassis hardware<br><clei-models><br><detail   extensive><br><models>                                                                     |
| <b>Syntax (TX Matrix Router)</b>                            | show chassis hardware<br><clei-models><br><detail   extensive><br><models><br><lcc <i>number</i>   scc>                                        |
| <b>Syntax (TX Matrix Plus Router)</b>                       | show chassis hardware<br><clei-models><br><detail   extensive><br><models><br><lcc <i>number</i>   sfc <i>number</i> >                         |
| <b>Syntax (MX Series Routers and EX Series Switches)</b>    | show chassis hardware<br><detail   extensive><br><clei-models><br><models><br><all-members><br><local><br><member <i>member-id</i> >           |
| <b>Syntax (MX2010 and MX2020 3D Universal Edge Routers)</b> | show chassis hardware<br><clei-models><br><detail   extensive><br><models>                                                                     |
| <b>Syntax (QFX Series)</b>                                  | show chassis hardware<br><detail   extensive><br><clei-models><br><interconnect-device <i>name</i> ><br><node-device <i>name</i> ><br><models> |
| <b>Syntax (PTX Series Packet Transport Switches)</b>        | show chassis hardware<br><detail   extensive><br><clei-models><br><models>                                                                     |

**Syntax (ACX Series  
Universal Access  
Routers)**    `show chassis hardware`  
                  `<detail | extensive>`  
                  `<clei-models>`  
                  `<models>`

**Release Information**    Command introduced before Junos OS Release 7.4.  
                              **models** option introduced in Junos OS Release 8.2.  
                              Command introduced in Junos OS Release 9.0 for EX Series switches.  
                              **sfc** option introduced for the TX Matrix Plus router in Junos OS Release 9.6.  
                              Command introduced in Junos OS Release 11.1 for QFX Series.  
                              Command introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches.  
                              Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.  
                              Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.  
                              Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.

**Description**    Display a list of all Flexible PIC Concentrators (FPCs) and PICs installed in the router or switch chassis, including the hardware version level and serial number.

In the EX Series switch command output, FPC refers to the following:

- On EX2200 switches, EX3200 switches, EX4200 standalone switches, and EX4500 switches—Refers to the switch; FPC **number** is always 0.
- On EX4200 switches in a Virtual Chassis configuration—Refers to the member of a Virtual Chassis; FPC **number** equals the member ID, from 0 through 9.
- On EX8208 and EX8216 switches—Refers to a line card; FPC **number** equals the slot number for the line card.

On a QFX3500 standalone switch, both the FPC and FPC **number** are always 0.

On Type 5 FPC on T4000 routers, there are no **top temperature sensor** or **bottom temperature sensor** parameters. Instead, **fan intake temperature sensor** and **fan exhaust temperature sensors** parameters are displayed.

**Options**    **none**—Display information about hardware. For a TX Matrix router, display information about the TX Matrix router and its attached T640 routers. For a TX Matrix Plus router, display information about the TX Matrix Plus router and its attached T1600 routers.

**clei-models**—(Optional) Display Common Language Equipment Identifier (CLEI) barcode and model number for orderable field-replaceable units (FRUs).

**detail**—(Optional) Include RAM and disk information in output.

**extensive**—(Optional) Display ID EEPROM information.

**all-members**—(MX Series routers and EX Series switches only) (Optional) Display hardware-specific information for all the members of the Virtual Chassis configuration.

**interconnect-device name**—(QFabric systems only) (Optional) Display hardware-specific information for the Interconnect device.

**lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display hardware information for a specified T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display hardware information for a specified T1600 router (or line-card chassis) that is connected to the TX Matrix Plus router. Replace *number* with a value from 0 through 3.

**local**—(MX Series routers and EX Series switches only) (Optional) Display hardware-specific information for the local Virtual Chassis members.

**member *member-id***—(MX Series routers and EX Series switches only) (Optional) Display hardware-specific information for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**models**—(Optional) Display model numbers and part numbers for orderable FRUs and, for components that use ID EEPROM format v2, the CLEI code.

**node-device *name***—(QFabric systems only) (Optional) Display hardware-specific information for the Node device.

**scc**—(TX Matrix router only) (Optional) Display hardware information for the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus router only) (Optional) Display hardware information for the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

**Additional Information** The **show chassis hardware detail** command now displays DIMM information for the following Routing Engines:

**Table 319: Routing Engines Displaying DIMM Information**

| Routing Engines             | Routers                         |
|-----------------------------|---------------------------------|
| RE-S-1800x2 and RE-S-1800x4 | MX240, MX480, and MX960 routers |
| RE-A-1800x2                 | M120 and M320 routers           |

**Required Privilege Level** view

**Related Documentation**

- *show chassis power*

**List of Sample Output**

- [show chassis hardware \(EX8216 Switch\) on page 4334](#)
- [show chassis hardware clei-models \(EX8216 Switch\) on page 4335](#)
- [show chassis hardware clei-models \(T1600 Router\) on page 4336](#)
- [show chassis hardware detail \(EX4200 Switch\) on page 4337](#)
- [show chassis hardware models \(EX4500 Switch\) on page 4337](#)
- [show chassis hardware \(J6350 Router\) on page 4337](#)
- [show chassis hardware \(J6300 Router\) on page 4337](#)
- [show chassis hardware \(M7i Router\) on page 4338](#)
- [show chassis hardware \(M10 Router\) on page 4338](#)
- [show chassis hardware models \(M10 Router\) on page 4339](#)

[show chassis hardware \(M20 Router\) on page 4339](#)  
[show chassis hardware models \(M20 Router\) on page 4340](#)  
[show chassis hardware \(M40 Router\) on page 4340](#)  
[show chassis hardware \(M40e Router\) on page 4341](#)  
[show chassis hardware \(M120 Router\) on page 4341](#)  
[show chassis hardware detail \(M120 Router\) on page 4342](#)  
[show chassis hardware models \(M120 Router\) on page 4343](#)  
[show chassis hardware \(M160 Router\) on page 4344](#)  
[show chassis hardware models \(M160 Router\) on page 4344](#)  
[show chassis hardware detail \(M160 Router\) on page 4345](#)  
[show chassis hardware \(M320 Router\) on page 4346](#)  
[show chassis hardware models \(M320 Router\) on page 4347](#)  
[show chassis hardware \(MX5 Router\) on page 4347](#)  
[show chassis hardware \(MX10 Router\) on page 4348](#)  
[show chassis hardware \(MX40 Router\) on page 4349](#)  
[show chassis hardware \(Fixed MX80 Router\) on page 4349](#)  
[show chassis hardware \(Modular MX80 Router\) on page 4350](#)  
[show chassis hardware \(MX240 Router\) on page 4350](#)  
[show chassis hardware detail \(MX240 Router with RE Displaying DIMM information\) on page 4351](#)  
[show chassis hardware \(MX240 Router with Enhanced MX SCB\) on page 4351](#)  
[show chassis hardware \(MX480 Router\) on page 4352](#)  
[show chassis hardware \(MX480 Router with Enhanced MX SCB\) on page 4353](#)  
[show chassis hardware \(MX960 Router\) on page 4353](#)  
[show chassis hardware \(MX960 Router with Bidirectional Optics\) on page 4353](#)  
[show chassis hardware \(MX960 Router with Enhanced MX SCB\) on page 4354](#)  
[show chassis hardware models \(MX960 Router with Enhanced MX SCB\) on page 4356](#)  
[show chassis hardware detail \(MX960 Router\) on page 4357](#)  
[show chassis hardware \(MX2010 Router\) on page 4357](#)  
[show chassis hardware detail \(MX2010 Router\) on page 4359](#)  
[show chassis hardware extensive \(MX2010 Router\) on page 4364](#)  
[show chassis hardware models \(MX2010 Router\) on page 4370](#)  
[show chassis hardware clei-models \(MX2010 Routers\) on page 4370](#)  
[show chassis hardware \(MX2020 Router\) on page 4371](#)  
[show chassis hardware detail \(MX2020 Router\) on page 4380](#)  
[show chassis hardware models \(MX2020 Router\) on page 4388](#)  
[show chassis hardware clei-models \(MX2020 Router\) on page 4390](#)  
[show chassis hardware \(MX Series Routers with ATM MIC\) on page 4391](#)  
[show chassis hardware \(MX240, MX480, MX960 Routers with Application Services Modular Line Card\) on page 4392](#)  
[show chassis hardware extensive \(MX240, MX480, MX960 Routers with Application Services Modular Line Card\) on page 4392](#)  
[show chassis hardware \(MX480 Router with MPC4E\) on page 4393](#)  
[show chassis hardware \(MX2020 Router with MPC4E\) on page 4394](#)  
[show chassis hardware \(T320 Router\) on page 4396](#)  
[show chassis hardware \(T640 Router\) on page 4397](#)  
[show chassis hardware models \(T640 Router\) on page 4397](#)  
[show chassis hardware extensive \(T640 Router\) on page 4398](#)  
[show chassis hardware \(T4000 Router\) on page 4399](#)



[show chassis hardware \(T4000 Router with 16 GB line card chassis \(LCC\) Routing Engine\) on page 4401](#)  
[show chassis hardware clei-models \(T4000 Router\) on page 4401](#)  
[show chassis hardware detail \(T4000 Router\) on page 4402](#)  
[show chassis hardware models \(T4000 Router\) on page 4404](#)  
[show chassis hardware lcc \(TX Matrix Router\) on page 4404](#)  
[show chassis hardware scc \(TX Matrix Router\) on page 4405](#)  
[show chassis hardware \(T1600 Router\) on page 4405](#)  
[show chassis hardware \(TX Matrix Plus Router\) on page 4408](#)  
[show chassis hardware sfc \(TX Matrix Plus Router\) on page 4412](#)  
[show chassis hardware extensive \(TX Matrix Plus Router\) on page 4414](#)  
[show chassis hardware clei-models \(TX Matrix Plus Router\) on page 4415](#)  
[show chassis hardware detail \(TX Matrix Plus Router\) on page 4417](#)  
[show chassis hardware models \(TX Matrix Plus Router\) on page 4419](#)  
[show chassis hardware \(16-Port 10-Gigabit Ethernet MPC with SFP+ Optics \[MX Series Routers\]\) on page 4422](#)  
[show chassis hardware \(MPC3E \[MX Series Routers\]\) on page 4422](#)  
[show chassis hardware \(QFX3500 Switches\) on page 4424](#)  
[show chassis hardware detail \(QFX3500 Switches\) on page 4424](#)  
[show chassis hardware models \(QFX3500 Switches\) on page 4425](#)  
[show chassis hardware clei-models \(QFX3500 Switches\) on page 4425](#)  
[show chassis hardware interconnect-device \(QFabric Systems\) on page 4425](#)  
[show chassis hardware node-device \(QFabric Systems\) on page 4426](#)  
[show chassis hardware \(PTX5000 Packet Transport Switch\) on page 4426](#)  
[show chassis hardware clei-models \(PTX5000 Packet Transport Switch\) on page 4427](#)  
[show chassis hardware detail \(PTX5000 Packet Transport Switch\) on page 4428](#)  
[show chassis hardware models \(PTX5000 Packet Transport Switch\) on page 4429](#)  
[show chassis hardware extensive \(PTX5000 Packet Transport Switch\) on page 4430](#)  
[show chassis hardware \(MX Routers with Media Services Blade \[MSB\]\) on page 4431](#)  
[show chassis hardware extensive \(MX Routers with Media Services Blade \[MSB\]\) on page 4431](#)

**Output Fields** [Table 60 on page 649](#) lists the output fields for the **show chassis hardware** command. Output fields are listed in the approximate order in which they appear.

Table 320: show chassis hardware Output Fields

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output  |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Item</b>                   | <p>Chassis component:</p> <ul style="list-style-type: none"> <li>(EX Series switches)—Information about the chassis, Routing Engine (SRE and Routing Engine modules in EX8200 switches), power supplies, fan trays, and LCD panel. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs). Information about the backplane, midplane, and SIBs (SF modules) is displayed for EX8200 switches. See <i>EX Series Switches Hardware and CLI Terminology Mapping</i>.</li> <li>(MX Series routers and EX Series switches)—Information about the backplane, Routing Engine, Power Entry Modules (PEMs), and fan trays. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs), Modular Port Concentrators (MPCs) and associated Modular Interface Cards (MICs), or Dense Port Concentrators (DPCs). MX80 routers have a single Routing Engine and a built-in Packet Forwarding Engine that attaches directly to MICs. The Packet Forwarding Engine has two “pseudo” FPCs (FPC 0 and FPC1). MX80 routers also have a Forwarding Engine Board (FEB).</li> <li>(M Series routers, except for the M320 router)—Information about the backplane; power supplies; fan trays; Routing Engine; maxicab (the connection between the Routing Engine and the backplane, for the M40 router only); SCB, SSB, SFM, or FEB; MCS and PCG (for the M160 router only); each FPC and PIC; and each fan, blower, and impeller.</li> <li>(M120, M320, and T Series routers)—Information about the backplane, power supplies, fan trays, midplane, FPM (craft interface), CIP, PEM, SCG, CB, FPC, PIC, SFP, SPMB, and SIB.</li> <li>(QFX Series)—Information about the chassis, Routing Engine, power supplies, fan trays, Interconnect devices, and Node devices. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs).</li> <li>(PTX Series)—Information about the chassis, midplane, craft interface (FPM), power distribution units (PDUs) and Power Supply Modules (PSMs), Centralized Clock Generators (CCGs), Routing Engines, Control Boards (CBs) and Switch Processor Mezzanine Boards (SPMBs), Flexible PIC Concentrators (FPCs), PICs, Switch Interface Boards (SIBs), and fan trays (vertical and horizontal).</li> <li>(MX2010 and MX2020 routers)—Information about the chassis, midplane, craft interface (FPM), power midplane (PMP), Power Supply Modules (PSMs), Power Distribution Modules (PDMs), Routing Engines, Control Boards (CBs) and Switch Processor Mezzanine Boards (SPMBs), Switch Fabric Boards (SFBs), Flexible PIC Concentrators (FPCs), PICs, adapter cards (ADCs) and fan trays.</li> </ul> | All levels       |
| <b>Version</b>                | Revision level of the chassis component.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | All levels       |
| <b>Part number</b>            | Part number of the chassis component.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels       |
| <b>Serial number</b>          | Serial number of the chassis component. The serial number of the backplane is also the serial number of the router chassis. Use this serial number when you need to contact Juniper Networks Customer Support about the router or switch chassis.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels       |
| <b>Assb ID or Assembly ID</b> | ( <b>extensive</b> keyword only) Identification number that describes the FRU hardware.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>extensive</b> |

Table 320: show chassis hardware Output Fields (*continued*)

| Field Name       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output  |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| Assembly Version | ( <b>extensive</b> keyword only) Version number of the FRU hardware.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>extensive</b> |
| Assembly Flags   | ( <b>extensive</b> keyword only) Flags.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>extensive</b> |
| FRU model number | ( <b>clei-models</b> , <b>extensive</b> , and <b>models</b> keyword only) Model number of the FRU hardware component.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | none specified   |
| CLEI code        | ( <b>clei-models</b> and <b>extensive</b> keyword only) Common Language Equipment Identifier code. This value is displayed only for hardware components that use ID EEPROM format v2. This value is not displayed for components that use ID EEPROM format v1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | none specified   |
| EEPROM Version   | ID EEPROM version used by the hardware component: <b>0x00</b> (version 0), <b>0x01</b> (version 1), or <b>0x02</b> (version 2).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>extensive</b> |
| Description      | <p>Brief description of the hardware item:</p> <ul style="list-style-type: none"> <li>• Type of power supply.</li> <li>• Type of PIC. If the PIC type is not supported on the current software release, the output states <b>Hardware Not Supported</b>.</li> <li>• Type of FPC: <b>FPC Type 1</b>, <b>FPC Type 2</b>, <b>FPC Type 3</b>, <b>FPC Type 4</b>, or <b>FPC TypeOC192</b>.</li> </ul> <p>On EX Series switches, a brief description of the FPC.</p> <p>On the J Series routers, the FPC type corresponds to the Physical Interface Module (PIM). The following list shows the PIM abbreviation in the output and the corresponding PIM name.</p> <ul style="list-style-type: none"> <li>• <b>2x FE</b>—Either two built-in Fast Ethernet interfaces (fixed PIM) or dual-port Fast Ethernet PIM</li> <li>• <b>4x FE</b>—4-port Fast Ethernet ePIM</li> <li>• <b>1x GE Copper</b>—Copper Gigabit Ethernet ePIM (one 10-Mbps, 100-Mbps, or 1000-Mbps port)</li> <li>• <b>1x GE SFP</b>—SFP Gigabit Ethernet ePIM (one fiber port)</li> <li>• <b>4x GE Base PIC</b>—Four built-in Gigabit Ethernet ports on a J4350 or J6350 chassis (fixed PIM)</li> <li>• <b>2x Serial</b>—Dual-port serial PIM</li> <li>• <b>2x T1</b>—Dual-port T1 PIM</li> <li>• <b>2x E1</b>—Dual-port E1 PIM</li> <li>• <b>2x CTIE1</b>—Dual-port channelized T1/E1 PIM</li> <li>• <b>1x T3</b>—T3 PIM (one port)</li> <li>• <b>1x E3</b>—E3 PIM (one port)</li> <li>• <b>4x BRI S/T</b>—4-port ISDN BRI S/T PIM</li> <li>• <b>4x BRI U</b>—4-port ISDN BRI U PIM</li> <li>• <b>1x ADSL Annex A</b>—ADSL 2/2+ Annex A PIM (one port, for POTS)</li> <li>• <b>1x ADSL Annex B</b>—ADSL 2/2+ Annex B PIM (one port, for ISDN)</li> <li>• <b>1x TGM550</b>—TGM550 Telephony Gateway Module (Avaya VoIP gateway module with one console port, two analog <b>LINE</b> ports, and two analog <b>TRUNK</b> ports)</li> </ul> | All levels       |

Table 320: show chassis hardware Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
|            | <ul style="list-style-type: none"> <li>• <b>1x DS1 TIM510</b>—TIM510 E1/T1 Telephony Interface Module (Avaya VoIP media module with one E1 or T1 trunk termination port and ISDN PRI backup)</li> <li>• <b>4x FXS, 4x FX0, TIM514</b>—TIM514 Analog Telephony Interface Module (Avaya VoIP media module with four analog <b>LINE</b> ports and four analog <b>TRUNK</b> ports)</li> <li>• <b>4x BRI TIM521</b>—TIM521 BRI Telephony Interface Module (Avaya VoIP media module with four ISDN BRI ports)</li> <li>• <b>Crypto Accelerator Module</b>—For enhanced performance of cryptographic algorithms used in IP Security (IPsec) services</li> <li>• <b>MPC M 16x 10GE</b>—16-port 10-Gigabit Module Port Concentrator that supports SFP+ optical transceivers. (Not on EX Series switches.)</li> <li>• For hosts, the Routing Engine type.</li> <li>• For small form-factor pluggable transceiver (SFP) modules, the type of fiber: <b>LX</b>, <b>SX</b>, <b>LH</b>, or <b>T</b>.</li> <li>• LCD description for EX Series switches (except EX2200 switches).</li> <li>• <b>MPC2</b>—1-port MPC2 that supports two separate slots for MICs.</li> <li>• <b>MPC3E</b>—1-port MPC3E that supports two separate slots for MICs (MIC-3D-1X100GE-CFP and MIC-3D-20GE-SFP) on MX960, MX480, and MX240 routers. The MPC3E maps one MIC to one PIC (1 MIC, 1 PIC), which differs from the mapping of legacy MPCs.</li> <li>• 100GBASE-LR4, pluggable CFP optics</li> <li>• Supports the Enhanced MX Switch Control Board with fabric redundancy and existing SCBs without fabric redundancy.</li> <li>• Interoperates with existing MX Series line cards, including Flexible Port Concentrators (FPC), Dense Port Concentrators (DPCs), and Modular Port Concentrators (MPCs).</li> <li>• <b>MPC4E</b>—Fixed configuration MPC4E that is available in two flavors: MPC4E-3D-32XGE-SFP and MPC4E-3D-2CGE-8XGE) on MX2020, MX960, MX480, and MX240 routers.</li> <li>• LCD description for MX Series routers and EX Series switches</li> </ul> |                 |

## Sample Output

### show chassis hardware (EX8216 Switch)

```

user@host> show chassis hardware
Hardware inventory:
Item Version Part number Serial number Description
Chassis REV 06 CY0109220035 EX8216
Midplane REV 06 710-016845 BA0909120112 EX8216-MP
CB 0 REV 22 710-020771 AX0109197723 EX8216-RE320
CB 1 REV 22 710-020771 AX0109197726 EX8216-RE320
Routing Engine 1 BUILTIN BUILTIN RE-EX8216
FPC 3 REV 19 710-020683 BC0109083125 EX8200-48F
CPU REV 13 710-020598 BF0109144549 EX8200-CPU
FPC 4 REV 17 710-020683 BC0108500127 EX8200-48F
CPU REV 10 710-020598 BF0108460510 EX8200-CPU
PIC 0 BUILTIN BUILTIN 48x 100 Base-QFX/1000
Base-X
Xcvr 1 REV 01 740-011613 PE70V89 SFP-SX

```

|                 |        |            |              |              |
|-----------------|--------|------------|--------------|--------------|
| Xcvr 11         | REV 01 | 740-011613 | PE70YCE      | SFP-SX       |
| Xcvr 12         | REV 01 | 740-011613 | PE70VSH      | SFP-SX       |
| Xcvr 13         | REV 01 | 740-011613 | E08C02063    | SFP-SX       |
| Xcvr 14         | REV 01 | 740-011613 | PE70VKU      | SFP-SX       |
| Xcvr 15         | REV 01 | 740-011613 | E08E03372    | SFP-SX       |
| Xcvr 21         | REV 01 | 740-011613 | PE70VAD      | SFP-SX       |
| Xcvr 22         | REV 01 | 740-011613 | E08E01228    | SFP-SX       |
| Xcvr 23         | REV 01 | 740-011613 | PE70VSL      | SFP-SX       |
| Xcvr 24         | REV 01 | 740-011613 | E08E03409    | SFP-SX       |
| Xcvr 25         | REV 01 | 740-011613 | PE70VL4      | SFP-SX       |
| Xcvr 26         | REV 01 | 740-011613 | PDQ4L2Z      | SFP-SX       |
| Xcvr 27         | REV 01 | 740-011613 | PE70WFK      | SFP-SX       |
| Xcvr 28         | REV 01 | 740-011782 | PBD2B5U      | SFP-SX       |
| Xcvr 29         | REV 01 | 740-011613 | PE70UQX      | SFP-SX       |
| Xcvr 30         | REV 01 | 740-011613 | PE70VL5      | SFP-SX       |
| Xcvr 31         | REV 01 | 740-011613 | PE70V0F      | SFP-SX       |
| Xcvr 32         | REV 01 | 740-011613 | E08C02052    | SFP-SX       |
| Xcvr 33         | REV 01 | 740-011613 | E08C02197    | SFP-SX       |
| Xcvr 34         | REV 01 | 740-011613 | PE70V0L      | SFP-SX       |
| Xcvr 35         | REV 01 | 740-011613 | E08E03390    | SFP-SX       |
| Xcvr 36         | REV 01 | 740-011613 | PDQ4VL9      | SFP-SX       |
| Xcvr 37         | REV 01 | 740-011613 | E08E03370    | SFP-SX       |
| Xcvr 38         | REV 01 | 740-011613 | E08E03362    | SFP-SX       |
| Xcvr 39         | REV 01 | 740-011613 | E08C02065    | SFP-SX       |
| Xcvr 40         | REV 01 | 740-011613 | E08E03405    | SFP-SX       |
| Xcvr 41         | REV 01 | 740-011613 | E08E03411    | SFP-SX       |
| Xcvr 43         | REV 01 | 740-011613 | E08C02171    | SFP-SX       |
| Xcvr 45         | REV 01 | 740-011613 | E08E03410    | SFP-SX       |
| FPC 13          | REV 16 | 710-016837 | BB0109051344 | EX8200-8XS   |
| CPU             |        |            |              |              |
| SIB 0           | REV 10 | 710-021613 | AY0109166244 | EX8216-SF320 |
| SIB 1           | REV 10 | 710-021613 | AY0109166357 | EX8216-SF320 |
| SIB 2           | REV 10 | 710-021613 | AY0109166362 | EX8216-SF320 |
| SIB 3           | REV 10 | 710-021613 | AY0109166338 | EX8216-SF320 |
| SIB 4           | REV 10 | 710-021613 | AY0109166350 | EX8216-SF320 |
| SIB 5           | REV 10 | 710-021613 | AY0109166365 | EX8216-SF320 |
| SIB 6           | REV 10 | 710-021613 | AY0109166361 | EX8216-SF320 |
| SIB 7           | REV 10 | 710-021613 | AY0109166399 | EX8216-SF320 |
| PSU 0           | REV 17 | 740-021466 | BG0709170003 | EX8200-AC2K  |
| PSU 1           | REV 17 | 740-021466 | BG0709170004 | EX8200-AC2K  |
| PSU 2           | REV 17 | 740-021466 | BG0709170020 | EX8200-AC2K  |
| PSU 3           | REV 17 | 740-021466 | BG0709170017 | EX8200-AC2K  |
| PSU 4           | REV 17 | 740-021466 | BG0709170008 | EX8200-AC2K  |
| PSU 5           | REV 17 | 740-021466 | BG0709170018 | EX8200-AC2K  |
| Top Fan Tray    |        |            |              |              |
| FTC 0           | REV 4  | 760-022620 | CX1209140212 | EX8216-FT    |
| FTC 1           | REV 4  | 760-022620 | CX1209140212 | EX8216-FT    |
| Bottom Fan Tray |        |            |              |              |
| FTC 0           | REV 4  | 760-022620 | CX1209140211 | EX8216-FT    |
| FTC 1           | REV 4  | 760-022620 | CX1209140211 | EX8216-FT    |
| LCD 0           | REV 04 | 710-025742 | CE0109186919 | EX8200 LCD   |

### show chassis hardware clei-models (EX8216 Switch)

```

user@host> show chassis hardware clei-models
Hardware inventory:
Item Version Part number CLEI code FRU model number
Midplane REV 08 710-016845
PSU 0 REV 05 740-023002 COUPAEAEAA EX8200-PWR-AC3KR
PSU 1 REV 05 740-023002 COUPAEAEAA EX8200-PWR-AC3KR
PSU 2 REV 05 740-023002 COUPAEAEAA EX8200-PWR-AC3KR

```

```

PSU 3 REV 05 740-023002 COUPAEAEAA EX8200-PWR-AC3KR
PSU 4 REV 05 740-023002 COUPAEAEAA EX8200-PWR-AC3KR
PSU 5 REV 05 740-023002 COUPAEAEAA EX8200-PWR-AC3KR
Top Fan Tray
Bottom Fan Tray

```

### show chassis hardware clei-models (T1600 Router)

```

user@host> show chassis hardware clei-models
Hardware inventory:
Item Version Part number CLEI code FRU model number
Midplane REV 03 710-005608
FPM Display REV 05 710-002897
CIP REV 06 710-002895
PEM 0 Rev 07 740-017906 IPUPAC7KTA PWR-T1600-3-80-DC-S
PEM 1 Rev 18 740-002595 PWR-T-DC-S
SCG 0 REV 15 710-003423 SCG-T-S
Routing Engine 0 REV 08 740-014082 RE-A-2000-4096-S
Routing Engine 1 REV 07 740-014082 RE-A-2000-4096-S
CB 0 REV 05 710-007655 CB-T-S
CB 1 REV 03 710-017707 CB-T-S
FPC 0 REV 07 710-013558 T640-FPC2-E2
 PIC 0 REV 01 750-010618 PB-4GE-SFP
 PIC 1 REV 06 750-001900 PB-10C48-SON-SMSR
 PIC 2 REV 14 750-001901 PB-40C12-SON-SMIR
 PIC 3 REV 07 750-001900 PB-10C48-SON-SMSR
FPC 1 REV 06 710-013553 T640-FPC1-E2
 PIC 0 REV 08 750-001072 P-1GE-SX
 PIC 1 REV 10 750-012266 PB-4GE-TYPE1-SFP-IQ2
 PIC 2 REV 22 750-005634 PB-1CHOC12SMIR-QPP
FPC 2
 PIC 0 REV 16 750-007141 PC-10GE-SFP
 PIC 1 REV 06 750-015217 PC-8GE-TYPE3-SFP-IQ2
 PIC 2 REV 05 750-004695 PC-TUNNEL
 PIC 3 REV 17 750-009553 PC-40C48-SON-SFP
FPC 3 REV 01 710-010154 T640-FPC3-E
 PIC 0 REV 07 750-012793 PC-1XGE-TYPE3-XFP-IQ2
 PIC 1 REV 25 750-007141 PC-10GE-SFP
 PIC 2 REV 17 750-009553 PC-40C48-SON-SFP
 PIC 3 REV 32 750-003700 PC-10C192-SON-VSR
FPC 4 REV 16 710-013037 T1600-FPC4-ES
 PIC 1 REV 06 750-034781 PD-1CE-CFP
FPC 5 REV 02 710-013037 T1600-FPC4-ES
 PIC 0 REV 16 750-012518 PD-40C192-SON-XFP
 PIC 1 REV 01 750-010850 PD-10C768-SON-SR
FPC 6 REV 14 710-013037 T1600-FPC4-ES
 PIC 0 REV 11 750-017405 PD-4XGE-XFP
 PIC 1 REV 13 750-017405 PD-4XGE-XFP
FPC 7 REV 09 710-007529 T640-FPC3
 PIC 0 REV 10 750-012793 PC-1XGE-TYPE3-XFP-IQ2
 PIC 1 REV 01 750-015217 PC-8GE-TYPE3-SFP-IQ2
 PIC 2 REV 01 750-015217 PC-8GE-TYPE3-SFP-IQ2
 PIC 3 REV 15 750-009450 PC-10C192-SON-SR2
SIB 0 REV 07 710-013074 SIB-I-T1600-S
SIB 1 REV 07 710-013074 SIB-I-T1600-S
SIB 2 REV 07 710-013074 SIB-I-T1600-S
SIB 3 REV 07 710-013074 SIB-I-T1600-S
SIB 4 REV 07 710-013074 SIB-I-T1600-S
Fan Tray 0 FANTRAY-T-S

```

```

Fan Tray 1
Fan Tray 2
FANTRAY-T-S
FAN-REAR-TX-T640-S

```

#### show chassis hardware detail (EX4200 Switch)

```

user@host> show chassis hardware detail
Hardware inventory:
Item Version Part number Serial number Description
Chassis
Routing Engine 0 REV 11 750-021256 BM0208327733 EX4200-24T, 8 POE
Routing Engine 0 REV 11 750-021256 BM0208327733 EX4200-24T, 8 POE
FPC 0
CPU BUILTIN BUILTIN FPC CPU
PIC 0
PIC 1 REV 03B 711-021270 AR0208162285 4x GE SFP
BRD REV 08 711-021264 AK0208328289 EX4200-24T, 8 POE
Power Supply 0 REV 03 740-020957 AT0508346354 PS 320W AC
Fan Tray

```

#### show chassis hardware models (EX4500 Switch)

```

user@host> show chassis hardware models
Hardware inventory:
Item Version Part number Serial number FRU model number
Routing Engine 0 REV 01 750-035700 GG0210271867 EX4500-40F-FB-C
FPC 0
PIC 0
Power Supply 1 REV 01 740-029654 H884FS00JC09 EX4500-PWR1-AC-FB

```

#### show chassis hardware (J6350 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item Version Part number Serial number Description
Chassis
Midplane REV 03 710-014593 NP1265 JSR6350
System IO REV 01 710-016210 NN9950 JX350 System IO
Crypto Module
Routing Engine REV 08 710-015273 NM6509 RE-J6350-3400
ad0 248 MB 256MB CKS 00102006C24A00000039 Compact
Flash
FPC 0
PIC 0
FPC 1 REV 06 750-010355 AI07030023 FPC
PIC 0
FPC 3 REV 06 750-011148 AJ06520151 FPC
PIC 0
FPC 6 REV 06 750-013492 NC4170 FPC
PIC 0
Power Supply 0

```

#### show chassis hardware (J6300 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item Version Part number Serial number Description
Chassis
Midplane REV 02.04 710-010001 CORE99570 J6300
System IO REV 02.00 710-010003 CORE100848 System IO board
Routing Engine RevX2.6 750-010006 IWGS40735390 RE-J.3
FPC 0
PIC 0

```

|       |         |            |          |                    |
|-------|---------|------------|----------|--------------------|
| FPC 1 | RevX2.0 | 750-011380 | N3960005 | FPC                |
| PIC 0 |         |            |          | 1xADSL pic Annex A |
| FPC 2 | RevX2.0 | 750-011380 | N3960002 | FPC                |
| PIC 0 |         |            |          | 1xADSL pic Annex B |
| FPC 3 | REV 03  | 750-010354 | N0780028 | FPC                |
| PIC 0 |         |            |          | 1x T3              |

**show chassis hardware (M7i Router)**

user@host&gt; show chassis hardware

Hardware inventory:

| Item           | Version | Part number | Serial number | Description             |
|----------------|---------|-------------|---------------|-------------------------|
| Chassis        |         |             | 31959         | M7i                     |
| Midplane       | REV 02  | 710-008761  | CA0209        | M7i Midplane            |
| Power Supply 0 | Rev 04  | 740-008537  | PD10272       | AC Power Supply         |
| Routing Engine | REV 01  | 740-008846  | 1000396803    | RE-5.0                  |
| CFEB           | REV 02  | 750-009492  | CA0166        | Internet Processor IIv1 |
| FPC 0          |         |             |               | E-FPC                   |
| PIC 0          | REV 04  | 750-003163  | HJ6416        | 1x G/E, 1000 BASE-SX    |
| PIC 1          | REV 04  | 750-003163  | HJ6423        | 1x G/E, 1000 BASE-SX    |
| PIC 2          | REV 04  | 750-003163  | HJ6421        | 1x G/E, 1000 BASE-SX    |
| PIC 3          | REV 02  | 750-003163  | HJ0425        | 1x G/E, 1000 BASE-SX    |
| FPC 1          |         |             |               | E-FPC                   |
| PIC 2          | REV 01  | 750-009487  | HM2275        | ASP - Integrated        |
| PIC 3          | REV 01  | 750-009098  | CA0142        | 2x F/E, 100 BASE-TX     |

Hardware inventory:

| Item           | Version | Part number | Serial number | Description           |
|----------------|---------|-------------|---------------|-----------------------|
| Chassis        |         |             | B1157         | M7i                   |
| Midplane       | REV 05  | 710-008761  | DM0840        | M7i Midplane          |
| Power Supply 0 | Rev 08  | 740-008537  | TE53755       | AC Power Supply       |
| Routing Engine | REV 07  | 740-011202  | 1000736567    | RE-850                |
| CFEB           | REV 09  | 750-010463  | DK6952        | Internet Processor II |
| FPC 0          |         |             |               | E-FPC                 |
| PIC 0          | REV 12  | 750-012838  | DL7993        | 4x 1GE(LAN), IQ2      |
| Xcvr 0         | REV 01  | 740-011614  | PD94TDJ       | SFP-LX10              |
| Xcvr 1         | REV 01  | 740-011615  | PAD5EER       | UNSUPPORTED           |
| Xcvr 2         | REV 01  | 740-011614  | PD94THU       | SFP-LX10              |
| Xcvr 3         |         | NON-JNPR    | PDC2E7A       | SFP-LX10              |
| PIC 1          | REV 03  | 750-023116  | JT0203        | 4x CHSTM1 SDH CE SFP  |
| Xcvr 0         | REV 01  | 740-012434  | AGT063832PS   | SFP-SR                |
| Xcvr 1         | REV 01  | 740-012434  | AGT063832LY   | SFP-SR                |
| Xcvr 3         | REV 01  | 740-016064  | C06J19018     | SFP-LR                |
| PIC 2          | REV 15  | 750-014895  | DM5757        | MultiServices 100     |
| PIC 3          | REV 01  | 750-025390  | JW9448        | 12x T1/E1 CE          |
| FPC 1          |         |             |               | E-FPC                 |
| PIC 2          |         | BUILTIN     | BUILTIN       | 1x Tunnel             |
| PIC 3          | REV 09  | 750-009099  | DM0899        | 1x G/E, 1000 BASE     |
| Xcvr 0         | REV 01  | 740-012434  | AGT07150HGJ   | UNSUPPORTED           |
| Fan Tray       |         |             |               | Rear Fan Tray         |

**show chassis hardware (M10 Router)**

user@host&gt; show chassis hardware

Hardware inventory:

| Item           | Version | Part number | Serial number    | Description |
|----------------|---------|-------------|------------------|-------------|
| Chassis        |         |             | 1122             | M10         |
| Midplane       | REV 1.1 | 710-001950  | S/N AC6626       |             |
| Power supply A | Rev 01  | 740-002497  | S/N LC36095      | AC          |
| Power supply B | Rev 01  | 740-002497  | S/N LC36100      | AC          |
| Display        | REV 1.2 | 710-001995  | S/N AC6656       |             |
| Host           |         |             | 18000005dfb3fb01 | teknon      |



|            |        |            |            |                       |
|------------|--------|------------|------------|-----------------------|
| FEB        | REV 01 | 710-001948 | S/N AC6632 | Internet Processor II |
| FPC 0      |        |            |            |                       |
| PIC 0      | REV 08 | 750-001072 | S/N AB2485 | 1x G/E, 1000 BASE-SX  |
| PIC 1      | REV 01 | 750-000613 | S/N AA1048 | 1x OC-12 SONET, SMIR  |
| FPC 1      |        |            |            |                       |
| Fan Tray 0 |        |            |            | FANTRAY-M10I-S        |
| Fan Tray 1 |        |            |            | FANTRAY-M10I-S        |

#### show chassis hardware models (M10 Router)

```
user@host> show chassis hardware models
```

Hardware inventory:

| Item             | Version | Part number | CLEI code | FRU model number  |
|------------------|---------|-------------|-----------|-------------------|
| Midplane         | REV 04  | 710-008920  |           | CHAS-MP-M10i-S    |
| Power Supply 0   | Rev 06  | 740-008537  |           | PWR-M10i-M7i-AC-S |
| Power Supply 1   | Rev 06  | 740-008537  |           | PWR-M10i-M7i-AC-S |
| HCM 0            | REV 03  | 710-010580  |           | HCM-M10i-S        |
| HCM 1            | REV 03  | 710-010580  |           | HCM-M10i-S        |
| Routing Engine 0 | REV 09  | 740-009459  |           | RE-400-256-S      |
| CFEB 0           | REV 05  | 750-010465  |           | FEB-M10i-M7i-S    |
| FPC 0            |         |             |           |                   |
| PIC 0            | REV 10  | 750-002971  |           | PE-40C3-SON-MM    |
| PIC 1            | REV 11  | 750-002992  |           | PE-4FE-TX         |
| PIC 2            | REV 03  | 750-002977  |           | PE-20C3-ATM-MM    |
| PIC 3            | REV 08  | 750-005724  |           | PE-20C3-ATM2-MM   |
| FPC 1            |         |             |           |                   |
| PIC 2            | REV 12  | 750-008425  |           | PE-AS             |
| PIC 3            | REV 13  | 750-005636  |           | PE-4CHDS3-QPP     |
| Fan Tray 0       |         |             |           | FANTRAY-M10I-S    |
| Fan Tray 1       |         |             |           | FANTRAY-M10I-S    |

#### show chassis hardware (M20 Router)

```
user@host> show chassis hardware
```

Hardware inventory:

| Item           | Version | Part number | Serial number    | Description           |
|----------------|---------|-------------|------------------|-----------------------|
| Chassis        |         |             | 20033            | M20                   |
| Backplane      | REV 07  | 710-001517  | S/N AA7940       |                       |
| Power supply B | Rev 01  | 740-001465  | S/N 000001       | AC                    |
| Display        | REV 02  | 710-001519  | S/N AA9704       |                       |
| Host 0         |         |             | 98000004f8f27501 | teknor                |
| SSB slot 0     | REV 01  | 710-001951  | S/N AD5905       | Internet Processor II |
| SSRAM bank 0   | REV 01  | 710-001385  | S00480           | 2 MB                  |
| SSRAM bank 1   | REV 01  | 710-001385  | S00490           | 2 MB                  |
| SSRAM bank 2   | REV 01  | 710-001385  | S001:?           | 2 MB                  |
| SSRAM bank 3   | REV 01  | 710-001385  | S00483           | 2 MB                  |
| SSB slot 1     | N/A     | N/A         | N/A              | Backup                |
| FPC 1          | REV 01  | 710-001292  | S/N AB7528       |                       |
| SSRAM          | REV 01  | 710-000077  | S/N 304209       | 1 MB                  |
| SDRAM bank 0   | REV 01  | 710-000099  | S/N 000603       | 64 MB                 |
| SDRAM bank 1   | REV 01  | 710-000099  | S/N 000414       | 64 MB                 |
| PIC 0          | REV 03  | 750-000612  | S/N AB8433       | 2x OC-3 ATM, MM       |
| PIC 1          | REV 01  | 750-000616  | S/N AA1168       | 1x OC-12 ATM, MM      |
| PIC 2          | REV 01  | 750-000613  | S/N AA1008       | 1x OC-12 SONET, SMIR  |
| PIC 3          | REV 01  | 750-002501  | S/N AD5810       | 4x E3                 |
| FPC 2          | REV 01  | 710-001292  | S/N AC0119       |                       |
| SSRAM          | REV 01  | 710-000077  | S/N 503241       | 1 MB                  |
| SDRAM bank 0   | REV 01  | 710-000099  | S/N 306835       | 64 MB                 |
| SDRAM bank 1   | REV 01  | 710-000099  | S/N 306832       | 64 MB                 |
| Fan Tray 0     |         |             |                  | Front Upper Fan Tray  |
| Fan Tray 1     |         |             |                  | Front Middle Fan Tray |

Fan Tray 2  
Fan Tray 3

Front Bottom Fan Tray  
Rear Fan Tray

### show chassis hardware models (M20 Router)

user@host> show chassis hardware models

Hardware inventory:

| Item             | Version | Part number | CLEI code | FRU model number |
|------------------|---------|-------------|-----------|------------------|
| Backplane        | REV 03  | 710-002334  |           | CHAS-MP-M20-S    |
| Power Supply A   | REV 06  | 740-001465  |           | PWR-M20-AC-S     |
| Display          | REV 04  | 710-001519  |           | CRAFT-M20-S      |
| Routing Engine 0 | REV 06  | 740-003239  |           | RE-333-768-S     |
| Routing Engine 1 | REV 06  | 740-003239  |           | RE-333-768-S     |
| SSB 0            | REV 02  | 710-001951  |           | SSB-E-M20        |
| SSB 1            | N/A     | N/A         |           |                  |
| FPC 0            | REV 03  | 710-003308  |           | FPC-E            |
| PIC 0            | REV 08  | 750-002303  |           | P-4FE-TX         |
| PIC 1            | REV 07  | 750-004745  |           | P-2MCD53         |
| PIC 2            | REV 03  | 750-002965  |           | PE-4CHDS3        |
| FPC 1            | REV 03  | 710-003308  |           | FPC-E            |
| PIC 0            | REV 03  | 750-002914  |           | P-20C3-ATM-MM    |
| Fan Tray 0       |         |             |           | FANTRAY-F-M20-S  |
| Fan Tray 1       |         |             |           | FANTRAY-F-M20-S  |
| Fan Tray 2       |         |             |           | FANTRAY-F-M20-S  |
| Fan Tray 3       |         |             |           | FANTRAY-R-M20-S  |

### show chassis hardware (M40 Router)

user@host> show chassis hardware

Hardware inventory:

| Item           | Version | Part number | Serial number | Description          |
|----------------|---------|-------------|---------------|----------------------|
| Backplane      | REV 02  | 710-000073  | S/N AA0053    |                      |
| Power supply A | Rev 2   | 740-000235  | S/N 000042    | DC                   |
| Maxicab        | REV X1  | 710-000229  | S/N AA0139    |                      |
| Minicab        | REV X1  | 710-000482  | S/N AA0201    |                      |
| Display        | REV 06  | 710-000150  | S/N AA0905    |                      |
| Host           |         |             |               | cpv5000              |
| SCB            | REV X1  | 710-000075  | S/N AA0158    | Internet Processor I |
| SSRAM bank 0   | REV 02  | 710-000077  | S/N AA2267    | 1 MB                 |
| SSRAM bank 1   | REV 02  | 710-000077  | S/N AA2270    | 1 MB                 |
| SSRAM bank 2   | REV 02  | 710-000077  | S/N AA2269    | 1 MB                 |
| SSRAM bank 3   | REV 02  | 710-000077  | S/N AA2268    | 1 MB                 |
| FPC 0          | REV 01  | 710-000175  | S/N AA0048    |                      |
| SSRAM          | REV 01  | 710-000077  | S/N AA2333    | 1 MB                 |
| SDRAM bank 0   | REV 01  | 710-000099  | S/N AA2332    | 64 MB                |
| SDRAM bank 1   | REV X1  | 710-000099  | S/N AA2337    | 64 MB                |
| PIC 0          | REV 04  | 750-000613  | S/N aa0343    | 1x OC-12 SONET, SMIR |
| PIC 1          | REV 04  | 750-000613  | S/N AA0379    | 1x OC-12 SONET, SMIR |
| PIC 2          | REV 04  | 750-000613  | S/N AA0377    | 1x OC-12 SONET, SMIR |
| PIC 3          | REV 04  | 750-000613  | S/N AA0378    | 1x Tunnel            |
| FPC 2          | REV 01  | 710-000175  | S/N AA0042    |                      |
| SSRAM          | REV 02  | 710-000077  | S/N AA2288    | 1 MB                 |
| SDRAM bank 0   | REV 01  | 710-000099  | S/N AA2331    | 64 MB                |
| SDRAM bank 1   | REV 01  | 710-000099  | S/N AA2330    | 64 MB                |
| PIC 0          | REV X1  | 750-000603  | S/N AA0143    | 4x OC-3 SONET, SMIR  |
| PIC 1          | REV X1  | 750-000615  | S/N AA0149    | 4x OC-3 SONET, MM    |
| PIC 2          | REV X1  | 750-000611  | S/N AA0148    | 4x OC-3 SONET, MM    |
| PIC 3          | REV 04  | 750-000613  | S/N AA0330    | 1x OC-12 SONET, SMIR |
| FPC 4          | REV 01  | 710-000175  | S/N AA0050    |                      |
| SSRAM          | REV 01  | 710-000077  | S/N AA2327    | 1 MB                 |
| SDRAM bank 0   | REV 01  | 710-000099  | S/N AA2329    | 64 MB                |

|              |        |            |            |                      |
|--------------|--------|------------|------------|----------------------|
| SDRAM bank 1 | REV 01 | 710-000099 | S/N AA2328 | 64 MB                |
| PIC 0        | REV 04 | 750-000613 | S/N AA0320 | 1x OC-12 SONET, SMIR |
| PIC 2        | REV 05 | 750-000616 | S/N AA1341 | 1x OC-12 ATM, MM     |
| PIC 3        | REV 08 | 750-001072 | S/N AB2462 | 1x G/E, 1000 BASE-SX |
| FPC 5        | REV 10 | 710-000175 | S/N AA7663 |                      |
| SSRAM        | REV 01 | 710-000077 | S/N 501590 | 1 MB                 |
| SDRAM bank 0 | REV 01 | 710-000099 | S/N 300949 | 64 MB                |
| SDRAM bank 1 | REV 01 | 710-000099 | S/N 300868 | 64 MB                |
| PIC 1        | REV 01 | 750-001323 | S/N AB1670 | 1x Tunnel            |

#### show chassis hardware (M40e Router)

```
user@host> show chassis hardware
```

Hardware inventory:

| Item        | Version | Part number | Serial number    | Description           |
|-------------|---------|-------------|------------------|-----------------------|
| Chassis     |         |             |                  | m40e                  |
| Midplane    | REV 01  | 710-005071  | AX3671           |                       |
| FPM CMB     | REV 03  | 710-001642  | AR9074           |                       |
| FPM Display | REV 03  | 710-001647  | AR7331           |                       |
| CIP         | REV 04  | 710-002649  | BB4449           |                       |
| PEM 0       | Rev 01  | 740-003787  | MC12364          | Power Entry Module    |
| PEM 1       | Rev 01  | 740-003787  | MC12383          | Power Entry Module    |
| PCG 0       | REV 07  | 710-001568  | AG1332           |                       |
| PCG 1       | REV 07  | 710-001568  | AR3789           |                       |
| Host 0      |         |             | 3e000007c8176601 | Present               |
| MCS 0       | REV 11  | 710-001226  | AN5813           |                       |
| SFM 0 SPP   | REV 07  | 710-001228  | AG4676           |                       |
| SFM 0 SPR   | REV 05  | 710-002189  | AE4735           | Internet Processor II |
| SFM 1 SPP   | REV 07  | 710-001228  | AP1347           |                       |
| SFM 1 SPR   | REV 05  | 710-002189  | BE0063           | Internet Processor II |
| FPC 0       | REV 01  | 710-011725  | BE0669           | M40e-EP-FPC Type 1    |
| CPU         | REV 01  | 710-004600  | BD9504           |                       |
| PIC 0       | REV 03  | 750-003737  | AY3991           | 4x G/E, 1000 BASE-SX  |
| FPC 1       | REV 01  | 710-005197  | BD9842           | M40e-FPC Type 2       |
| CPU         | REV 01  | 710-004600  | BB4869           |                       |
| PIC 0       | REV 07  | 750-001900  | AR8278           | 1x OC-48 SONET, SMSR  |
| FPC 2       | REV 02  | 710-005197  | BD9824           | M40e-FPC Type 2       |
| CPU         | REV 01  | 710-004600  | BD9531           |                       |
| PIC 0       | REV 03  | 750-003737  | AY3986           | 4x G/E, 1000 BASE-SX  |
| FPC 4       | REV 02  | 710-005078  | BE0664           | M40e-FPC Type 1       |
| CPU         | REV 01  | 710-004600  | BD9559           |                       |
| PIC 0       | REV 03  | 750-001894  | AG7963           | 1x G/E, 1000 BASE-SX  |
| PIC 2       | REV 01  | 750-002575  | AF2472           | 4x OC-3 SONET, SMIR   |
| FPC 6       | REV 02  | 710-005078  | BE0652           | M40e-FPC Type 1       |
| CPU         | REV 01  | 710-004600  | BD9607           |                       |
| PIC 0       | REV 02  | 750-002911  | AN2286           | 4x F/E, 100 BASE-TX   |
| PIC 2       | REV 01  | 750-002577  | AP6345           | 4x OC-3 SONET, MM     |

#### show chassis hardware (M120 Router)

```
user@host> show chassis hardware
```

Hardware inventory:

| Item             | Version | Part number | Serial number | Description           |
|------------------|---------|-------------|---------------|-----------------------|
| Chassis          |         |             | JN000054AC    | M120                  |
| Midplane         | REV 01  | 710-013667  | RB4170        | M120 Midplane         |
| FPM Board        | REV 02  | 710-011407  | CJ9186        | M120 FPM Board        |
| FPM Display      | REV 02  | 710-011405  | CJ9173        | M120 FPM Display      |
| FPM CIP          | REV 02  | 710-011410  | CJ9221        | M120 FPM CIP          |
| PEM 0            | Rev 05  | 740-011936  | RM28320       | AC Power Entry Module |
| PEM 1            | Rev 05  | 740-011936  | RM28321       | AC Power Entry Module |
| Routing Engine 0 | REV 03  | 740-014080  | 1000642883    | RE-A-1000             |

|            |                                |            |              |                         |
|------------|--------------------------------|------------|--------------|-------------------------|
| CB 0       | REV 03                         | 710-011403 | CM8346       | M120 Control Board      |
| CB 1       | REV 06                         | 710-011403 | CP6728       | M120 Control Board      |
| FPC 1      | REV 02                         | 710-015908 | CP6925       | M120 CFPC 10GE          |
| PIC 0      |                                | BUILTIN    | BUILTIN      | 1x 10GE(LAN/WAN) XFP    |
| Xcvr 0     | REV 01                         | 740-014279 | 62E204N00007 | XFP-10G-LR              |
| FPC 3      | REV 03                         | 710-011393 | CJ9234       | M120 FPC Type 2         |
| PIC 0      | REV 16                         | 750-008155 | NB5229       | 2x G/E IQ, 1000 BASE    |
| Xcvr 0     | REV 01                         | 740-011613 | P9F15JB      | SFP-SX                  |
| Xcvr 1     | REV 01                         | 740-007326 | P4Q0R9G      | SFP-SX                  |
| PIC 1      | REV 09                         | 750-007745 | CG4360       | 4x OC-3 SONET, SMIR     |
| PIC 2      | REV 16                         | 750-008155 | ND7787       | 2x G/E IQ, 1000 BASE    |
| Xcvr 0     | REV 01                         | 740-011613 | P9F12AS      | SFP-SX                  |
| Xcvr 1     | REV 01                         | 740-011613 | P9F1ALU      | SFP-SX                  |
| PIC 3      | REV 07                         | 750-011800 | JW1284       | 8x 1GE(LAN), IQ2        |
| Xcvr 0     | REV 01                         | 740-011613 | P9F1AM6      | SFP-SX                  |
| Xcvr 6     | REV 01                         | 740-011613 | P9F16NN      | SFP-SX                  |
| Xcvr 7     | REV 01                         | 740-011782 | P8C29Y7      | SFP-SX                  |
| Board B    | REV 02                         | 710-011395 | CN3754       | M120 FPC Mezz           |
| FPC 4      | REV 02                         | 710-011398 | CP6741       | M120 FPC Type 3         |
| PIC 0      | REV 16                         | 750-007141 | NB2855       | 10x 1GE(LAN), 1000 BASE |
| Xcvr 0     | REV 01                         | 740-011782 | P922A1F      | SFP-SX                  |
| Xcvr 1     | REV 01                         | 740-011782 | P922A16      | SFP-SX                  |
| Xcvr 2     | REV 01                         | 740-011782 | P922A0U      | SFP-SX                  |
| Xcvr 3     | REV 01                         | 740-011782 | P9229UZ      | SFP-SX                  |
| Xcvr 4     | REV 01                         | 740-009029 | P11JXWP      | SFP-LX                  |
| Xcvr 6     | REV 01                         | 740-011613 | P9F1ALW      | SFP-SX                  |
| FPC 5      | REV 01                         | 710-011388 | CJ9088       | M120 FPC Type 1         |
| PIC 0      | *** Hardware Not Supported *** |            |              |                         |
| PIC 1      | REV 05                         | 750-012052 | NB0410       | 1x CHOC3 IQ SONET, SMLR |
| PIC 2      | REV 01                         | 750-013167 | CM3824       | 4x CHDS3 IQ             |
| PIC 3      | REV 01                         | 750-010240 | CB5366       | 1x G/E SFP, 1000 BASE   |
| Board B    | REV 01                         | 710-011390 | CJ9103       | M120 FPC Mezz Board     |
| FEB 3      | REV 04                         | 710-011663 | CP6673       | M120 FEB                |
| FEB 4      | REV 04                         | 710-011663 | CJ9368       | M120 FEB                |
| FEB 5      | REV 04                         | 710-011663 | CJ9386       | M120 FEB                |
| Fan Tray 0 |                                |            |              | Front Top Fan Tray      |
| Fan Tray 1 |                                |            |              | Front Bottom Fan Tray   |
| Fan Tray 2 |                                |            |              | Rear Top Fan Tray       |
| Fan Tray 3 |                                |            |              | Rear Bottom Fan Tray    |

### show chassis hardware detail (M120 Router)

```

user@host> show chassis hardware detail
Hardware inventory:
Item Version Part number Serial number Description
Chassis JN000054AC M120
Midplane REV 01 710-013667 RB4170 M120 Midplane
FPM Board REV 02 710-011407 CJ9186 M120 FPM Board
FPM Display REV 02 710-011405 CJ9173 M120 FPM Display
FPM CIP REV 02 710-011410 CJ9221 M120 FPM CIP
PEM 0 Rev 05 740-011936 RM28320 AC Power Entry Module
PEM 1 Rev 05 740-011936 RM28321 AC Power Entry Module
Routing Engine 0 REV 03 740-014080 1000642883 RE-A-1000
 ad0 248 MB SILICONSYSTEMS INC 256M 126CT505S0763SC00110 Compact Flash
 ad2 38154 MB HTE541040G9SA00 MPBBTOX2HS2E3M Hard Disk
CB 0 REV 03 710-011403 CM8346 M120 Control Board
CB 1 REV 06 710-011403 CP6728 M120 Control Board
FPC 1 REV 02 710-015908 CP6925 M120 CFPC 10GE
PIC 0 BUILTIN BUILTIN BUILTIN 1x 10GE(LAN/WAN) XFP

```

|            |                                |            |              |                         |
|------------|--------------------------------|------------|--------------|-------------------------|
| Xcvr 0     | REV 01                         | 740-014279 | 62E204N00007 | XFP-10G-LR              |
| FPC 3      | REV 03                         | 710-011393 | CJ9234       | M120 FPC Type 2         |
| PIC 0      | REV 16                         | 750-008155 | NB5229       | 2x G/E IQ, 1000 BASE    |
| Xcvr 0     | REV 01                         | 740-011613 | P9F15JB      | SFP-SX                  |
| Xcvr 1     | REV 01                         | 740-007326 | P4Q0R9G      | SFP-SX                  |
| PIC 1      | REV 09                         | 750-007745 | CG4360       | 4x OC-3 SONET, SMIR     |
| PIC 2      | REV 16                         | 750-008155 | ND7787       | 2x G/E IQ, 1000 BASE    |
| Xcvr 0     | REV 01                         | 740-011613 | P9F12AS      | SFP-SX                  |
| Xcvr 1     | REV 01                         | 740-011613 | P9F1ALU      | SFP-SX                  |
| PIC 3      | REV 07                         | 750-011800 | JW1284       | 8x 1GE(LAN), IQ2        |
| Xcvr 0     | REV 01                         | 740-011613 | P9F1AM6      | SFP-SX                  |
| Xcvr 6     | REV 01                         | 740-011613 | P9F16NN      | SFP-SX                  |
| Xcvr 7     | REV 01                         | 740-011782 | P8C29Y7      | SFP-SX                  |
| Board B    | REV 02                         | 710-011395 | CN3754       | M120 FPC Mezz           |
| FPC 4      | REV 02                         | 710-011398 | CP6741       | M120 FPC Type 3         |
| PIC 0      | REV 16                         | 750-007141 | NB2855       | 10x 1GE(LAN), 1000 BASE |
| Xcvr 0     | REV 01                         | 740-011782 | P922A1F      | SFP-SX                  |
| Xcvr 1     | REV 01                         | 740-011782 | P922A16      | SFP-SX                  |
| Xcvr 2     | REV 01                         | 740-011782 | P922A0U      | SFP-SX                  |
| Xcvr 3     | REV 01                         | 740-011782 | P9229UZ      | SFP-SX                  |
| Xcvr 4     | REV 01                         | 740-009029 | P11JXWP      | SFP-LX                  |
| Xcvr 6     | REV 01                         | 740-011613 | P9F1ALW      | SFP-SX                  |
| FPC 5      | REV 01                         | 710-011388 | CJ9088       | M120 FPC Type 1         |
| PIC 0      | *** Hardware Not Supported *** |            |              |                         |
| PIC 1      | REV 05                         | 750-012052 | NB0410       | 1x CHOC3 IQ SONET, SMLR |
| PIC 2      | REV 01                         | 750-013167 | CM3824       | 4x CHDS3 IQ             |
| PIC 3      | REV 01                         | 750-010240 | CB5366       | 1x G/E SFP, 1000 BASE   |
| Board B    | REV 01                         | 710-011390 | CJ9103       | M120 FPC Mezz Board     |
| FEB 3      | REV 04                         | 710-011663 | CP6673       | M120 FEB                |
| FEB 4      | REV 04                         | 710-011663 | CJ9368       | M120 FEB                |
| FEB 5      | REV 04                         | 710-011663 | CJ9386       | M120 FEB                |
| Fan Tray 0 |                                |            |              | Front Top Fan Tray      |
| Fan Tray 1 |                                |            |              | Front Bottom Fan Tray   |
| Fan Tray 2 |                                |            |              | Rear Top Fan Tray       |
| Fan Tray 3 |                                |            |              | Rear Bottom Fan Tray    |

### show chassis hardware models (M120 Router)

```
user@host> show chassis hardware models
Hardware inventory:
```

| Item             | Version | Part number | CLEI code | FRU model number     |
|------------------|---------|-------------|-----------|----------------------|
| Midplane         | REV 01  | 710-013667  |           |                      |
| FPM CIP          | REV 02  | 710-011410  |           | CRAFT-M120-S         |
| PEM 0            | Rev 05  | 740-011936  |           | PWR-M120-AC-S        |
| PEM 1            | Rev 05  | 740-011936  |           | PWR-M120-AC-S        |
| Routing Engine 0 | REV 03  | 740-014080  |           | RE-A-1000-2048-S     |
| CB 0             | REV 03  | 710-011403  |           | CB-M120-S            |
| CB 1             | REV 06  | 710-011403  |           | CB-M120-S            |
| FPC 1            | REV 02  | 710-015908  |           | M120-cFPC-1XGE-XFP   |
| FPC 3            |         |             |           |                      |
| PIC 0            | REV 16  | 750-008155  |           | PB-2GE-SFP-QPP       |
| PIC 1            | REV 09  | 750-007745  |           | PC-40C3-SON-SMIR     |
| PIC 2            | REV 16  | 750-008155  |           | PB-2GE-SFP-QPP       |
| PIC 3            | REV 07  | 750-011800  |           | PB-8GE-TYPE2-SFP-IQ2 |
| FPC 4            |         |             |           |                      |
| PIC 0            | REV 16  | 750-007141  |           | PC-10GE-SFP          |
| FPC 5            |         |             |           |                      |
| PIC 1            | REV 05  | 750-012052  |           | PB-1CHOC3-SMIR-QPP   |
| PIC 2            | REV 01  | 750-013167  |           | PE-4CHDS3-QPP        |

|            |        |            |                 |
|------------|--------|------------|-----------------|
| PIC 3      | REV 01 | 750-010240 | PB-1GE-SFP      |
| Fan Tray 0 |        |            | FFANTRAY-M120-S |
| Fan Tray 1 |        |            | FFANTRAY-M120-S |
| Fan Tray 2 |        |            | RFANTRAY-M120-S |
| Fan Tray 3 |        |            | RFANTRAY-M120-S |

### show chassis hardware (M160 Router)

```
user@host> show chassis hardware
```

| Item        | Version | Part number | Serial number    | Description          |
|-------------|---------|-------------|------------------|----------------------|
| Chassis     |         |             | 101              | M160                 |
| Midplane    | REV 02  | 710-001245  | S/N AB4107       |                      |
| FPM CMB     | REV 01  | 710-001642  | S/N AA2911       |                      |
| FPM Display | REV 01  | 710-001647  | S/N AA2999       |                      |
| CIP         | REV 02  | 710-001593  | S/N AA9563       |                      |
| PEM 0       | Rev 01  | 740-001243  | S/N KJ35769      | DC                   |
| PEM 1       | Rev 01  | 740-001243  | S/N KJ35765      | DC                   |
| PCG 0       | REV 01  | 710-001568  | S/N AA9794       |                      |
| PCG 1       | REV 01  | 710-001568  | S/N AA9804       |                      |
| Host 1      |         |             | da000004f8d57001 | teknor               |
| MCS 1       | REV 03  | 710-001226  | S/N AA9777       |                      |
| SFM 0 SPP   | REV 04  | 710-001228  | S/N AA2975       |                      |
| SFM 0 SPR   | REV 02  | 710-001224  | S/N AA9838       | Internet Processor I |
| SFM 1 SPP   | REV 04  | 710-001228  | S/N AA2860       |                      |
| SFM 1 SPR   | REV 01  | 710-001224  | S/N AB0139       | Internet Processor I |
| FPC 0       | REV 03  | 710-001255  | S/N AA9806       | FPC Type 1           |
| CPU         | REV 02  | 710-001217  | S/N AA9590       |                      |
| PIC 1       | REV 05  | 750-000616  | S/N AA1527       | 1x OC-12 ATM, MM     |
| PIC 2       | REV 05  | 750-000616  | S/N AA1535       | 1x OC-12 ATM, MM     |
| PIC 3       | REV 01  | 750-000616  | S/N AA1519       | 1x OC-12 ATM, MM     |
| FPC 1       | REV 02  | 710-001611  | S/N AA9523       | FPC Type 2           |
| CPU         | REV 02  | 710-001217  | S/N AA9571       |                      |
| PIC 0       | REV 03  | 750-001900  | S/N AA9626       | 1x STM-16 SDH, SMIR  |
| PIC 1       | REV 01  | 710-002381  | S/N AD3633       | 2x G/E, 1000 BASE-SX |
| FPC 2       |         |             |                  | FPC Type OC192       |
| CPU         | REV 03  | 710-001217  | S/N AB3329       |                      |
| PIC 0       | REV 01  |             |                  | 1x OC-192 SM SR-2    |
| Fan Tray 0  |         |             |                  | Rear Bottom Blower   |
| Fan Tray 1  |         |             |                  | Rear Top Blower      |
| Fan Tray 2  |         |             |                  | Front Top Blower     |
| Fan Tray 3  |         |             |                  | Front Fan Tray       |

### show chassis hardware models (M160 Router)

```
user@host> show chassis hardware models
```

Hardware inventory:

| Item             | Version | Part number | CLEI code | FRU model number  |
|------------------|---------|-------------|-----------|-------------------|
| Midplane         | REV 03  | 710-009120  |           | CHAS-BP-M320-S    |
| FPM Display      | REV 02  | 710-009351  |           | CRAFT-M320-S      |
| CIP              | REV 03  | 710-005926  |           | CIP-M320-S        |
| PEM 2            | Rev X4  | 740-009148  |           | PWR-M-DC-S        |
| PEM 3            | Rev X4  | 740-009148  |           | PWR-M-DC-S        |
| Routing Engine 0 | REV 02  | 740-008883  |           | RE-1600-2048-S    |
| Routing Engine 1 | REV 02  | 740-008883  |           | RE-1600-2048-S    |
| FPC 0            | REV 02  | 710-010419  |           | M320-FPC1         |
| PIC 0            | REV 01  | 750-001323  |           | P-TUNNEL          |
| PIC 1            | REV 02  | 750-002987  |           | PE-10C12-SON-SMIR |
| PIC 2            | REV 04  | 750-001894  |           | PB-1GE-SX         |
| PIC 3            | REV 04  | 750-001896  |           | PB-10C12-SON-SMIR |
| FPC 1            | REV 02  | 710-010419  |           | M320-FPC1         |
| PIC 0            | REV 04  | 750-001894  |           | PB-1GE-SX         |

|            |        |            |                    |
|------------|--------|------------|--------------------|
| PIC 1      | REV 04 | 750-001894 | PB-1GE-SX          |
| PIC 3      | REV 03 | 750-001894 | PB-1GE-SX          |
| FPC 2      | REV 02 | 710-010419 | M320-FPC1          |
| PIC 0      | REV 10 | 750-005634 | PB-1CHOC12SMIR-QPP |
| PIC 1      | REV 10 | 750-005634 | PB-1CHOC12SMIR-QPP |
| PIC 2      | REV 07 | 750-005634 | PB-1CHOC12SMIR-QPP |
| PIC 3      | REV 07 | 750-005634 | PB-1CHOC12SMIR-QPP |
| PIC 1      | REV 10 | 750-005634 | PB-1CHOC12SMIR-QPP |
| PIC 2      | REV 07 | 750-005634 | PB-1CHOC12SMIR-QPP |
| PIC 3      | REV 07 | 750-005634 | PB-1CHOC12SMIR-QPP |
| FPC 3      |        |            |                    |
| PIC 0      | REV 03 | 750-001895 | PB-10C12-SON-MM    |
| PIC 1      | REV 04 | 750-001894 | PB-1GE-SX          |
| PIC 3      | REV 04 | 750-003141 | PB-1GE-SX-B        |
| FPC 4      | REV 02 | 710-010419 | M320-FPC1          |
| FPC 5      | REV 02 | 710-010419 | M320-FPC1          |
| FPC 6      | REV 02 | 710-010419 | M320-FPC1          |
| FPC 7      |        |            |                    |
| PIC 0      | REV 15 | 750-001901 | PB-40C12-SON-SMIR  |
| PIC 1      | REV 06 | 750-001900 | PB-10C48-SON-SMSR  |
| PIC 2      | REV 07 | 750-001900 | PB-10C48-SON-SMSR  |
| PIC 3      | REV 05 | 750-003737 | PB-4GE-SX          |
| SIB 0      | REV 03 | 710-009184 | SIB-M-S            |
| SIB 1      | REV 03 | 710-009184 | SIB-M-S            |
| SIB 2      | REV 03 | 710-009184 | SIB-M-S            |
| SIB 3      | REV 03 | 710-009184 | SIB-M-S            |
| Fan Tray 0 |        |            | FFANTRAY-M320-S    |
| Fan Tray 1 |        |            | FFANTRAY-M320-S    |
| Fan Tray 2 |        |            | RFANTRAY-M320-S    |

### show chassis hardware detail (M160 Router)

```

user@host> show chassis hardware detail
Hardware inventory:

```

| Item         | Version | Part number | Serial number    | Description          |
|--------------|---------|-------------|------------------|----------------------|
| Chassis      |         |             | 101              | M160                 |
| Midplane     | REV 02  | 710-001245  | S/N AB4107       |                      |
| FPM CMB      | REV 01  | 710-001642  | S/N AA2911       |                      |
| FPM Display  | REV 01  | 710-001647  | S/N AA2999       |                      |
| CIP          | REV 02  | 710-001593  | S/N AA9563       |                      |
| PEM 0        | Rev 01  | 740-001243  | S/N KJ35769      | DC                   |
| PEM 1        | Rev 01  | 740-001243  | S/N KJ35765      | DC                   |
| PCG 0        | REV 01  | 710-001568  | S/N AA9794       |                      |
| PCG 1        | REV 01  | 710-001568  | S/N AA9804       |                      |
| Host 1       |         |             | da000004f8d57001 | teknor               |
| MCS 1        | REV 03  | 710-001226  | S/N AA9777       |                      |
| SFM 0 SPP    | REV 04  | 710-001228  | S/N AA2975       |                      |
| SFM 0 SPR    | REV 02  | 710-001224  | S/N AA9838       | Internet Processor I |
| SSRAM bank 0 | REV 01  | 710-000077  | S/N 306456       | 1 MB                 |
| SSRAM bank 1 | REV 01  | 710-000077  | S/N 306474       | 1 MB                 |
| SSRAM bank 2 | REV 01  | 710-000077  | S/N 306388       | 1 MB                 |
| SSRAM bank 3 | REV 01  | 710-000077  | S/N 306392       | 1 MB                 |
| SFM 1 SPP    | REV 04  | 710-001228  | S/N AA2860       |                      |
| SFM 1 SPR    | REV 01  | 710-001224  | S/N AB0139       | Internet Processor I |
| SSRAM bank 0 | REV 01  | 710-000077  | S/N 302917       | 1 MB                 |
| SSRAM bank 1 | REV 01  | 710-000077  | S/N 302662       | 1 MB                 |
| SSRAM bank 2 | REV 01  | 710-000077  | S/N 302593       | 1 MB                 |
| SSRAM bank 3 | REV 01  | 710-000077  | S/N 100160       | 1 MB                 |
| FPC 0        | REV 03  | 710-001255  | S/N AA9806       | FPC Type 1           |
| CPU          | REV 02  | 710-001217  | S/N AA9590       |                      |
| SSRAM        | REV 01  | 710-000077  | S/N 302836       | 1 MB                 |

|           |        |            |            |                      |
|-----------|--------|------------|------------|----------------------|
| SDRAM 0   | REV 01 | 710-001196 | S00141     | 32 MB                |
| SDRAM 1   | REV 01 | 710-001196 | S0010;     | 32 MB                |
| SSRAM     | REV 01 | 710-000077 | S/N 302633 | 1 MB                 |
| SDRAM 0   | REV 01 | 710-001196 | S00143     | 32 MB                |
| SDRAM 1   | REV 01 | 710-001196 | S00115     | 32 MB                |
| SSRAM     | REV 01 | 710-000077 | S/N 302952 | 1 MB                 |
| SDRAM 0   | REV 01 | 710-001196 | S00135     | 32 MB                |
| SDRAM 1   | REV 01 | 710-001196 | S001=3     | 32 MB                |
| SSRAM     | REV 01 | 710-000077 | S/N 302892 | 1 MB                 |
| SDRAM 0   | REV 01 | 710-001196 | S000?6     | 32 MB                |
| SDRAM 1   | REV 01 | 710-001196 | S001=5     | 32 MB                |
| PIC 1     | REV 05 | 750-000616 | S/N AA1527 | 1x OC-12 ATM, MM     |
| PIC 2     | REV 05 | 750-000616 | S/N AA1535 | 1x OC-12 ATM, MM     |
| PIC 3     | REV 01 | 750-000616 | S/N AA1519 | 1x OC-12 ATM, MM     |
| FPC 1     | REV 02 | 710-001611 | S/N AA9523 | FPC Type 2           |
| CPU       | REV 02 | 710-001217 | S/N AA9571 |                      |
| SSRAM     | REV 01 | 710-000077 | S/N 306340 | 1 MB                 |
| SDRAM 0   | REV 01 | 710-001196 | S00012     | 32 MB                |
| SDRAM 1   | REV 01 | 710-001196 | S0001?     | 32 MB                |
| SSRAM     | REV 01 | 710-000077 | S/N 306454 | 1 MB                 |
| SDRAM 0   | REV 01 | 710-001196 | S00028     | 32 MB                |
| SDRAM 1   | REV 01 | 710-001196 | S0002?     | 32 MB                |
| SSRAM     | REV 01 | 710-000077 | S/N 306492 | 1 MB                 |
| SDRAM 0   | REV 01 | 710-001196 | S00015     | 32 MB                |
| SDRAM 1   | REV 01 | 710-001196 | S00031     | 32 MB                |
| SSRAM     | REV 01 | 710-000077 | S/N 306363 | 1 MB                 |
| SDRAM 0   | REV 01 | 710-001196 | S00013     | 32 MB                |
| SDRAM 1   | REV 01 | 710-001196 | S00032     | 32 MB                |
| PIC 0     | REV 03 | 750-001900 | S/N AA9626 | 1x STM-16 SDH, SMIR  |
| PIC 1     | REV 01 | 710-002381 | S/N AD3633 | 2x G/E, 1000 BASE-SX |
| FPC 2     |        |            |            | FPC Type OC192       |
| ... SSRAM | REV 01 | 710-000077 | S/N 306466 | 1 MB                 |

**show chassis hardware (M320 Router)**

```

user@host> show chassis hardware
Hardware inventory:
Item Version Part number Serial number Description
Chassis 67245 M320
Midplane REV 05 710-009120 RB1202 M320 Midplane
FPM GBUS REV 04 710-005928 HZ5697 M320 Board
FPM Display REV 05 710-009351 HR1464 M320 FPM Display
CIP REV 04 710-005926 HT8672 M320 CIP
PEM 0 Rev 05 740-009148 QK34208 DC Power Entry Module
PEM 1 Rev 05 740-009148 QK34262 DC Power Entry Module
PEM 2 Rev 05 740-009148 QF10449 DC Power Entry Module
PEM 3 Rev 05 740-009148 QJ18257 DC Power Entry Module
Routing Engine 0 REV 06 740-008883 P11123901185 RE-4.0
CB 0 REV 07 710-009115 JB2382 M320 Control Board
FPC 0 REV 02 710-005017 CD9926 M320 FPC Type 2
CPU REV 01 710-011659 CJ6940 M320 PCA SCPU
PIC 0 REV 07 750-001900 AT1594 1x OC-48 SONET, SMSR
PIC 1 REV 03 750-001850 HS2746 1x Tunnel
PIC 2 REV 05 750-010618 JE7117 4x G/E SFP, 1000 BASE
PIC 3 REV 06 750-001900 HE6083 1x OC-48 SONET, SMSR
FPC 2 REV 02 710-005017 CH0319 M320 FPC Type 1
CPU REV 01 710-011659 CJ6942 M320 PCA SCPU
PIC 0 REV 05 750-003034 BD8705 4x OC-3 SONET, SMIR
FPC 5 REV 02 710-005017 CD9938 M320 FPC Type 2
CPU
FPC 7 REV 02 710-005017 CD9934 M320 FPC Type 2

```



|            |        |            |        |                       |
|------------|--------|------------|--------|-----------------------|
| CPU        |        |            |        |                       |
| SIB 0      | REV 09 | 710-009184 | JA6540 | M320 SIB              |
| SIB 1      | REV 09 | 710-009184 | HV9511 | M320 SIB              |
| SIB 2      | REV 09 | 710-009184 | HW2057 | M320 SIB              |
| SIB 3      | REV 09 | 710-009184 | JA6687 | M320 SIB              |
| Fan Tray 0 |        |            |        | Front Top Fan Tray    |
| Fan Tray 1 |        |            |        | Front Bottom Fan Tray |
| Fan Tray 2 |        |            |        | Rear Fan Tray         |

### show chassis hardware models (M320 Router)

```
user@host> show chassis hardware models
Hardware inventory:
```

| Item             | Version | Part number | CLEI code | FRU model number   |
|------------------|---------|-------------|-----------|--------------------|
| Midplane         | REV 03  | 710-009120  |           | CHAS-BP-M320-S     |
| FPM Display      | REV 02  | 710-009351  |           | CRAFT-M320-S       |
| CIP              | REV 03  | 710-005926  |           | CIP-M320-S         |
| PEM 2            | Rev X4  | 740-009148  |           | PWR-M-DC-S         |
| PEM 3            | Rev X4  | 740-009148  |           | PWR-M-DC-S         |
| Routing Engine 0 | REV 02  | 740-008883  |           | RE-1600-2048-S     |
| Routing Engine 1 | REV 02  | 740-008883  |           | RE-1600-2048-S     |
| FPC 0            | REV 02  | 710-010419  |           | M320-FPC1          |
| PIC 0            | REV 01  | 750-001323  |           | P-TUNNEL           |
| PIC 1            | REV 02  | 750-002987  |           | PE-10C12-SON-SMIR  |
| PIC 2            | REV 04  | 750-001894  |           | PB-1GE-SX          |
| PIC 3            | REV 04  | 750-001896  |           | PB-10C12-SON-SMIR  |
| FPC 1            | REV 02  | 710-010419  |           | M320-FPC1          |
| PIC 0            | REV 04  | 750-001894  |           | PB-1GE-SX          |
| PIC 1            | REV 04  | 750-001894  |           | PB-1GE-SX          |
| PIC 3            | REV 03  | 750-001894  |           | PB-1GE-SX          |
| FPC 2            | REV 02  | 710-010419  |           | M320-FPC1          |
| PIC 0            | REV 10  | 750-005634  |           | PB-1CHOC12SMIR-QPP |
| PIC 1            | REV 10  | 750-005634  |           | PB-1CHOC12SMIR-QPP |
| PIC 2            | REV 07  | 750-005634  |           | PB-1CHOC12SMIR-QPP |
| PIC 3            | REV 07  | 750-005634  |           | PB-1CHOC12SMIR-QPP |
| PIC 1            | REV 10  | 750-005634  |           | PB-1CHOC12SMIR-QPP |
| PIC 2            | REV 07  | 750-005634  |           | PB-1CHOC12SMIR-QPP |
| PIC 3            | REV 07  | 750-005634  |           | PB-1CHOC12SMIR-QPP |
| FPC 3            |         |             |           |                    |
| PIC 0            | REV 03  | 750-001895  |           | PB-10C12-SON-MM    |
| PIC 1            | REV 04  | 750-001894  |           | PB-1GE-SX          |
| PIC 3            | REV 04  | 750-003141  |           | PB-1GE-SX-B        |
| FPC 4            | REV 02  | 710-010419  |           | M320-FPC1          |
| FPC 5            | REV 02  | 710-010419  |           | M320-FPC1          |
| FPC 6            | REV 02  | 710-010419  |           | M320-FPC1          |
| FPC 7            |         |             |           |                    |
| PIC 0            | REV 15  | 750-001901  |           | PB-40C12-SON-SMIR  |
| PIC 1            | REV 06  | 750-001900  |           | PB-10C48-SON-SMSR  |
| PIC 2            | REV 07  | 750-001900  |           | PB-10C48-SON-SMSR  |
| PIC 3            | REV 05  | 750-003737  |           | PB-4GE-SX          |
| SIB 0            | REV 03  | 710-009184  |           | SIB-M-S            |
| SIB 1            | REV 03  | 710-009184  |           | SIB-M-S            |
| SIB 2            | REV 03  | 710-009184  |           | SIB-M-S            |
| SIB 3            | REV 03  | 710-009184  |           | SIB-M-S            |
| Fan Tray 0       |         |             |           | FFANTRAY-M320-S    |
| Fan Tray 1       |         |             |           | FFANTRAY-M320-S    |
| Fan Tray 2       |         |             |           | RFANTRAY-M320-S    |

### show chassis hardware (MX5 Router)

```
user@host> show chassis hardware
```

## Hardware inventory:

| Item           | Version | Part number | Serial number | Description           |
|----------------|---------|-------------|---------------|-----------------------|
| Chassis        |         |             | E1368         | MX5-T                 |
| Midplane       | REV 01  | 711-038215  | YF5288        | MX5-T                 |
| PEM 0          | Rev 04  | 740-028288  | VA01215       | AC Power Entry Module |
| PEM 1          | Rev 04  | 740-028288  | VA01218       | AC Power Entry Module |
| Routing Engine |         | BUILTIN     | BUILTIN       | Routing Engine        |
| TFEB 0         |         | BUILTIN     | BUILTIN       | Forwarding Engine     |
| Processor      |         |             |               |                       |
| QXM 0          | REV 05  | 711-028408  | ZA9136        | MPC QXM               |
| FPC 0          |         | BUILTIN     | BUILTIN       | MPC BUILTIN           |
| MIC 0          |         | BUILTIN     | BUILTIN       | 4x 10GE XFP           |
| PIC 0          |         | BUILTIN     | BUILTIN       | 4x 10GE XFP           |
| FPC 1          |         | BUILTIN     | BUILTIN       | MPC BUILTIN           |
| MIC 0          | REV 24  | 750-028392  | YX9820        | 3D 20x 1GE(LAN) SFP   |
| PIC 0          |         | BUILTIN     | BUILTIN       | 10x 1GE(LAN) SFP      |
| Xcvr 0         | REV 01  | 740-031851  | AM1045SUAQ3   | SFP-SX                |
| Xcvr 1         | REV 01  | 740-031851  | AM1045SUAPA   | SFP-SX                |
| Xcvr 2         | REV 01  | 740-031851  | AM1045SUAN7   | SFP-SX                |
| Xcvr 3         | REV 01  | 740-031851  | AM1045SU91Q   | SFP-SX                |
| Xcvr 4         | REV 01  | 740-031851  | AM1045SUDDR   | SFP-SX                |
| Xcvr 9         | REV 01  | 740-011613  | AM0848SB6A1   | SFP-SX                |
| PIC 1          |         | BUILTIN     | BUILTIN       | 10x 1GE(LAN) SFP      |
| Xcvr 0         | REV 01  | 740-031851  | AM1045SUANO   | SFP-SX                |
| Xcvr 1         | REV 01  | 740-011613  | AS0812S0719   | SFP-SX                |
| Xcvr 2         | REV 01  | 740-011613  | AM0821SA121   | SFP-SX                |
| Xcvr 3         | REV 01  | 740-011613  | PF21K21       | SFP-SX                |
| Xcvr 4         | REV 01  | 740-011613  | AM0848SB69Z   | SFP-SX                |
| Xcvr 5         | REV 01  | 740-011782  | P9POXV3       | SFP-SX                |
| Xcvr 6         | REV 01  | 740-011613  | AM0812S8WJN   | SFP-SX                |
| Xcvr 7         | REV 01  | 740-011613  | PAM3G9Q       | SFP-SX                |
| Xcvr 8         | REV 01  | 740-011613  | AM0848SB4A6   | SFP-SX                |
| Xcvr 9         | REV 01  | 740-011782  | P9MOU37       | SFP-SX                |
| MIC 1          | REV 20  | 750-028380  | ZG2657        | 3D 2x 10GE XFP        |
| PIC 2          |         | BUILTIN     | BUILTIN       | 1x 10GE XFP           |
| PIC 3          |         | BUILTIN     | BUILTIN       | 1x 10GE XFP           |
| Fan Tray       |         |             |               | Fan Tray              |

## show chassis hardware (MX10 Router)

user@host&gt; show chassis hardware

## Hardware inventory:

| Item           | Version | Part number | Serial number | Description           |
|----------------|---------|-------------|---------------|-----------------------|
| Chassis        |         |             | E1372         | MX10-T                |
| Midplane       | REV 01  | 711-038211  | YF5285        | MX10-T                |
| PEM 0          | Rev 04  | 740-028288  | VB01678       | AC Power Entry Module |
| Routing Engine |         | BUILTIN     | BUILTIN       | Routing Engine        |
| TFEB 0         |         | BUILTIN     | BUILTIN       | Forwarding Engine     |
| Processor      |         |             |               |                       |
| QXM 0          | REV 05  | 711-028408  | ZA9053        | MPC QXM               |
| FPC 0          |         | BUILTIN     | BUILTIN       | MPC BUILTIN           |
| MIC 0          |         | BUILTIN     | BUILTIN       | 4x 10GE XFP           |
| PIC 0          |         | BUILTIN     | BUILTIN       | 4x 10GE XFP           |
| FPC 1          |         | BUILTIN     | BUILTIN       | MPC BUILTIN           |
| MIC 0          | REV 24  | 750-028392  | YX9436        | 3D 20x 1GE(LAN) SFP   |
| PIC 0          |         | BUILTIN     | BUILTIN       | 10x 1GE(LAN) SFP      |
| Xcvr 0         | REV 01  | 740-031851  | AM1107SUFQW   | SFP-SX                |
| PIC 1          |         | BUILTIN     | BUILTIN       | 10x 1GE(LAN) SFP      |
| Fan Tray       |         |             |               | Fan Tray              |

## show chassis hardware (MX40 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item Version Part number Serial number Description
Chassis E1367 MX40-T
Midplane REV 01 711-038211 YF5284 MX40-T
PEM 0 Rev 04 740-028288 VB01680 AC Power Entry Module
PEM 1 Rev 04 740-028288 VB01700 AC Power Entry Module
Routing Engine BUILTIN BUILTIN BUILTIN Routing Engine
TFEB 0 BUILTIN BUILTIN BUILTIN Forwarding Engine
Processor
 QXM 0 REV 05 711-028408 ZA9048 MPC QXM
 FPC 0 BUILTIN BUILTIN BUILTIN MPC BUILTIN
 MIC 0 BUILTIN BUILTIN BUILTIN 4x 10GE XFP
 PIC 0 BUILTIN BUILTIN BUILTIN 4x 10GE XFP
 Xcvr 0 REV 01 740-014279 M7067UPP XFP-10G-LR
 Xcvr 1 NON-JNPR K9J02UN XFP-10G-LR
 FPC 1 BUILTIN BUILTIN BUILTIN MPC BUILTIN
 MIC 0 REV 24 750-028392 YX3504 3D 20x 1GE(LAN) SFP
 PIC 0 BUILTIN BUILTIN BUILTIN 10x 1GE(LAN) SFP
 Xcvr 0 REV 01 740-011613 AM0812S8WTE SFP-SX
 Xcvr 1 REV 01 740-011613 PFA6KV2 SFP-SX
 Xcvr 2 REV 01 740-031851 AM1045SUDDM SFP-SX
 Xcvr 3 REV 01 740-011613 PD63C7M SFP-SX
 Xcvr 4 REV 01 740-011613 PD63DJY SFP-SX
 Xcvr 5 REV 02 740-011613 AA0950STLL9 SFP-SX
 Xcvr 6 REV 01 740-011782 PAR1YHC SFP-SX
 Xcvr 7 REV 01 740-011782 P9P0XXL SFP-SX
 Xcvr 8 REV 01 740-011613 PD63D95 SFP-SX
 Xcvr 9 REV 01 740-031851 AM1045SU9B8 SFP-SX
 PIC 1 BUILTIN BUILTIN BUILTIN 10x 1GE(LAN) SFP
 Xcvr 0 REV 01 740-011613 PF21L3Z SFP-SX
 Xcvr 1 REV 01 740-031851 AM1045SU7M9 SFP-SX
 Xcvr 2 REV 01 740-031851 AM1045SUAPT SFP-SX
 Xcvr 3 REV 01 740-011613 PFF2BZH SFP-SX
 Xcvr 4 REV 01 740-031851 AM1045SUDDN SFP-SX
 Xcvr 5 REV 01 740-031851 AM1039S00ZR SFP-SX
 Xcvr 6 REV 01 740-031851 AM1045SUD6Y SFP-SX
 Xcvr 8 REV 01 740-011613 PFM1QBS SFP-SX
 Xcvr 9 REV 01 740-011613 PFF2E25 SFP-SX
 MIC 1 REV 01 750-021130 KG4391 3D 2x 10GE XFP
 PIC 2 BUILTIN BUILTIN BUILTIN 1x 10GE XFP
 Xcvr 0 REV 01 740-011571 C645XJ04G XFP-10G-SR
 PIC 3 BUILTIN BUILTIN BUILTIN 1x 10GE XFP
 Xcvr 0 NON-JNPR CA49BK0AE XFP-10G-SR
Fan Tray

```

## show chassis hardware (Fixed MX80 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item Version Part number Serial number Description
Chassis MX80-48T
Midplane REV 01 711-031603 KF9250 MX80-48T
Routing Engine BUILTIN BUILTIN BUILTIN Routing Engine
FEB 0 BUILTIN BUILTIN BUILTIN Forwarding Engine Board
FPC 0 BUILTIN BUILTIN BUILTIN MPC BUILTIN
 MIC 0 BUILTIN BUILTIN BUILTIN 4x 10GE XFP
 PIC 0 BUILTIN BUILTIN BUILTIN 4x 10GE XFP
 Xcvr 0 NON-JNPR M6439D41 XFP-10G-LR

```

|          |        |            |              |                   |
|----------|--------|------------|--------------|-------------------|
| Xcvr 1   | REV 01 | 740-014279 | 6XE931N00202 | XFP-10G-LR        |
| Xcvr 2   | REV 01 | 740-014289 | C715XU05F    | XFP-10G-SR        |
| Xcvr 3   | REV 01 | 740-014289 | C650XU0EP    | XFP-10G-SR        |
| FPC 1    |        | BUILTIN    | BUILTIN      | MPC BUILTIN       |
| MIC 0    | REV 01 | 711-029399 | JR6981       | 12x 1GE(LAN) RJ45 |
| PIC 0    |        | BUILTIN    | BUILTIN      | 12x 1GE(LAN) RJ45 |
| PIC 1    |        | BUILTIN    | BUILTIN      | 12x 1GE(LAN) RJ45 |
| MIC 1    | REV 01 | BUILTIN    | BUILTIN      | 12x 1GE(LAN) RJ45 |
| PIC 2    |        | BUILTIN    | BUILTIN      | 12x 1GE(LAN) RJ45 |
| PIC 3    |        | BUILTIN    | BUILTIN      | 12x 1GE(LAN) RJ45 |
| Fan Tray |        |            |              | Fan Tray          |

### show chassis hardware (Modular MX80 Router)

```
user@host> show chassis hardware
```

```
Hardware inventory:
Item Version Part number Serial number Description
Chassis MX80
Midplane REV 02 711-031594 JR7084 MX80
PEM 0 Rev 01 740-028288 000018 AC Power Entry Module
Routing Engine Routing Engine
FEB 0 BUILTIN BUILTIN BUILTIN Forwarding Engine Board

QXM 0 REV 05 711-028408 JR7041 MPC QXM
FPC 0 BUILTIN BUILTIN BUILTIN MPC BUILTIN
MIC 0 BUILTIN BUILTIN BUILTIN 4x 10GE XFP
PIC 0 BUILTIN BUILTIN BUILTIN 4x 10GE XFP
FPC 1 BUILTIN BUILTIN BUILTIN MPC BUILTIN
MIC 0 REV 02 750-028380 JR6598 3D 2x 10GE XFP
PIC 0 BUILTIN BUILTIN BUILTIN 1x 10GE XFP
Xcvr 0 REV 01 740-014289 T07M86365 XFP-10G-SR
PIC 1 BUILTIN BUILTIN BUILTIN 1x 10GE XFP
Xcvr 0 REV 01 740-014289 T07M71094 XFP-10G-SR
MIC 1 REV 02 750-028380 JG8548 3D 2x 10GE XFP
PIC 2 BUILTIN BUILTIN BUILTIN 1x 10GE XFP
Xcvr 0 REV 02 740-014289 T08L86302 XFP-10G-SR
PIC 3 BUILTIN BUILTIN BUILTIN 1x 10GE XFP
Xcvr 0 REV 02 740-014289 C810XU0BA XFP-10G-SR
Fan Tray Fan Tray
```

### show chassis hardware (MX240 Router)

```
user@host> show chassis hardware
```

```
Hardware inventory:
Item Version Part number Serial number Description
Chassis MX240
Midplane REV 01 710-021041 TR1502 MX240 Backplane
FPM Board REV 01 710-017254 KD4017 Front Panel Display
PEM 0 Rev 02 740-017330 000332 PS 1.2-1.7kW; 100-240V
AC in
PEM 1 Rev 02 740-017330 000226 PS 1.2-1.7kW; 100-240V
AC in
Routing Engine 0 REV 06 740-013063 1000703522 RE-S-2000
Routing Engine 1 REV 06 740-015113 1000687625 RE-S-1300
CB 0 REV 07 710-013385 KC9057 MX SCB
CB 1 REV 05 710-013385 JY4760 MX SCB
FPC 1 REV 01 750-021679 KC7340 DPCE 40x 1GE R
CPU REV 06 710-013713 KD4078 DPC PMB
PIC 0 BUILTIN BUILTIN BUILTIN 10x 1GE(LAN)
Xcvr 0 REV 01 740-011613 P9F18ME SFP-SX
```

|            |        |            |         |                   |
|------------|--------|------------|---------|-------------------|
| PIC 1      |        | BUILTIN    | BUILTIN | 10x 1GE(LAN)      |
| PIC 2      |        | BUILTIN    | BUILTIN | 10x 1GE(LAN)      |
| PIC 3      |        | BUILTIN    | BUILTIN | 10x 1GE(LAN)      |
| FPC 2      | REV 04 | 710-016669 | JS4529  | DPCE 40x 1GE R EQ |
| CPU        | REV 06 | 710-013713 | KB3969  | DPC PMB           |
| PIC 0      |        | BUILTIN    | BUILTIN | 10x 1GE(LAN) EQ   |
| Xcvr 0     | REV 01 | 740-011613 | PBG3Y79 | SFP-SX            |
| Xcvr 1     | REV 01 | 740-011613 | PBG3XU8 | SFP-SX            |
| Xcvr 2     | REV 01 | 740-011613 | PBG3YG6 | SFP-SX            |
| Xcvr 3     | REV 01 | 740-011613 | PBG3XUG | SFP-SX            |
| Xcvr 4     | REV 01 | 740-011613 | PBG3XTJ | SFP-SX            |
| PIC 1      |        | BUILTIN    | BUILTIN | 10x 1GE(LAN) EQ   |
| Xcvr 0     | REV 01 | 740-011613 | PBG3ZUM | SFP-SX            |
| Xcvr 1     | REV 01 | 740-011613 | PBG3Y5H | SFP-SX            |
| Xcvr 2     | REV 01 | 740-011613 | PBG3UZT | SFP-SX            |
| Xcvr 3     | REV 01 | 740-011613 | PBG3US1 | SFP-SX            |
| PIC 2      |        | BUILTIN    | BUILTIN | 10x 1GE(LAN) EQ   |
| Xcvr 0     | REV 01 | 740-011613 | PBG3YG7 | SFP-SX            |
| Xcvr 1     | REV 01 | 740-011613 | PBG3XZ9 | SFP-SX            |
| Xcvr 2     | REV 01 | 740-011613 | PBG3XTY | SFP-SX            |
| Xcvr 3     | REV 01 | 740-011613 | PBG3UZG | SFP-SX            |
| PIC 3      |        | BUILTIN    | BUILTIN | 10x 1GE(LAN) EQ   |
| Xcvr 0     | REV 01 | 740-011613 | PBG3Y8W | SFP-SX            |
| Xcvr 1     | REV 01 | 740-011613 | PBG3YVX | SFP-SX            |
| Xcvr 2     | REV 01 | 740-011613 | PBG3YB3 | SFP-SX            |
| Xcvr 3     | REV 01 | 740-011613 | PBG43VQ | SFP-SX            |
| Fan Tray 0 | REV 01 | 710-021113 | JS4642  | MX240 Fan Tray    |

#### show chassis hardware detail (MX240 Router with RE Displaying DIMM information)

```
user@host> show chassis hardware detail
```

| Item             | Version                            | Part number | Serial number        | Description             |
|------------------|------------------------------------|-------------|----------------------|-------------------------|
| Chassis          |                                    |             | JN11279B4AFC         | MX240 Backplane         |
| Midplane         | REV 07                             | 760-021404  | TS2474               | MX240 Backplane         |
| FPM Board        | REV 03                             | 760-021392  | XC2643               | Front Panel Display     |
| PEM 0            | Rev 03                             | 740-017343  | QCS0908A068          | DC Power Entry Module   |
| Routing Engine 0 | REV 01                             | 740-031117  | AARCH00              | RE-S-1800x4             |
| ad0 3764 MB      | STEC M2+                           | CF 9.0.2    | STIM2Q3209239145303  | Removable Compact Flash |
| ad1 28626 MB     | WDC SSD-F0030S-5000                |             | C933Z036237215548S00 | Compact Flash           |
| usb0 (addr 1)    | EHCI root hub 0                    |             | Intel                | uhub0                   |
| usb0 (addr 2)    | product 0x0020 32                  |             | vendor 0x8087        | uhub1                   |
| DIMM 0           | VL31B5263E-F8S DIE REV-0 PCB REV-0 |             |                      | MFR ID-ce80             |
| DIMM 1           | VL31B5263E-F8S DIE REV-0 PCB REV-0 |             |                      | MFR ID-ce80             |
| DIMM 2           | VL31B5263E-F8S DIE REV-0 PCB REV-0 |             |                      | MFR ID-ce80             |
| DIMM 3           | SL31B5263E-F8S DIE REV-0 PCB REV-0 |             |                      | MFR ID-ce80             |
| CB 0             | REV 03                             | 710-021523  | XD7225               | MX SCB                  |
| Fan Tray 0       | REV 01                             | 710-021113  | WZ4986               | MX240 Fan Tray          |

#### show chassis hardware (MX240 Router with Enhanced MX SCB)

```
user@host> show chassis hardware
```

Hardware inventory:

| Item      | Version | Part number | Serial number | Description            |
|-----------|---------|-------------|---------------|------------------------|
| Chassis   |         |             | JN10C7F7EAFC  | MX240                  |
| Midplane  | REV 01  | 710-021041  | TR1502        | MX240 Backplane        |
| FPM Board | REV 01  | 710-017254  | KD4017        | Front Panel Display    |
| PEM 0     | Rev 02  | 740-017330  | 000332        | PS 1.2-1.7kW; 100-240V |
| AC in     |         |             |               |                        |
| PEM 1     | Rev 02  | 740-017330  | 000226        | PS 1.2-1.7kW; 100-240V |
| AC in     |         |             |               |                        |

|                  |        |            |            |                   |
|------------------|--------|------------|------------|-------------------|
| Routing Engine 0 | REV 06 | 740-013063 | 1000703522 | RE-S-2000         |
| Routing Engine 1 | REV 06 | 740-015113 | 1000687625 | RE-S-1300         |
| CB 0             | REV 02 | 710-031391 | YE8494     | Enhanced MX SCB   |
| CB 1             | REV 05 | 710-031391 | YOP5764    | Enhanced MX SCB   |
| FPC 1            | REV 01 | 750-021679 | KC7340     | DPCE 40x 1GE R    |
| CPU              | REV 06 | 710-013713 | KD4078     | DPC PMB           |
| PIC 0            |        | BUILTIN    | BUILTIN    | 10x 1GE(LAN)      |
| Xcvr 0           | REV 01 | 740-011613 | P9F18ME    | SFP-SX            |
| PIC 1            |        | BUILTIN    | BUILTIN    | 10x 1GE(LAN)      |
| PIC 2            |        | BUILTIN    | BUILTIN    | 10x 1GE(LAN)      |
| PIC 3            |        | BUILTIN    | BUILTIN    | 10x 1GE(LAN)      |
| FPC 2            | REV 04 | 710-016669 | JS4529     | DPCE 40x 1GE R EQ |
| CPU              | REV 06 | 710-013713 | KB3969     | DPC PMB           |
| PIC 0            |        | BUILTIN    | BUILTIN    | 10x 1GE(LAN) EQ   |
| Xcvr 0           | REV 01 | 740-011613 | PBG3Y79    | SFP-SX            |
| Xcvr 1           | REV 01 | 740-011613 | PBG3XU8    | SFP-SX            |
| Xcvr 2           | REV 01 | 740-011613 | PBG3YG6    | SFP-SX            |
| Xcvr 3           | REV 01 | 740-011613 | PBG3XUG    | SFP-SX            |
| Xcvr 4           | REV 01 | 740-011613 | PBG3XTJ    | SFP-SX            |
| PIC 1            |        | BUILTIN    | BUILTIN    | 10x 1GE(LAN) EQ   |
| Xcvr 0           | REV 01 | 740-011613 | PBG3ZUM    | SFP-SX            |
| Xcvr 1           | REV 01 | 740-011613 | PBG3Y5H    | SFP-SX            |
| Xcvr 2           | REV 01 | 740-011613 | PBG3UZT    | SFP-SX            |
| Xcvr 3           | REV 01 | 740-011613 | PBG3US1    | SFP-SX            |
| PIC 2            |        | BUILTIN    | BUILTIN    | 10x 1GE(LAN) EQ   |
| Xcvr 0           | REV 01 | 740-011613 | PBG3YG7    | SFP-SX            |
| Xcvr 1           | REV 01 | 740-011613 | PBG3XZ9    | SFP-SX            |
| Xcvr 2           | REV 01 | 740-011613 | PBG3XTY    | SFP-SX            |
| Xcvr 3           | REV 01 | 740-011613 | PBG3UZG    | SFP-SX            |
| PIC 3            |        | BUILTIN    | BUILTIN    | 10x 1GE(LAN) EQ   |
| Xcvr 0           | REV 01 | 740-011613 | PBG3Y8W    | SFP-SX            |
| Xcvr 1           | REV 01 | 740-011613 | PBG3YVX    | SFP-SX            |
| Xcvr 2           | REV 01 | 740-011613 | PBG3YB3    | SFP-SX            |
| Xcvr 3           | REV 01 | 740-011613 | PBG43VQ    | SFP-SX            |
| Fan Tray 0       | REV 01 | 710-021113 | JS4642     | MX240 Fan Tray    |

#### show chassis hardware (MX480 Router)

```
user@host> show chassis hardware
```

| Hardware inventory: |         |             |               |                        |
|---------------------|---------|-------------|---------------|------------------------|
| Item                | Version | Part number | Serial number | Description            |
| Chassis             |         |             | JN10C7F7FAFB  | MX480                  |
| Midplane            | REV 04  | 710-017414  | TR2071        | MX480 Midplane         |
| FPM Board           | REV 02  | 710-017254  | KB8459        | Front Panel Display    |
| PEM 0               | Rev 02  | 740-017330  | QCS07519029   | PS 1.2-1.7kW; 100-240V |
| AC in               |         |             |               |                        |
| PEM 1               | Rev 02  | 740-017330  | QCS07519041   | PS 1.2-1.7kW; 100-240V |
| AC in               |         |             |               |                        |
| PEM 2               | Rev 02  | 740-017330  | QCS07519097   | PS 1.2-1.7kW; 100-240V |
| AC in               |         |             |               |                        |
| Routing Engine 0    | REV 07  | 740-013063  | 1000733381    | RE-S-2000              |
| Routing Engine 1    | REV 07  | 740-013063  | 1000733540    | RE-S-2000              |
| CB 0                | REV 07  | 710-013385  | KA8022        | MX SCB                 |
| CB 1                | REV 07  | 710-013385  | KA8303        | MX SCB                 |
| FPC 0               | REV 09  | 750-020452  | KA8660        | DPCE 40x 1GE X EQ      |
| CPU                 | REV 06  | 710-013713  | KA8185        | DPC PMB                |
| PIC 0               |         | BUILTIN     | BUILTIN       | 10x 1GE(LAN) EQ        |
| PIC 1               |         | BUILTIN     | BUILTIN       | 10x 1GE(LAN) EQ        |
| PIC 2               |         | BUILTIN     | BUILTIN       | 10x 1GE(LAN) EQ        |

|          |         |         |                 |
|----------|---------|---------|-----------------|
| PIC 3    | BUILTIN | BUILTIN | 10x 1GE(LAN) EQ |
| Fan Tray |         |         | Left Fan Tray   |

#### show chassis hardware (MX480 Router with Enhanced MX SCB)

```
user@host> show chassis hardware
Hardware inventory:
Item Version Part number Serial number Description
Chassis JN10C7F7FAFB MX480
Midplane REV 04 710-017414 TR2071 MX480 Midplane
FPM Board REV 02 710-017254 KB8459 Front Panel Display
PEM 0 Rev 02 740-017330 QCS07519029 PS 1.2-1.7kW; 100-240V
AC in
PEM 1 Rev 02 740-017330 QCS07519041 PS 1.2-1.7kW; 100-240V
AC in
PEM 2 Rev 02 740-017330 QCS07519097 PS 1.2-1.7kW; 100-240V
AC in
Routing Engine 0 REV 07 740-013063 1000733381 RE-S-2000
Routing Engine 1 REV 07 740-013063 1000733540 RE-S-2000
CB 0 REV 07 710-013385 KA8022 Enhanced MX SCB
CB 1 REV 07 710-013385 KA8303 Enhanced MX SCB
FPC 0 REV 09 750-020452 KA8660 DPCE 40x 1GE X EQ
CPU REV 06 710-013713 KA8185 DPC PMB
PIC 0 BUILTIN BUILTIN 10x 1GE(LAN) EQ
PIC 1 BUILTIN BUILTIN 10x 1GE(LAN) EQ
PIC 2 BUILTIN BUILTIN 10x 1GE(LAN) EQ
PIC 3 BUILTIN BUILTIN 10x 1GE(LAN) EQ
Fan Tray
```

#### show chassis hardware (MX960 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item Version Part number Serial number Description
Chassis MX960
Midplane REV 01 710-013698 AA6082 MX960 Midplane
PIM Rev 01 740-013110 000008 Power Inlet Module
PEM 2
PEM 3 Rev 01 740-013682 000038 PS 1.7kW; 200-240VAC in
Routing Engine 0 REV 00 740-015113 1000617944 RE-S-1300
CB 0 REV 05 710-013725 JK6947 MX960 Test SCB
FPC 4 REV 01 710-013305 JM7617 MX960 Test DPC
CPU
PIC 0 BUILTIN BUILTIN 1x 10GE(LAN/WAN)
PIC 1 BUILTIN BUILTIN 10x 1GE
FPC 7 REV 01 710-013305 JL9634 MX960 Test DPC
CPU
PIC 0 BUILTIN BUILTIN 1x 10GE(LAN/WAN)
Xcvr 0 NON-JNPR MYBG65I82C XFP-10G-SR
PIC 1 BUILTIN BUILTIN 10x 1GE
Xcvr 1 REV 01 740-011782 P7N0368 SFP-SX
Xcvr 4 REV 01 740-011782 P8J1W27 SFP-SX
Xcvr 6 REV 01 740-011782 P8J1VSD SFP-SX
Xcvr 9 REV 01 740-011782 P8J1W25 SFP-SX
Fan Tray 0
Fan Tray 1
```

#### show chassis hardware (MX960 Router with Bidirectional Optics)

```
user@host> show chassis hardware
Hardware inventory:
Item Version Part number Serial number Description
```

|                  |        |            |              |                           |
|------------------|--------|------------|--------------|---------------------------|
| Chassis          |        |            | JN10BA5B9AFA | MX960                     |
| Midplane         | REV 03 | 710-013698 | TR0234       | MX960 Backplane           |
| FPM Board        | REV 03 | 710-014974 | JA0878       | Front Panel Display       |
| PDM              | Rev 03 | 740-013110 | QCS11135028  | Power Distribution Module |
| PEM 0            | Rev 03 | 740-013682 | QCS11154036  | PS 1.7kW; 200-240VAC in   |
| PEM 1            | Rev 03 | 740-013682 | QCS11154010  | PS 1.7kW; 200-240VAC in   |
| PEM 2            | Rev 03 | 740-013682 | QCS11154022  | PS 1.7kW; 200-240VAC in   |
| Routing Engine 0 | REV 06 | 740-013063 | 1000691458   | RE-S-2000                 |
| CB 0             | REV 07 | 710-013385 | KA2190       | MX SCB                    |
| CB 1             | REV 07 | 710-013385 | KA0837       | MX SCB                    |
| FPC 3            | REV 02 | 750-018122 | KB3890       | DPCE 40x 1GE R            |
| CPU              |        |            |              |                           |
| FPC 4            | REV 01 | 750-018122 | KB3889       | DPCE 40x 1GE R            |
| CPU              | REV 06 | 710-013713 | KB3976       | DPC PMB                   |
| PIC 0            |        | BUILTIN    | BUILTIN      | 10x 1GE(LAN)              |
| Xcvr 1           | REV 01 | 740-020426 | 4910549      | SFP-1000BASE-BX40-D       |
| Xcvr 2           | REV 01 | 740-020426 | 4910551      | SFP-1000BASE-BX40-D       |
| Xcvr 5           | REV 01 | 740-021340 | 77E245N00006 | SFP-1000BASE-BX10-U       |
| Xcvr 6           | REV 01 | 740-020425 | 4882821      | SFP-1000BASE-BX40-U       |
| Xcvr 8           | REV 01 | 740-020425 | 4882820      | SFP-1000BASE-BX40-U       |
| PIC 1            |        | BUILTIN    | BUILTIN      | 10x 1GE(LAN)              |
| Xcvr 0           | REV 01 | 740-020465 | 77E555N00894 | SFP-1000BASE-BX10-D       |
| Xcvr 1           | REV 01 | 740-020465 | 75E467X00818 | SFP-1000BASE-BX10-D       |
| Xcvr 2           | REV 01 | 740-020465 | 75E467X00573 | SFP-1000BASE-BX10-D       |
| Xcvr 3           | REV 01 | 740-020465 | 4888227      | SFP-1000BASE-BX10-D       |
| Xcvr 4           | REV 01 | 740-020465 | 4888241      | SFP-1000BASE-BX10-D       |
| Xcvr 5           | REV 01 | 740-021340 | 77E245N00005 | SFP-1000BASE-BX10-U       |
| Xcvr 6           | REV 01 | 740-021340 | 76E245X00487 | SFP-1000BASE-BX10-U       |
| Xcvr 7           | REV 01 | 740-021341 | 5255889      | SFP-1000BASE-BX10-U       |
| Xcvr 8           | REV 01 | 740-021341 | 5255887      | SFP-1000BASE-BX10-U       |
| Xcvr 9           | REV 01 | 740-021340 | 77E245N00004 | SFP-1000BASE-BX10-U       |
| PIC 2            |        | BUILTIN    | BUILTIN      | 10x 1GE(LAN)              |
| Xcvr 0           | REV 01 | 740-020424 | 5007582      | SFP-1000BASE-BX10-D       |
| Xcvr 1           | REV 01 | 740-020424 | 4888187      | SFP-1000BASE-BX10-D       |
| Xcvr 2           | REV 01 | 740-020424 | 4656500      | SFP-1000BASE-BX10-D       |
| Xcvr 5           | REV 01 | 740-021341 | 5255886      | SFP-1000BASE-BX10-U       |
| Xcvr 7           | REV 01 | 740-021340 | 77E245N00003 | SFP-1000BASE-BX10-U       |
| Xcvr 8           | REV 01 | 740-021341 | 5255888      | SFP-1000BASE-BX10-U       |
| PIC 3            |        | BUILTIN    | BUILTIN      | 10x 1GE(LAN)              |
| Xcvr 0           | REV 01 | 740-017726 | 74S184H30341 | SFP-EX                    |
| Xcvr 1           | REV 01 | 740-017726 | 4814061      | SFP-EX                    |
| Xcvr 5           | REV 01 | 740-017726 | 6ZS184H31108 | SFP-EX                    |
| Xcvr 9           | REV 01 | 740-021340 | 76E245X00486 | SFP-1000BASE-BX10-U       |
| Fan Tray 0       |        |            |              |                           |
| Fan Tray 1       | REV 03 | 740-014971 | TP0850       | Fan Tray                  |

### show chassis hardware (MX960 Router with Enhanced MX SCB)

```
user@host> show chassis hardware
```

| Hardware inventory: |         |             |               |                           |
|---------------------|---------|-------------|---------------|---------------------------|
| Item                | Version | Part number | Serial number | Description               |
| Chassis             |         |             | JN1096805AFA  | MX960                     |
| Midplane            | REV 03  | 710-013698  | TR0183        | MX960 Backplane           |
| Fan Extender        | REV 02  | 710-018051  | JY5227        | Extended Cable Manager    |
| FPM Board           | REV 03  | 710-014974  | JZ6876        | Front Panel Display       |
| PDM                 | Rev 03  | 740-013110  | QCS11035023   | Power Distribution Module |
| PEM 1               | Rev 03  | 740-013682  | QCS1109400L   | PS 1.7kW; 200-240VAC in   |
| PEM 2               | Rev 03  | 740-013682  | QCS11094015   | PS 1.7kW; 200-240VAC in   |
| PEM 3               | Rev 03  | 740-013682  | QCS11094012   | PS 1.7kW; 200-240VAC in   |
| Routing Engine 0    | REV 06  | 740-013063  | 1000687969    | RE-S-2000                 |
| Routing Engine 1    | REV 06  | 740-013063  | 1000687955    | RE-S-2000                 |



|        |        |            |                 |                     |
|--------|--------|------------|-----------------|---------------------|
| CB 0   | REV 11 | 750-031391 | YZ6072          | Enhanced MX SCB     |
| CB 1   | REV 11 | 750-031391 | YZ6068          | Enhanced MX SCB     |
| CB 2   | REV 11 | 750-031391 | YZ6081          | Enhanced MX SCB     |
| FPC 0  | REV 01 | 750-018122 | KA5576          | DPCE 40x 1GE R      |
| CPU    | REV 06 | 710-013713 | KB3961          | DPC PMB             |
| PIC 0  |        | BUILTIN    | BUILTIN         | 10x 1GE(LAN)        |
| Xcvr 0 | REV 01 | 740-011613 | P9F18GF         | SFP-SX              |
| Xcvr 2 | REV 01 | 740-011782 | P9M0TL9         | SFP-SX              |
| Xcvr 7 | REV 01 | 740-011782 | P9P0XXH         | SFP-SX              |
| Xcvr 9 | REV 01 | 740-011782 | P9M0TN1         | SFP-SX              |
| PIC 1  |        | BUILTIN    | BUILTIN         | 10x 1GE(LAN)        |
| Xcvr 0 | REV 01 | 740-011613 | PAJ4UHC         | SFP-SX              |
| PIC 2  |        | BUILTIN    | BUILTIN         | 10x 1GE(LAN)        |
| Xcvr 0 | REV 01 | 740-011613 | PFF2CD0         | SFP-SX              |
| Xcvr 1 | REV 01 | 740-011613 | PBG3ZUT         | SFP-SX              |
| Xcvr 2 | REV 01 | 740-011613 | PFF2DDV         | SFP-SX              |
| Xcvr 5 | REV 01 | 740-011613 | P8E2SST         | SFP-SX              |
| Xcvr 9 | REV 01 | 740-011782 | PB8329N         | SFP-SX              |
| PIC 3  |        | BUILTIN    | BUILTIN         | 10x 1GE(LAN)        |
| Xcvr 0 | REV 01 | 740-026192 | 1U0201084503342 | SFP-100BASE-BX10-U  |
| Xcvr 1 | REV 01 | 740-026193 | 1U1201084503313 | SFP-100BASE-BX10-D  |
| Xcvr 2 | REV 01 | 740-011613 | PAJ4Y5B         | SFP-SX              |
| Xcvr 6 | REV 01 | 740-011782 | P9M0U3M         | SFP-SX              |
| Xcvr 7 | REV 01 | 740-011782 | P9M0TLA         | SFP-SX              |
| FPC 1  | REV 16 | 750-031089 | YL0719          | MPC Type 2 3D       |
| CPU    | REV 06 | 711-030884 | YL1463          | MPC PMB 2G          |
| MIC 0  | REV 07 | 750-028387 | JR6500          | 3D 4x 10GE XFP      |
| PIC 0  |        | BUILTIN    | BUILTIN         | 2x 10GE XFP         |
| Xcvr 0 | REV 01 | 740-014279 | 733019A00154    | XFP-10G-LR          |
| Xcvr 1 | REV 02 | 740-014289 | T09F55034       | XFP-10G-SR          |
| PIC 1  |        | BUILTIN    | BUILTIN         | 2x 10GE XFP         |
| Xcvr 0 | REV 01 | 740-014279 | 913019B00791    | XFP-10G-LR          |
| Xcvr 1 | REV 01 | 740-014289 | 98S803A90384    | XFP-10G-SR          |
| MIC 1  | REV 24 | 750-028387 | YJ3950          | 3D 4x 10GE XFP      |
| PIC 2  |        | BUILTIN    | BUILTIN         | 2x 10GE XFP         |
| Xcvr 0 | REV 02 | 740-014279 | T10B36134       | XFP-10G-LR          |
| Xcvr 1 | REV 01 | 740-014289 | T07M86354       | XFP-10G-SR          |
| PIC 3  |        | BUILTIN    | BUILTIN         | 2x 10GE XFP         |
| FPC 2  | REV 08 | 710-014219 | JY9654          | DPCE 4x 10GE R      |
| CPU    | REV 06 | 710-013713 | JZ6549          | DPC PMB             |
| PIC 0  |        | BUILTIN    | BUILTIN         | 1x 10GE(LAN/WAN)    |
| PIC 1  |        | BUILTIN    | BUILTIN         | 1x 10GE(LAN/WAN)    |
| PIC 2  |        | BUILTIN    | BUILTIN         | 1x 10GE(LAN/WAN)    |
| Xcvr 0 | REV 03 | 740-011571 | C931BK028       | XFP-10G-SR          |
| PIC 3  |        | BUILTIN    | BUILTIN         | 1x 10GE(LAN/WAN)    |
| FPC 3  | REV 10 | 750-024199 | XJ6692          | MX FPC Type 3       |
| CPU    | REV 03 | 710-022351 | XF5182          | DPC PMB             |
| PIC 0  | REV 17 | 750-009553 | RJ2945          | 4x OC-48 SONET      |
| Xcvr 1 | REV 01 | 740-011785 | PCP3YLL         | SFP-SR              |
| Xcvr 3 | REV 01 | 740-011785 | PDSOMRY         | SFP-SR              |
| PIC 1  | REV 32 | 750-003700 | DP2113          | 1x OC-192 12xMM VSR |
| FPC 5  | REV 25 | 750-028467 | YM8256          | MPC 3D 16x 10GE     |
| CPU    | REV 10 | 711-029089 | YL3029          | AMPC PMB            |
| PIC 0  |        | BUILTIN    | BUILTIN         | 4x 10GE(LAN) SFP+   |
| Xcvr 1 | REV 01 | 740-031980 | AHNOX1Z         | SFP+-10G-SR         |
| PIC 1  |        | BUILTIN    | BUILTIN         | 4x 10GE(LAN) SFP+   |
| PIC 2  |        | BUILTIN    | BUILTIN         | 4x 10GE(LAN) SFP+   |
| PIC 3  |        | BUILTIN    | BUILTIN         | 4x 10GE(LAN) SFP+   |
| FPC 7  | REV 02 | 750-031092 | JR6658          | MPC Type 1 3D Q     |
| CPU    | REV 01 | 711-030884 | JZ9038          | MPC PMB 2G          |
| MIC 0  | REV 08 | 750-028392 | JZ8737          | 3D 20x 1GE(LAN) SFP |

|            |        |            |              |                      |
|------------|--------|------------|--------------|----------------------|
| PIC 0      |        | BUILTIN    | BUILTIN      | 10x 1GE(LAN) SFP     |
| Xcvr 0     | REV 01 | 740-011782 | PBE2C6Y      | SFP-SX               |
| Xcvr 2     |        | NON-JNPR   | U8105N8      | SFP-SX               |
| Xcvr 4     | REV 01 | 740-011613 | PFM18EF      | SFP-SX               |
| Xcvr 7     | REV 01 | 740-011613 | PFF2AM8      | SFP-SX               |
| Xcvr 8     | REV 01 | 740-011613 | PFF2CT6      | SFP-SX               |
| PIC 1      |        | BUILTIN    | BUILTIN      | 10x 1GE(LAN) SFP     |
| Xcvr 0     | REV 01 | 740-011782 | PB82VHH      | SFP-SX               |
| Xcvr 1     | REV 01 | 740-011613 | PFF2CSW      | SFP-SX               |
| Xcvr 9     | REV 01 | 740-011613 | PFF2BY0      | SFP-SX               |
| QXM 0      | REV 04 | 711-028408 | JR6372       | MPC QXM              |
| FPC 8      | REV 05 | 750-024387 | JW9754       | MX FPC Type 2        |
| CPU        | REV 03 | 710-022351 | KF1651       | DPC PMB              |
| PIC 0      | REV 08 | 750-014730 | DM3664       | 4x OC-3 1x OC-12 SFP |
| Xcvr 0     | REV 01 | 740-016065 | 81S290N00077 | SFP-SR               |
| Xcvr 1     |        | NON-JNPR   | 2191844      | SFP-SR               |
| Xcvr 2     | REV 01 | 740-011618 | PD81EE5      | SFP-IR               |
| PIC 1      | REV 08 | 750-014637 | DM3671       | 4x OC-12-3 SFP       |
| Xcvr 0     | REV 01 | 740-011785 | PCK3UNK      | SFP-SR               |
| Xcvr 3     | REV 01 | 740-011785 | PDSOMPZ      | SFP-SR               |
| FPC 10     | REV 04 | 710-013699 | JY4654       | DPCE 40x 1GE R       |
| CPU        | REV 05 | 710-013713 | JS9717       | DPC PMB              |
| PIC 0      |        | BUILTIN    | BUILTIN      | 10x 1GE(LAN)         |
| Xcvr 5     | REV 01 | 740-011782 | PAR1L72      | SFP-SX               |
| Xcvr 6     | REV 01 | 740-011782 | P8N1YQ4      | SFP-SX               |
| PIC 1      |        | BUILTIN    | BUILTIN      | 10x 1GE(LAN)         |
| PIC 2      |        | BUILTIN    | BUILTIN      | 10x 1GE(LAN)         |
| Xcvr 0     | REV 01 | 740-011782 | P8Q2AVL      | SFP-SX               |
| Xcvr 5     | REV 01 | 740-011782 | PAR1L7B      | SFP-SX               |
| Xcvr 6     | REV 01 | 740-011782 | PAR1L2J      | SFP-SX               |
| Xcvr 8     | REV 01 | 740-011782 | P8N1YMY      | SFP-SX               |
| PIC 3      |        | BUILTIN    | BUILTIN      | 10x 1GE(LAN)         |
| Fan Tray 0 | REV 03 | 740-014971 | TP0567       | Fan Tray             |
| Fan Tray 1 | REV 03 | 740-014971 | TP0702       | Fan Tray             |

### show chassis hardware models (MX960 Router with Enhanced MX SCB)

```
user@host> show chassis hardware models
```

| Hardware inventory: |         |             |               |                        |
|---------------------|---------|-------------|---------------|------------------------|
| Item                | Version | Part number | Serial number | FRU model number       |
| Midplane            | REV 03  | 710-013698  | TR0183        | CHAS-BP-MX960-S        |
| Fan Extender        | REV 02  | 710-018051  | JY5227        | ECM-MX960              |
| FPM Board           | REV 03  | 710-014974  | JZ6876        | CRAFT-MX960-S          |
| Routing Engine 0    | REV 06  | 740-013063  | 1000687969    | RE-S-2000-4096-S       |
| Routing Engine 1    | REV 06  | 740-013063  | 1000687955    | RE-S-2000-4096-S       |
| CB 0                | REV 11  | 750-031391  | YZ6072        | SCBE-MX-S              |
| CB 1                | REV 11  | 750-031391  | YZ6068        | SCBE-MX-S              |
| CB 2                | REV 11  | 750-031391  | YZ6081        | SCBE-MX-S              |
| FPC 0               | REV 01  | 750-018122  | KA5576        | DPCE-R-40GE-SFP        |
| FPC 1               | REV 16  | 750-031089  | YL0719        | MX-MPC2-3D             |
| MIC 0               | REV 07  | 750-028387  | JR6500        | MIC-3D-4XGE-XFP        |
| MIC 1               | REV 24  | 750-028387  | YJ3950        | MIC-3D-4XGE-XFP        |
| FPC 2               | REV 08  | 710-014219  | JY9654        | DPC-R-4XGE-XFP         |
| FPC 3               | REV 10  | 750-024199  | XJ6692        | MX-FPC3                |
| PIC 0               | REV 17  | 750-009553  | RJ2945        | PC-40C48-SON-SFP       |
| PIC 1               | REV 32  | 750-003700  | DP2113        | PC-10C192-SON-VSR      |
| FPC 5               | REV 25  | 750-028467  | YM8256        | MPC-3D-16XGE-SFPP      |
| FPC 7               | REV 02  | 750-031092  | JR6658        | MX-MPC1-3D-Q           |
| MIC 0               | REV 08  | 750-028392  | JZ8737        | MIC-3D-20GE-SFP        |
| FPC 8               | REV 05  | 750-024387  | JW9754        | MX-FPC2                |
| PIC 0               | REV 08  | 750-014730  | DM3664        | PB-40C3-10C12-SON2-SFP |

|            |        |            |        |                       |
|------------|--------|------------|--------|-----------------------|
| PIC 1      | REV 08 | 750-014637 | DM3671 | PB-40C3-40C12-SON-SFP |
| FPC 10     | REV 04 | 710-013699 | JY4654 | DPC-R-40GE-SFP        |
| Fan Tray 0 | REV 03 | 740-014971 | TP0567 | FFANTRAY-MX960-S      |
| Fan Tray 1 | REV 03 | 740-014971 | TP0702 | FFANTRAY-MX960-S      |

### show chassis hardware detail (MX960 Router)

```

user@host> show chassis hardware detail
Hardware inventory:

```

| Item             | Version  | Part number       | Serial number    | Description             |
|------------------|----------|-------------------|------------------|-------------------------|
| Chassis          |          |                   |                  | MX960                   |
| Midplane         | REV 01   | 710-013698        | AA6082           | MX960 Midplane          |
| PIM              | Rev 01   | 740-013110        | 000008           | Power Inlet Module      |
| PEM 2            |          |                   |                  |                         |
| PEM 3            | Rev 01   | 740-013682        | 000038           | PS 1.7kW; 200-240VAC in |
| Routing Engine 0 | REV 00   | 740-015113        | 1000617944       | RE-S-1300               |
| ad0              | 245 MB   | SanDisk SDCFB-256 | 111419E1805T1141 | Compact Flash           |
| ad2              | 38154 MB | FUJITSU MHT2040BH | NROWT5925N77     | Hard Disk               |
| CB 0             | REV 05   | 710-013725        | JK6947           | MX960 Test SCB          |
| FPC 4            | REV 01   | 710-013305        | JM7617           | MX960 Test DPC          |
| CPU              |          |                   |                  |                         |
| PIC 0            |          | BUILTIN           | BUILTIN          | 1x 10GE(LAN/WAN)        |
| PIC 1            |          | BUILTIN           | BUILTIN          | 10x 1GE                 |
| FPC 7            | REV 01   | 710-013305        | JL9634           | MX960 Test DPC          |
| CPU              |          |                   |                  |                         |
| PIC 0            |          | BUILTIN           | BUILTIN          | 1x 10GE(LAN/WAN)        |
| Xcvr 0           |          | NON-JNPR          | MYBG65I82C       | XFP-10G-SR              |
| PIC 1            |          | BUILTIN           | BUILTIN          | 10x 1GE                 |
| Xcvr 1           | REV 01   | 740-011782        | P7N0368          | SFP-SX                  |
| Xcvr 4           | REV 01   | 740-011782        | P8J1W27          | SFP-SX                  |
| Xcvr 6           | REV 01   | 740-011782        | P8J1VSD          | SFP-SX                  |
| Xcvr 9           | REV 01   | 740-011782        | P8J1W25          | SFP-SX                  |
| Fan Tray 0       |          |                   |                  |                         |
| Fan Tray 1       |          |                   |                  |                         |

### show chassis hardware (MX2010 Router)

```

user@host > show chassis hardware
Hardware inventory:

```

| Item             | Version | Part number | Serial number | Description          |
|------------------|---------|-------------|---------------|----------------------|
| Chassis          |         |             | JN11E3217AFK  | MX2010               |
| Midplane         | REV 01  | 750-044636  | ABAB8506      | Lower Backplane      |
| Midplane 1       | REV 01  | 711-044557  | ZY8296        | Upper Backplane      |
| PMP              | REV 03  | 711-032426  | ACAJ1388      | Power Midplane       |
| FPM Board        | REV 06  | 711-032349  | ZX8744        | Front Panel Display  |
| PSM 4            | REV 0C  | 740-033727  | VK00254       | DC 52V Power Supply  |
| Module           |         |             |               |                      |
| PSM 5            | REV 0B  | 740-033727  | VG00015       | DC 52V Power Supply  |
| Module           |         |             |               |                      |
| PSM 6            | REV 0B  | 740-033727  | VH00097       | DC 52V Power Supply  |
| Module           |         |             |               |                      |
| PSM 7            | REV 0C  | 740-033727  | VJ00151       | DC 52V Power Supply  |
| Module           |         |             |               |                      |
| PSM 8            | REV 0C  | 740-033727  | VJ00149       | DC 52V Power Supply  |
| Module           |         |             |               |                      |
| PDM 0            | REV 0B  | 740-038109  | WA00008       | DC Power Dist Module |
| PDM 1            | REV 0B  | 740-038109  | WA00014       | DC Power Dist Module |
| Routing Engine 0 | REV 02  | 740-041821  | 9009094134    | RE-S-1800x4          |
| Routing Engine 1 | REV 02  | 740-041821  | 9009094141    | RE-S-1800x4          |
| CB 0             | REV 08  | 750-040257  | CAAB3491      | Control Board        |
| CB 1             | REV 08  | 750-040257  | CAAB3489      | Control Board        |

|        |        |            |              |                     |
|--------|--------|------------|--------------|---------------------|
| SPMB 0 | REV 02 | 711-041855 | CAAA6135     | PMB Board           |
| SPMB 1 | REV 02 | 711-041855 | CAAA6137     | PMB Board           |
| SFB 0  | REV 06 | 711-032385 | ZV1828       | Switch Fabric Board |
| SFB 1  | REV 07 | 711-032385 | ZZ2568       | Switch Fabric Board |
| SFB 2  | REV 07 | 711-032385 | ZZ2563       | Switch Fabric Board |
| SFB 3  | REV 07 | 711-032385 | ZZ2564       | Switch Fabric Board |
| SFB 4  | REV 07 | 711-032385 | ZZ2580       | Switch Fabric Board |
| SFB 5  | REV 07 | 711-032385 | ZZ2579       | Switch Fabric Board |
| SFB 6  | REV 07 | 711-032385 | CAAB4882     | Switch Fabric Board |
| SFB 7  | REV 07 | 711-032385 | CAAB4898     | Switch Fabric Board |
| FPC 0  | REV 33 | 750-028467 | CAAB1919     | MPC 3D 16x 10GE     |
| CPU    | REV 11 | 711-029089 | CAAB7174     | AMPC PMB            |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+   |
| Xcvr 0 | REV 01 | 740-021308 | AMH02RE      | SFP+-10G-SR         |
| Xcvr 1 | REV 01 | 740-021308 | AMH038C      | SFP+-10G-SR         |
| Xcvr 2 | REV 01 | 740-021308 | AMH0390      | SFP+-10G-SR         |
| Xcvr 3 | REV 01 | 740-021308 | AMG0SUA      | SFP+-10G-SR         |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+   |
| Xcvr 0 | REV 01 | 740-021308 | AMH0579      | SFP+-10G-SR         |
| Xcvr 1 | REV 01 | 740-021308 | AMG0SGP      | SFP+-10G-SR         |
| Xcvr 2 | REV 01 | 740-021308 | AMH04SV      | SFP+-10G-SR         |
| Xcvr 3 | REV 01 | 740-021308 | AMH04X3      | SFP+-10G-SR         |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+   |
| Xcvr 0 | REV 01 | 740-021308 | AMH0135      | SFP+-10G-SR         |
| Xcvr 1 | REV 01 | 740-021308 | AMH02NC      | SFP+-10G-SR         |
| Xcvr 2 | REV 01 | 740-021308 | AMH02XB      | SFP+-10G-SR         |
| Xcvr 3 | REV 01 | 740-021308 | AMH02PN      | SFP+-10G-SR         |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+   |
| Xcvr 0 | REV 01 | 740-021308 | AMH057Y      | SFP+-10G-SR         |
| Xcvr 1 | REV 01 | 740-021308 | AMG0JHE      | SFP+-10G-SR         |
| Xcvr 2 | REV 01 | 740-021308 | AMH02HT      | SFP+-10G-SR         |
| Xcvr 3 | REV 01 | 740-021308 | AMH04V4      | SFP+-10G-SR         |
| FPC 1  | REV 21 | 750-033205 | ZG5027       | MPCE Type 3D        |
| CPU    | REV 04 | 711-035209 | YT4780       | HMPC PMB 2G         |
| MIC 0  | REV 03 | 750-033307 | ZV6299       | 10X10GE SFPP        |
| PIC 0  |        | BUILTIN    | BUILTIN      | 10X10GE SFPP        |
| Xcvr 0 | REV 01 | 740-031980 | 083363A00410 | SFP+-10G-SR         |
| Xcvr 1 | REV 01 | 740-031980 | 083363A00334 | SFP+-10G-SR         |
| Xcvr 2 | REV 01 | 740-031980 | 113363A00125 | SFP+-10G-SR         |
| Xcvr 3 | REV 01 | 740-031980 | 083363A00953 | SFP+-10G-SR         |
| Xcvr 4 | REV 01 | 740-031980 | AHR013D      | SFP+-10G-SR         |
| Xcvr 5 | REV 01 | 740-031980 | AJ40JUR      | SFP+-10G-SR         |
| Xcvr 6 | REV 01 | 740-031980 | AJ40JKL      | SFP+-10G-SR         |
| Xcvr 7 | REV 01 | 740-031980 | AJ30ECK      | SFP+-10G-SR         |
| Xcvr 8 | REV 01 | 740-021308 | 19T511100864 | SFP+-10G-SR         |
| Xcvr 9 | REV 01 | 740-021308 | 19T511100868 | SFP+-10G-SR         |
| MIC 1  | REV 03 | 750-033307 | ZV6268       | 10X10GE SFPP        |
| PIC 2  |        | BUILTIN    | BUILTIN      | 10X10GE SFPP        |
| Xcvr 0 | REV 01 | 740-031980 | AJCOJML      | SFP+-10G-SR         |
| Xcvr 1 | REV 01 | 740-031980 | AJ403PC      | SFP+-10G-SR         |
| Xcvr 2 | REV 01 | 740-031980 | AJ10N25      | SFP+-10G-SR         |
| Xcvr 3 | REV 01 | 740-031980 | AJ40JF4      | SFP+-10G-SR         |
| Xcvr 4 | REV 01 | 740-031980 | AJ40JSJ      | SFP+-10G-SR         |
| Xcvr 5 | REV 01 | 740-031980 | AJ403V7      | SFP+-10G-SR         |
| Xcvr 6 | REV 01 | 740-031980 | AJ40JN3      | SFP+-10G-SR         |
| Xcvr 7 | REV 01 | 740-031980 | AJ40JSU      | SFP+-10G-SR         |
| Xcvr 8 | REV 01 | 740-021308 | 19T511100468 | SFP+-10G-SR         |
| Xcvr 9 | REV 01 | 740-021308 | 19T511101363 | SFP+-10G-SR         |
| FPC 8  | REV 22 | 750-031089 | ZT9746       | MPC Type 2 3D       |
| CPU    | REV 06 | 711-030884 | ZS1271       | MPC PMB 2G          |
| MIC 0  | REV 26 | 750-028392 | ABBS1150     | 3D 20x 1GE(LAN) SFP |

|            |        |            |              |                        |
|------------|--------|------------|--------------|------------------------|
| PIC 0      |        | BUILTIN    | BUILTIN      | 10x 1GE(LAN) SFP       |
| Xcvr 0     | REV 01 | 740-031851 | PLG023C      | SFP-SX                 |
| Xcvr 1     | REV 01 | 740-031851 | PLG09C6      | SFP-SX                 |
| Xcvr 2     | REV 02 | 740-011613 | AM0950SF9L7  | SFP-SX                 |
| Xcvr 3     | REV 02 | 740-011613 | AM1001SFN1H  | SFP-SX                 |
| Xcvr 4     | REV 02 | 740-011613 | AM1001SFM9D  | SFP-SX                 |
| Xcvr 5     | REV 02 | 740-011613 | AM1001SFLTJ  | SFP-SX                 |
| Xcvr 6     | REV 01 | 740-031851 | AC1108S03L9  | SFP-SX                 |
| Xcvr 7     | REV 01 | 740-031851 | AC1102S00NC  | SFP-SX                 |
| Xcvr 8     | REV 01 | 740-031851 | AC1102S00MX  | SFP-SX                 |
| Xcvr 9     | REV 01 | 740-031851 | AC1102S0085  | SFP-SX                 |
| PIC 1      |        | BUILTIN    | BUILTIN      | 10x 1GE(LAN) SFP       |
| Xcvr 0     | REV 01 | 740-031851 | AC1102S00KU  | SFP-SX                 |
| Xcvr 1     | REV 01 | 740-031851 | AC1102S00NG  | SFP-SX                 |
| Xcvr 2     | REV 01 | 740-031851 | AC1102S00K3  | SFP-SX                 |
| Xcvr 3     | REV 01 | 740-031851 | AC1102S008R  | SFP-SX                 |
| Xcvr 4     | REV 01 | 740-031851 | AM1107SUFVJ  | SFP-SX                 |
| Xcvr 5     | REV 01 | 740-031851 | AC1108S03LG  | SFP-SX                 |
| MIC 1      | REV 26 | 750-028387 | ABBR9582     | 3D 4x 10GE XFP         |
| PIC 2      |        | BUILTIN    | BUILTIN      | 2x 10GE XFP            |
| Xcvr 0     |        | NON-JNPR   | T10A91703    | XFP-10G-SR             |
| Xcvr 1     |        | NON-JNPR   | T09L42604    | XFP-10G-SR             |
| PIC 3      |        | BUILTIN    | BUILTIN      | 2x 10GE XFP            |
| FPC 9      | REV 11 | 750-036284 | ZL3591       | MPC 3D 16x 10GE EM     |
| CPU        | REV 10 | 711-029089 | ZL0513       | AMPC PMB               |
| PIC 0      |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+      |
| Xcvr 0     | REV 01 | 740-031980 | 1YT517101825 | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | 1YT517101821 | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | 1YT517101682 | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-031980 | ALQ13R6      | SFP+-10G-SR            |
| PIC 1      |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+      |
| Xcvr 0     | REV 01 | 740-031980 | 1YT517101828 | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | 1YT517101716 | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | 1YT517101732 | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-031980 | ALP0TR1      | SFP+-10G-SR            |
| PIC 2      |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+      |
| Xcvr 0     | REV 01 | 740-031980 | 1YT517101741 | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | 1YT517101829 | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | 1YT517101669 | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-031980 | ALQ14E3      | SFP+-10G-SR            |
| PIC 3      |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+      |
| Xcvr 0     | REV 01 | 740-031980 | 1YT517101826 | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | 1YT517101817 | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | 1YT517101735 | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-031980 | ALQ159A      | SFP+-10G-SR            |
| ADC 0      | REV 05 | 750-043596 | CAAC2073     | Adapter Card           |
| ADC 1      | REV 01 | 750-043596 | ZV4117       | Adapter Card           |
| ADC 8      | REV 01 | 750-043596 | ZV4107       | Adapter Card           |
| ADC 9      | REV 02 | 750-043596 | ZW1555       | Adapter Card           |
| Fan Tray 0 | REV 2A | 760-046960 | ACAY0015     | 172mm FanTray - 6 Fans |
| Fan Tray 1 | REV 2A | 760-046960 | ACAY0019     | 172mm FanTray - 6 Fans |
| Fan Tray 2 | REV 2A | 760-046960 | ACAY0020     | 172mm FanTray - 6 Fans |
| Fan Tray 3 | REV 2A | 760-046960 | ACAY0021     | 172mm FanTray - 6 Fans |

#### show chassis hardware detail (MX2010 Router)

```
user@host > show chassis hardware detail
```

```
Hardware inventory:
```

| Item     | Version | Part number | Serial number | Description     |
|----------|---------|-------------|---------------|-----------------|
| Chassis  |         |             | JN11E233DAFK  | MX2010          |
| Midplane | REV 26  | 750-044636  | ABAB9357      | Lower Backplane |

|                  |        |                                          |                          |                      |
|------------------|--------|------------------------------------------|--------------------------|----------------------|
| Midplane 1       | REV 01 | 711-044557                               | ABAB8643                 | Upper Backplane      |
| PMP              | REV 04 | 711-032426                               | ACAJ1677                 | Power Midplane       |
| FPM Board        | REV 08 | 760-044634                               | ABBV9726                 | Front Panel Display  |
| PSM 0            | REV 01 | 740-045050                               | 1E02224000P              | DC 52V Power Supply  |
| Module           |        |                                          |                          |                      |
| PSM 1            | REV 01 | 740-045050                               | 1E02224000M              | DC 52V Power Supply  |
| Module           |        |                                          |                          |                      |
| PSM 2            | REV 01 | 740-045050                               | 1E022240010              | DC 52V Power Supply  |
| Module           |        |                                          |                          |                      |
| PSM 3            | REV 01 | 740-045050                               | 1E02224000G              | DC 52V Power Supply  |
| Module           |        |                                          |                          |                      |
| PSM 4            | REV 01 | 740-045050                               | 1E022240013              | DC 52V Power Supply  |
| Module           |        |                                          |                          |                      |
| PSM 5            | REV 01 | 740-045050                               | 1E022240007              | DC 52V Power Supply  |
| Module           |        |                                          |                          |                      |
| PSM 6            | REV 01 | 740-045050                               | 1E02224001C              | DC 52V Power Supply  |
| Module           |        |                                          |                          |                      |
| PSM 7            | REV 01 | 740-045050                               | 1E02224001D              | DC 52V Power Supply  |
| Module           |        |                                          |                          |                      |
| PSM 8            | REV 01 | 740-045050                               | 1E02224001B              | DC 52V Power Supply  |
| Module           |        |                                          |                          |                      |
| PDM 0            | REV 01 | 740-045234                               | 1E262250067              | DC Power Dist Module |
| Routing Engine 0 | REV 02 | 740-041821                               | 9009099704               | RE-S-1800x4          |
| ad0 3831 MB      |        | UGB30SFA4000T1                           | SFA4000T1 00000651       | Compact Flash        |
| ad1 30533 MB     |        | UGB94BPH32H0S1-KCI                       | 11000019592              | Disk 1               |
| usb0 (addr 1)    |        | EHCI root hub 0                          | Intel                    | uhub0                |
| usb0 (addr 2)    |        | product 0x0020 32                        | vendor 0x8087            | uhub1                |
| DIMM 0           |        | SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 | MFR ID-ce80              |                      |
| DIMM 1           |        | SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 | MFR ID-ce80              |                      |
| DIMM 2           |        | SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 | MFR ID-ce80              |                      |
| DIMM 3           |        | SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 | MFR ID-ce80              |                      |
| Routing Engine 1 | REV 02 | 740-041821                               | 9009099706               | RE-S-1800x4          |
| ad0 3998 MB      |        | Virtium - TuffDrive                      | VCF P1T0200262860208 114 | Compact Flash        |
| ad1 30533 MB     |        | UGB94ARF32H0S3-KC                        | UNIGEN-499551-000404     | Disk 1               |
| CB 0             | REV 13 | 750-040257                               | CAAF8436                 | Control Board        |
| CB 1             | REV 13 | 750-040257                               | CAAF8434                 | Control Board        |
| SPMB 0           | REV 02 | 711-041855                               | ABBV3825                 | PMB Board            |
| SPMB 1           | REV 02 | 711-041855                               | ABBV3833                 | PMB Board            |
| SFB 0            | REV 05 | 711-044466                               | ABBX5682                 | Switch Fabric Board  |
| SFB 1            | REV 05 | 711-044466                               | ABBX5676                 | Switch Fabric Board  |
| SFB 2            | REV 05 | 711-044466                               | ABBX5665                 | Switch Fabric Board  |
| SFB 3            | REV 05 | 711-044466                               | ABBX5699                 | Switch Fabric Board  |
| SFB 4            | REV 05 | 711-044466                               | ABBX5603                 | Switch Fabric Board  |
| SFB 5            | REV 05 | 711-044466                               | ABBX5587                 | Switch Fabric Board  |
| SFB 6            | REV 05 | 711-044466                               | ABBX5607                 | Switch Fabric Board  |
| SFB 7            | REV 05 | 711-044466                               | ABBX5669                 | Switch Fabric Board  |
| FPC 0            | REV 09 | 750-037355                               | CAAF0924                 | MPC Type 4-2         |
| CPU              | REV 08 | 711-035209                               | CAAB9842                 | HMPC PMB 2G          |
| PIC 0            |        | BUILTIN                                  | BUILTIN                  | 4x10GE SFPP          |
| Xcvr 0           | REV 01 | 740-021308                               | 19T511101656             | SFP+-10G-SR          |
| Xcvr 1           | REV 01 | 740-031980                               | AMA04RU                  | SFP+-10G-SR          |
| Xcvr 2           | REV 01 | 740-031980                               | 193363A00558             | SFP+-10G-SR          |
| Xcvr 3           | REV 01 | 740-031980                               | B10M00202                | SFP+-10G-SR          |
| PIC 1            |        | BUILTIN                                  | BUILTIN                  | 1X100GE CFP          |
| Xcvr 0           |        | NON-JNPR                                 | X12J00328                | CFP-100G-SR10        |
| PIC 2            |        | BUILTIN                                  | BUILTIN                  | 4x10GE SFPP          |
| Xcvr 0           | REV 01 | 740-031980                               | AMA088W                  | SFP+-10G-SR          |
| Xcvr 1           | REV 01 | 740-031980                               | B10L04211                | SFP+-10G-SR          |
| Xcvr 2           | REV 01 | 740-021308                               | 19T511101602             | SFP+-10G-SR          |
| Xcvr 3           | REV 01 | 740-031980                               | B10L04151                | SFP+-10G-SR          |
| PIC 3            |        | BUILTIN                                  | BUILTIN                  | 1X100GE CFP          |

|        |        |            |              |                   |
|--------|--------|------------|--------------|-------------------|
| Xcvr 0 |        | NON-JNPR   | X12J00332    | CFP-100G-SR10     |
| FPC 1  | REV 18 | 750-033205 | ZE0128       | MPCE Type 3D      |
| CPU    | REV 06 | 711-035209 | ZG5431       | HMPC PMB 2G       |
| MIC 0  | REV 15 | 750-033199 | ZP6435       | 1X100GE CFP       |
| PIC 0  |        | BUILTIN    | BUILTIN      | 1X100GE CFP       |
| Xcvr 0 | REV 01 | 740-032210 | J11E46118    | CFP-100G-LR4      |
| MIC 1  | REV 15 | 750-033199 | ZP6442       | 1X100GE CFP       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 1X100GE CFP       |
| Xcvr 0 | REV 01 | 740-032210 | UMN03T4      | CFP-100G-LR4      |
| FPC 2  | REV 16 | 750-037358 | CAAL1001     | MPC Type 4-1      |
| CPU    | REV 08 | 711-035209 | CAAK7927     | HMPC PMB 2G       |
| PIC 0  |        | BUILTIN    | BUILTIN      | 8X10GE SFPP       |
| Xcvr 0 | REV 01 | 740-031980 | 193363A00589 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-021308 | 973152A00028 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 193363A00376 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-021308 | 973152A00016 | SFP+-10G-SR       |
| Xcvr 4 | REV 01 | 740-031980 | 193363A00499 | SFP+-10G-SR       |
| Xcvr 5 | REV 01 | 740-021308 | 973152A00039 | SFP+-10G-SR       |
| Xcvr 6 | REV 01 | 740-031980 | B11E01239    | SFP+-10G-SR       |
| Xcvr 7 | REV 01 | 740-021308 | 973152A00058 | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 8X10GE SFPP       |
| Xcvr 0 | REV 01 | 740-031980 | B10M00075    | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-021308 | 973152A00014 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AMA0638      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-021308 | 973152A00063 | SFP+-10G-SR       |
| Xcvr 4 | REV 01 | 740-031980 | AMA0629      | SFP+-10G-SR       |
| Xcvr 5 | REV 01 | 740-021308 | 973152A00053 | SFP+-10G-SR       |
| Xcvr 6 | REV 01 | 740-031980 | 193363A00344 | SFP+-10G-SR       |
| Xcvr 7 | REV 01 | 740-021308 | 973152A00046 | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 8X10GE SFPP       |
| Xcvr 0 | REV 01 | 740-031980 | AMA062M      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-021308 | 973152A00080 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 193363A00580 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-021308 | 973152A00064 | SFP+-10G-SR       |
| Xcvr 4 | REV 01 | 740-031980 | 093363A01494 | SFP+-10G-SR       |
| Xcvr 5 | REV 01 | 740-021308 | 973152A00020 | SFP+-10G-SR       |
| Xcvr 6 | REV 01 | 740-031980 | 123363A00047 | SFP+-10G-SR       |
| Xcvr 7 | REV 01 | 740-021308 | 973152A00072 | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 8X10GE SFPP       |
| Xcvr 0 | REV 01 | 740-021308 | 03DZ06A01033 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-021308 | 973152A00022 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-021308 | 03DZ06A01026 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-021308 | 973152A00013 | SFP+-10G-SR       |
| Xcvr 4 | REV 01 | 740-021308 | 03DZ06A01028 | SFP+-10G-SR       |
| Xcvr 5 | REV 01 | 740-021308 | 973152A00079 | SFP+-10G-SR       |
| Xcvr 6 | REV 01 | 740-021308 | 03DZ06A01018 | SFP+-10G-SR       |
| Xcvr 7 | REV 01 | 740-021308 | 973152A00025 | SFP+-10G-SR       |
| FPC 3  | REV 33 | 750-028467 | CAAF5400     | MPC 3D 16x 10GE   |
| CPU    | REV 11 | 711-029089 | CAAH7626     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-021308 | 973152A00066 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-021308 | 973152A00021 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-021308 | 973152A00062 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-021308 | 973152A00027 | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-021308 | 973152A00065 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-021308 | 973152A00069 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-021308 | 973152A00026 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-021308 | 973152A00003 | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-021308 | 973152A00035 | SFP+-10G-SR       |

|        |        |            |              |                   |
|--------|--------|------------|--------------|-------------------|
| Xcvr 1 | REV 01 | 740-021308 | 973152A00004 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-021308 | 973152A00049 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-021308 | 973152A00055 | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-021308 | 973152A00010 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-021308 | 973152A00001 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-021308 | 973152A00073 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-021308 | 973152A00012 | SFP+-10G-SR       |
| FPC 4  | REV 21 | 750-033205 | ZG5028       | MPCE Type 3D      |
| CPU    | REV 05 | 711-035209 | YX3911       | HMPC PMB 2G       |
| MIC 0  | REV 03 | 750-036233 | ZL2036       | 2X40GE QSFP       |
| PIC 0  |        | BUILTIN    | BUILTIN      | 2X40GE QSFP       |
| Xcvr 0 | REV 01 | 740-032986 | QB220708     | QSFP+-40G-SR4     |
| Xcvr 1 | REV 01 | 740-032986 | QB220735     | QSFP+-40G-SR4     |
| MIC 1  | REV 03 | 750-036233 | ZL2028       | 2X40GE QSFP       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 2X40GE QSFP       |
| Xcvr 0 | REV 01 | 740-032986 | QB220727     | QSFP+-40G-SR4     |
| Xcvr 1 | REV 01 | 740-032986 | QB220715     | QSFP+-40G-SR4     |
| FPC 5  | REV 11 | 750-037358 | CAAE2196     | MPC Type 4-1      |
| CPU    | REV 08 | 711-035209 | CAAD9074     | HMPC PMB 2G       |
| PIC 0  |        | BUILTIN    | BUILTIN      | 8X10GE SFPP       |
| Xcvr 0 | REV 01 | 740-031980 | AMA062S      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AMA062P      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AMA052R      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AMA0632      | SFP+-10G-SR       |
| Xcvr 4 | REV 01 | 740-031980 | 193363A00564 | SFP+-10G-SR       |
| Xcvr 5 | REV 01 | 740-031980 | 193363A00229 | SFP+-10G-SR       |
| Xcvr 6 | REV 01 | 740-031980 | 193363A00363 | SFP+-10G-SR       |
| Xcvr 7 | REV 01 | 740-031980 | 193363A00278 | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 8X10GE SFPP       |
| Xcvr 0 | REV 01 | 740-031980 | AMA04CC      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-021308 | AD0927A001W  | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AMA04N2      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AMA062U      | SFP+-10G-SR       |
| Xcvr 4 | REV 01 | 740-031980 | 193363A00491 | SFP+-10G-SR       |
| Xcvr 5 | REV 01 | 740-031980 | 183363A01511 | SFP+-10G-SR       |
| Xcvr 6 | REV 01 | 740-031980 | 193363A00565 | SFP+-10G-SR       |
| Xcvr 7 | REV 01 | 740-031980 | 193363A00405 | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 8X10GE SFPP       |
| Xcvr 0 | REV 01 | 740-031980 | AMA07QX      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AMA06MS      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 193363A00318 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 193363A00402 | SFP+-10G-SR       |
| Xcvr 4 | REV 01 | 740-031980 | 193363A00174 | SFP+-10G-SR       |
| Xcvr 5 | REV 01 | 740-031980 | 193363A00388 | SFP+-10G-SR       |
| Xcvr 6 | REV 01 | 740-031980 | 193363A00377 | SFP+-10G-SR       |
| Xcvr 7 | REV 01 | 740-031980 | 193363A00234 | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 8X10GE SFPP       |
| Xcvr 0 | REV 01 | 740-031980 | AMA062T      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 193363A00550 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 193363A00364 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AMA0630      | SFP+-10G-SR       |
| Xcvr 4 | REV 01 | 740-031980 | 193363A00509 | SFP+-10G-SR       |
| Xcvr 5 | REV 01 | 740-031980 | 193363A00459 | SFP+-10G-SR       |
| Xcvr 6 | REV 01 | 740-031980 | 113363A00191 | SFP+-10G-SR       |
| Xcvr 7 | REV 01 | 740-031980 | 193363A00352 | SFP+-10G-SR       |
| FPC 6  | REV 33 | 750-028467 | CAAF5552     | MPC 3D 16x 10GE   |
| CPU    | REV 11 | 711-029089 | CAAH7601     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-021308 | AD0927A0036  | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-021308 | AD0927A003M  | SFP+-10G-SR       |



|        |        |            |              |                   |
|--------|--------|------------|--------------|-------------------|
| Xcvr 2 | REV 01 | 740-021308 | AD0927A003G  | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-021308 | AD0927A0031  | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 193363A00331 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 193363A00325 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 193363A00417 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 183363A02509 | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-021308 | T09K75140    | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | B11A04356    | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | B11K01952    | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | B11K01914    | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-021308 | T09K75157    | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-021308 | T09K75194    | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | B11K01926    | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | B11K01936    | SFP+-10G-SR       |
| FPC 7  | REV 16 | 750-037358 | CAAL1012     | MPC Type 4-1      |
| CPU    | REV 08 | 711-035209 | CAAJ3851     | HMPC PMB 2G       |
| PIC 0  |        | BUILTIN    | BUILTIN      | 8X10GE SFPP       |
| Xcvr 0 | REV 01 | 740-031980 | AMA04NK      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | B11F00260    | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | B11E02192    | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AMA04CP      | SFP+-10G-SR       |
| Xcvr 4 | REV 01 | 740-031980 | AJ40JJK      | SFP+-10G-SR       |
| Xcvr 5 | REV 01 | 740-031980 | B11F00238    | SFP+-10G-SR       |
| Xcvr 6 | REV 01 | 740-031980 | B10M00275    | SFP+-10G-SR       |
| Xcvr 7 | REV 01 | 740-031980 | 193363A00211 | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 8X10GE SFPP       |
| Xcvr 0 | REV 01 | 740-031980 | B11D05577    | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | B11G00586    | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AMA08B7      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AMA04Q0      | SFP+-10G-SR       |
| Xcvr 4 | REV 01 | 740-031980 | B11D05840    | SFP+-10G-SR       |
| Xcvr 5 | REV 01 | 740-031980 | B11E00467    | SFP+-10G-SR       |
| Xcvr 6 | REV 01 | 740-031980 | B11E00029    | SFP+-10G-SR       |
| Xcvr 7 | REV 01 | 740-021308 | 19T511101712 | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 8X10GE SFPP       |
| Xcvr 0 | REV 01 | 740-031980 | 193363A00568 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | B10M00166    | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | B10M00212    | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | B11D05823    | SFP+-10G-SR       |
| Xcvr 4 | REV 01 | 740-021308 | 03DZ06A01005 | SFP+-10G-SR       |
| Xcvr 5 | REV 01 | 740-021308 | 03DZ06A01003 | SFP+-10G-SR       |
| Xcvr 6 | REV 01 | 740-021308 | 03DZ06A01009 | SFP+-10G-SR       |
| Xcvr 7 | REV 01 | 740-021308 | 03DZ06A01004 | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 8X10GE SFPP       |
| Xcvr 0 | REV 01 | 740-021308 | 03DZ06A01017 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-021308 | 03DZ06A01016 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-021308 | 03DZ06A01024 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-021308 | 03DZ06A01008 | SFP+-10G-SR       |
| Xcvr 4 | REV 01 | 740-030658 | AD0946A02UH  | SFP+-10G-USR      |
| Xcvr 5 | REV 01 | 740-021308 | T09J67913    | SFP+-10G-SR       |
| Xcvr 6 | REV 01 | 740-021308 | AD0837ES09G  | SFP+-10G-SR       |
| Xcvr 7 | REV 01 | 740-021308 | 03DZ06A01015 | SFP+-10G-SR       |
| FPC 8  | REV 03 | 750-045372 | CAAD3111     | MPCE Type 3D      |
| CPU    | REV 08 | 711-035209 | CAAD8033     | HMPC PMB 2G       |
| MIC 0  | REV 03 | 750-036233 | ZL2032       | 2X40GE QSFP       |
| PIC 0  |        | BUILTIN    | BUILTIN      | 2X40GE QSFP       |
| Xcvr 0 | REV 01 | 740-032986 | QB230273     | QSFP+-40G-SR4     |
| Xcvr 1 | REV 01 | 740-032986 | QB230254     | QSFP+-40G-SR4     |

|            |        |            |              |                        |
|------------|--------|------------|--------------|------------------------|
| MIC 1      | REV 03 | 750-036233 | ZL2021       | 2X40GE QSFP            |
| PIC 2      |        | BUILTIN    | BUILTIN      | 2X40GE QSFP            |
| Xcvr 0     | REV 01 | 740-032986 | QB390962     | QSFP+-40G-SR4          |
| Xcvr 1     | REV 01 | 740-032986 | QB390960     | QSFP+-40G-SR4          |
| FPC 9      | REV 09 | 750-037355 | CAAF1531     | MPC Type 4-2           |
| CPU        | REV 08 | 711-035209 | CAAB9927     | HMPC PMB 2G            |
| PIC 0      |        | BUILTIN    | BUILTIN      | 4x10GE SFP             |
| Xcvr 0     | REV 01 | 740-031980 | 193363A00525 | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | 193363A00504 | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | 193363A00368 | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-031980 | AJ40JSS      | SFP+-10G-SR            |
| PIC 1      |        | BUILTIN    | BUILTIN      | 1X100GE CFP            |
| PIC 2      |        | BUILTIN    | BUILTIN      | 4x10GE SFP             |
| Xcvr 0     | REV 01 | 740-031980 | 123363A00042 | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | B10M00023    | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | AJ802EM      | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-031980 | B11E02348    | SFP+-10G-SR            |
| PIC 3      |        | BUILTIN    | BUILTIN      | 1X100GE CFP            |
| ADC 0      | REV 13 | 750-043596 | ABBX5532     | Adapter Card           |
| ADC 1      | REV 13 | 750-043596 | ABBX5550     | Adapter Card           |
| ADC 2      | REV 13 | 750-043596 | ABBX5571     | Adapter Card           |
| ADC 3      | REV 13 | 750-043596 | ABBX5568     | Adapter Card           |
| ADC 4      | REV 13 | 750-043596 | ABBX5556     | Adapter Card           |
| ADC 5      | REV 13 | 750-043596 | ABBX5553     | Adapter Card           |
| ADC 6      | REV 13 | 750-043596 | ABBX5541     | Adapter Card           |
| ADC 7      | REV 13 | 750-043596 | ABBX5578     | Adapter Card           |
| ADC 8      | REV 13 | 750-043596 | ABBX5560     | Adapter Card           |
| ADC 9      | REV 07 | 750-043596 | ABBV7188     | Adapter Card           |
| Fan Tray 0 | REV 03 | 760-046960 | ACAY0127     | 172mm FanTray - 6 Fans |
| Fan Tray 1 | REV 2A | 760-046960 | ACAY0068     | 172mm FanTray - 6 Fans |
| Fan Tray 2 | REV 2A | 760-046960 | ACAY0072     | 172mm FanTray - 6 Fans |
| Fan Tray 3 | REV 2A | 760-046960 | ACAY0070     | 172mm FanTray - 6 Fans |

### show chassis hardware extensive (MX2010 Router)

```

user@host > show chassis hardware extensive
Hardware inventory:
Item Version Part number Serial number Description
Chassis JN11E233DAFK MX2010
Jedec Code: 0x7fb0 EEPROM Version: 0x02
 S/N: JN11E233DAFK
Assembly ID: 0x0557 Assembly Version: 00.00
Date: 00-00-0000 Assembly Flags: 0x00
ID: MX2010
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 05 57 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: 4a 4e 31 31 45 32 33 33 44 41 46 4b 00 00 00 00
Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane REV 26 750-044636 ABAB9357 Lower Backplane
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 750-044636 S/N: S/N ABAB9357
Assembly ID: 0x0b66 Assembly Version: 01.26
Date: 08-28-2012 Assembly Flags: 0x00
Version: REV 26 CLEI Code: PROTOXCLEI

```

```

ID: Lower Backplane FRU Model Number: PROTO-ASSEMBLY
Board Information Record:
 Address 0x00: ad 01 08 00 2c 21 72 70 a0 00 ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 02 ff 0b 66 01 1a 52 45 56 20 32 36 00 00
 Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 36 33 36 00 00
 Address 0x20: 53 2f 4e 20 41 42 41 42 39 33 35 37 00 1c 08 07
 Address 0x30: dc ff ff ff ad 01 08 00 2c 21 72 70 a0 00 ff ff
 Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
 Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
 Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
 Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff
Midplane 1 REV 01 711-044557 ABAB8643 Upper Backplane
Jedec Code: 0x7fb0 EEPROM Version: 0x01
P/N: 711-044557 S/N: S/N ABAB8643
Assembly ID: 0x0b65 Assembly Version: 01.01
Date: 07-27-2012 Assembly Flags: 0x00
Version: REV 01
ID: Upper Backplane
Board Information Record:
 Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 01 ff 0b 65 01 01 52 45 56 20 30 31 00 00
 Address 0x10: 00 00 00 00 37 31 31 2d 30 34 34 35 35 37 00 00
 Address 0x20: 53 2f 4e 20 41 42 41 42 38 36 34 33 00 1b 07 07
 Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
 Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
PMP REV 04 711-032426 ACAJ1677 Power Midplane
Jedec Code: 0x7fb0 EEPROM Version: 0x01
P/N: 711-032426 S/N: S/N ACAJ1677
Assembly ID: 0x045d Assembly Version: 01.04
Date: 07-20-2012 Assembly Flags: 0x00
Version: REV 04
ID: Power Midplane
Board Information Record:
 Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 01 ff 04 5d 01 04 52 45 56 20 30 34 00 00
 Address 0x10: 00 00 00 00 37 31 31 2d 30 33 32 34 32 36 00 00
 Address 0x20: 53 2f 4e 20 41 43 41 4a 31 36 37 37 00 14 07 07
 Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
 Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
 Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM Board REV 08 760-044634 ABBV9726 Front Panel Display
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 760-044634 S/N: S/N ABBV9726
Assembly ID: 0x0b64 Assembly Version: 01.08
Date: 09-10-2012 Assembly Flags: 0x00
Version: REV 08 CLEI Code: IPMYA4EJRA
ID: Front Panel Display FRU Model Number: MX2010-CRAFT-S
Board Information Record:
 Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
 Address 0x00: 7f b0 02 ff 0b 64 01 08 52 45 56 20 30 38 00 00
 Address 0x10: 00 00 00 00 37 36 30 2d 30 34 34 36 33 34 00 00
 Address 0x20: 53 2f 4e 20 41 42 42 56 39 37 32 36 00 0a 09 07

```

```

Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 4d 59 41 34 45 4a 52 41 4d
Address 0x50: 58 32 30 31 30 2d 43 52 41 46 54 2d 53 00 00 00
Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 93 ff ff ff ff ff ff ff ff ff ff ff ff
PSM 0 REV 01 740-045050 1E02224000P DC 52V Power Supply
Module
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 740-045050 S/N: 1E02224000P
Assembly ID: 0x0478 Assembly Version: 01.01
Date: 12-06-2012 Assembly Flags: 0x00
Version: REV 01 CLEI Code: XXXXXXXXXX
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-HC-DC-S-A
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 30 35 30 00 00
Address 0x20: 31 45 30 32 32 32 34 30 30 30 50 00 00 06 0c 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 58 58 58 58 58 58 58 58 58 58 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 48 43 2d 44 43 2d
Address 0x60: 53 2d 41 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 4a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 1 REV 01 740-045050 1E02224000M DC 52V Power Supply
Module
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 740-045050 S/N: 1E02224000M
Assembly ID: 0x0478 Assembly Version: 01.01
Date: 12-06-2012 Assembly Flags: 0x00
Version: REV 01 CLEI Code: XXXXXXXXXX
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-HC-DC-S-A
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 30 35 30 00 00
Address 0x20: 31 45 30 32 32 32 34 30 30 30 4d 00 00 06 0c 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 58 58 58 58 58 58 58 58 58 58 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 48 43 2d 44 43 2d
Address 0x60: 53 2d 41 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 4a 00 00 00 00 00 00 00 00 00 00 00 00
...
PDM 0 REV 01 740-045234 1E262250067 DC Power Dist Module
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 740-045234 S/N: 1E262250067
Assembly ID: 0x047b Assembly Version: 01.01
Date: 06-28-2012 Assembly Flags: 0x00
Version: REV 01 CLEI Code: IPUPAJSKAA
ID: DC Power Dist Module FRU Model Number: MX2000-PDM-DC-S-A
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 7b 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 32 33 34 00 00
Address 0x20: 31 45 32 36 32 32 35 30 30 36 37 00 00 1c 06 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4a 53 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 44 4d 2d 44 43 2d 53 2d 41
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff

```

```

Address 0x70: ff ff ff 89 00 00 00 00 00 00 00 00 00 00 00 00
Routing Engine 0 REV 02 740-041821 9009099704 RE-S-1800x4
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 740-041821 S/N: 9009099704
Assembly ID: 0x09c0 Assembly Version: 01.02
Date: 03-15-2012 Assembly Flags: 0x00
Version: REV 02
ID: RE-S-1800x4 FRU Model Number: RE-S-1800X4-16G-S
Board Information Record:
Address 0x00: 54 32 30 32 37 44 41 2d 34 34 47 42 23 41 23 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 c0 01 02 52 45 56 20 30 32 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 31 38 32 31 00 00
Address 0x20: 39 30 30 39 30 39 39 37 30 34 00 00 00 0f 03 07
Address 0x30: dc ff ff ff 54 32 30 32 37 44 41 2d 34 34 47 42
Address 0x40: 23 41 23 00 01 00 00 00 00 00 00 00 00 00 00 52
Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 8c ff ff ff ff ff ff ff ff ff ff ff ff
ad0 3831 MB UGB30SFA4000T1 SFA4000T1 00000651 Compact Flash
ad1 30533 MB UGB94BPH32H0S1-KCI 11000019592 Disk 1
usb0 (addr 1) EHCI root hub 0 Intel uhub0
usb0 (addr 2) product 0x0020 32 vendor 0x8087 uhub1
DIMM 0 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 1 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 2 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 3 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1 REV 02 740-041821 9009099706 RE-S-1800x4
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 740-041821 S/N: 9009099706
Assembly ID: 0x09c0 Assembly Version: 01.02
Date: 02-23-2012 Assembly Flags: 0x00
Version: REV 02
ID: RE-S-1800x4 FRU Model Number: RE-S-1800X4-16G-S
Board Information Record:
Address 0x00: 54 32 30 32 37 44 41 2d 34 34 47 42 23 41 23 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 c0 01 02 52 45 56 20 30 32 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 31 38 32 31 00 00
Address 0x20: 39 30 30 39 30 39 39 37 30 36 00 00 00 17 02 07
Address 0x30: dc ff ff ff 54 32 30 32 37 44 41 2d 34 34 47 42
Address 0x40: 23 41 23 00 01 00 00 00 00 00 00 00 00 00 00 52
Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 8c ff ff ff ff ff ff ff ff ff ff ff ff
ad0 3998 MB Virtium - TuffDrive VCF P1T0200262860208 114 Compact Flash
ad1 30533 MB UGB94ARF32H0S3-KC UNIGEN-499551-000404 Disk 1
CB 0 REV 13 750-040257 CAAF8436 Control Board
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 750-040257 S/N: S/N CAAF8436
Assembly ID: 0x0b26 Assembly Version: 01.13
Date: 08-29-2012 Assembly Flags: 0x00
Version: REV 13 CLEI Code: PROTOXCLEI
ID: Control Board FRU Model Number: PROTO-ASSEMBLY
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 26 01 0d 52 45 56 20 31 33 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 34 30 32 35 37 00 00
Address 0x20: 53 2f 4e 20 43 41 41 46 38 34 33 36 00 1d 08 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

```

```

Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff

...
SPMB 0 REV 02 711-041855 ABBV3825 PMB Board
Jedec Code: 0x7fb0 EEPROM Version: 0x01
P/N: 711-041855 S/N: S/N ABBV3825
Assembly ID: 0x0b29 Assembly Version: 01.02
Date: 08-14-2012 Assembly Flags: 0x00
Version: REV 02
ID: PMB Board
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 0b 29 01 02 52 45 56 20 30 32 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 34 31 38 35 35 00 00
Address 0x20: 53 2f 4e 20 41 42 42 56 33 38 32 35 00 0e 08 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00

...
SFB 0 REV 05 711-044466 ABBX5682 Switch Fabric Board
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 711-044466 S/N: S/N ABBX5682
Assembly ID: 0x0b25 Assembly Version: 01.05
Date: 09-07-2012 Assembly Flags: 0x00
Version: REV 05 CLEI Code: PROTOXCLEI
ID: Switch Fabric Board FRU Model Number: PROTO-ASSEMBLY
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 25 01 05 52 45 56 20 30 35 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 34 34 36 36 00 00
Address 0x20: 53 2f 4e 20 41 42 42 58 35 36 38 32 00 07 09 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c2 00 00 00 01 00 00 00 00 00 00 48 00

...
FPC 0 REV 09 750-037355 CAAF0924 MPC Type 4-2
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 750-037355 S/N: S/N CAAF0924
Assembly ID: 0x0b4e Assembly Version: 01.09
Date: 05-21-2012 Assembly Flags: 0x00
Version: REV 09 CLEI Code: PROTOXCLEI
ID: MPC Type 4-2 FRU Model Number: MPC4E-2CGE-8XGE
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 4e 01 09 52 45 56 20 30 39 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 33 35 35 00 00
Address 0x20: 53 2f 4e 20 43 41 41 46 30 39 32 34 00 15 05 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 4d
Address 0x50: 50 43 34 45 2d 32 43 47 45 2d 38 58 47 45 00 00
Address 0x60: 00 00 00 00 00 00 30 39 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c6 ff ff ff ff ff ff ff ff ff ff ff ff

```

```

CPU REV 08 711-035209 CAAB9842 HMPC PMB 2G
Jedec Code: 0x7fb0 EEPROM Version: 0x01
P/N: 711-035209 S/N: S/N CAAB9842
Assembly ID: 0x0b04 Assembly Version: 01.08
Date: 05-17-2012 Assembly Flags: 0x00
Version: REV 08
ID: HMPC PMB 2G
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 0b 04 01 08 52 45 56 20 30 38 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 35 32 30 39 00 00
Address 0x20: 53 2f 4e 20 43 41 41 42 39 38 34 32 00 11 05 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
PIC 0 BUILTIN BUILTIN 4x10GE SFPP
Jedec Code: 0x0000 EEPROM Version: 0x00
P/N: BUILTIN S/N: BUILTIN
Assembly ID: 0x0a53 Assembly Version: 00.00
Date: 00-00-0000 Assembly Flags: 0x00
ID: 4x10GE SFPP
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 00 00 00 00 0a 53 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 4d 58 43 00
Address 0x20: 42 55 49 4c 54 49 4e 00 4d 58 43 00 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 c0 02 ae 64 00 00 00 00 0a 52 00 00
Xcvr 0 REV 01 740-021308 19T511101656 SFP+-10G-SR
Xcvr 1 REV 01 740-031980 AMA04RU SFP+-10G-SR
Xcvr 2 REV 01 740-031980 193363A00558 SFP+-10G-SR
Xcvr 3 REV 01 740-031980 B10M00202 SFP+-10G-SR
...
ADC 0 REV 13 750-043596 ABBX5532 Adapter Card
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 750-043596 S/N: S/N ABBX5532
Assembly ID: 0x0b3d Assembly Version: 01.13
Date: 09-12-2012 Assembly Flags: 0x00
Version: REV 13 CLEI Code: IPUCBA8CAA
ID: Adapter Card FRU Model Number: MX2000-LC-ADAPTER
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 3d 01 0d 52 45 56 20 31 33 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 34 33 35 39 36 00 00
Address 0x20: 53 2f 4e 20 41 42 42 58 35 35 33 32 00 0c 09 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 43 42 41 38 43 41 41 4d
Address 0x50: 58 32 30 30 30 2d 4c 43 2d 41 44 41 50 54 45 52
Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 3a 00 00 00 00 00 00 00 00 00 00 00 00
...

```

## show chassis hardware models (MX2010 Router)

user@host &gt; show chassis hardware models

Hardware inventory:

| Item             | Version | Part number | Serial number | FRU model number         |
|------------------|---------|-------------|---------------|--------------------------|
| FPM Board        | REV 06  | 711-032349  | ZX8744        | 711-032349               |
| PSM 4            | REV 0C  | 740-033727  | VK00254       | 000000000000000000000000 |
| PSM 5            | REV 0B  | 740-033727  | VG00015       | 000000000000000000000000 |
| PSM 6            | REV 0B  | 740-033727  | VH00097       | 000000000000000000000000 |
| PSM 7            | REV 0C  | 740-033727  | VJ00151       | 000000000000000000000000 |
| PSM 8            | REV 0C  | 740-033727  | VJ00149       | 000000000000000000000000 |
| PDM 0            | REV 0B  | 740-038109  | WA00008       |                          |
| PDM 1            | REV 0B  | 740-038109  | WA00014       |                          |
| Routing Engine 0 | REV 02  | 740-041821  | 9009094134    | RE-S-1800X4-16G-S        |
| Routing Engine 1 | REV 02  | 740-041821  | 9009094141    | RE-S-1800X4-16G-S        |
| CB 0             | REV 08  | 750-040257  | CAAB3491      | 750-040257               |
| CB 1             | REV 08  | 750-040257  | CAAB3489      | 750-040257               |
| SFB 0            | REV 06  | 711-032385  | ZV1828        | 711-032385               |
| SFB 1            | REV 07  | 711-032385  | ZZ2568        | 711-032385               |
| SFB 2            | REV 07  | 711-032385  | ZZ2563        | 711-032385               |
| SFB 3            | REV 07  | 711-032385  | ZZ2564        | 711-032385               |
| SFB 4            | REV 07  | 711-032385  | ZZ2580        | 711-032385               |
| SFB 5            | REV 07  | 711-032385  | ZZ2579        | 711-0323856              |
| SFB 6            | REV 07  | 711-032385  | CAAB4882      | 711-044170               |
| SFB 7            | REV 07  | 711-032385  | CAAB4898      | 711-044170               |
| FPC 0            | REV 33  | 750-028467  | CAAB1919      | MPC-3D-16XGE-SFPP        |
| FPC 1            | REV 21  | 750-033205  | ZG5027        | MX-MPC3-3D               |
| MIC 0            | REV 03  | 750-033307  | ZV6299        | MIC3-3D-10XGE-SFPP       |
| MIC 1            | REV 03  | 750-033307  | ZV6268        | MIC3-3D-10XGE-SFPP       |
| FPC 8            | REV 22  | 750-031089  | ZT9746        | MX-MPC2-3D               |
| MIC 0            | REV 26  | 750-028392  | ABBS1150      | MIC-3D-20GE-SFP          |
| MIC 1            | REV 26  | 750-028387  | ABBR9582      | MIC-3D-4XGE-XFP          |
| FPC 9            | REV 11  | 750-036284  | ZL3591        | MPCE-3D-16XGE-SFPP       |
| ADC 0            | REV 05  | 750-043596  | CAAC2073      | 750-043596               |
| ADC 1            | REV 01  | 750-043596  | ZV4117        | 750-043596               |
| ADC 8            | REV 01  | 750-043596  | ZV4107        | 750-043596               |
| ADC 9            | REV 02  | 750-043596  | ZW1555        | 750-043596               |
| Fan Tray 0       | REV 2A  | 760-046960  | ACAY0015      |                          |
| Fan Tray 1       | REV 2A  | 760-046960  | ACAY0019      |                          |
| Fan Tray 2       | REV 2A  | 760-046960  | ACAY0020      |                          |
| Fan Tray 3       | REV 2A  | 760-046960  | ACAY0021      |                          |

## show chassis hardware clei-models (MX2010 Routers)

user@host &gt; show chassis hardware clei-models

Hardware inventory:

| Item             | Version | Part number | CLEI code  | FRU model number         |
|------------------|---------|-------------|------------|--------------------------|
| FPM Board        | REV 06  | 711-032349  | PROTOXCLEI | 711-032349               |
| PSM 4            | REV 0C  | 740-033727  | 0000000000 | 000000000000000000000000 |
| PSM 5            | REV 0B  | 740-033727  | 0000000000 | 000000000000000000000000 |
| PSM 6            | REV 0B  | 740-033727  | 0000000000 | 000000000000000000000000 |
| PSM 7            | REV 0C  | 740-033727  | 0000000000 | 000000000000000000000000 |
| PSM 8            | REV 0C  | 740-033727  | 0000000000 | 000000000000000000000000 |
| PDM 0            | REV 0B  | 740-038109  |            |                          |
| PDM 1            | REV 0B  | 740-038109  |            |                          |
| Routing Engine 0 | REV 02  | 740-041821  |            | RE-S-1800X4-16G-S        |
| Routing Engine 1 | REV 02  | 740-041821  |            | RE-S-1800X4-16G-S        |
| CB 0             | REV 08  | 750-040257  | PROTOXCLEI | 750-040257               |
| CB 1             | REV 08  | 750-040257  | PROTOXCLEI | 750-040257               |
| SFB 0            | REV 06  | 711-032385  | PROTOXCLEI | 711-032385               |
| SFB 1            | REV 07  | 711-032385  | PROTOXCLEI | 711-032385               |



|            |        |            |            |                    |
|------------|--------|------------|------------|--------------------|
| SFB 2      | REV 07 | 711-032385 | PROTOXCLEI | 711-032385         |
| SFB 3      | REV 07 | 711-032385 | PROTOXCLEI | 711-032385         |
| SFB 4      | REV 07 | 711-032385 | PROTOXCLEI | 711-032385         |
| SFB 5      | REV 07 | 711-032385 | PROTOXCLEI | 711-0323856        |
| SFB 6      | REV 07 | 711-032385 | PROTOXCLEI | 711-044170         |
| SFB 7      | REV 07 | 711-032385 | PROTOXCLEI | 711-044170         |
| FPC 0      | REV 33 | 750-028467 |            | MPC-3D-16XGE-SFPP  |
| FPC 1      | REV 21 | 750-033205 |            | MX-MPC3-3D         |
| MIC 0      | REV 03 | 750-033307 | PROTOXCLEI | MIC3-3D-10XGE-SFPP |
| MIC 1      | REV 03 | 750-033307 | PROTOXCLEI | MIC3-3D-10XGE-SFPP |
| FPC 8      | REV 22 | 750-031089 | COUIBAYBAA | MX-MPC2-3D         |
| MIC 0      | REV 26 | 750-028392 | COUIA15BAA | MIC-3D-20GE-SFP    |
| MIC 1      | REV 26 | 750-028387 | COUIA16BAA | MIC-3D-4XGE-XFP    |
| FPC 9      | REV 11 | 750-036284 | CMUIACGBAA | MPCE-3D-16XGE-SFPP |
| ADC 0      | REV 05 | 750-043596 | PROTOXCLEI | 750-043596         |
| ADC 1      | REV 01 | 750-043596 | PROTOXCLEI | 750-043596         |
| ADC 8      | REV 01 | 750-043596 | PROTOXCLEI | 750-043596         |
| ADC 9      | REV 02 | 750-043596 | PROTOXCLEI | 750-043596         |
| Fan Tray 0 | REV 2A | 760-046960 |            |                    |
| Fan Tray 1 | REV 2A | 760-046960 |            |                    |
| Fan Tray 2 | REV 2A | 760-046960 |            |                    |
| Fan Tray 3 | REV 2A | 760-046960 |            |                    |

#### show chassis hardware (MX2020 Router)

```
user@host > show chassis hardware
```

```
Hardware inventory:
```

| Item       | Version | Part number | Serial number | Description          |
|------------|---------|-------------|---------------|----------------------|
| Chassis    |         |             | JN11E2227AFJ  | MX2020               |
| Midplane   | REV 27  | 750-040240  | ABAB9384      | Lower Power Midplane |
| Midplane 1 | REV 04  | 711-032386  | ABAB9386      | Upper Backplane      |
| PMP 1      | REV 05  | 711-032428  | ACAJ1579      | Upper Power Midplane |
| PMP 0      | REV 04  | 711-032426  | ACAJ1524      | Lower Power Midplane |
| FPM Board  | REV 06  | 760-040242  | ABBT8837      | Front Panel Display  |
| PSM 0      | REV 01  | 740-045050  | 1E022240056   | DC 52V Power Supply  |
| Module     |         |             |               |                      |
| PSM 1      | REV 01  | 740-045050  | 1E022240054   | DC 52V Power Supply  |
| Module     |         |             |               |                      |
| PSM 2      | REV 01  | 740-045050  | 1E02224005H   | DC 52V Power Supply  |
| Module     |         |             |               |                      |
| PSM 3      | REV 01  | 740-045050  | 1E022240053   | DC 52V Power Supply  |
| Module     |         |             |               |                      |
| PSM 4      | REV 01  | 740-045050  | 1E02224004K   | DC 52V Power Supply  |
| Module     |         |             |               |                      |
| PSM 7      | REV 01  | 740-045050  | 1E02224006W   | DC 52V Power Supply  |
| Module     |         |             |               |                      |
| PSM 8      | REV 01  | 740-045050  | 1E022240062   | DC 52V Power Supply  |
| Module     |         |             |               |                      |
| PSM 9      | REV 01  | 740-045050  | 1E02224005B   | DC 52V Power Supply  |
| Module     |         |             |               |                      |
| PSM 10     | REV 01  | 740-045050  | 1E02224005A   | DC 52V Power Supply  |
| Module     |         |             |               |                      |
| PSM 11     | REV 01  | 740-045050  | 1E022240052   | DC 52V Power Supply  |
| Module     |         |             |               |                      |
| PSM 12     | REV 01  | 740-045050  | 1E022240051   | DC 52V Power Supply  |
| Module     |         |             |               |                      |
| PSM 13     | REV 01  | 740-045050  | 1E022240058   | DC 52V Power Supply  |
| Module     |         |             |               |                      |
| PSM 14     | REV 01  | 740-045050  | 1E02224004L   | DC 52V Power Supply  |
| Module     |         |             |               |                      |
| PSM 15     | REV 01  | 740-045050  | 1E02224005M   | DC 52V Power Supply  |

|                  |        |            |             |                      |  |
|------------------|--------|------------|-------------|----------------------|--|
| Module           |        |            |             |                      |  |
| PSM 16           | REV 01 | 740-045050 | 1E02224006S | DC 52V Power Supply  |  |
| Module           |        |            |             |                      |  |
| PSM 17           | REV 01 | 740-045050 | 1E02224005Z | DC 52V Power Supply  |  |
| Module           |        |            |             |                      |  |
| PDM 0            | REV 01 | 740-045234 | 1E012150033 | DC Power Dist Module |  |
| PDM 1            | REV 01 | 740-045234 | 1E012150027 | DC Power Dist Module |  |
| PDM 2            | REV 01 | 740-045234 | 1E012150028 | DC Power Dist Module |  |
| PDM 3            | REV 01 | 740-045234 | 1E012150045 | DC Power Dist Module |  |
| Routing Engine 0 | REV 02 | 740-041821 | 9009089704  | RE-S-1800x4          |  |
| Routing Engine 1 | REV 02 | 740-041821 | 9009094138  | RE-S-1800x4          |  |
| CB 0             | REV 14 | 750-040257 | CAAF8430    | Control Board        |  |
| CB 1             | REV 08 | 750-040257 | CAAB3482    | Control Board        |  |
| SPMB 0           | REV 01 | 711-041855 | ZS2290      | PMB Board            |  |
| SPMB 1           | REV 02 | 711-041855 | CAA6141     | PMB Board            |  |
| SFB 0            | REV 03 | 711-044466 | ABBV6789    | Switch Fabric Board  |  |
| SFB 1            | REV 05 | 711-044466 | ABBX5666    | Switch Fabric Board  |  |
| SFB 2            | REV 05 | 711-044466 | ABBX5678    | Switch Fabric Board  |  |
| SFB 3            | REV 05 | 711-044466 | ABBX5687    | Switch Fabric Board  |  |
| SFB 4            | REV 05 | 711-044466 | ABBX5609    | Switch Fabric Board  |  |
| SFB 5            | REV 05 | 711-044466 | ABBX5675    | Switch Fabric Board  |  |
| SFB 6            | REV 03 | 711-044466 | ABBV6805    | Switch Fabric Board  |  |
| SFB 7            | REV 05 | 711-044466 | ABBX5701    | Switch Fabric Board  |  |
| FPC 0            | REV 30 | 750-028467 | ABBN0284    | MPC 3D 16x 10GE      |  |
| CPU              | REV 10 | 711-029089 | ABBN0507    | AMPC PMB             |  |
| PIC 0            |        | BUILTIN    | BUILTIN     | 4x 10GE(LAN) SFP+    |  |
| Xcvr 0           | REV 01 | 740-030658 | B11E00990   | SFP+-10G-USR         |  |
| Xcvr 1           | REV 01 | 740-030658 | B11E04357   | SFP+-10G-USR         |  |
| Xcvr 2           | REV 01 | 740-030658 | B11F01327   | SFP+-10G-USR         |  |
| Xcvr 3           | REV 01 | 740-030658 | B11E04375   | SFP+-10G-USR         |  |
| PIC 1            |        | BUILTIN    | BUILTIN     | 4x 10GE(LAN) SFP+    |  |
| Xcvr 0           | REV 01 | 740-030658 | B11E02760   | SFP+-10G-USR         |  |
| Xcvr 1           | REV 01 | 740-030658 | B11E02904   | SFP+-10G-USR         |  |
| Xcvr 2           | REV 01 | 740-030658 | B11E03963   | SFP+-10G-USR         |  |
| Xcvr 3           | REV 01 | 740-030658 | B11E00756   | SFP+-10G-USR         |  |
| PIC 2            |        | BUILTIN    | BUILTIN     | 4x 10GE(LAN) SFP+    |  |
| Xcvr 0           | REV 01 | 740-030658 | B11E04418   | SFP+-10G-USR         |  |
| Xcvr 1           | REV 01 | 740-030658 | B11E01077   | SFP+-10G-USR         |  |
| Xcvr 2           | REV 01 | 740-030658 | B11E01128   | SFP+-10G-USR         |  |
| Xcvr 3           | REV 01 | 740-030658 | B11F01253   | SFP+-10G-USR         |  |
| PIC 3            |        | BUILTIN    | BUILTIN     | 4x 10GE(LAN) SFP+    |  |
| Xcvr 0           | REV 01 | 740-030658 | B11E01140   | SFP+-10G-USR         |  |
| Xcvr 1           | REV 01 | 740-030658 | B11F01626   | SFP+-10G-USR         |  |
| Xcvr 2           | REV 01 | 740-030658 | B11E01075   | SFP+-10G-USR         |  |
| Xcvr 3           | REV 01 | 740-030658 | B11E01177   | SFP+-10G-USR         |  |
| FPC 1            | REV 30 | 750-028467 | ABBN0208    | MPC 3D 16x 10GE      |  |
| CPU              | REV 10 | 711-029089 | ABBJ1084    | AMPC PMB             |  |
| PIC 0            |        | BUILTIN    | BUILTIN     | 4x 10GE(LAN) SFP+    |  |
| Xcvr 0           | REV 01 | 740-030658 | B11E04745   | SFP+-10G-USR         |  |
| Xcvr 1           | REV 01 | 740-030658 | B11F01570   | SFP+-10G-USR         |  |
| Xcvr 2           | REV 01 | 740-030658 | B11E04388   | SFP+-10G-USR         |  |
| Xcvr 3           | REV 01 | 740-030658 | B11F01439   | SFP+-10G-USR         |  |
| PIC 1            |        | BUILTIN    | BUILTIN     | 4x 10GE(LAN) SFP+    |  |
| Xcvr 0           | REV 01 | 740-030658 | B11E04739   | SFP+-10G-USR         |  |
| Xcvr 1           | REV 01 | 740-030658 | B11F01869   | SFP+-10G-USR         |  |
| Xcvr 2           | REV 01 | 740-030658 | B11F01675   | SFP+-10G-USR         |  |
| Xcvr 3           | REV 01 | 740-030658 | B11F01901   | SFP+-10G-USR         |  |
| PIC 2            |        | BUILTIN    | BUILTIN     | 4x 10GE(LAN) SFP+    |  |
| Xcvr 0           | REV 01 | 740-030658 | B11F01346   | SFP+-10G-USR         |  |
| Xcvr 1           | REV 01 | 740-030658 | B11F01288   | SFP+-10G-USR         |  |
| Xcvr 2           | REV 01 | 740-030658 | B11F01824   | SFP+-10G-USR         |  |

|        |        |            |           |                   |
|--------|--------|------------|-----------|-------------------|
| Xcvr 3 | REV 01 | 740-030658 | B11E04312 | SFP+-10G-USR      |
| PIC 3  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E02811 | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E03847 | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11F01495 | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11F01265 | SFP+-10G-USR      |
| FPC 2  | REV 30 | 750-028467 | ZM5111    | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ZP6607    | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80LJA   | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80MFZ   | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80NKL   | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80KF4   | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80FBJ   | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80MM2   | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80LJV   | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80NXV   | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80N1H   | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80NLS   | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80FL5   | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80NL9   | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80NG2   | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80KDU   | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80MG1   | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80MM0   | SFP+-10G-SR       |
| FPC 3  | REV 30 | 750-028467 | ABB0302   | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABB0495   | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11F01581 | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E01176 | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11F01251 | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E02752 | SFP+-10G-USR      |
| PIC 1  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E00786 | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E01020 | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E01023 | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E02819 | SFP+-10G-USR      |
| PIC 2  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E02812 | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11D04437 | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11F01279 | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11F01333 | SFP+-10G-USR      |
| PIC 3  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E00978 | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E01018 | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11F01784 | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-031980 | AK80NKP   | SFP+-10G-SR       |
| FPC 4  | REV 30 | 750-028467 | ABB0308   | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABB11095  | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E04305 | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E01147 | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E01195 | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11F01743 | SFP+-10G-USR      |
| PIC 1  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11F01892 | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E02880 | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E00725 | SFP+-10G-USR      |

|        |        |            |           |                   |
|--------|--------|------------|-----------|-------------------|
| Xcvr 3 | REV 01 | 740-030658 | B11E01057 | SFP+-10G-USR      |
| PIC 2  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E02816 | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11C04501 | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E02764 | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E00789 | SFP+-10G-USR      |
| PIC 3  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11F01250 | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E02847 | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E00787 | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E03803 | SFP+-10G-USR      |
| FPC 5  | REV 30 | 750-028467 | ABBN0316  | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABB11082  | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | B11K00523 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | B11K01848 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | B11K01865 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | B11K00540 | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | B11K00422 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | B11K00428 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | B11K00423 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | B11K01855 | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | B11K01847 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | B11K00526 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | B11K00529 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | B11K00525 | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | B11K00425 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | B11K00530 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | B11K01851 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | B11K00528 | SFP+-10G-SR       |
| FPC 6  | REV 32 | 750-028467 | ABBN6832  | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBN6534  | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80MB4   | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80FQ6   | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80N1F   | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80NLQ   | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80KDR   | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80FGJ   | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80N5G   | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80KD8   | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80LET   | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80N1X   | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80NRF   | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80NL2   | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80N3D   | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80MRB   | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80LEQ   | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80LER   | SFP+-10G-SR       |
| FPC 7  | REV 32 | 750-028467 | ABBN6811  | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBN7288  | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80NK8   | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80LJG   | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80LBU   | SFP+-10G-SR       |

|        |        |            |              |                   |
|--------|--------|------------|--------------|-------------------|
| Xcvr 3 | REV 01 | 740-031980 | AK80N21      | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80LEU      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80NLM      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80NL6      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80LES      | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80LEN      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80ME0      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80LMG      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80MM1      | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80MG7      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80KF9      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80NRQ      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80NLE      | SFP+-10G-SR       |
| FPC 8  | REV 23 | 750-028467 | YN2977       | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | YP1856       | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 183363A00875 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 183363A00851 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 183363A00772 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 183363A00882 | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 183363A00735 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 183363A00169 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 183363A00726 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 183363A00077 | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 183363A00168 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 183363A00676 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 183363A00732 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 183363A00091 | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 183363A00725 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 183363A00642 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 183363A00871 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 183363A00853 | SFP+-10G-SR       |
| FPC 9  | REV 32 | 750-028467 | ABBN6798     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBK6556     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-021308 | 9ZDZ06A00055 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 183363A00239 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-021308 | AD0915E003K  | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-021308 | AD0915E003A  | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80MRC      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80NL5      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80NKN      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80N3U      | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80N1T      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AJ808DJ      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80NG4      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80FND      | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80FKQ      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80NLT      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80NKR      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80LKM      | SFP+-10G-SR       |
| FPC 10 | REV 32 | 750-028467 | ABBN6813     | MPC 3D 16x 10GE   |

|        |        |            |              |                   |
|--------|--------|------------|--------------|-------------------|
| CPU    | REV 10 | 711-029089 | ABBK6542     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80NA3      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80NLF      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80MRH      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80KE4      | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-021308 | 973152A00030 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80L9H      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80ME8      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80NLR      | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80NG1      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80MCA      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80LFC      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80LEM      | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80N9X      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80LAC      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80LF2      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80N8T      | SFP+-10G-SR       |
| FPC 11 | REV 30 | 750-028467 | ABBN0281     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBN0526     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11F01326    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E03973    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E00950    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E00674    | SFP+-10G-USR      |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E00775    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E04461    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E01074    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E02821    | SFP+-10G-USR      |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E04501    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E00757    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11F01623    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E01022    | SFP+-10G-USR      |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E04359    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E02751    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E02736    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E01178    | SFP+-10G-USR      |
| FPC 12 | REV 32 | 750-028467 | ABBN6796     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBN7259     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | B11K01856    | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | B11K01853    | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | B11K01863    | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A02863 | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A02668 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02881 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A01671 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A02627 | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A02725 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02692 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A02730 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A03081 | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |

|        |        |            |              |                   |
|--------|--------|------------|--------------|-------------------|
| Xcvr 0 | REV 01 | 740-031980 | 163363A02736 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02568 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A02747 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A02579 | SFP+-10G-SR       |
| FPC 13 | REV 30 | 750-028467 | ABBN0270     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABB0966      | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80NL1      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80NXW      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80KD2      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80FMD      | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80NKQ      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80MGH      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80N38      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80NL7      | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80LEL      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80NKD      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80KCY      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80LHK      | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80M5J      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80MBE      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80NLG      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80LFH      | SFP+-10G-SR       |
| FPC 14 | REV 32 | 750-028467 | ABBN6790     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBK6515     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80LZM      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80MCC      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80KCM      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80KE0      | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-021310 | C10F99155    | SFP+-10G-LRM      |
| Xcvr 1 | REV 01 | 740-021310 | C10F99049    | SFP+-10G-LRM      |
| Xcvr 2 | REV 01 | 740-021310 | C10F99128    | SFP+-10G-LRM      |
| Xcvr 3 | REV 01 | 740-021310 | C10F99169    | SFP+-10G-LRM      |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80LF3      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02597 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A03060 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A03057 | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80LEX      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80FEU      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80FNM      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-021308 | AJQQQ5G      | SFP+-10G-SR       |
| FPC 15 | REV 32 | 750-028467 | ABBN6791     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBN7289     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | B11K00424    | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | B11K01849    | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | B11K01862    | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | B11K01852    | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | B11K00427    | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | B11K00430    | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | B11K01854    | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | B11K00426    | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |

|        |        |            |              |                   |
|--------|--------|------------|--------------|-------------------|
| Xcvr 0 | REV 01 | 740-031980 | B11K00429    | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | B11K01864    | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | B11K01850    | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | B11K00522    | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E01144    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E00985    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E00796    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-031980 | B11K01866    | SFP+-10G-SR       |
| FPC 16 | REV 30 | 750-028467 | ABBM4592     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBN0465     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11F01435    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E01052    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11F01328    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11F01254    | SFP+-10G-USR      |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E02738    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E02881    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11F01624    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E00889    | SFP+-10G-USR      |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E02883    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E00681    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E04306    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E02813    | SFP+-10G-USR      |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11F01801    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E02753    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E01156    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E04324    | SFP+-10G-USR      |
| FPC 17 | REV 32 | 750-028467 | ABBN6810     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBN7237     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A02638 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02082 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A01674 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A03058 | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A03048 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02729 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A02566 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A02567 | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A02878 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02739 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A01959 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A02660 | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A02731 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02588 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A02673 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A02654 | SFP+-10G-SR       |
| FPC 18 | REV 30 | 750-028467 | ABBM4739     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBN0487     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A02569 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02886 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A03082 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 133363A00297 | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |



|            |        |            |              |                        |
|------------|--------|------------|--------------|------------------------|
| Xcvr 0     | REV 01 | 740-031980 | 163363A02726 | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | 163363A03050 | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | 163363A02884 | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-031980 | 163363A03076 | SFP+-10G-SR            |
| PIC 2      |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+      |
| Xcvr 0     | REV 01 | 740-031980 | 163363A02581 | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | 163363A02873 | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | 163363A02582 | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-031980 | 163363A03083 | SFP+-10G-SR            |
| PIC 3      |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+      |
| Xcvr 0     | REV 01 | 740-031981 | UL70BU6      | SFP+-10G-LR            |
| Xcvr 1     | REV 01 | 740-031981 | UL50QC6      | SFP+-10G-LR            |
| Xcvr 2     | REV 01 | 740-031981 | UL708N6      | SFP+-10G-LR            |
| Xcvr 3     | REV 01 | 740-031981 | UL603KK      | SFP+-10G-LR            |
| FPC 19     | REV 32 | 750-028467 | ABBN6827     | MPC 3D 16x 10GE        |
| CPU        | REV 10 | 711-029089 | ABBK6508     | AMPC PMB               |
| PIC 0      |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+      |
| Xcvr 0     | REV 01 | 740-031980 | 163363A01688 | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | 163363A01724 | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | 163363A01773 | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-031980 | 163363A02593 | SFP+-10G-SR            |
| PIC 1      |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+      |
| Xcvr 0     | REV 01 | 740-031980 | 163363A03061 | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | 163363A03056 | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | 163363A02669 | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-031980 | 163363A03070 | SFP+-10G-SR            |
| PIC 2      |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+      |
| Xcvr 0     | REV 01 | 740-031980 | 163363A02572 | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | 163363A02697 | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | 163363A02585 | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-031980 | 163363A03052 | SFP+-10G-SR            |
| PIC 3      |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+      |
| Xcvr 0     | REV 01 | 740-031980 | 163363A02591 | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | 163363A02649 | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | 163363A02577 | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-031980 | 163363A02698 | SFP+-10G-SR            |
| ADC 0      | REV 13 | 750-043596 | ABBX5561     | Adapter Card           |
| ADC 1      | REV 13 | 750-043596 | ABBX5546     | Adapter Card           |
| ADC 2      | REV 13 | 750-043596 | ABBX5535     | Adapter Card           |
| ADC 3      | REV 13 | 750-043596 | ABBX5552     | Adapter Card           |
| ADC 4      | REV 13 | 750-043596 | ABBX5581     | Adapter Card           |
| ADC 5      | REV 13 | 750-043596 | ABBX5545     | Adapter Card           |
| ADC 6      | REV 13 | 750-043596 | ABBX5554     | Adapter Card           |
| ADC 7      | REV 07 | 750-043596 | ABBV7194     | Adapter Card           |
| ADC 8      | REV 07 | 750-043596 | ABBV7251     | Adapter Card           |
| ADC 9      | REV 07 | 750-043596 | ABBV7202     | Adapter Card           |
| ADC 10     | REV 13 | 750-043596 | ABBX5538     | Adapter Card           |
| ADC 11     | REV 13 | 750-043596 | ABBX5566     | Adapter Card           |
| ADC 12     | REV 13 | 750-043596 | ABBX5542     | Adapter Card           |
| ADC 13     | REV 13 | 750-043596 | ABBX5539     | Adapter Card           |
| ADC 14     | REV 13 | 750-043596 | ABBX5555     | Adapter Card           |
| ADC 15     | REV 13 | 750-043596 | ABBX5557     | Adapter Card           |
| ADC 16     | REV 13 | 750-043596 | ABBX5536     | Adapter Card           |
| ADC 17     | REV 13 | 750-043596 | ABBX5559     | Adapter Card           |
| ADC 18     | REV 13 | 750-043596 | ABBX5537     | Adapter Card           |
| ADC 19     | REV 11 | 750-043596 | ABBW5685     | Adapter Card           |
| Fan Tray 0 | REV 2A | 760-046960 | ACAY0030     | 172mm FanTray - 6 Fans |
| Fan Tray 1 | REV 2A | 760-046960 | ACAY0039     | 172mm FanTray - 6 Fans |
| Fan Tray 2 | REV 2A | 760-046960 | ACAY0033     | 172mm FanTray - 6 Fans |
| Fan Tray 3 | REV 2A | 760-046960 | ACAY0062     | 172mm FanTray - 6 Fans |

## show chassis hardware detail (MX2020 Router)

```

user@host> show chassis hardware detail
Hardware inventory:
Item Version Part number Serial number Description
Chassis JN11E2227AFJ MX2020
Midplane REV 27 750-040240 ABAB9384 Lower Power Midplane
Midplane 1 REV 04 711-032386 ABAB9386 Upper Backplane
PMP 1 REV 05 711-032428 ACAJ1821 Upper Power Midplane
PMP 0 REV 04 711-032426 ACAJ1524 Lower Power Midplane
FPM Board REV 06 760-040242 ABBT8837 Front Panel Display
PSM 0 REV 01 740-045050 1E02224006G DC 52V Power Supply
Module
PSM 1 REV 01 740-045050 1E022240053 DC 52V Power Supply
Module
PSM 2 REV 01 740-045050 1E02224004K DC 52V Power Supply
Module
PSM 3 REV 01 740-045050 1E022240056 DC 52V Power Supply
Module
PSM 4 REV 01 740-045050 1E022240054 DC 52V Power Supply
Module
PSM 5 REV 01 740-045050 1E02224005H DC 52V Power Supply
Module
PSM 6 REV 01 740-045050 1E02224006S DC 52V Power Supply
Module
PSM 7 REV 01 740-045050 1E02224005M DC 52V Power Supply
Module
PSM 8 REV 01 740-045050 1E022240062 DC 52V Power Supply
Module
PSM 9 REV 03 740-045050 1EDB2350095 DC 52V Power Supply
Module
PSM 10 REV 03 740-045050 1EDB235009L DC 52V Power Supply
Module
PSM 11 REV 03 740-045050 1EDB2350092 DC 52V Power Supply
Module
PSM 12 REV 03 740-045050 1EDB23500AT DC 52V Power Supply
Module
PSM 13 REV 03 740-045050 1EDB2350094 DC 52V Power Supply
Module
PSM 15 REV 03 740-045050 1EDB235008X DC 52V Power Supply
Module
PDM 0 REV 01 740-045234 1E012150033 DC Power Dist Module
PDM 1 REV 01 740-045234 1E012150027 DC Power Dist Module
PDM 2 REV 01 740-045234 1E262250072 DC Power Dist Module
Routing Engine 0 REV 02 740-041821 9009094138 RE-S-1800x4
 ad0 3998 MB Virtium - TuffDisk VCF3 20110825A021D0000064 Compact Flash
 ad1 30533 MB UGB94ARF32H0S3-KC UNIGEN-499551-000347 Disk 1
 usb0 (addr 1) EHCI root hub 0 Intel uhub0
 usb0 (addr 2) product 0x0020 32 vendor 0x8087 uhub1
 DIMM 0 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
 DIMM 1 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
 DIMM 2 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
 DIMM 3 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1 REV 02 740-041821 9009089709 RE-S-1800x4
 ad0 3831 MB UGB30SFA4000T1 SFA4000T1 00000113 Compact Flash
 ad1 30533 MB UGB94ARF32H0S3-KC UNIGEN-478612-001044 Disk 1
CB 0 REV 08 750-040257 CAAB3482 Control Board
CB 1 REV 04 750-040257 ZT2864 Control Board
SPMB 0 REV 02 711-041855 CAAA6141 PMB Board
SPMB 1 REV 01 711-041855 ZS2275 PMB Board
SFB 0 REV 05 711-044466 ABBT2161 Switch Fabric Board

```

|        |        |            |           |                     |
|--------|--------|------------|-----------|---------------------|
| SFB 1  | REV 05 | 711-044466 | ABBT2159  | Switch Fabric Board |
| SFB 2  | REV 05 | 711-044466 | ABBX3718  | Switch Fabric Board |
| SFB 3  | REV 05 | 711-044466 | ABBT2152  | Switch Fabric Board |
| SFB 4  | REV 05 | 711-044466 | ABBT2160  | Switch Fabric Board |
| SFB 5  | REV 05 | 711-044466 | ABBT2145  | Switch Fabric Board |
| SFB 6  | REV 05 | 711-044466 | ABBT2150  | Switch Fabric Board |
| SFB 7  | REV 05 | 711-044466 | ABBT2163  | Switch Fabric Board |
| FPC 0  | REV 30 | 750-028467 | ABBN0284  | MPC 3D 16x 10GE     |
| CPU    | REV 10 | 711-029089 | ABBN0507  | AMPC PMB            |
| PIC 0  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+   |
| Xcvr 0 | REV 01 | 740-030658 | B11E00990 | SFP+-10G-USR        |
| Xcvr 1 | REV 01 | 740-030658 | B11E04357 | SFP+-10G-USR        |
| Xcvr 2 | REV 01 | 740-030658 | B11F01327 | SFP+-10G-USR        |
| Xcvr 3 | REV 01 | 740-030658 | B11E04375 | SFP+-10G-USR        |
| PIC 1  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+   |
| Xcvr 0 | REV 01 | 740-030658 | B11E02760 | SFP+-10G-USR        |
| Xcvr 1 | REV 01 | 740-030658 | B11E02904 | SFP+-10G-USR        |
| Xcvr 2 | REV 01 | 740-030658 | B11E03963 | SFP+-10G-USR        |
| Xcvr 3 | REV 01 | 740-030658 | B11E00756 | SFP+-10G-USR        |
| PIC 2  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+   |
| Xcvr 0 | REV 01 | 740-030658 | B11E04418 | SFP+-10G-USR        |
| Xcvr 1 | REV 01 | 740-030658 | B11E01077 | SFP+-10G-USR        |
| Xcvr 2 | REV 01 | 740-030658 | B11E01128 | SFP+-10G-USR        |
| Xcvr 3 | REV 01 | 740-030658 | B11F01253 | SFP+-10G-USR        |
| PIC 3  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+   |
| Xcvr 0 | REV 01 | 740-030658 | B11E01140 | SFP+-10G-USR        |
| Xcvr 1 | REV 01 | 740-030658 | B11F01626 | SFP+-10G-USR        |
| Xcvr 2 | REV 01 | 740-030658 | B11E01075 | SFP+-10G-USR        |
| Xcvr 3 | REV 01 | 740-030658 | B11E01177 | SFP+-10G-USR        |
| FPC 1  | REV 30 | 750-028467 | ABBN0308  | MPC 3D 16x 10GE     |
| CPU    | REV 10 | 711-029089 | ABBJ1095  | AMPC PMB            |
| PIC 0  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+   |
| Xcvr 0 | REV 01 | 740-030658 | B11E04305 | SFP+-10G-USR        |
| Xcvr 1 | REV 01 | 740-030658 | B11E01147 | SFP+-10G-USR        |
| Xcvr 2 | REV 01 | 740-030658 | B11E01195 | SFP+-10G-USR        |
| Xcvr 3 | REV 01 | 740-030658 | B11F01743 | SFP+-10G-USR        |
| PIC 1  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+   |
| Xcvr 0 | REV 01 | 740-030658 | B11F01892 | SFP+-10G-USR        |
| Xcvr 1 | REV 01 | 740-030658 | B11E02880 | SFP+-10G-USR        |
| Xcvr 2 | REV 01 | 740-030658 | B11E00725 | SFP+-10G-USR        |
| Xcvr 3 | REV 01 | 740-030658 | B11E01057 | SFP+-10G-USR        |
| PIC 2  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+   |
| Xcvr 0 | REV 01 | 740-030658 | B11E02816 | SFP+-10G-USR        |
| Xcvr 1 | REV 01 | 740-030658 | B11C04501 | SFP+-10G-USR        |
| Xcvr 2 | REV 01 | 740-030658 | B11E02764 | SFP+-10G-USR        |
| Xcvr 3 | REV 01 | 740-030658 | B11E00789 | SFP+-10G-USR        |
| PIC 3  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+   |
| Xcvr 0 | REV 01 | 740-030658 | B11F01250 | SFP+-10G-USR        |
| Xcvr 1 | REV 01 | 740-030658 | B11E02847 | SFP+-10G-USR        |
| Xcvr 2 | REV 01 | 740-030658 | B11E00787 | SFP+-10G-USR        |
| Xcvr 3 | REV 01 | 740-030658 | B11E03803 | SFP+-10G-USR        |
| FPC 2  | REV 30 | 750-028467 | ABBN0316  | MPC 3D 16x 10GE     |
| CPU    | REV 10 | 711-029089 | ABBJ1082  | AMPC PMB            |
| PIC 0  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+   |
| Xcvr 0 | REV 01 | 740-031980 | B11K00523 | SFP+-10G-SR         |
| Xcvr 1 | REV 01 | 740-031980 | B11K01848 | SFP+-10G-SR         |
| Xcvr 2 | REV 01 | 740-031980 | B11K01865 | SFP+-10G-SR         |
| Xcvr 3 | REV 01 | 740-031980 | B11K00540 | SFP+-10G-SR         |
| PIC 1  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+   |
| Xcvr 0 | REV 01 | 740-031980 | B11K00422 | SFP+-10G-SR         |
| Xcvr 1 | REV 01 | 740-031980 | B11K00428 | SFP+-10G-SR         |

|        |        |            |           |                   |
|--------|--------|------------|-----------|-------------------|
| Xcvr 2 | REV 01 | 740-031980 | B11K00423 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | B11K01855 | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | B11K01847 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | B11K00526 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | B11K00529 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | B11K00525 | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | B11K00425 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | B11K00530 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | B11K01851 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | B11K00528 | SFP+-10G-SR       |
| FPC 3  | REV 32 | 750-028467 | ABBN6832  | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBK6534  | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80MB4   | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80FQ6   | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80N1F   | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80NLQ   | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80KDR   | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80FGJ   | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80N5G   | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80KD8   | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80LET   | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80N1X   | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80NRF   | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80NL2   | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80N3D   | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80MRB   | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80LEQ   | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80LER   | SFP+-10G-SR       |
| FPC 4  | REV 32 | 750-028467 | ABBN6811  | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBN7288  | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80NK8   | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80LJG   | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80LBU   | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80N21   | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80LEU   | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80NLM   | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80NL6   | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80LES   | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80LEN   | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80ME0   | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80LMG   | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80MM1   | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80MG7   | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80KF9   | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80NRQ   | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80NLE   | SFP+-10G-SR       |
| FPC 5  | REV 32 | 750-028467 | ABBN6791  | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBN7289  | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN   | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | B11K00424 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | B11K01849 | SFP+-10G-SR       |

|        |        |            |              |                   |
|--------|--------|------------|--------------|-------------------|
| Xcvr 2 | REV 01 | 740-031980 | B11K01862    | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | B11K01852    | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP  |
| Xcvr 0 | REV 01 | 740-031980 | B11K00427    | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | B11K00430    | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | B11K01854    | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | B11K00426    | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | B11K00429    | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | B11K01864    | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | B11K01850    | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | B11K00522    | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E01144    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E00985    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E00796    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-031980 | B11K01866    | SFP+-10G-SR       |
| FPC 6  | REV 30 | 750-028467 | ABBM4592     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBN0465     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11F01435    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E01052    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11F01328    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11F01254    | SFP+-10G-USR      |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E02738    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E02881    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11F01624    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E00889    | SFP+-10G-USR      |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E02883    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E00681    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E04306    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E02813    | SFP+-10G-USR      |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11F01801    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E02753    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E01156    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E04324    | SFP+-10G-USR      |
| FPC 7  | REV 32 | 750-028467 | ABBN6810     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBN7237     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A03058 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02082 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A01674 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A02638 | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A03048 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02729 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A02566 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A02567 | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A02878 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02739 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A01959 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A02660 | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A02731 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02588 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A02673 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A02654 | SFP+-10G-SR       |

|        |        |            |              |                   |
|--------|--------|------------|--------------|-------------------|
| FPC 8  | REV 30 | 750-028467 | ABBM4739     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBN0487     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A02569 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02886 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A03082 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 133363A00297 | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A02726 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A03050 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A02884 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A03076 | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A02581 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02873 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A02582 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A03083 | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031981 | UL70BU6      | SFP+-10G-LR       |
| Xcvr 1 | REV 01 | 740-031981 | UL50QC6      | SFP+-10G-LR       |
| Xcvr 2 | REV 01 | 740-031981 | UL708N6      | SFP+-10G-LR       |
| Xcvr 3 | REV 01 | 740-031981 | UL603KK      | SFP+-10G-LR       |
| FPC 9  | REV 32 | 750-028467 | ABBN6827     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBK6508     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A01688 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A01724 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A01773 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A02593 | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A03061 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A03056 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A02669 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A03070 | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A02572 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02697 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A02585 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A03052 | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A02591 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02649 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A02577 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A02698 | SFP+-10G-SR       |
| FPC 10 | REV 30 | 750-028467 | ABBN0302     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBN0495     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11F01581    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E01176    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11F01251    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E02752    | SFP+-10G-USR      |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E00786    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E01020    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E01023    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E02819    | SFP+-10G-USR      |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E02812    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11D04437    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11F01279    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11F01333    | SFP+-10G-USR      |

|        |        |            |              |                   |
|--------|--------|------------|--------------|-------------------|
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E00978    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E01018    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11F01784    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-031980 | AK80NKP      | SFP+-10G-SR       |
| FPC 11 | REV 32 | 750-028467 | ABBN6790     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBK6515     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80LZM      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80MCC      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80KCM      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80KE0      | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-021310 | C10F99155    | SFP+-10G-LRM      |
| Xcvr 1 | REV 01 | 740-021310 | C10F99049    | SFP+-10G-LRM      |
| Xcvr 2 | REV 01 | 740-021310 | C10F99128    | SFP+-10G-LRM      |
| Xcvr 3 | REV 01 | 740-021310 | C10F99169    | SFP+-10G-LRM      |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80LF3      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02597 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A03060 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A03057 | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80LEX      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80FEU      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80FNM      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-021308 | AJQQQ5G      | SFP+-10G-SR       |
| FPC 12 | REV 30 | 750-028467 | ZM5111       | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ZP6607       | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80LJA      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80MFZ      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80NKL      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80KF4      | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80FBJ      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80MM2      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80LJV      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80NXV      | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80N1H      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80NLS      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80FL5      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80NL9      | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80NG2      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80KDU      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80MG1      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80MM0      | SFP+-10G-SR       |
| FPC 13 | REV 30 | 750-028467 | ABBN0208     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABB11084     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E04745    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11F01570    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E04388    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11F01439    | SFP+-10G-USR      |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E04739    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11F01869    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11F01675    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11F01901    | SFP+-10G-USR      |

|        |        |            |              |                 |                   |
|--------|--------|------------|--------------|-----------------|-------------------|
| PIC 2  |        |            | BUILTIN      | BUILTIN         | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11F01346    | SFP+-10G-USR    |                   |
| Xcvr 1 | REV 01 | 740-030658 | B11F01288    | SFP+-10G-USR    |                   |
| Xcvr 2 | REV 01 | 740-030658 | B11F01824    | SFP+-10G-USR    |                   |
| Xcvr 3 | REV 01 | 740-030658 | B11E04312    | SFP+-10G-USR    |                   |
| PIC 3  |        |            | BUILTIN      | BUILTIN         | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E02811    | SFP+-10G-USR    |                   |
| Xcvr 1 | REV 01 | 740-030658 | B11E03847    | SFP+-10G-USR    |                   |
| Xcvr 2 | REV 01 | 740-030658 | B11F01495    | SFP+-10G-USR    |                   |
| Xcvr 3 | REV 01 | 740-030658 | B11F01265    | SFP+-10G-USR    |                   |
| FPC 14 | REV 23 | 750-028467 | YN2977       | MPC 3D 16x 10GE |                   |
| CPU    | REV 10 | 711-029089 | YP1856       | AMPC PMB        |                   |
| PIC 0  |        |            | BUILTIN      | BUILTIN         | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 183363A00875 | SFP+-10G-SR     |                   |
| Xcvr 1 | REV 01 | 740-031980 | 183363A00851 | SFP+-10G-SR     |                   |
| Xcvr 2 | REV 01 | 740-031980 | 183363A00772 | SFP+-10G-SR     |                   |
| Xcvr 3 | REV 01 | 740-031980 | 183363A00882 | SFP+-10G-SR     |                   |
| PIC 1  |        |            | BUILTIN      | BUILTIN         | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 183363A00735 | SFP+-10G-SR     |                   |
| Xcvr 1 | REV 01 | 740-031980 | 183363A00169 | SFP+-10G-SR     |                   |
| Xcvr 2 | REV 01 | 740-031980 | 183363A00726 | SFP+-10G-SR     |                   |
| Xcvr 3 | REV 01 | 740-031980 | 183363A00077 | SFP+-10G-SR     |                   |
| PIC 2  |        |            | BUILTIN      | BUILTIN         | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 183363A00168 | SFP+-10G-SR     |                   |
| Xcvr 1 | REV 01 | 740-031980 | 183363A00676 | SFP+-10G-SR     |                   |
| Xcvr 2 | REV 01 | 740-031980 | 183363A00732 | SFP+-10G-SR     |                   |
| Xcvr 3 | REV 01 | 740-031980 | 183363A00091 | SFP+-10G-SR     |                   |
| PIC 3  |        |            | BUILTIN      | BUILTIN         | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 183363A00725 | SFP+-10G-SR     |                   |
| Xcvr 1 | REV 01 | 740-031980 | 183363A00642 | SFP+-10G-SR     |                   |
| Xcvr 2 | REV 01 | 740-031980 | 183363A00871 | SFP+-10G-SR     |                   |
| Xcvr 3 | REV 01 | 740-031980 | 183363A00853 | SFP+-10G-SR     |                   |
| FPC 15 | REV 32 | 750-028467 | ABBN6798     | MPC 3D 16x 10GE |                   |
| CPU    | REV 10 | 711-029089 | ABBK6556     | AMPC PMB        |                   |
| PIC 0  |        |            | BUILTIN      | BUILTIN         | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-021308 | 9ZDZ06A00055 | SFP+-10G-SR     |                   |
| Xcvr 1 | REV 01 | 740-031980 | 183363A00239 | SFP+-10G-SR     |                   |
| Xcvr 2 | REV 01 | 740-021308 | AD0915E003K  | SFP+-10G-SR     |                   |
| Xcvr 3 | REV 01 | 740-021308 | AD0915E003A  | SFP+-10G-SR     |                   |
| PIC 1  |        |            | BUILTIN      | BUILTIN         | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80MRC      | SFP+-10G-SR     |                   |
| Xcvr 1 | REV 01 | 740-031980 | AK80NL5      | SFP+-10G-SR     |                   |
| Xcvr 2 | REV 01 | 740-031980 | AK80NKN      | SFP+-10G-SR     |                   |
| Xcvr 3 | REV 01 | 740-031980 | AK80N3U      | SFP+-10G-SR     |                   |
| PIC 2  |        |            | BUILTIN      | BUILTIN         | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80N1T      | SFP+-10G-SR     |                   |
| Xcvr 1 | REV 01 | 740-031980 | AJ808DJ      | SFP+-10G-SR     |                   |
| Xcvr 2 | REV 01 | 740-031980 | AK80NG4      | SFP+-10G-SR     |                   |
| Xcvr 3 | REV 01 | 740-031980 | AK80FND      | SFP+-10G-SR     |                   |
| PIC 3  |        |            | BUILTIN      | BUILTIN         | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80FKQ      | SFP+-10G-SR     |                   |
| Xcvr 1 | REV 01 | 740-031980 | AK80NLT      | SFP+-10G-SR     |                   |
| Xcvr 2 | REV 01 | 740-031980 | AK80NKR      | SFP+-10G-SR     |                   |
| Xcvr 3 | REV 01 | 740-031980 | AK80LKM      | SFP+-10G-SR     |                   |
| FPC 16 | REV 30 | 750-028467 | ABBN0270     | MPC 3D 16x 10GE |                   |
| CPU    | REV 10 | 711-029089 | ABBJ0966     | AMPC PMB        |                   |
| PIC 0  |        |            | BUILTIN      | BUILTIN         | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80NL1      | SFP+-10G-SR     |                   |
| Xcvr 1 | REV 01 | 740-031980 | AK80NXW      | SFP+-10G-SR     |                   |
| Xcvr 2 | REV 01 | 740-031980 | AK80KD2      | SFP+-10G-SR     |                   |
| Xcvr 3 | REV 01 | 740-031980 | AK80FMD      | SFP+-10G-SR     |                   |



|        |        |            |              |                   |
|--------|--------|------------|--------------|-------------------|
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80NKQ      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80MGH      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80N38      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80NL7      | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80M5J      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80NKD      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80KCY      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80LHK      | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | AK80LEL      | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | AK80MBE      | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | AK80NLG      | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | AK80LFH      | SFP+-10G-SR       |
| FPC 17 | REV 32 | 750-028467 | ABBN6796     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBN7259     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | B11K01856    | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | B11K01853    | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | B11K01863    | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A02863 | SFP+-10G-SR       |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A02668 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02881 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A01671 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A02627 | SFP+-10G-SR       |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A02725 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02692 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A02730 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A03081 | SFP+-10G-SR       |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-031980 | 163363A02736 | SFP+-10G-SR       |
| Xcvr 1 | REV 01 | 740-031980 | 163363A02568 | SFP+-10G-SR       |
| Xcvr 2 | REV 01 | 740-031980 | 163363A02747 | SFP+-10G-SR       |
| Xcvr 3 | REV 01 | 740-031980 | 163363A02579 | SFP+-10G-SR       |
| FPC 18 | REV 30 | 750-028467 | ABBN0281     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBN0526     | AMPC PMB          |
| PIC 0  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11F01326    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E03973    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E00950    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E00674    | SFP+-10G-USR      |
| PIC 1  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E00775    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E04461    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E01074    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E02821    | SFP+-10G-USR      |
| PIC 2  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E04501    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E00757    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11F01623    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E01022    | SFP+-10G-USR      |
| PIC 3  |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+ |
| Xcvr 0 | REV 01 | 740-030658 | B11E04359    | SFP+-10G-USR      |
| Xcvr 1 | REV 01 | 740-030658 | B11E02751    | SFP+-10G-USR      |
| Xcvr 2 | REV 01 | 740-030658 | B11E02736    | SFP+-10G-USR      |
| Xcvr 3 | REV 01 | 740-030658 | B11E01178    | SFP+-10G-USR      |
| FPC 19 | REV 32 | 750-028467 | ABBN6813     | MPC 3D 16x 10GE   |
| CPU    | REV 10 | 711-029089 | ABBK6542     | AMPC PMB          |

|            |        |            |              |                        |
|------------|--------|------------|--------------|------------------------|
| PIC 0      |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+      |
| Xcvr 0     | REV 01 | 740-031980 | AK80NA3      | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | AK80NLF      | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | AK80MRH      | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-031980 | AK80KE4      | SFP+-10G-SR            |
| PIC 1      |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+      |
| Xcvr 0     | REV 01 | 740-021308 | 973152A00030 | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | AK80L9H      | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | AK80ME8      | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-031980 | AK80NLR      | SFP+-10G-SR            |
| PIC 2      |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+      |
| Xcvr 0     | REV 01 | 740-031980 | AK80NG1      | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | AK80MCA      | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | AK80LFC      | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-031980 | AK80LEM      | SFP+-10G-SR            |
| PIC 3      |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+      |
| Xcvr 0     | REV 01 | 740-031980 | AK80N9X      | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | AK80LAC      | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | AK80LFC      | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-031980 | AK80N8T      | SFP+-10G-SR            |
| ADC 0      | REV 13 | 750-043596 | ABBX5561     | Adapter Card           |
| ADC 1      | REV 13 | 750-043596 | ABBX5546     | Adapter Card           |
| ADC 2      | REV 13 | 750-043596 | ABBX5535     | Adapter Card           |
| ADC 3      | REV 13 | 750-043596 | ABBX5552     | Adapter Card           |
| ADC 4      | REV 13 | 750-043596 | ABBX5581     | Adapter Card           |
| ADC 5      | REV 13 | 750-043596 | ABBX5545     | Adapter Card           |
| ADC 6      | REV 13 | 750-043596 | ABBX5554     | Adapter Card           |
| ADC 7      | REV 07 | 750-043596 | ABBV7194     | Adapter Card           |
| ADC 8      | REV 07 | 750-043596 | ABBV7251     | Adapter Card           |
| ADC 9      | REV 07 | 750-043596 | ABBV7202     | Adapter Card           |
| ADC 10     | REV 13 | 750-043596 | ABBX5579     | Adapter Card           |
| ADC 11     | REV 13 | 750-043596 | ABBX5548     | Adapter Card           |
| ADC 12     | REV 13 | 750-043596 | ABBX5575     | Adapter Card           |
| ADC 13     | REV 13 | 750-043596 | ABBX5539     | Adapter Card           |
| ADC 14     | REV 13 | 750-043596 | ABBX5555     | Adapter Card           |
| ADC 15     | REV 13 | 750-043596 | ABBX5557     | Adapter Card           |
| ADC 16     | REV 13 | 750-043596 | ABBX5536     | Adapter Card           |
| ADC 17     | REV 13 | 750-043596 | ABBX5559     | Adapter Card           |
| ADC 18     | REV 13 | 750-043596 | ABBX5537     | Adapter Card           |
| ADC 19     | REV 11 | 750-043596 | ABBW5685     | Adapter Card           |
| Fan Tray 0 | REV 04 | 760-046960 | ACAY0090     | 172mm FanTray - 6 Fans |
| Fan Tray 1 | REV 04 | 760-046960 | ACAY0088     | 172mm FanTray - 6 Fans |
| Fan Tray 2 | REV 04 | 760-046960 | ACAY0089     | 172mm FanTray - 6 Fans |
| Fan Tray 3 | REV 04 | 760-046960 | ACAY0108     | 172mm FanTray - 6 Fans |

#### show chassis hardware models (MX2020 Router)

```
user@host > show chassis hardware models
```

```
Hardware inventory:
```

| Item      | Version | Part number | Serial number | FRU model number     |
|-----------|---------|-------------|---------------|----------------------|
| Midplane  | REV 27  | 750-040240  | ABAB9384      | 750-040240           |
| FPM Board | REV 06  | 760-040242  | ABBT8837      | 760-040242           |
| PSM 0     | REV 01  | 740-045050  | 1E02224006G   | MX2000-PSM-HC-DC-S-A |
| PSM 1     | REV 01  | 740-045050  | 1E022240053   | MX2000-PSM-HC-DC-S-A |
| PSM 2     | REV 01  | 740-045050  | 1E02224004K   | MX2000-PSM-HC-DC-S-A |
| PSM 3     | REV 01  | 740-045050  | 1E022240056   | MX2000-PSM-HC-DC-S-A |
| PSM 4     | REV 01  | 740-045050  | 1E022240054   | MX2000-PSM-HC-DC-S-A |
| PSM 5     | REV 01  | 740-045050  | 1E02224005H   | MX2000-PSM-HC-DC-S-A |
| PSM 6     | REV 01  | 740-045050  | 1E02224006S   | MX2000-PSM-HC-DC-S-A |
| PSM 7     | REV 01  | 740-045050  | 1E02224005M   | MX2000-PSM-HC-DC-S-A |
| PSM 8     | REV 01  | 740-045050  | 1E022240062   | MX2000-PSM-HC-DC-S-A |

|                  |        |            |             |                   |
|------------------|--------|------------|-------------|-------------------|
| PSM 9            | REV 03 | 740-045050 | 1EDB2350095 | MX2000-PSM-DC-S-A |
| PSM 10           | REV 03 | 740-045050 | 1EDB235009L | MX2000-PSM-DC-S-A |
| PSM 11           | REV 03 | 740-045050 | 1EDB2350092 | MX2000-PSM-DC-S-A |
| PSM 12           | REV 03 | 740-045050 | 1EDB23500AT | MX2000-PSM-DC-S-A |
| PSM 13           | REV 03 | 740-045050 | 1EDB2350094 | MX2000-PSM-DC-S-A |
| PSM 15           | REV 03 | 740-045050 | 1EDB235008X | MX2000-PSM-DC-S-A |
| PDM 0            | REV 01 | 740-045234 | 1E012150033 |                   |
| PDM 1            | REV 01 | 740-045234 | 1E012150027 |                   |
| PDM 2            | REV 01 | 740-045234 | 1E262250072 | MX2000-PDM-DC-S-A |
| Routing Engine 0 | REV 02 | 740-041821 | 9009094138  | RE-S-1800X4-16G-S |
| Routing Engine 1 | REV 02 | 740-041821 | 9009089709  | RE-S-1800X4-16G-S |
| CB 0             | REV 08 | 750-040257 | CAAB3482    | 750-040257        |
| CB 1             | REV 04 | 750-040257 | ZT2864      | 750-040257        |
| SFB 0            | REV 05 | 711-044466 | ABBT2161    | MX2000-SFB-S      |
| SFB 1            | REV 05 | 711-044466 | ABBT2159    | MX2000-SFB-S      |
| SFB 2            | REV 05 | 711-044466 | ABBX3718    | MX2000-SFB-S      |
| SFB 4            | REV 05 | 711-044466 | ABBT2160    | MX2000-SFB-S      |
| SFB 5            | REV 05 | 711-044466 | ABBT2145    | MX2000-SFB-S      |
| SFB 7            | REV 05 | 711-044466 | ABBT2163    | MX2000-SFB-S      |
| FPC 0            | REV 30 | 750-028467 | ABBN0284    | MPC-3D-16XGE-SFPP |
| FPC 1            | REV 30 | 750-028467 | ABBN0308    | MPC-3D-16XGE-SFPP |
| FPC 2            | REV 30 | 750-028467 | ABBN0316    | MPC-3D-16XGE-SFPP |
| FPC 3            | REV 32 | 750-028467 | ABBN6832    | MPC-3D-16XGE-SFPP |
| FPC 4            | REV 32 | 750-028467 | ABBN6811    | MPC-3D-16XGE-SFPP |
| FPC 5            | REV 32 | 750-028467 | ABBN6791    | MPC-3D-16XGE-SFPP |
| FPC 6            | REV 30 | 750-028467 | ABBM4592    | MPC-3D-16XGE-SFPP |
| FPC 7            | REV 32 | 750-028467 | ABBN6810    | MPC-3D-16XGE-SFPP |
| FPC 8            | REV 30 | 750-028467 | ABBM4739    | MPC-3D-16XGE-SFPP |
| FPC 9            | REV 32 | 750-028467 | ABBN6827    | MPC-3D-16XGE-SFPP |
| FPC 10           | REV 30 | 750-028467 | ABBN0302    | MPC-3D-16XGE-SFPP |
| FPC 11           | REV 32 | 750-028467 | ABBN6790    | MPC-3D-16XGE-SFPP |
| FPC 12           | REV 30 | 750-028467 | ZM5111      | MPC-3D-16XGE-SFPP |
| FPC 13           | REV 30 | 750-028467 | ABBN0208    | MPC-3D-16XGE-SFPP |
| FPC 14           | REV 23 | 750-028467 | YN2977      | MPC-3D-16XGE-SFPP |
| FPC 15           | REV 32 | 750-028467 | ABBN6798    | MPC-3D-16XGE-SFPP |
| FPC 16           | REV 30 | 750-028467 | ABBN0270    | MPC-3D-16XGE-SFPP |
| FPC 17           | REV 32 | 750-028467 | ABBN6796    | MPC-3D-16XGE-SFPP |
| FPC 18           | REV 30 | 750-028467 | ABBN0281    | MPC-3D-16XGE-SFPP |
| FPC 19           | REV 32 | 750-028467 | ABBN6813    | MPC-3D-16XGE-SFPP |
| ADC 0            | REV 13 | 750-043596 | ABBX5561    | PROTO-ASSEMBLY    |
| ADC 1            | REV 13 | 750-043596 | ABBX5546    | PROTO-ASSEMBLY    |
| ADC 2            | REV 13 | 750-043596 | ABBX5535    | MX2000-LC-ADAPTER |
| ADC 3            | REV 13 | 750-043596 | ABBX5552    | MX2000-LC-ADAPTER |
| ADC 4            | REV 13 | 750-043596 | ABBX5581    | MX2000-LC-ADAPTER |
| ADC 5            | REV 13 | 750-043596 | ABBX5545    | PROTO-ASSEMBLY    |
| ADC 6            | REV 13 | 750-043596 | ABBX5554    | PROTO-ASSEMBLY    |
| ADC 7            | REV 07 | 750-043596 | ABBV7194    | MX2000-LC-ADAPTER |
| ADC 8            | REV 07 | 750-043596 | ABBV7251    | MX2000-LC-ADAPTER |
| ADC 9            | REV 07 | 750-043596 | ABBV7202    | MX2000-LC-ADAPTER |
| ADC 10           | REV 13 | 750-043596 | ABBX5579    | MX2000-LC-ADAPTER |
| ADC 12           | REV 13 | 750-043596 | ABBX5575    | MX2000-LC-ADAPTER |
| ADC 13           | REV 13 | 750-043596 | ABBX5539    | PROTO-ASSEMBLY    |
| ADC 14           | REV 13 | 750-043596 | ABBX5555    | PROTO-ASSEMBLY    |
| ADC 15           | REV 13 | 750-043596 | ABBX5557    | MX2000-LC-ADAPTER |
| ADC 16           | REV 13 | 750-043596 | ABBX5536    | PROTO-ASSEMBLY    |
| ADC 17           | REV 13 | 750-043596 | ABBX5559    | PROTO-ASSEMBLY    |
| ADC 18           | REV 13 | 750-043596 | ABBX5537    | PROTO-ASSEMBLY    |
| ADC 19           | REV 11 | 750-043596 | ABBW5685    | PROTO-ASSEMBLY    |
| Fan Tray 0       | REV 04 | 760-046960 | ACAY0090    |                   |
| Fan Tray 1       | REV 04 | 760-046960 | ACAY0088    |                   |

```

Fan Tray 2 REV 04 760-046960 ACAY0089
Fan Tray 3 REV 04 760-046960 ACAY0108

```

### show chassis hardware clei-models (MX2020 Router)

```

user@ host > show chassis hardware clei-models
Hardware inventory:
Item Version Part number CLEI code FRU model number
Midplane REV 27 750-040240 PROTOXCLEI 750-040240
FPM Board REV 06 760-040242 PROTOXCLEI 760-040242
PSM 0 REV 01 740-045050 IPUPAJMKAA MX2000-PSM-HC-DC-S-A
PSM 1 REV 01 740-045050 IPUPAJMKAA MX2000-PSM-HC-DC-S-A
PSM 2 REV 01 740-045050 IPUPAJMKAA MX2000-PSM-HC-DC-S-A
PSM 3 REV 01 740-045050 IPUPAJMKAA MX2000-PSM-HC-DC-S-A
PSM 4 REV 01 740-045050 IPUPAJMKAA MX2000-PSM-HC-DC-S-A
PSM 5 REV 01 740-045050 IPUPAJMKAA MX2000-PSM-HC-DC-S-A
PSM 6 REV 01 740-045050 IPUPAJMKAA MX2000-PSM-HC-DC-S-A
PSM 7 REV 01 740-045050 IPUPAJMKAA MX2000-PSM-HC-DC-S-A
PSM 8 REV 01 740-045050 IPUPAJMKAA MX2000-PSM-HC-DC-S-A
PSM 9 REV 03 740-045050 IPUPAJMKAA MX2000-PSM-DC-S-A
PSM 10 REV 03 740-045050 IPUPAJMKAA MX2000-PSM-DC-S-A
PSM 11 REV 03 740-045050 IPUPAJMKAA MX2000-PSM-DC-S-A
PSM 12 REV 03 740-045050 IPUPAJMKAA MX2000-PSM-DC-S-A
PSM 13 REV 03 740-045050 IPUPAJMKAA MX2000-PSM-DC-S-A
PSM 15 REV 03 740-045050 IPUPAJMKAA MX2000-PSM-DC-S-A
PDM 0 REV 01 740-045234
PDM 1 REV 01 740-045234
PDM 2 REV 01 740-045234 IPUPAJSKAA MX2000-PDM-DC-S-A
Routing Engine 0 REV 02 740-041821 RE-S-1800X4-16G-S
Routing Engine 1 REV 02 740-041821 RE-S-1800X4-16G-S
CB 0 REV 08 750-040257 PROTOXCLEI 750-040257
CB 1 REV 04 750-040257 PROTOXCLEI 750-040257
SFB 0 REV 05 711-044466 IPUCBA6CAA MX2000-SFB-S
SFB 1 REV 05 711-044466 IPUCBA6CAA MX2000-SFB-S
SFB 2 REV 05 711-044466 IPUCBA6CAA MX2000-SFB-S
SFB 4 REV 05 711-044466 IPUCBA6CAA MX2000-SFB-S
SFB 5 REV 05 711-044466 IPUCBA6CAA MX2000-SFB-S
SFB 7 REV 05 711-044466 IPUCBA6CAA MX2000-SFB-S
FPC 0 REV 30 750-028467 MPC-3D-16XGE-SFPP
FPC 1 REV 30 750-028467 MPC-3D-16XGE-SFPP
FPC 2 REV 30 750-028467 MPC-3D-16XGE-SFPP
FPC 3 REV 32 750-028467 MPC-3D-16XGE-SFPP
FPC 4 REV 32 750-028467 MPC-3D-16XGE-SFPP
FPC 5 REV 32 750-028467 MPC-3D-16XGE-SFPP
FPC 6 REV 30 750-028467 MPC-3D-16XGE-SFPP
FPC 7 REV 32 750-028467 MPC-3D-16XGE-SFPP
FPC 8 REV 30 750-028467 MPC-3D-16XGE-SFPP
FPC 9 REV 32 750-028467 MPC-3D-16XGE-SFPP
FPC 10 REV 30 750-028467 MPC-3D-16XGE-SFPP
FPC 11 REV 32 750-028467 MPC-3D-16XGE-SFPP
FPC 12 REV 30 750-028467 MPC-3D-16XGE-SFPP
FPC 13 REV 30 750-028467 MPC-3D-16XGE-SFPP
FPC 14 REV 23 750-028467 MPC-3D-16XGE-SFPP
FPC 15 REV 32 750-028467 MPC-3D-16XGE-SFPP
FPC 16 REV 30 750-028467 MPC-3D-16XGE-SFPP
FPC 17 REV 32 750-028467 MPC-3D-16XGE-SFPP
FPC 18 REV 30 750-028467 MPC-3D-16XGE-SFPP
FPC 19 REV 32 750-028467 MPC-3D-16XGE-SFPP
ADC 0 REV 13 750-043596 PROTOXCLEI PROTO-ASSEMBLY
ADC 1 REV 13 750-043596 PROTOXCLEI PROTO-ASSEMBLY
ADC 2 REV 13 750-043596 IPUCBA8CAA MX2000-LC-ADAPTER

```

|            |        |            |            |                   |
|------------|--------|------------|------------|-------------------|
| ADC 3      | REV 13 | 750-043596 | IPUCBA8CAA | MX2000-LC-ADAPTER |
| ADC 4      | REV 13 | 750-043596 | IPUCBA8CAA | MX2000-LC-ADAPTER |
| ADC 5      | REV 13 | 750-043596 | PROTOXCLEI | PROTO-ASSEMBLY    |
| ADC 6      | REV 13 | 750-043596 | PROTOXCLEI | PROTO-ASSEMBLY    |
| ADC 7      | REV 07 | 750-043596 | PROTOXCLEI | MX2000-LC-ADAPTER |
| ADC 8      | REV 07 | 750-043596 | PROTOXCLEI | MX2000-LC-ADAPTER |
| ADC 9      | REV 07 | 750-043596 | PROTOXCLEI | MX2000-LC-ADAPTER |
| ADC 10     | REV 13 | 750-043596 | IPUCBA8CAA | MX2000-LC-ADAPTER |
| ADC 12     | REV 13 | 750-043596 | IPUCBA8CAA | MX2000-LC-ADAPTER |
| ADC 13     | REV 13 | 750-043596 | PROTOXCLEI | PROTO-ASSEMBLY    |
| ADC 14     | REV 13 | 750-043596 | PROTOXCLEI | PROTO-ASSEMBLY    |
| ADC 15     | REV 13 | 750-043596 | IPUCBA8CAA | MX2000-LC-ADAPTER |
| ADC 16     | REV 13 | 750-043596 | PROTOXCLEI | PROTO-ASSEMBLY    |
| ADC 17     | REV 13 | 750-043596 | PROTOXCLEI | PROTO-ASSEMBLY    |
| ADC 18     | REV 13 | 750-043596 | PROTOXCLEI | PROTO-ASSEMBLY    |
| ADC 19     | REV 11 | 750-043596 | PROTOXCLEI | PROTO-ASSEMBLY    |
| Fan Tray 0 | REV 04 | 760-046960 |            |                   |
| Fan Tray 1 | REV 04 | 760-046960 |            |                   |
| Fan Tray 2 | REV 04 | 760-046960 |            |                   |
| Fan Tray 3 | REV 04 | 760-046960 |            |                   |

### show chassis hardware (MX Series Routers with ATM MIC)

Can I retain this output by just deleting "ATM MIC"?

```
user@host> show chassis hardware
```

Hardware inventory:

| Item             | Version | Part number | Serial number | Description            |
|------------------|---------|-------------|---------------|------------------------|
| Chassis          |         |             | JN115736EAFc  | MX240                  |
| Midplane         | REV 07  | 760-021404  | ABAA5038      | MX240 Backplane        |
| FPM Board        | REV 03  | 760-021392  | ABBA2758      | Front Panel Display    |
| PEM 0            | Rev 01  | 740-022697  | QCS0937C07K   | PS 1.2-1.7kW; 100-240V |
| AC in            |         |             |               |                        |
| PEM 1            | Rev 01  | 740-022697  | QCS0939C04X   | PS 1.2-1.7kW; 100-240V |
| AC in            |         |             |               |                        |
| PEM 2            | Rev 01  | 740-022697  | QCS0937C06B   | PS 1.2-1.7kW; 100-240V |
| AC in            |         |             |               |                        |
| PEM 3            | Rev 01  | 740-022697  | QCS0937C07U   | PS 1.2-1.7kW; 100-240V |
| AC in            |         |             |               |                        |
| Routing Engine 0 | REV 12  | 740-013063  | 9009042291    | RE-S-2000              |
| Routing Engine 1 | REV 12  | 740-013063  | 9009042266    | RE-S-2000              |
| CB 0             | REV 06  | 710-021523  | ABBC1435      | MX SCB                 |
| CB 1             | REV 06  | 710-021523  | ABBC1497      | MX SCB                 |
| FPC 2            | REV 14  | 750-031088  | YH8446        | MPC Type 2 3D Q        |
| CPU              | REV 06  | 711-030884  | YH9612        | MPC PMB 2G             |
| MIC 0            |         |             |               |                        |
| MIC 1            | REV 10  | 750-036132  | ZP7062        | 2xOC12/8xOC3 CC-CE     |
| PIC 2            |         | BUILTIN     | BUILTIN       | 2xOC12/8xOC3 CC-CE     |
| Xcvr 0           |         | NON-JNPR    | 23393-00492   | UNKNOWN                |
| Xcvr 1           |         | NON-JNPR    | 23393-00500   | UNKNOWN                |
| Xcvr 2           |         | NON-JNPR    | 23393-00912   | UNKNOWN                |
| Xcvr 3           | REV 01  | 740-015638  | 22216-00575   | Load SFP               |
| Xcvr 4           | REV 01  | 740-015638  | 24145-00110   | Load SFP               |
| Xcvr 5           | REV 01  | 740-015638  | 24145-00016   | Load SFP               |
| Xcvr 6           | REV 01  | 740-015638  | 24145-00175   | Load SFP               |
| Xcvr 7           |         | NON-JNPR    | 23393-00627   | UNKNOWN                |
| QXM 0            | REV 05  | 711-028408  | YF4681        | MPC QXM                |
| QXM 1            | REV 05  | 711-028408  | YF4817        | MPC QXM                |
| Fan Tray 0       | REV 01  | 710-021113  | XL3645        | MX240 Fan Tray         |

**show chassis hardware (MX240, MX480, MX960 Routers with Application Services Modular Line Card)**

```
user@host>show chassis hardware
```

```
Hardware inventory:
```

| Item             | Version | Part number | Serial number | Description               |
|------------------|---------|-------------|---------------|---------------------------|
| Chassis          |         |             | JN11D969BAFA  | MX960                     |
| Midplane         | REV 03  | 710-013698  | ACAA2362      | MX960 Backplane           |
| FPM Board        | REV 03  | 710-014974  | ZR0639        | Front Panel Display       |
| PDM              | Rev 03  | 740-013110  | QCS152250SX   | Power Distribution Module |
| PEM 0            | Rev 10  | 740-013683  | QCS1512718W   | DC Power Entry Module     |
| PEM 1            | Rev 10  | 740-013683  | QCS1512702Y   | DC Power Entry Module     |
| Routing Engine 0 | REV 15  | 740-013063  | 9012024667    | RE-S-2000                 |
| Routing Engine 1 | REV 15  | 740-013063  | 9012024649    | RE-S-2000                 |
| CB 0             | REV 14  | 750-031391  | ZJ7749        | Enhanced MX SCB           |
| CB 1             | REV 14  | 750-031391  | ZJ7750        | Enhanced MX SCB           |
| CB 2             | REV 14  | 750-031391  | ZY9233        | Enhanced MX SCB           |
| FPC 0            | REV 17  | 750-031089  | YR7434        | MPC Type 2 3D             |
| CPU              |         |             |               |                           |
| FPC 1            | REV 11  | 750-037207  | ZW9727        | AS-MCC                    |
| CPU              | REV 04  | 711-038173  | ZW4817        | AS-MCC-PMB                |
| MIC 0            | REV 01  | 750-037214  | ZH3764        | AS-MSC                    |
| PIC 0            |         | BUILTIN     | BUILTIN       | AS-MSC                    |
| MIC 1            | REV 01  | 711-028408  | JZ9200        | AS-MXC                    |
| PIC 2            |         | BUILTIN     | BUILTIN       | AS-MXC                    |
| FPC 4            | REV 30  | 750-028467  | ABBN0232      | MPC 3D 16x 10GE           |
| CPU              |         |             |               |                           |
| FPC 5            | REV 04  | 750-037207  | ZK9074        | AS-MCC                    |
| CPU              |         |             |               |                           |
| Fan Tray 0       | REV 05  | 740-014971  | VT5683        | Fan Tray                  |
| Fan Tray 1       | REV 05  | 740-014971  | VT5684        | Fan Tray                  |

**show chassis hardware extensive (MX240, MX480, MX960 Routers with Application Services Modular Line Card)**

```
user@host> show chassis hardware extensive
```

```
ID: AS-MCC FRU Model Number: 750-037207
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 37 01 0b 52 45 56 20 31 31 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 30 37 00 00
Address 0x20: 53 2f 4e 20 5a 57 39 37 32 37 00 00 00 11 02 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 35 30 2d 30 33 37 32 30 37 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 31 31 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 5e ff ff ff ff ff ff ff ff ff ff ff ff
CPU REV 04 711-038173 ZW4817 AS-MCC-PMB
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 711-038173 S/N: S/N ZW4817
Assembly ID: 0x0b38 Assembly Version: 01.04
Date: 12-30-2011 Assembly Flags: 0x00
Version: REV 04
ID: AS-MCC-PMB
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 38 01 04 52 45 56 20 30 34 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 38 31 37 33 00 00
Address 0x20: 53 2f 4e 20 5a 57 34 38 31 37 00 00 00 1e 0c 07
Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
```

```

Address 0x40: ff ff ff ff 00 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 31 31 2d 30 33 38 31 37 33 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 30 34 00 ff ff ff ff ff ff
Address 0x70: ff ff ff 60 00 00 00 00 00 00 00 00 00 00 00
MIC 0 REV 01 750-037214 ZH3764 AS-MS
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 750-037214 S/N: S/N ZH3764
Assembly ID: 0x0a44 Assembly Version: 01.01
Date: 07-04-2011 Assembly Flags: 0x00
Version: REV 01
ID: AS-MS
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff I2C Hex Data:
Address 0x00: 7f b0 02 ff 0a 44 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 31 34 00 00
Address 0x20: 53 2f 4e 20 5a 48 33 37 36 34 00 00 00 04 07 07
Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff f6 c0 03 e1 bc 00 00 00 00 00 00 00 00
PIC 0 BUILTIN BUILTIN AS-MS
FPC 4 REV 30 750-028467 ABBN0232 MPC 3D 16x 10GE
Jedec Code: 0x7fb0 EEPROM Version: 0x01

```

#### show chassis hardware (MX480 Router with MPC4E)

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

| Item             | Version     | Part number | Serial number | Description            |
|------------------|-------------|-------------|---------------|------------------------|
| Chassis          |             |             | JN10FF57BAFB  | MX480                  |
| Midplane         | REV 05      | 750-047849  | Good          | MX480 Midplane         |
| FPM Board        | REV 02      | 710-017254  | KG2066        | Front Panel Display    |
| PEM 0            | Rev 03      | 740-017330  | QCS081590BJ   | PS 1.2-1.7kW; 100-240V |
| AC in            |             |             |               |                        |
| PEM 1            | Rev 03      | 740-017330  | QCS0815908Z   | PS 1.2-1.7kW; 100-240V |
| AC in            |             |             |               |                        |
| PEM 2            | Rev 03      | 740-029970  | QCS1001U001   | PS 1.4-2.52kW; 90-264V |
| AC in            |             |             |               |                        |
| Routing Engine 0 | REV 05      | 740-031116  | 9009089502    | RE-S-1800x4            |
| Routing Engine 1 | REV 05      | 740-031116  | 9009089624    | RE-S-1800x4            |
| CB 0             | REV 02      | 750-031391  | YE8506        | Enhanced MX SCB        |
| CB 1             | REV 14      | 750-031391  | ZK8265        | Enhanced MX SCB        |
| FPC 2            | REV 05      | 750-037358  | ZT0638        | MPC4E 3D 32XGE         |
| CPU              | REV 07      | 711-035209  | ZK3187        | HMPC PMB 2G            |
| PIC 0            |             | BUILTIN     | BUILTIN       | 8X10GE SFPP            |
| PIC 1            |             | BUILTIN     | BUILTIN       | 8X10GE SFPP            |
| PIC 2            |             | BUILTIN     | BUILTIN       | 8X10GE SFPP            |
| PIC 3            |             | BUILTIN     | BUILTIN       | 8X10GE SFPP            |
| FPC 3            | REV 06      | 750-037355  | CAAB1144      | MPC4E 3D 2CGE+8XGE     |
| CPU              | REV 08      | 711-035209  | CAAB1278      | HMPC PMB 2G            |
| PIC 0            |             | BUILTIN     | BUILTIN       | 4x10GE SFPP            |
| Xcvr 0           | REV 01      | 740-031980  | B11E01439     | SFP+-10G-SR            |
| Xcvr 1           | REV 01      | 740-031980  | B11D05809     | SFP+-10G-SR            |
| PIC 1            |             | BUILTIN     | BUILTIN       | 1X100GE CFP            |
| Xcvr 0           |             | NON-JNPR    | D5418         | UNKNOWN                |
| PIC 2            |             | BUILTIN     | BUILTIN       | 4x10GE SFPP            |
| PIC 3            |             | BUILTIN     | BUILTIN       | 1X100GE CFP            |
| Xcvr 0           |             | NON-JNPR    | X12J00362     | CFP-100G-SR10          |
| FPC 4            | REV 12.3.10 | 750-033205  | YR9445        | MPCE Type 3 3D         |

CPU  
Fan Tray

Enhanced Left Fan Tray

### show chassis hardware (MX2020 Router with MPC4E)

user@host> show chassis hardware

Hardware inventory:

| Item             | Version | Part number | Serial number | Description           |
|------------------|---------|-------------|---------------|-----------------------|
| Chassis          |         |             | JN11E188CAFJ  | MX2020                |
| Midplane         | REV 04  | 711-032387  | ABAC7474      | Lower Backplane       |
| Midplane 1       | REV 04  | 711-032386  | ABAC7408      | Upper Backplane       |
| PMP 1            | REV 03  | 711-032428  | ACAJ1137      | Upper Power Midplane  |
| PMP 0            | REV 03  | 711-032426  | ACAJ1016      | Lower Power Midplane  |
| FPM Board        | REV 06  | 760-040242  | ABBT8832      | Front Panel Display   |
| PSM 3            | REV 0C  | 740-033727  | VK00255       | DC 52V Power Supply   |
| Module           |         |             |               |                       |
| PSM 4            | REV 0C  | 740-033727  | VJ00148       | DC 52V Power Supply   |
| Module           |         |             |               |                       |
| PSM 5            | REV 0C  | 740-033727  | VK00207       | DC 52V Power Supply   |
| Module           |         |             |               |                       |
| PSM 6            | REV 0C  | 740-033727  | VK00319       | DC 52V Power Supply   |
| Module           |         |             |               |                       |
| PSM 7            | REV 0C  | 740-033727  | VK00264       | DC 52V Power Supply   |
| Module           |         |             |               |                       |
| PSM 8            | REV 0B  | 740-033727  | VG00025       | DC 52V Power Supply   |
| Module           |         |             |               |                       |
| PSM 13           | REV 0C  | 740-033727  | VK00274       | DC 52V Power Supply   |
| Module           |         |             |               |                       |
| PSM 14           | REV 0C  | 740-033727  | VJ00167       | DC 52V Power Supply   |
| Module           |         |             |               |                       |
| PSM 15           | REV 0C  | 740-033727  | VK00299       | DC 52V Power Supply   |
| Module           |         |             |               |                       |
| PSM 16           | REV 0C  | 740-033727  | VK00213       | DC 52V Power Supply   |
| Module           |         |             |               |                       |
| PSM 17           | REV 0C  | 740-033727  | VK00253       | DC 52V Power Supply   |
| Module           |         |             |               |                       |
| PDM 0            | REV 0B  | 740-038109  | VJ00040       | DC Power Dist Module  |
| PDM 2            | REV 0B  | 740-038109  | VJ00025       | DC Power Dist Module  |
| Routing Engine 0 | REV 02  | 740-041821  | 9009089735    | RE-S-1800x4           |
| Routing Engine 1 | REV 02  | 740-041821  | 9009089731    | RE-S-1800x4           |
| CB 0             | REV 04  | 750-040257  | ZT2846        | Control Board         |
| CB 1             | REV 04  | 750-040257  | ZT2877        | Control Board         |
| SPMB 0           | REV 01  | 711-041855  | ZS2282        | PMB Board             |
| SPMB 1           | REV 01  | 711-041855  | ZS2261        | PMB Board             |
| SFB 0            | REV 07  | 711-032385  | ZZ2582        | Switch Fabric Board   |
| SFB 1            | REV 04  | 711-032385  | ZV4229        | Switch Fabric Board   |
| SFB 2            | REV 07  | 711-032385  | CAAB4902      | Switch Fabric Board   |
| SFB 3            | REV 07  | 711-032385  | CAAB4891      | Switch Fabric Board   |
| SFB 4            | REV 07  | 711-032385  | CAAB4883      | Switch Fabric Board   |
| SFB 5            | REV 07  | 711-032385  | CAAB4889      | Switch Fabric Board   |
| SFB 6            | REV 06  | 711-032385  | ZV1818        | Switch Fabric Board   |
| SFB 7            | REV 07  | 711-032385  | CAAB4897      | Switch Fabric Board   |
| FPC 0            | REV 34  | 750-031090  | ZT9799        | MPC Type 2 3D EQ      |
| CPU              | REV 06  | 711-030884  | ZS1122        | MPC PMB 2G            |
| MIC 0            | REV 11  | 750-033535  | CAAD7674      | MIC-3D-10C192-XFP     |
| PIC 0            |         | BUILTIN     | BUILTIN       | MIC-3D-10C192-XFP     |
| Xcvr 0           | REV 01  | 740-014279  | 753019A00404  | XFP-0C192-SR          |
| MIC 1            | REV 14  | 750-031967  | ZM6103        | MIC-3D-80C30C12-40C48 |
| PIC 2            |         | BUILTIN     | BUILTIN       | MIC-3D-80C30C12-40C48 |
| Xcvr 0           | REV 01  | 740-011615  | PEF1AZP       | SFP-IR                |
| Xcvr 1           | REV 01  | 740-011615  | PEF1AZN       | SFP-IR                |



|            |        |            |              |                        |
|------------|--------|------------|--------------|------------------------|
| Xcvr 2     | REV 01 | 740-021308 | ANA0N8S      | SFP+-10G-SR            |
| QXM 0      | REV 06 | 711-028408 | ZT9339       | MPC QXM                |
| QXM 1      | REV 06 | 711-028408 | ZT9237       | MPC QXM                |
| FPC 9      | REV 34 | 750-031090 | ZT9770       | MPC Type 2 3D EQ       |
| CPU        | REV 06 | 711-030884 | ZS1302       | MPC PMB 2G             |
| MIC 0      | REV 24 | 750-028387 | YJ3950       | 3D 4x 10GE XFP         |
| PIC 0      |        | BUILTIN    | BUILTIN      | 2x 10GE XFP            |
| Xcvr 0     |        | NON-JNPR   | T09M52516    | XFP-10G-SR             |
| Xcvr 1     |        | NON-JNPR   | CA49BK095    | XFP-10G-SR             |
| PIC 1      |        | BUILTIN    | BUILTIN      | 2x 10GE XFP            |
| Xcvr 0     | REV 02 | 740-014289 | C834XU01T    | XFP-10G-SR             |
| Xcvr 1     |        | NON-JNPR   | T09M52515    | XFP-10G-SR             |
| MIC 1      | REV 11 | 750-033535 | CAAD7681     | MIC-3D-10C192-XFP      |
| PIC 2      |        | BUILTIN    | BUILTIN      | MIC-3D-10C192-XFP      |
| Xcvr 0     | REV 01 | 740-014279 | KBQ02BE      | XFP-OC192-SR           |
| QXM 0      | REV 06 | 711-028408 | ZT9151       | MPC QXM                |
| QXM 1      | REV 06 | 711-028408 | ZT9116       | MPC QXM                |
| FPC 10     | REV 27 | 750-033205 | ZL6215       | MPCE Type 3 3D         |
| CPU        | REV 07 | 711-035209 | ZK9038       | HMPC PMB 2G            |
| MIC 0      | REV 18 | 750-028380 | YG6885       | 3D 2x 10GE XFP         |
| PIC 0      |        | BUILTIN    | BUILTIN      | 1x 10GE XFP            |
| Xcvr 0     | REV 01 | 740-014289 | C706XU0AG    | XFP-10G-SR             |
| PIC 1      |        | BUILTIN    | BUILTIN      | 1x 10GE XFP            |
| Xcvr 0     | REV 02 | 740-014289 | T08L84366    | XFP-10G-SR             |
| FPC 14     | REV 09 | 750-037355 | CAAF1534     | MPC4E 3D 2CGE+8XGE     |
| CPU        | REV 08 | 711-035209 | CAAB9879     | HMPC PMB 2G            |
| PIC 0      |        | BUILTIN    | BUILTIN      | 4x10GE SFPP            |
| Xcvr 0     | REV 01 | 740-021308 | 21T511100436 | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | AHPOGPM      | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | 123363A00032 | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-021308 | 19T511100477 | SFP+-10G-SR            |
| PIC 1      |        | BUILTIN    | BUILTIN      | 1X100GE CFP            |
| Xcvr 0     |        | NON-JNPR   | X12J00260    | CFP-100G-SR10          |
| PIC 2      |        | BUILTIN    | BUILTIN      | 4x10GE SFPP            |
| Xcvr 0     | REV 01 | 740-021308 | 21T511104086 | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-021308 | 21T511104627 | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-021308 | 21T511104644 | SFP+-10G-SR            |
| PIC 3      |        | BUILTIN    | BUILTIN      | 1X100GE CFP            |
| FPC 19     | REV 32 | 750-028467 | ZR2008       | MPC 3D 16x 10GE        |
| CPU        | REV 10 | 711-029089 | ZT6933       | AMPC PMB               |
| PIC 0      |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+      |
| Xcvr 0     | REV 01 | 740-021308 | 19T511100291 | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-021308 | AMH02VE      | SFP+-10G-SR            |
| PIC 1      |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+      |
| Xcvr 0     | REV 01 | 740-021308 | 23T511102128 | SFP+-10G-SR            |
| PIC 2      |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+      |
| Xcvr 0     | REV 01 | 740-021308 | AMS15PP      | SFP+-10G-SR            |
| PIC 3      |        | BUILTIN    | BUILTIN      | 4x 10GE(LAN) SFP+      |
| Xcvr 0     | REV 01 | 740-031980 | 123363A00716 | SFP+-10G-SR            |
| ADC 0      | REV 05 | 750-043596 | CAAC2072     | Adapter Card           |
| ADC 9      | REV 01 | 750-043596 | ZV4111       | Adapter Card           |
| ADC 10     | REV 05 | 750-043596 | CAAC2058     | Adapter Card           |
| ADC 14     | REV 02 | 750-043596 | ZW1561       | Adapter Card           |
| ADC 19     | REV 01 | 750-043596 | ZV4127       | Adapter Card           |
| Fan Tray 0 | REV 03 | 760-046960 | ACAY0124     | 172mm FanTray - 6 Fans |
| Fan Tray 1 | REV 2A | 760-046960 | ACAY0022     | 172mm FanTray - 6 Fans |
| Fan Tray 2 | REV 2A | 760-046960 | ACAY0023     | 172mm FanTray - 6 Fans |
| Fan Tray 3 | REV 2A | 760-046960 | ACAY0025     | 172mm FanTray - 6 Fans |

## show chassis hardware (T320 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item Version Part number Serial number Description
Chassis 19093 T320
Midplane REV 04 710-004339 BC1436 T320 Backplane
FPM GBUS REV 03 710-004461 BC1407 T320 FPM Board
FPM Display REV 04 710-002897 BE0763 FPM Display
CIP REV 05 710-002895 BB2311 T Series CIP
PEM 0 Rev 01 740-004359 NB12546 Power Entry Module
SCG 0 REV 06 710-004455 AY4522 T320 Sonet
Clock Gen.
Routing Engine 0
CB 0 REV 13 710-002728 BC1577 unknown
 T Series
Control Board
CB 1 REV 13 710-002728 BC1595 T Series
Control Board
FPC 1 REV 09 710-007531 HS1572 FPC Type 2
CPU REV 15 710-001726 HR8763 FPC CPU
PIC 0 REV 01 750-010618 CB5579 4x G/E SFP,
1000 BASE
SFP 0 REV 01 740-007326 P5809Z1 SFP-SX
SFP 1 REV 01 740-007326 P4Q10XU SFP-SX
SFP 2 NON-JNPR RA45020031 SFP-SX
SFP 3 NON-JNPR RA45020032 SFP-SX
PIC 1 REV 01 750-010618 CD9587 4x G/E SFP,
1000 BASE
SFP 0 NON-JNPR P5A08QZ SFP-T
SFP 1 REV 01 740-007326 P4Q133K SFP-SX
SFP 2 REV 01 740-007326 P5809YY SFP-SX
SFP 3 REV 01 740-007327 4C81704 SFP-LX
MMB 1 REV 03 710-005555 HR9401 MMB-288mbit
PPB 0 REV 04 710-003758 HR2886 PPB Type 2
FPC 2 REV 07 710-005860 HP2392 FPC Type 1
CPU REV 14 710-001726 HP7797 FPC CPU
PIC 0 REV 02 750-007643 HM0853 1x G/E QPP,
1000 BASE
SFP 0 REV 01 740-007326 P11E9JJ SFP-SX
MMB 1 REV 02 710-005555 HN2379 MMB-288mbit
PPB 0 REV 04 710-003758 HP8092 PPB Type 2
FPC 3 REV 07 710-005860 HP2393 FPC Type 1
CPU REV 14 710-001726 HP0968 FPC CPU
PIC 0 REV 01 750-010240 CB5363 1x G/E SFP,
1000 BASE
SFP 0 REV 01 740-007326 P4R0PNH SFP-SX
PIC 1 REV 03 750-003034 HD2832 4x OC-3 SONET,
SMIR
MMB 1 REV 02 710-005555 HN6307 MMB-288mbit
PPB 0 REV 04 710-003758 HP5051 PPB Type 2
FPC 4 REV 01 710-010845 JD3872 FPC Type 4
CPU REV 02 710-011481 JB6042 FPC CPU
5 REV 01 710-005802 BC1566 FPC Type 2
CPU REV 09 710-001726 AY4922 FPC CPU
PIC 0 REV 02 750-008155 BE2114 2x G/E QPP,
1000 BASE
SFP 0 REV 01 740-007326 P4R0PMQ SFP-SX
SFP 1 REV 01 740-007326 P4R0PN9 SFP-SX
PIC 1 REV 01 750-008155 BE2116 2x G/E QPP,
1000 BASE
SFP 0 REV 01 740-007326 P4R0PNZ SFP-SX

```

|       |        |            |        |             |
|-------|--------|------------|--------|-------------|
| SFP 1 |        | NON-JNPR   | 2908   | SFP-T       |
| MMB 1 | REV 01 | 710-005555 | AZ2246 | MMB-288mbit |
| PPB 0 | REV 03 | 710-003758 | AY4839 | PPB Type 2  |
| FPC 7 | REV 01 | 710-005803 | AZ2123 | FPC Type 3  |
| ...   |        |            |        |             |

### show chassis hardware (T640 Router)

```
user@host> show chassis hardware
Hardware inventory:
```

| Item             | Version | Part number | Serial number | Description            |
|------------------|---------|-------------|---------------|------------------------|
| Chassis          |         |             | 19182         | T640                   |
| Midplane         | REV 04  | 710-002726  | AX5608        | T640 Backplane         |
| FPM GBUS         | REV 02  | 710-002901  | HE3064        | T640 FPM Board         |
| FPM Display      | REV 02  | 710-002897  | HE7864        | FPM Display            |
| CIP              | REV 05  | 710-002895  | HA5024        | T Series CIP           |
| PEM 0            | Rev 02  | 740-029522  | VH26235       | AC PEM 10kW US         |
| PEM 1            | Rev 02  | 740-029522  | VH26230       | AC PEM 10kW US         |
| SCG 0            | REV 03  | 710-003423  | HA4508        | T640 Sonet Clock Gen.  |
| Routing Engine 0 | REV 02  | 740-005022  | 210865700483  | RE-3.0 (RE-600)        |
| CB 0             | REV 01  | 710-002728  | HD3044        | T Series Control Board |
| FPC 2            | REV 04  | 710-001721  | HD5572        | FPC Type 3             |
| CPU              | REV 06  | 710-001726  | HA4712        | FPC CPU                |
| PIC 1            | REV 03  | 750-009567  | HV2331        | 1x 10GE(LAN),XENPAK    |
| SFP 0            | REV 01  | 740-009898  | USC202R103    | XENPAK-SR              |
| PIC 2            | REV 03  | 750-009567  | HV2332        | 1x 10GE(LAN),XENPAK    |
| SFP 0            | REV 01  | 740-011268  | USC202R112    | XENPAK-ZR              |
| PIC 3            | REV 03  | 750-009567  | HX4416        | 1x 10GE(LAN),XENPAK    |
| SFP 0            | REV 01  | 740-012056  | 434TC004      | XENPAK-CX4             |
| PIC 4            | REV 03  | 750-009567  | HX4420        | 1x 10GE(LAN),XENPAK    |
| SFP 0            | REV 01  | 740-012058  | 434TC124      | XENPAK-LX4             |
| FPC 5            | REV 01  | 710-013553  | JE4839        | E2-FPC Type 1          |
| CPU              | REV 01  | 710-013569  | JW9163        | FPC CPU                |
| PIC 0            | REV 01  | 750-009567  | HX4419        | 1x 10GE(LAN),XENPAK    |
| SFP 0            | REV 01  | 740-009898  | USC202RT05    | XENPAK-LR              |
| PIC 1            | REV 03  | 750-009567  | HN7426        | 1x 10GE(LAN),XENPAK    |
| SFP 0            | REV 01  | 740-009550  | 03L90051      | XENPAK-ER              |
| PIC 2            | REV 03  | 750-009467  | HT7423        | 1x 10GE(LAN),XENPAK    |
| SFP 0            |         | NON-JNPR    |               | UNKNOWN                |
| PIC 3            | REV 04  | 750-005100  | AY4850        | 1x 10GE(LAN),DWDM      |
| FPC 4            | REV 01  | 710-010845  | JD3872        | FPC Type 4             |
| CPU              | REV 02  | 710-011481  | JB6042        | FPC CPU                |
| Fan Tray 0       |         |             |               | Front Top Fan Tray     |
| Fan Tray 1       |         |             |               | Front Bottom Fan Tray  |
| Fan Tray 2       |         |             |               | Rear Fan Tray          |

### show chassis hardware models (T640 Router)

```
user@host> show chassis hardware models
Hardware inventory:
```

| Item             | Version | Part number | CLEI code | FRU model number |
|------------------|---------|-------------|-----------|------------------|
| Midplane         | REV 04  | 710-002726  |           | CHAS-BP-T640-S   |
| FPM Display      | REV 02  | 710-002897  |           | CRAFT-T640-S     |
| CIP              | REV 05  | 710-002895  |           | CIP-L-T640-S     |
| PEM 0            | Rev 01  | 740-002595  |           | PWR-T-DC-S       |
| SCG 0            | REV 04  | 710-003423  |           | SCG-T-S          |
| SCG 1            | REV 04  | 710-003423  |           | SCG-T-S          |
| Routing Engine 0 | REV 01  | 740-005022  |           | RE-600-2048-S    |
| Routing Engine 1 | REV 07  | 740-005022  |           | RE-600-2048-S    |
| CB 0             | REV 06  | 710-002726  |           | CHAS-BP-T640-S   |
| CB 1             | REV 06  | 710-002728  |           | CB-L-T-S         |

|            |        |            |                    |
|------------|--------|------------|--------------------|
| FPC 5      | REV 05 | 710-007527 | T640-FPC2          |
| PIC 0      | REV 05 | 750-002510 | PB-2GE-SX          |
| PIC 1      | REV 05 | 750-001901 | PB-40C12-SON-SMIR  |
| FPC 6      | REV 03 | 710-001721 | T640-FPC3          |
| PIC 1      | REV 01 | 750-009553 | PC-40C48-SON-SFP   |
| SIB 4      | REV 02 | 750-005486 | SIB-I-T640-S       |
| Fan Tray 0 |        |            | FANTRAY-T-S        |
| Fan Tray 1 |        |            | FANTRAY-T-S        |
| Fan Tray 2 |        |            | FAN-REAR-TX-T640-S |

### show chassis hardware extensive (T640 Router)

```

user@host> show chassis hardware extensive
Hardware inventory:
Item Version Part number Serial number Description
Chassis
Jedec Code: 0x7fb0 EEPROM Version: 0x01
P/N: S/N:
Assembly ID: 0x0507 Assembly Version: 00.00
Date: 00-00-0000 Assembly Flags: 0x00
Version:
ID: Gibson LCC Chassis
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 01 ff 05 07 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane REV 04 710-002726 AX5633
Jedec Code: 0x7fb0 EEPROM Version: 0x01
P/N: 710-002726. S/N: S/N AX5633.
Assembly ID: 0x0127 Assembly Version: 01.04
Date: 06-27-2001 Assembly Flags: 0x00
Version: REV 04.....
ID: Gibson Backplane
Board Information Record:
Address 0x00: ad 01 08 00 00 90 69 0e f8 00 ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 01 27 01 04 52 45 56 20 30 34 00 00
Address 0x10: 00 00 00 00 00 37 31 30 2d 30 30 32 37 32 36 00 00
Address 0x20: 53 2f 4e 20 41 58 35 36 33 33 00 00 00 1b 06 07
Address 0x30: d1 ff ff ff ad 01 08 00 00 90 69 0e f8 00 ff ff
Address 0x40: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM GBUS REV 02 710-002901 HE3245
...
FPM Display REV 02 710-002897 HA4873
...
CIP REV 05 710-002895 HA4729
...
PEM 1 RevX02 740-002595 MD21815 Power Entry Module
...
SCG 0 REV 04 710-003423 HF6023
...
SCG 1 REV 04 710-003423 HF6061
...
Routing Engine 0 REV 01 740-005022 210865700292 RE-3.0
...
CB 0 REV 06 710-002728 HE3614
...

```

|       |        |            |        |            |
|-------|--------|------------|--------|------------|
| FPC 1 | REV 01 | 710-002385 | HE3009 | FPC Type 1 |
| ...   | REV 06 | 710-001726 | HC0010 |            |

### show chassis hardware (T4000 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item Version Part number Serial number Description
Chassis JN1172F25AHA T4000
Midplane REV 01 710-027486 RC8355 T-series Backplane
FPM GBUS REV 13 710-002901 BBAE0927 T640 FPM Board
FPM Display REV 01 710-021387 EF6764 T1600 FPM Display
CIP REV 06 710-002895 BBAD9210 T-series CIP
PEM 0 REV 01 740-036442 VA00016 Power Entry Module 6x60
SCG 0 REV 18 710-003423 BBAD7248 T640 Sonet Clock Gen.
SCG 1 REV 18 710-003423 BBAE3874 T640 Sonet Clock Gen.
Routing Engine 0 REV 05 740-026941 P737F-002248 RE-DUO-1800
Routing Engine 1 REV 06 740-026941 P737F-002653 RE-DUO-1800
CB 0 REV 09 710-022597 ED0295 LCC Control Board
CB 1 REV 09 710-022597 EA6050 LCC Control Board
FPC 0 REV 26 750-032819 EK1173 FPC Type 5-3D
 CPU REV 12 711-030686 EJ8584 SNG PMB
 PIC 0 REV 07 750-034624 EF6837 12x10GE (LAN/WAN) SFPP
 Xcvr 0 REV 01 740-031980 123363A01145 SFP+-10G-SR
 Xcvr 1 REV 01 740-031980 123363A01147 SFP+-10G-SR
 Xcvr 2 REV 01 740-031980 AJJ01P3 SFP+-10G-SR
 Xcvr 3 REV 01 740-031980 B10M03256 SFP+-10G-SR
 Xcvr 4 REV 01 740-031980 AJJ01M2 SFP+-10G-SR
 Xcvr 5 REV 01 740-031980 123363A01137 SFP+-10G-SR
 Xcvr 6 REV 01 740-031980 AJJ01PN SFP+-10G-SR
 Xcvr 7 REV 01 740-031980 AJJ01NW SFP+-10G-SR
 Xcvr 8 REV 01 740-031980 123363A01139 SFP+-10G-SR
 Xcvr 9 REV 01 740-031980 AJJ01KE SFP+-10G-SR
 Xcvr 10 REV 01 740-031980 123363A01336 SFP+-10G-SR
 Xcvr 11 REV 01 740-031980 B10M01325 SFP+-10G-SR
 PIC 1 REV 07 750-034624 EF6800 12x10GE (LAN/WAN) SFPP
 Xcvr 0 REV 01 740-031980 AJJ01SA SFP+-10G-SR
 Xcvr 1 REV 01 740-031980 AJJ01QZ SFP+-10G-SR
 Xcvr 2 REV 01 740-031980 AJH0217 SFP+-10G-SR
 Xcvr 3 REV 01 740-031980 AJJ01TE SFP+-10G-SR
 Xcvr 4 REV 01 740-031980 AJJ01KV SFP+-10G-SR
 Xcvr 5 REV 01 740-031980 AJJ01MU SFP+-10G-SR
 Xcvr 6 REV 01 740-031980 AJJ01R0 SFP+-10G-SR
 Xcvr 7 REV 01 740-031980 AJJ01TC SFP+-10G-SR
 Xcvr 8 REV 01 740-031980 AJJ0364 SFP+-10G-SR
 Xcvr 9 REV 01 740-031980 AJD0GV3 SFP+-10G-SR
 Xcvr 10 REV 01 740-031980 B10M03343 SFP+-10G-SR
 Xcvr 11 REV 01 740-031980 AJJ01QJ SFP+-10G-SR
 LMB 0 REV 05 711-034381 EJ8490 Type-0 LMB
 LMB 1 REV 04 711-035774 EJ8517 Type-1 LMB
 LMB 2 REV 05 711-034381 EJ8489 Type-0 LMB
FPC 3 REV 07 750-032819 EG3637 FPC Type 5-3D
 CPU REV 09 711-030686 EG0150 SNG PMB
 PIC 0 REV 08 750-035293 EF3657 1x100GE
 Xcvr 0 REV 01 740-032210 C22CQNJ CFP-100G-LR4
 PIC 1 REV 10 750-034624 BBAN4098 12x10GE (LAN/WAN) SFPP
 Xcvr 0 REV 01 740-031980 B11J04902 SFP+-10G-SR
 Xcvr 1 REV 01 740-031980 B11J04891 SFP+-10G-SR
 Xcvr 2 REV 01 740-031980 AJJ01MX SFP+-10G-SR
 Xcvr 3 REV 01 740-031980 B11J04183 SFP+-10G-SR
 Xcvr 4 REV 01 740-031980 B11J04894 SFP+-10G-SR

```

|            |        |            |              |                        |
|------------|--------|------------|--------------|------------------------|
| Xcvr 5     | REV 01 | 740-031980 | B11J04184    | SFP+-10G-SR            |
| Xcvr 6     | REV 01 | 740-031980 | B11J04897    | SFP+-10G-SR            |
| Xcvr 7     | REV 01 | 740-031980 | B11J04899    | SFP+-10G-SR            |
| Xcvr 8     | REV 01 | 740-031980 | AJJ01TV      | SFP+-10G-SR            |
| Xcvr 9     | REV 01 | 740-031980 | B11J04057    | SFP+-10G-SR            |
| Xcvr 10    | REV 01 | 740-031980 | AJJ01M4      | SFP+-10G-SR            |
| Xcvr 11    | REV 01 | 740-031980 | B11J04905    | SFP+-10G-SR            |
| LMB 0      | REV 04 | 711-034381 | EG1524       | Type-0 LMB             |
| LMB 1      | REV 03 | 711-035774 | EG0345       | Type-1 LMB             |
| LMB 2      | REV 04 | 711-034381 | EG1522       | Type-0 LMB             |
| FPC 5      | REV 03 | 710-033871 | BBAJ0768     | FPC Type 4-ES          |
| CPU        | REV 11 | 710-016744 | BBAH9342     | ST-PMB2                |
| PIC 0      | REV 09 | 750-029262 | EE6789       | 100GE                  |
| PIC 1      | REV 03 | 750-034781 | EE6655       | 100GE CFP              |
| Xcvr 0     | REV 01 | 740-032210 | J11A22334    | CFP-100G-LR4           |
| BRIDGE 0   | REV 03 | 711-029995 | EE6572       | 100GE Bridge Board     |
| MMB 0      | REV 07 | 710-025563 | BBAJ4657     | ST-MMB2                |
| MMB 1      | REV 07 | 710-025563 | BBAJ3073     | ST-MMB2                |
| FPC 6      | REV 05 | 750-010153 | EF4936       | FPC Type 5-3D          |
| CPU        | REV 06 | 711-030686 | EF4189       | SNG PMB                |
| PIC 0      | REV 10 | 750-034624 | BBAN4109     | 12x10GE (LAN/WAN) SFPP |
| Xcvr 0     | REV 01 | 740-031980 | B11J04895    | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | B11J04898    | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | B11J04021    | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-031980 | B11J04903    | SFP+-10G-SR            |
| Xcvr 4     | REV 01 | 740-031980 | B11J04311    | SFP+-10G-SR            |
| Xcvr 5     | REV 01 | 740-031980 | B11J04059    | SFP+-10G-SR            |
| Xcvr 6     | REV 01 | 740-031980 | B11J04016    | SFP+-10G-SR            |
| Xcvr 7     | REV 01 | 740-031980 | B11J04017    | SFP+-10G-SR            |
| Xcvr 8     | REV 01 | 740-031980 | B11J04887    | SFP+-10G-SR            |
| Xcvr 9     | REV 01 | 740-031980 | B11J04297    | SFP+-10G-SR            |
| Xcvr 10    | REV 01 | 740-031980 | B11J04893    | SFP+-10G-SR            |
| Xcvr 11    | REV 01 | 740-031980 | B11J04022    | SFP+-10G-SR            |
| PIC 1      | REV 02 | 750-034624 | EE3711       | 12x10GE (LAN/WAN) SFPP |
| Xcvr 0     | REV 01 | 740-031980 | AJH033X      | SFP+-10G-SR            |
| Xcvr 1     | REV 01 | 740-031980 | AJJ01N0      | SFP+-10G-SR            |
| Xcvr 2     | REV 01 | 740-031980 | AJJ01SV      | SFP+-10G-SR            |
| Xcvr 3     | REV 01 | 740-031980 | AJJ032L      | SFP+-10G-SR            |
| Xcvr 4     | REV 01 | 740-031980 | B10M01593    | SFP+-10G-SR            |
| Xcvr 5     | REV 01 | 740-031980 | AJD0FF1      | SFP+-10G-SR            |
| Xcvr 6     | REV 01 | 740-031980 | AJJ01NU      | SFP+-10G-SR            |
| Xcvr 7     | REV 01 | 740-031980 | 123363A01305 | SFP+-10G-SR            |
| Xcvr 8     | REV 01 | 740-031980 | B10M00361    | SFP+-10G-SR            |
| Xcvr 9     | REV 01 | 740-031980 | AJJ01M7      | SFP+-10G-SR            |
| Xcvr 10    | REV 01 | 740-031980 | AJJ032X      | SFP+-10G-SR            |
| Xcvr 11    | REV 01 | 740-031980 | AJJ01PG      | SFP+-10G-SR            |
| LMB 0      | REV 04 | 711-034381 | EF3838       | Type-0 LMB             |
| LMB 1      | REV 03 | 711-035774 | EF3821       | Type-1 LMB             |
| LMB 2      | REV 04 | 711-034381 | EF3834       | Type-0 LMB             |
| SPMB 0     | REV 05 | 710-023321 | ED1990       | LCC Switch CPU         |
| SPMB 1     | REV 05 | 710-023321 | EA2768       | LCC Switch CPU         |
| SIB 0      | REV 02 | 711-036340 | EF8802       | SIB-HC-3D              |
| SIB 1      | REV 07 | 711-036340 | EG2286       | SIB-HC-3D              |
| SIB 2      | REV 07 | 711-036340 | EG2252       | SIB-HC-3D              |
| SIB 3      | REV 02 | 711-036340 | EF1358       | SIB-HC-3D              |
| SIB 4      | REV 02 | 711-036340 | EF8806       | SIB-HC-3D              |
| Fan Tray 0 |        |            |              | Front Top Fan Tray     |
| Fan Tray 1 |        |            |              | Front Bottom Fan Tray  |
| -- Rev 2   |        |            |              |                        |
| Fan Tray 2 |        |            |              | Rear Fan Tray -- Rev 3 |

## show chassis hardware (T4000 Router with 16 GB line card chassis (LCC) Routing Engine)

```

user@host> show chassis hardware
Hardware inventory:
Item Version Part number Serial number Description
Chassis JN11BDF2CAHA T1600
Midplane REV 01 710-027486 ACAJ0774 T640 Backplane
FPM GBUS REV 13 710-002901 BBAL6812 T640 FPM Board
FPM Display REV 04 710-021387 BBAP2679 T1600 FPM Display
CIP REV 06 710-002895 BBAP4758 T-series CIP
PEM 0 Rev 03 740-026384 XF86421 Power Entry Module 3x80
PEM 1 Rev 03 740-026384 XF86429 Power Entry Module 3x80
SCG 0 REV 18 710-003423 BBAP1896 T640 Sonet Clock Gen.
SCG 1 REV 18 710-003423 BBAN8659 T640 Sonet Clock Gen.
Routing Engine 0 REV 01 740-042243 737F-002238 RE-DUO-1800-16G
Routing Engine 1 REV 01 740-042243 737F-002403 RE-DUO-1800-16G
CB 1 REV 11 710-022597 EK4526 LCC Control Board
CB 1 REV 11 710-022597 EK4527 LCC Control Board
FPC 0 REV 05 710-033871 EK5644 FPC Type 4-ES
CPU REV 11 710-016744 EK3428 ST-PMB2
PIC 0 REV 20 750-017405 EJ3041 4x 10GE (LAN/WAN) XFP
PIC 1 REV 17 750-026962 EH7536 10x10GE(LAN/WAN) SFPP
MMB 0 REV 07 710-025563 EK6039 ST-MMB2
MMB 1 REV 07 710-025563 EK6086 ST-MMB2
FPC 1 REV 05 710-033871 EK6583 FPC Type 4-ES
CPU REV 11 710-016744 EK3401 ST-PMB2
PIC 0 REV 17 750-026962 EJ8948 10x10GE(LAN/WAN) SFPP
MMB 0 REV 07 710-025563 EK6202 ST-MMB2
MMB 1 REV 07 710-025563 EK6112 ST-MMB2
SPMB 1 REV 05 710-023321 EK4900 LCC Switch CPU
SIB 0 REV 11 710-013074 EK5958 SIB-I8-SF
SIB 1 REV 11 710-013074 EK4606 SIB-I8-SF
SIB 2 REV 11 710-013074 EK5971 SIB-I8-SF
SIB 3 REV 11 710-013074 EK4609 SIB-I8-SF
SIB 4 REV 11 710-013074 EK4602 SIB-I8-SF
Fan Tray 0 Front Top Fan Tray
Fan Tray 1 Front Bottom Fan Tray
Fan Tray 2 Rear Fan Tray -- Rev 2

```

## show chassis hardware clei-models (T4000 Router)

```

user@host> show chassis hardware clei-models
Hardware inventory:
Item Version Part number CLEI code FRU model number
Midplane REV 01 710-027486 IPMJ700DRD CHAS-BP-T1600-S
FPM Display REV 01 710-021387 IPUPAG6KAA CRAFT-T1600-S
CIP REV 06 710-002895 IPUPAG6KAA CIP-L-T640-S
PEM 0 REV 01 740-036442 IPUPAG6KAA PWR-T-6-60-DC
SCG 0 REV 18 710-003423 IPUPAG6KAA SCG-T-S
SCG 1 REV 18 710-003423 IPUPAG6KAA SCG-T-S
Routing Engine 0 REV 05 740-026941 IPUPAG6KAA RE-DUO-C1800-8G-S
Routing Engine 1 REV 06 740-026941 IPUPAG6KAA RE-DUO-C1800-8G-S
CB 0 REV 09 710-022597 IPUPAG6KAA CB-LCC-S
CB 1 REV 09 710-022597 IPUPAG6KAA CB-LCC-S
FPC 3
PIC 0 REV 08 750-035293 XXXXXXXXBB PF-1CGE-CFP
PIC 1 REV 10 750-034624 XXXXXXXXCX PF-12XGE-SFPP
FPC 5 REV 03 710-033871 IPUCAMBCTD T1600-FPC4-ES
PIC 1 REV 03 750-034781 IPUIBKLMMA PD-1CE-CFP-FPC4
FPC 6
PIC 0 REV 10 750-034624 XXXXXXXXCX PF-12XGE-SFPP

```

|            |                 |
|------------|-----------------|
| Fan Tray 0 | FANTRAY-T-S     |
| Fan Tray 1 | FANTRAY-T4000-S |
| Fan Tray 2 | FANTRAY-TXP-R-S |

**show chassis hardware detail (T4000 Router)**

```

user@host> show chassis hardware detail
Hardware inventory:
Item Version Part number Serial number Description
Chassis REV 01 710-027486 JN1172F25AHA T4000
Midplane REV 13 710-002901 RC8355 T-series Backplane
FPM GBUS REV 01 710-021387 EF6764 T640 FPM Board
FPM Display REV 06 710-002895 BBAD9210 T1600 FPM Display
CIP REV 01 740-036442 VA00016 Power Entry Module 6x60
PEM 0 REV 18 710-003423 BBAD7248 T640 Sonet Clock Gen.
SCG 0 REV 18 710-003423 BBAD7248 T640 Sonet Clock Gen.
SCG 1 REV 18 710-003423 BBAD7248 T640 Sonet Clock Gen.
Routing Engine 0 REV 05 740-026941 P737F-002248 RE-DUO-1800
 ad0 3823 MB SMART CF 2009121602A661576157 Compact Flash
 ad1 59690 MB STEC MACH-8 SSD STM000103FDB Disk 1
Routing Engine 1 REV 06 740-026941 P737F-002653 RE-DUO-1800
 ad0 3823 MB SMART CF 201011150153F52CF52C Compact Flash
 ad1 62720 MB SMART Lite SATA Drive 2010110900150A880A88 Disk 1
CB 0 REV 09 710-022597 ED0295 LCC Control Board
CB 1 REV 09 710-022597 EA6050 LCC Control Board
FPC 0 REV 26 750-032819 EK1173 FPC Type 5-3D
CPU REV 12 711-030686 EJ8584 SNG PMB
PIC 0 REV 07 750-034624 EF6837 12x10GE (LAN/WAN) SFPP
 Xcvr 0 REV 01 740-031980 123363A01145 SFP+-10G-SR
 Xcvr 1 REV 01 740-031980 123363A01147 SFP+-10G-SR
 Xcvr 2 REV 01 740-031980 AJJ01P3 SFP+-10G-SR
 Xcvr 3 REV 01 740-031980 B10M03256 SFP+-10G-SR
 Xcvr 4 REV 01 740-031980 AJJ01M2 SFP+-10G-SR
 Xcvr 5 REV 01 740-031980 123363A01137 SFP+-10G-SR
 Xcvr 6 REV 01 740-031980 AJJ01PN SFP+-10G-SR
 Xcvr 7 REV 01 740-031980 AJJ01NW SFP+-10G-SR
 Xcvr 8 REV 01 740-031980 123363A01139 SFP+-10G-SR
 Xcvr 9 REV 01 740-031980 AJJ01KE SFP+-10G-SR
 Xcvr 10 REV 01 740-031980 123363A01336 SFP+-10G-SR
 Xcvr 11 REV 01 740-031980 B10M01325 SFP+-10G-SR
PIC 1 REV 07 750-034624 EF6800 12x10GE (LAN/WAN) SFPP
 Xcvr 0 REV 01 740-031980 AJJ01SA SFP+-10G-SR
 Xcvr 1 REV 01 740-031980 AJJ01QZ SFP+-10G-SR
 Xcvr 2 REV 01 740-031980 AJH0217 SFP+-10G-SR
 Xcvr 3 REV 01 740-031980 AJJ01TE SFP+-10G-SR
 Xcvr 4 REV 01 740-031980 AJJ01KV SFP+-10G-SR
 Xcvr 5 REV 01 740-031980 AJJ01MU SFP+-10G-SR
 Xcvr 6 REV 01 740-031980 AJJ01R0 SFP+-10G-SR
 Xcvr 7 REV 01 740-031980 AJJ01TC SFP+-10G-SR
 Xcvr 8 REV 01 740-031980 AJJ0364 SFP+-10G-SR
 Xcvr 9 REV 01 740-031980 AJD0GV3 SFP+-10G-SR
 Xcvr 10 REV 01 740-031980 B10M03343 SFP+-10G-SR
 Xcvr 11 REV 01 740-031980 AJJ01QJ SFP+-10G-SR
LMB 0 REV 05 711-034381 EJ8490 Type-0 LMB
LMB 1 REV 04 711-035774 EJ8517 Type-1 LMB
LMB 2 REV 05 711-034381 EJ8489 Type-0 LMB
FPC 3 REV 07 750-032819 EG3637 FPC Type 5-3D
CPU REV 09 711-030686 EG0150 SNG PMB
PIC 0 REV 08 750-035293 EF3657 1x100GE
 Xcvr 0 REV 01 740-032210 C22CQNJ CFP-100G-LR4
PIC 1 REV 10 750-034624 BBAN4098 12x10GE (LAN/WAN) SFPP

```



|          |        |            |              |                        |
|----------|--------|------------|--------------|------------------------|
| Xcvr 0   | REV 01 | 740-031980 | B11J04902    | SFP+-10G-SR            |
| Xcvr 1   | REV 01 | 740-031980 | B11J04891    | SFP+-10G-SR            |
| Xcvr 2   | REV 01 | 740-031980 | AJJ01MX      | SFP+-10G-SR            |
| Xcvr 3   | REV 01 | 740-031980 | B11J04183    | SFP+-10G-SR            |
| Xcvr 4   | REV 01 | 740-031980 | B11J04894    | SFP+-10G-SR            |
| Xcvr 5   | REV 01 | 740-031980 | B11J04184    | SFP+-10G-SR            |
| Xcvr 6   | REV 01 | 740-031980 | B11J04897    | SFP+-10G-SR            |
| Xcvr 7   | REV 01 | 740-031980 | B11J04899    | SFP+-10G-SR            |
| Xcvr 8   | REV 01 | 740-031980 | AJJ01TV      | SFP+-10G-SR            |
| Xcvr 9   | REV 01 | 740-031980 | B11J04057    | SFP+-10G-SR            |
| Xcvr 10  | REV 01 | 740-031980 | AJJ01M4      | SFP+-10G-SR            |
| Xcvr 11  | REV 01 | 740-031980 | B11J04905    | SFP+-10G-SR            |
| LMB 0    | REV 04 | 711-034381 | EG1524       | Type-0 LMB             |
| LMB 1    | REV 03 | 711-035774 | EG0345       | Type-1 LMB             |
| LMB 2    | REV 04 | 711-034381 | EG1522       | Type-0 LMB             |
| FPC 5    | REV 03 | 710-033871 | BBAJ0768     | FPC Type 4-ES          |
| CPU      | REV 11 | 710-016744 | BBAH9342     | ST-PMB2                |
| PIC 0    | REV 09 | 750-029262 | EE6789       | 100GE                  |
| PIC 1    | REV 03 | 750-034781 | EE6655       | 100GE CFP              |
| Xcvr 0   | REV 01 | 740-032210 | J11A22334    | CFP-100G-LR4           |
| BRIDGE 0 | REV 03 | 711-029995 | EE6572       | 100GE Bridge Board     |
| MMB 0    | REV 07 | 710-025563 | BBAJ4657     | ST-MMB2                |
| MMB 1    | REV 07 | 710-025563 | BBAJ3073     | ST-MMB2                |
| FPC 6    | REV 05 | 750-010153 | EF4936       | FPC Type 5-3D          |
| CPU      | REV 06 | 711-030686 | EF4189       | SNG PMB                |
| PIC 0    | REV 10 | 750-034624 | BBAN4109     | 12x10GE (LAN/WAN) SFPP |
| Xcvr 0   | REV 01 | 740-031980 | B11J04895    | SFP+-10G-SR            |
| Xcvr 1   | REV 01 | 740-031980 | B11J04898    | SFP+-10G-SR            |
| Xcvr 2   | REV 01 | 740-031980 | B11J04021    | SFP+-10G-SR            |
| Xcvr 3   | REV 01 | 740-031980 | B11J04903    | SFP+-10G-SR            |
| Xcvr 4   | REV 01 | 740-031980 | B11J04311    | SFP+-10G-SR            |
| Xcvr 5   | REV 01 | 740-031980 | B11J04059    | SFP+-10G-SR            |
| Xcvr 6   | REV 01 | 740-031980 | B11J04016    | SFP+-10G-SR            |
| Xcvr 7   | REV 01 | 740-031980 | B11J04017    | SFP+-10G-SR            |
| Xcvr 8   | REV 01 | 740-031980 | B11J04887    | SFP+-10G-SR            |
| Xcvr 9   | REV 01 | 740-031980 | B11J04297    | SFP+-10G-SR            |
| Xcvr 10  | REV 01 | 740-031980 | B11J04893    | SFP+-10G-SR            |
| Xcvr 11  | REV 01 | 740-031980 | B11J04022    | SFP+-10G-SR            |
| PIC 1    | REV 02 | 750-034624 | EE3711       | 12x10GE (LAN/WAN) SFPP |
| Xcvr 0   | REV 01 | 740-031980 | AJH033X      | SFP+-10G-SR            |
| Xcvr 1   | REV 01 | 740-031980 | AJJ01N0      | SFP+-10G-SR            |
| Xcvr 2   | REV 01 | 740-031980 | AJJ01SV      | SFP+-10G-SR            |
| Xcvr 3   | REV 01 | 740-031980 | AJJ032L      | SFP+-10G-SR            |
| Xcvr 4   | REV 01 | 740-031980 | B10M01593    | SFP+-10G-SR            |
| Xcvr 5   | REV 01 | 740-031980 | AJD0FF1      | SFP+-10G-SR            |
| Xcvr 6   | REV 01 | 740-031980 | AJJ01NU      | SFP+-10G-SR            |
| Xcvr 7   | REV 01 | 740-031980 | 123363A01305 | SFP+-10G-SR            |
| Xcvr 8   | REV 01 | 740-031980 | B10M00361    | SFP+-10G-SR            |
| Xcvr 9   | REV 01 | 740-031980 | AJJ01M7      | SFP+-10G-SR            |
| Xcvr 10  | REV 01 | 740-031980 | AJJ032X      | SFP+-10G-SR            |
| Xcvr 11  | REV 01 | 740-031980 | AJJ01PG      | SFP+-10G-SR            |
| LMB 0    | REV 04 | 711-034381 | EF3838       | Type-0 LMB             |
| LMB 1    | REV 03 | 711-035774 | EF3821       | Type-1 LMB             |
| LMB 2    | REV 04 | 711-034381 | EF3834       | Type-0 LMB             |
| SPMB 0   | REV 05 | 710-023321 | ED1990       | LCC Switch CPU         |
| SPMB 1   | REV 05 | 710-023321 | EA2768       | LCC Switch CPU         |
| SIB 0    | REV 02 | 711-036340 | EF8802       | SIB-HC-3D              |
| SIB 1    | REV 07 | 711-036340 | EG2286       | SIB-HC-3D              |
| SIB 2    | REV 07 | 711-036340 | EG2252       | SIB-HC-3D              |
| SIB 3    | REV 02 | 711-036340 | EF1358       | SIB-HC-3D              |
| SIB 4    | REV 02 | 711-036340 | EF8806       | SIB-HC-3D              |

```

Fan Tray 0
Fan Tray 1
-- Rev 2
Fan Tray 2

Front Top Fan Tray
Front Bottom Fan Tray

Rear Fan Tray -- Rev 3

```

### show chassis hardware models (T4000 Router)

```
user@host> show chassis hardware models
```

```

Hardware inventory:
Item Version Part number Serial number FRU model number
Midplane REV 01 710-027486 RC8355 CHAS-BP-T1600-S
FPM Display REV 01 710-021387 EF6764 CRAFT-T1600-S
CIP REV 06 710-002895 BBAD9210 CIP-L-T640-S
PEM 0 REV 01 740-036442 VA00016 PWR-T-6-60-DC
SCG 0 REV 18 710-003423 BBAD7248 SCG-T-S
SCG 1 REV 18 710-003423 BBAE3874 SCG-T-S
Routing Engine 0 REV 05 740-026941 P737F-002248 RE-DUO-C1800-8G-S
Routing Engine 1 REV 06 740-026941 P737F-002653 RE-DUO-C1800-8G-S
CB 0 REV 09 710-022597 ED0295 CB-LCC-S
CB 1 REV 09 710-022597 EA6050 CB-LCC-S
FPC 3
 PIC 0 REV 08 750-035293 EF3657 PF-1CGE-CFP
 PIC 1 REV 10 750-034624 BBAN4098 PF-12XGE-SFPP
FPC 5 REV 03 710-033871 BBAJ0768 T1600-FPC4-ES
 PIC 1 REV 03 750-034781 EE6655 PD-1CE-CFP-FPC4
FPC 6
 PIC 0 REV 10 750-034624 BBAN4109 PF-12XGE-SFPP
Fan Tray 0 FANTRAY-T-S
Fan Tray 1 FANTRAY-T4000-S
Fan Tray 2 FAN-REAR-TXP-LCC

```

### show chassis hardware lcc (TX Matrix Router)

```
user@host> show chassis hardware lcc 0
lcc0-re0:
```

```

Hardware inventory:
Item Version Part number Serial number Description
Chassis 65751 T640
Midplane REV 03 710-005608 RA1408 T640 Backplane
FPM GBUS REV 09 710-002901 RA2784 T640 FPM Board
FPM Display REV 05 710-002897 RA2825 FPM Display
CIP REV 06 710-002895 HT0684 T Series CIP
PEM 0 Rev 11 740-002595 PM18483 Power Entry Module
PEM 1 Rev 11 740-002595 qb13984 Power Entry Module
SCG 0 REV 11 710-003423 HT0022 T640 Sonet Clock Gen.
Routing Engine 0 REV 13 740-005022 210865700363 RE-3.0 (RE-600)
CB 0 REV 03 710-007655 HW1195 Control Board (CB-T)
FPC 1 REV 05 710-007527 HM3245 FPC Type 2
 CPU REV 14 710-001726 HM1084 FPC CPU
 PIC 0 REV 02 750-007218 AZ1112 2x OC-12 ATM2 IQ, SMIR
 PIC 1 REV 02 750-007745 HG3462 4x OC-3 SONET, SMIR
 PIC 2 REV 14 750-001901 BA5390 4x OC-12 SONET, SMIR
 PIC 3 REV 09 750-008155 HS3012 2x G/E IQ, 1000 BASE
 SFP 0 NON-JNPR P1186TY SFP-S
 SFP 1 REV 01 740-007326 P11WLTF SFP-SX
 MMB 1 REV 02 710-005555 HL7514 MMB-288mbit
 PPB 0 REV 04 710-003758 HM4405 PPB Type 2
 PPB 1 REV 04 710-003758 AV1960 PPB Type 2
FPC 2 REV 08 710-010154 HZ3578 E-FPC Type 3

```

|         |        |            |            |                         |
|---------|--------|------------|------------|-------------------------|
| CPU     | REV 05 | 710-010169 | HZ3219     | FPC CPU-Enhanced        |
| PIC 0   | REV 02 | 750-009567 | HX2882     | 1x 10GE(LAN), XENPAK    |
| SFP 0   | REV 01 | 740-009898 | USC202U709 | XENPAK-LR               |
| PIC 1   | REV 03 | 750-003336 | HJ9954     | 4x OC-48 SONET, SMSR    |
| PIC 2   | REV 01 | 750-004535 | HC0235     | 1x OC-192 SM SR1        |
| PIC 3   | REV 07 | 750-007141 | HX1699     | 10x 1GE(LAN), 1000 BASE |
| SFP 0   | REV 01 | 740-007326 | 2441042    | SFP-SX                  |
| SFP 1   | REV 01 | 740-007326 | 2441027    | SFP-SX                  |
| MMB 0   | REV 03 | 710-010171 | HV2365     | MMB-5M3-288mbit         |
| MMB 1   | REV 03 | 710-010171 | HZ3888     | MMB-5M3-288mbit         |
| SPMB 0  | REV 09 | 710-003229 | HW5245     | T Series Switch CPU     |
| SIB 3   | REV 07 | 710-005781 | HR5927     | SIB-L8-F16              |
| B Board | REV 06 | 710-005782 | HR5971     | SIB-L8-F16 (B)          |
| SIB 4   | REV 07 | 710-005781 | HR5903     | SIB-L8-F16              |
| B Board | REV 06 | 710-005782 | HZ5275     | SIB-L8-F16 (B)          |

### show chassis hardware scc (TX Matrix Router)

```
user@host> show chassis hardware scc
scc-re0:
```

#### Hardware inventory:

| Item             | Version | Part number | Serial number | Description           |
|------------------|---------|-------------|---------------|-----------------------|
| Chassis          |         |             |               | TX Matrix             |
| Midplane         | REV 04  | 710-004396  | RB0014        | SCC Midplane          |
| FPM GBUS         | REV 04  | 710-004617  | HW9141        | SCC FPM Board         |
| FPM Display      | REV 04  | 710-004619  | HS5950        | SCC FPM               |
| CIP 0            | REV 01  | 710-010218  | HV9151        | SCC CIP               |
| CIP 1            | REV 01  | 710-010218  | HV9152        | SCC CIP               |
| PEM 1            | Rev 11  | 740-002595  | QB13977       | Power Entry Module    |
| Routing Engine 0 | REV 05  | 740-008883  | P11123900153  | RE-4.0 (RE-1600)      |
| CB 0             | REV 01  | 710-011709  | HR5964        | Control Board (CB-TX) |
| SPMB 0           | REV 09  | 710-003229  | HW5293        | T Series Switch CPU   |
| SIB 3            |         |             |               |                       |
| SIB 4            | REV 01  | 710-005839  | HW1177        | SIB-S8-F16            |
| B Board          | REV 01  | 710-005840  | HW1202        | SIB-S8-F16 (B)        |

### show chassis hardware (T1600 Router)

```
user@host> show chassis hardware
```

#### Hardware inventory:

| Item             | Version | Part number | Serial number | Description             |
|------------------|---------|-------------|---------------|-------------------------|
| Chassis          |         |             | B2703         | T1600                   |
| Midplane         | REV 03  | 710-005608  | RC4137        | T640 Backplane          |
| FPM GBUS         | REV 10  | 710-002901  | DT7062        | T640 FPM Board          |
| FPM Display      | REV 05  | 710-002897  | DS3067        | FPM Display             |
| CIP              | REV 06  | 710-002895  | DT3386        | T-series CIP            |
| PEM 0            | Rev 07  | 740-017906  | UA26344       | Power Entry Module 3x80 |
| PEM 1            | Rev 18  | 740-002595  | UF38441       | Power Entry Module      |
| SCG 0            | REV 15  | 710-003423  | DV0941        | T640 Sonet Clock Gen.   |
| Routing Engine 0 | REV 08  | 740-014082  | 9009014502    | RE-A-2000               |
| Routing Engine 1 | REV 07  | 740-014082  | 9009009591    | RE-A-2000               |
| CB 0             | REV 05  | 710-007655  | JA9360        | Control Board (CB-T)    |
| CB 1             | REV 03  | 710-017707  | DT3251        | Control Board (CB-T)    |
| FPC 0            | REV 07  | 710-013558  | DR4253        | E2-FPC Type 2           |
| CPU              | REV 05  | 710-013563  | DS3902        | FPC CPU-Enhanced        |
| PIC 0            | REV 01  | 750-010618  | CB5446        | 4x G/E SFP, 1000 BASE   |
| Xcvr 0           | REV 01  | 740-011613  | P9F11CW       | SFP-SX                  |
| Xcvr 1           | REV 01  | 740-011613  | P9F15C2       | SFP-SX                  |
| Xcvr 2           | REV 01  | 740-011782  | PB94K0L       | SFP-SX                  |

|        |        |            |             |                          |
|--------|--------|------------|-------------|--------------------------|
| PIC 1  | REV 06 | 750-001900 | HB6399      | 1x OC-48 SONET, SMSR     |
| PIC 2  | REV 14 | 750-001901 | AP1092      | 4x OC-12 SONET, SMIR     |
| PIC 3  | REV 07 | 750-001900 | AR8275      | 1x OC-48 SONET, SMSR     |
| MMB 1  | REV 07 | 710-010171 | DS1524      | MMB-5M3-288mbit          |
| FPC 1  | REV 06 | 710-013553 | DL9067      | E2-FPC Type 1            |
| CPU    | REV 04 | 710-013563 | DM1685      | FPC CPU-Enhanced         |
| PIC 0  | REV 08 | 750-001072 | AB1688      | 1x G/E, 1000 BASE-SX     |
| PIC 1  | REV 10 | 750-012266 | JX5519      | 4x 1GE(LAN), IQ2         |
| Xcvr 0 | REV 01 | 740-011613 | AM0812S8UK6 | SFP-SX                   |
| Xcvr 2 | REV 01 | 740-011613 | AM0812S8UK1 | SFP-SX                   |
| Xcvr 3 | REV 01 | 740-011782 | P8N1YHG     | SFP-SX                   |
| PIC 2  | REV 22 | 750-005634 | DP0083      | 1x CHOC12 IQ SONET, SMIR |
| MMB 1  | REV 07 | 710-008923 | DN1862      | MMB 3M 288-bit           |
| FPC 2  | REV 01 | 710-005548 | HJ9899      | FPC Type 3               |
| CPU    | REV 06 | 710-001726 | HC0586      | FPC CPU                  |
| PIC 0  | REV 16 | 750-007141 | NC9660      | 10x 1GE(LAN), 1000 BASE  |
| Xcvr 0 | REV 01 | 740-011613 | AM0812S8XAR | SFP-SX                   |
| Xcvr 1 | REV 01 | 740-011782 | P920E7B     | SFP-SX                   |
| Xcvr 2 | REV 01 | 740-011613 | AM0812S8XAU | SFP-SX                   |
| Xcvr 4 | REV 01 | 740-011613 | AM0812S8XAK | SFP-SX                   |
| Xcvr 5 | REV 01 | 740-011613 | AM0812S8XAA | SFP-SX                   |
| Xcvr 6 | REV 01 | 740-011613 | PAJ4NKY     | SFP-SX                   |
| Xcvr 7 | REV 01 | 740-011613 | AM0812S8UJW | SFP-SX                   |
| Xcvr 8 | REV 01 | 740-011782 | PB81X89     | SFP-SX                   |
| Xcvr 9 | REV 01 | 740-011613 | AM0812S8UJX | SFP-SX                   |
| PIC 1  | REV 06 | 750-015217 | DK3280      | 8x 1GE(TYPE3), IQ2       |
| Xcvr 0 | REV 01 | 740-011782 | P8P0A3T     | SFP-SX                   |
| Xcvr 1 | REV 01 | 740-013111 | 5090002     | SFP-T                    |
| Xcvr 2 | REV 01 | 740-011613 | AM0814S93BQ | SFP-SX                   |
| Xcvr 4 |        | NON-JNPR   | PDE0FAN     | SFP-SX                   |
| Xcvr 5 | REV 01 | 740-011782 | P8Q20XY     | SFP-SX                   |
| Xcvr 6 | REV 01 | 740-011613 | AM0812S8UJV | SFP-SX                   |
| Xcvr 7 | REV 01 | 740-011613 | AM0812S8UP7 | SFP-SX                   |
| PIC 2  | REV 05 | 750-004695 | HT4383      | 1x Tunnel                |
| PIC 3  | REV 17 | 750-009553 | RL0204      | 4x OC-48 SONET           |
| Xcvr 0 | REV 01 | 740-011785 | PDS3T23     | SFP-SR                   |
| Xcvr 1 | REV 01 | 740-011785 | P6Q0F3E     | SFP-SR                   |
| MMB 0  | REV 03 | 710-004047 | HD5843      | MMB-288mbit              |
| MMB 1  | REV 03 | 710-004047 | HE3208      | MMB-288mbit              |
| PPB 0  | REV 02 | 710-002845 | HA4524      | PPB Type 3               |
| PPB 1  | REV 02 | 710-002845 | HA4766      | PPB Type 3               |
| FPC 3  | REV 01 | 710-010154 | HR0863      | E-FPC Type 3             |
| CPU    | REV 01 | 710-010169 | HN3422      | FPC CPU-Enhanced         |
| PIC 0  | REV 07 | 750-012793 | WF5096      | 1x 10GE(LAN/WAN) IQ2     |
| Xcvr 0 |        | NON-JNPR   | M64294TP    | XFP-10G-LR               |
| PIC 1  | REV 25 | 750-007141 | DV2127      | 10x 1GE(LAN), 1000 BASE  |
| Xcvr 0 | REV 01 | 740-011613 | PFA6LTJ     | SFP-SX                   |
| Xcvr 1 | REV 01 | 740-011782 | P9P0XV4     | SFP-SX                   |
| Xcvr 2 | REV 01 | 740-011782 | P9M0TNX     | SFP-SX                   |
| Xcvr 4 | REV 01 | 740-011782 | P9B0TTP     | SFP-SX                   |
| Xcvr 5 |        | NON-JNPR   | PBS4LED     | SFP-SX                   |
| PIC 2  | REV 17 | 750-009553 | RL0212      | 4x OC-48 SONET           |
| Xcvr 0 | REV 01 | 740-011785 | PDS3T8G     | SFP-SR                   |
| PIC 3  | REV 32 | 750-003700 | DL1279      | 1x OC-192 12xMM VSR      |
| MMB 0  | REV 01 | 710-010171 | HR0821      | MMB-288mbit              |
| MMB 1  | REV 01 | 710-010171 | HR0818      | MMB-288mbit              |
| FPC 4  | REV 16 | 710-013037 | EB4919      | FPC Type 4-ES            |
| CPU    | REV 09 | 710-016744 | BBAA4382    | ST-PMB2                  |

|            |        |            |             |                       |
|------------|--------|------------|-------------|-----------------------|
| PIC 0      | REV 03 | 711-029996 | EB1569      | 100GE                 |
| PIC 1      | REV 05 | 711-029999 | EB9983      | 100GE CFP             |
| Xcvr 0     | REV 0  | 740-032210 | J10G80746   | CFP-100G-LR4          |
| BRIDGE 0   | REV 02 | 711-029995 | EB2235      | 100GE Bridge Board    |
| MMB 0      | REV 04 | 710-025563 | BBA7112     | ST-MMB2               |
| MMB 1      | REV 04 | 710-025563 | BBA7149     | ST-MMB2               |
| FPC 5      | REV 02 | 710-013037 | DE3407      | FPC Type 4-ES         |
| CPU        | REV 04 | 710-016744 | DA2124      | ST-PMB2               |
| PIC 0      | REV 16 | 750-012518 | DF2554      | 4x OC-192 SONET XFP   |
| Xcvr 0     | REV 01 | 740-014279 | AA0745N1FX8 | XFP-OC192-SR          |
| Xcvr 1     | REV 01 | 740-014279 | AA0748N1HN5 | XFP-OC192-SR          |
| Xcvr 2     | REV 01 | 740-014279 | AA0748N1HT6 | XFP-OC192-SR          |
| Xcvr 3     | REV 01 | 740-014279 | AA0744N1EC9 | XFP-OC192-SR          |
| PIC 1      | REV 01 | 750-010850 | JA0329      | 1x OC-768 SONET SR    |
| MMB 0      | REV 04 | 710-016036 | DE9577      | ST-MMB2               |
| MMB 1      | REV 04 | 710-016036 | DK4060      | ST-MMB2               |
| FPC 6      | REV 14 | 710-013037 | DV1431      | FPC Type 4-ES         |
| CPU        | REV 09 | 710-016744 | DT9020      | ST-PMB2               |
| PIC 0      | REV 11 | 750-017405 | DM6261      | 4x 10GE (LAN/WAN) XFP |
| Xcvr 0     | REV 01 | 740-014289 | C701XU05Q   | XFP-10G-SR            |
| Xcvr 1     | REV 01 | 740-014279 | AA0748N1HPT | XFP-10G-LR            |
| Xcvr 2     | REV 01 | 740-014289 | T08E19189   | XFP-10G-SR            |
| Xcvr 3     | REV 01 | 740-014289 | C715XU058   | XFP-10G-SR            |
| PIC 1      | REV 13 | 750-017405 | DP8772      | 4x 10GE (LAN/WAN) XFP |
| Xcvr 0     | REV 02 | 740-011571 | C850XJ037   | XFP-10G-SR            |
| Xcvr 1     | REV 02 | 740-014289 | C839XU0L9   | XFP-10G-SR            |
| Xcvr 2     | REV 02 | 740-014289 | C834XU05A   | XFP-10G-SR            |
| Xcvr 3     | REV 02 | 740-014289 | C810XU0CE   | XFP-10G-SR            |
| MMB 0      | REV 01 | 710-025563 | DT8454      | ST-MMB2               |
| MMB 1      | REV 01 | 710-025563 | DT8366      | ST-MMB2               |
| FPC 7      | REV 09 | 710-007529 | HZ7624      | FPC Type 3            |
| CPU        | REV 15 | 710-001726 | HZ1413      | FPC CPU               |
| PIC 0      | REV 10 | 750-012793 | DM5627      | 1x 10GE(LAN/WAN) IQ2  |
| Xcvr 0     | REV 02 | 740-011571 | C831XJ062   | XFP-10G-SR            |
| PIC 1      | REV 01 | 750-015217 | JT6762      | 8x 1GE(TYPE3), IQ2    |
| Xcvr 0     | REV 01 | 740-011782 | P8Q25JU     | SFP-SX                |
| Xcvr 1     | REV 01 | 740-011782 | P9B0U0K     | SFP-SX                |
| PIC 2      | REV 01 | 750-015217 | JS4268      | 8x 1GE(TYPE3), IQ2    |
| Xcvr 0     | REV 01 | 740-011613 | AM0812S8XBZ | SFP-SX                |
| Xcvr 1     | REV 01 | 740-011613 | AM0812S8XAP | SFP-SX                |
| Xcvr 2     | REV 01 | 740-011613 | AM0812S8XBY | SFP-SX                |
| Xcvr 3     | REV 01 | 740-011613 | AM0812S8XBX | SFP-SX                |
| Xcvr 4     | REV 01 | 740-011613 | P9F1652     | SFP-SX                |
| Xcvr 5     | REV 01 | 740-011782 | P8Q21YC     | SFP-SX                |
| Xcvr 6     | REV 01 | 740-011782 | P8Q27HQ     | SFP-SX                |
| Xcvr 7     | REV 01 | 740-011613 | P8E2SSU     | SFP-SX                |
| PIC 3      | REV 15 | 750-009450 | NB6790      | 1x OC-192 SM SR2      |
| MMB 0      | REV 03 | 710-005555 | HZ3450      | MMB-288mbit           |
| MMB 1      | REV 03 | 710-005555 | HZ3415      | MMB-288mbit           |
| PPB 0      | REV 04 | 710-002845 | HP0887      | PPB Type 3            |
| PPB 1      | REV 04 | 710-002845 | HW5255      | PPB Type 3            |
| SPMB 0     | REV 10 | 710-003229 | HX3699      | T-series Switch CPU   |
| SPMB 1     | REV 12 | 710-003229 | DT3091      | T-series Switch CPU   |
| SIB 0      | REV 07 | 710-013074 | DS4747      | SIB-I8-SF             |
| SIB 1      | REV 07 | 710-013074 | DS4942      | SIB-I8-SF             |
| SIB 2      | REV 07 | 710-013074 | DS4965      | SIB-I8-SF             |
| SIB 3      | REV 07 | 710-013074 | DS4990      | SIB-I8-SF             |
| SIB 4      | REV 07 | 710-013074 | DS4944      | SIB-I8-SF             |
| Fan Tray 0 |        |            |             | Front Top Fan Tray    |

Fan Tray 1  
Fan Tray 2

Front Bottom Fan Tray  
Rear Fan Tray -- Rev 2

### show chassis hardware (TX Matrix Plus Router)

```
user@host> show chassis hardware
sfc0-re0:
```

-----  
Hardware inventory:

| Item             | Version | Part number | Serial number | Description        |
|------------------|---------|-------------|---------------|--------------------|
| Chassis          |         |             | JN113186EAHB  | TXP                |
| Midplane         | REV 05  | 710-022574  | TS3822        | SFC Midplane       |
| FPM Display      | REV 03  | 710-024027  | DW4701        | TXP FPM Display    |
| CIP 0            | REV 05  | 710-023792  | DW7998        | TXP CIP            |
| CIP 1            | REV 05  | 710-023792  | DW7999        | TXP CIP            |
| PEM 0            | Rev 04  | 740-027463  | UM26367       | Power Entry Module |
| PEM 1            | Rev 04  | 740-027463  | UM26346       | Power Entry Module |
| Routing Engine 0 | REV 06  | 740-026942  | 737A-1081     | RE-DUO-2600        |
| Routing Engine 1 | REV 06  | 740-026942  | 737A-1043     | RE-DUO-2600        |
| CB 0             | REV 05  | 710-022606  | DW4435        | SFC Control Board  |
| CB 1             | REV 09  | 710-022606  | DW6100        | SFC Control Board  |
| SPMB 0           |         | BUILTIN     |               | SFC Switch CPU     |
| SPMB 1           |         | BUILTIN     |               | SFC Switch CPU     |
| SIB F13 0        | REV 04  | 750-024564  | DW5764        | F13 SIB            |
| B Board          | REV 03  | 710-023431  | DW9053        | F13 SIB Mezz       |
| SIB F13 3        | REV 04  | 750-024564  | DW5785        | F13 SIB            |
| B Board          | REV 03  | 710-023431  | DW9030        | F13 SIB Mezz       |
| SIB F13 6        |         |             |               |                    |
| SIB F13 8        | REV 04  | 750-024564  | DW5752        | F13 SIB            |
| B Board          | REV 03  | 710-023431  | DW9051        | F13 SIB Mezz       |
| SIB F13 11       | REV 04  | 750-024564  | DW5782        | F13 SIB            |
| B Board          | REV 03  | 710-023431  | DW9058        | F13 SIB Mezz       |
| SIB F13 12       | REV 03  | 750-024564  | DT9466        | F13 SIB            |
| B Board          | REV 02  | 710-023431  | DT6556        | F13 SIB Mezz       |
| SIB F2S 0/0      | REV 05  | 710-022603  | DW7898        | F2S SIB            |
| B Board          | REV 05  | 710-023787  | DW7625        | F2S SIB Mezz       |
| SIB F2S 0/2      | REV 05  | 710-022603  | DW7811        | F2S SIB            |
| B Board          | REV 05  | 710-023787  | DW7550        | F2S SIB Mezz       |
| SIB F2S 0/4      | REV 04  | 710-022603  | DW4873        | F2S SIB            |
| B Board          | REV 05  | 710-023787  | DW8509        | F2S SIB Mezz       |
| SIB F2S 0/6      | REV 04  | 710-022603  | DW4867        | F2S SIB            |
| B Board          | REV 05  | 710-023787  | DW8472        | F2S SIB Mezz       |
| SIB F2S 1/0      | REV 04  | 710-022603  | DW4871        | F2S SIB            |
| B Board          | REV 05  | 710-023787  | DW8497        | F2S SIB Mezz       |
| SIB F2S 1/2      | REV 05  | 710-022603  | DW7868        | F2S SIB            |
| B Board          | REV 05  | 710-023787  | DW7551        | F2S SIB Mezz       |
| SIB F2S 1/4      | REV 04  | 710-022603  | DW4854        | F2S SIB            |
| B Board          | REV 05  | 710-023787  | DW8496        | F2S SIB Mezz       |
| SIB F2S 1/6      | REV 05  | 710-022603  | DW7889        | F2S SIB            |
| B Board          | REV 05  | 710-023787  | DW7496        | F2S SIB Mezz       |
| SIB F2S 2/0      | REV 04  | 710-022603  | DW4852        | F2S SIB            |
| B Board          | REV 05  | 710-023787  | DW8498        | F2S SIB Mezz       |
| SIB F2S 2/2      | REV 04  | 710-022603  | DW4845        | F2S SIB            |
| B Board          | REV 05  | 710-023787  | DW8457        | F2S SIB Mezz       |
| SIB F2S 2/4      | REV 05  | 710-022603  | DW7802        | F2S SIB            |
| B Board          | REV 05  | 710-023787  | DW7562        | F2S SIB Mezz       |
| SIB F2S 2/6      | REV 04  | 710-022603  | DW4822        | F2S SIB            |
| B Board          | REV 05  | 710-023787  | DW8467        | F2S SIB Mezz       |
| SIB F2S 3/0      | REV 05  | 710-022603  | DW7815        | F2S SIB            |
| B Board          | REV 05  | 710-023787  | DW7518        | F2S SIB Mezz       |
| SIB F2S 3/2      | REV 03  | 710-022603  | DV0068        | F2S SIB            |

|             |        |            |        |                |
|-------------|--------|------------|--------|----------------|
| B Board     | REV 03 | 710-023787 | DT9974 | F2S SIB Mezz   |
| SIB F2S 3/4 | REV 05 | 710-022603 | DW7874 | F2S SIB        |
| B Board     | REV 05 | 710-023787 | DW7601 | F2S SIB Mezz   |
| SIB F2S 3/6 | REV 03 | 710-022603 | DV0033 | F2S SIB        |
| B Board     | REV 03 | 710-023787 | DT9969 | F2S SIB Mezz   |
| SIB F2S 4/0 | REV 03 | 710-022603 | DV0043 | F2S SIB        |
| B Board     | REV 03 | 710-023787 | DT9948 | F2S SIB Mezz   |
| SIB F2S 4/2 | REV 05 | 710-022603 | DW5446 | F2S SIB        |
| B Board     | REV 05 | 710-023787 | DW7611 | F2S SIB Mezz   |
| SIB F2S 4/4 | REV 04 | 710-022603 | DW4826 | F2S SIB        |
| B Board     | REV 05 | 710-023787 | DW8458 | F2S SIB Mezz   |
| SIB F2S 4/6 | REV 03 | 710-022603 | DV0026 | F2S SIB        |
| B Board     | REV 03 | 710-023787 | DT9963 | F2S SIB Mezz   |
| Fan Tray 0  | REV 02 | 760-024497 | DR8290 | Front Fan Tray |
| Fan Tray 1  | REV 02 | 760-024497 | DR8293 | Front Fan Tray |
| Fan Tray 2  | REV 05 | 760-024502 | DR8280 | Rear Fan Tray  |
| Fan Tray 3  |        |            |        |                |
| Fan Tray 4  | REV 05 | 760-024502 | DR8276 | Rear Fan Tray  |
| Fan Tray 5  | REV 02 | 760-024502 | DP5643 | Rear Fan Tray  |

lcc0-re0:

-----  
Hardware inventory:

| Item             | Version | Part number | Serial number | Description              |
|------------------|---------|-------------|---------------|--------------------------|
| Chassis          |         |             | JN11036F8AHA  | T1600                    |
| Midplane         | REV 03  | 710-017247  | RC3799        | T-series Backplane       |
| FPM GBUS         | REV 10  | 710-002901  | DP7009        | T640 FPM Board           |
| FPM Display      | REV 01  | 710-021387  | DN7026        | T1600 FPM Display        |
| CIP              | REV 06  | 710-002895  | DP6024        | T-series CIP             |
| PEM 1            | Rev 02  | 740-023211  | WA50019       | Power Entry Module 4x60A |
| SCG 0            | REV 15  | 710-003423  | DR6757        | T640 Sonet Clock Gen.    |
| SCG 1            | REV 15  | 710-003423  | DS2225        | T640 Sonet Clock Gen.    |
| Routing Engine 0 | REV 01  | 740-026941  | 737F-1040     | RE-DUO-1800              |
| Routing Engine 1 | REV 01  | 740-026941  | 737F-1016     | RE-DUO-1800              |
| CB 0             | REV 06  | 710-022597  | DX4011        | LCC Control Board        |
| CB 1             | REV 06  | 710-022597  | DX4017        | LCC Control Board        |
| FPC 1            | REV 07  | 710-013035  | DN5847        | FPC Type 3-ES            |
| CPU              | REV 08  | 710-016744  | DP2570        | ST-PMB2                  |
| PIC 0            | REV 05  | 750-015217  | DB0418        | 8x 1GE(TYPE3), IQ2       |
| Xcvr 0           | REV 01  | 740-011782  | P8Q27ZG       | SFP-SX                   |
| Xcvr 1           |         | NON-JNPR    | PDA1U0D       | SFP-SX                   |
| Xcvr 2           | REV 01  | 740-011613  | P9F1ALW       | SFP-SX                   |
| Xcvr 3           | REV 01  | 740-011782  | PBA403V       | SFP-SX                   |
| Xcvr 4           |         | NON-JNPR    | PDE09DP       | SFP-SX                   |
| Xcvr 5           | REV 01  | 740-011782  | PCH2P4K       | SFP-SX                   |
| Xcvr 6           | REV 01  | 740-011782  | PB94K0F       | SFP-SX                   |
| Xcvr 7           | REV 01  | 740-011782  | PBA2R2A       | SFP-SX                   |
| PIC 1            | REV 03  | 750-004424  | HJ4020        | 1x 10GE(LAN),DWDM        |
| PIC 2            | REV 01  | 750-003336  | HG6073        | 4x OC-48 SONET, SMSR     |
| MMB 0            | REV 04  | 710-016036  | DP3401        | ST-MMB2                  |
| FPC 3            | REV 12  | 710-013037  | DR1169        | FPC Type 4-ES            |
| CPU              | REV 08  | 710-016744  | DP9429        | ST-PMB2                  |
| PIC 0            | REV 02  | 750-010850  | JA0332        | 1x OC-768 SONET SR       |
| MMB 0            | REV 04  | 710-016036  | DR0628        | ST-MMB2                  |
| MMB 1            | REV 04  | 710-016036  | DR0592        | ST-MMB2                  |
| FPC 4            | REV 05  | 710-021534  | DR7350        | FPC Type 1-ES            |
| CPU              | REV 08  | 710-016744  | DP8096        | ST-PMB2                  |
| PIC 0            | REV 04  | 750-014627  | DP9171        | 4x OC-3 1x OC-12 SFP     |
| Xcvr 0           | REV 02  | 740-011615  | PDE2RVR       | SFP-SR                   |
| PIC 1            | REV 22  | 750-005634  | DS5815        | 1x CHOC12 IQ SONET, SMIR |

|            |        |            |             |                         |
|------------|--------|------------|-------------|-------------------------|
| PIC 2      | REV 09 | 750-002911 | CF4539      | 4x F/E, 100 BASE-TX     |
| PIC 3      | REV 08 | 750-021652 | DR2827      | 1x CHOC12 IQE SONET     |
| Xcvr 0     |        | NON-JNPR   | 8           | UNKNOWN                 |
| MMB 0      | REV 04 | 710-016036 | DR0809      | ST-MMB2                 |
| FPC 5      | REV 07 | 710-007529 | HS5608      | FPC Type 3              |
| CPU        | REV 15 | 710-001726 | HX4351      | FPC CPU                 |
| PIC 0      | REV 14 | 750-009567 | WJ8961      | 1x 10GE(LAN), XENPAK    |
| Xcvr 0     | REV 01 | 740-013170 | J05K05961   | XENPAK-LR               |
| PIC 1      | REV 16 | 750-007141 | JJ8146      | 10x 1GE(LAN), 1000 BASE |
| Xcvr 1     | REV 01 | 740-011613 | P9F117T     | SFP-SX                  |
| Xcvr 2     | REV 01 | 740-011782 | PBA2VCL     | SFP-SX                  |
| Xcvr 3     | REV 01 | 740-011782 | PB83DRB     | SFP-SX                  |
| Xcvr 4     | REV 01 | 740-011613 | AM0812S8UP8 | SFP-SX                  |
| PIC 2      | REV 12 | 750-009567 | WF3566      | 1x 10GE(LAN), XENPAK    |
| Xcvr 0     | REV 02 | 740-013170 | T07C94489   | XENPAK-LR               |
| MMB 0      | REV 03 | 710-005555 | HZ1907      | MMB-288mbit             |
| MMB 1      | REV 03 | 710-005555 | HW5283      | MMB-288mbit             |
| PPB 0      | REV 04 | 710-002845 | HZ7717      | PPB Type 3              |
| PPB 1      | REV 04 | 710-002845 | HS0110      | PPB Type 3              |
| FPC 6      | REV 07 | 710-013035 | DP7486      | FPC Type 3-ES           |
| CPU        | REV 08 | 710-016744 | DP2545      | ST-PMB2                 |
| PIC 0      | REV 09 | 750-009567 | NE6323      | 1x 10GE(LAN), XENPAK    |
| Xcvr 0     | REV 02 | 740-013170 | T09C71959   | XENPAK-LR               |
| PIC 1      | REV 06 | 750-015217 | DN4775      | 8x 1GE(TYPE3), IQ2      |
| Xcvr 0     | REV 01 | 740-011782 | P7E0T6M     | SFP-SX                  |
| Xcvr 1     | REV 01 | 740-011613 | AM0812S8XAY | SFP-SX                  |
| Xcvr 2     | REV 01 | 740-011782 | P7E0T6J     | SFP-SX                  |
| Xcvr 3     | REV 01 | 740-011782 | PCH2P7D     | SFP-SX                  |
| Xcvr 4     | REV 01 | 740-011782 | P9B0QYT     | SFP-SX                  |
| Xcvr 5     | REV 01 | 740-011613 | AM0812S8WQJ | SFP-SX                  |
| Xcvr 6     | REV 02 | 740-013111 | 9301220     | SFP-T                   |
| Xcvr 7     | REV 01 | 740-011782 | P9B0TZ5     | SFP-SX                  |
| PIC 2      | REV 06 | 750-015217 | DM6747      | 8x 1GE(TYPE3), IQ2      |
| Xcvr 0     | REV 01 | 740-011613 | PAP0ZB2     | SFP-SX                  |
| Xcvr 1     | REV 01 | 740-013111 | 70191002    | SFP-T                   |
| Xcvr 6     | REV 01 | 740-011782 | PBA29H8     | SFP-SX                  |
| Xcvr 7     | REV 01 | 740-011613 | AM0812S8WQG | SFP-SX                  |
| MMB 0      | REV 04 | 710-016036 | DP3238      | ST-MMB2                 |
| FPC 7      | REV 03 | 710-021540 | DV3154      | FPC Type 2-ES           |
| CPU        | REV 09 | 710-016744 | DT9053      | ST-PMB2                 |
| PIC 0      | REV 13 | 750-001901 | HB4225      | 4x OC-12 SONET, SMIR    |
| PIC 1      | REV 05 | 750-001900 | AD3644      | 1x OC-48 SONET, SMSR    |
| PIC 2      | REV 10 | 750-008155 | HV0335      | 2x G/E IQ, 1000 BASE    |
| Xcvr 0     | REV 01 | 740-011782 | PCH2UKF     | SFP-SX                  |
| Xcvr 1     | REV 01 | 740-011782 | PCH2V19     | SFP-SX                  |
| PIC 3      | REV 03 | 750-014638 | JS9493      | 1x OC-48-12-3 SFP       |
| Xcvr 0     | REV 01 | 740-011785 | P6Q0ENK     | SFP-SR                  |
| MMB 0      | REV 05 | 710-016036 | DP3323      | ST-MMB2                 |
| SPMB 0     | REV 04 | 710-023321 | DX3004      | LCC Switch CPU          |
| SPMB 1     | REV 04 | 710-023321 | DX3009      | LCC Switch CPU          |
| SIB 0      | REV 07 | 710-022594 | DW4195      | LCC SIB                 |
| B Board    | REV 07 | 710-023185 | DW3930      | LCC SIB Mezz            |
| SIB 1      | REV 07 | 710-022594 | DW4179      | LCC SIB                 |
| B Board    | REV 07 | 710-023185 | DW3919      | LCC SIB Mezz            |
| SIB 2      |        |            |             |                         |
| SIB 3      | REV 06 | 710-022594 | DT8251      | LCC SIB                 |
| B Board    | REV 06 | 710-023185 | DT5792      | LCC SIB Mezz            |
| SIB 4      | REV 08 | 710-022594 | DW8014      | LCC SIB                 |
| B Board    | REV 07 | 710-023185 | DW3917      | LCC SIB Mezz            |
| Fan Tray 0 |        |            |             | Front Top Fan Tray      |



Fan Tray 1  
Fan Tray 2

Front Bottom Fan Tray  
Rear Fan Tray -- Rev 3

lcc1-re0:

-----  
Hardware inventory:

| Item             | Version | Part number | Serial number | Description              |
|------------------|---------|-------------|---------------|--------------------------|
| Chassis          |         |             | JN1102270AHA  | T1600                    |
| Midplane         | REV 04  | 710-017247  | RC5358        | T-series Backplane       |
| FPM GBUS         | REV 10  | 710-002901  | DS3443        | T640 FPM Board           |
| FPM Display      | REV 01  | 710-021387  | DS6411        | T1600 FPM Display        |
| CIP              | REV 06  | 710-002895  | DS4235        | T-series CIP             |
| PEM 0            | Rev 02  | 740-023211  | VM82438       | Power Entry Module 4x60A |
| SCG 0            | REV 15  | 710-003423  | DS6649        | T640 Sonet Clock Gen.    |
| SCG 1            | REV 15  | 710-003423  | DR6775        | T640 Sonet Clock Gen.    |
| Routing Engine 0 | REV 01  | 740-026941  | 737F-1083     | RE-DUO-1800              |
| Routing Engine 1 | REV 01  | 740-026941  | 737F-1104     | RE-DUO-1800              |
| CB 0             | REV 06  | 710-022597  | DW8542        | LCC Control Board        |
| CB 1             | REV 06  | 710-014279  | DW8530        | LCC Control Board        |
| FPC 0            | REV 02  | 710-010845  | JE2392        | FPC Type 4               |
| CPU              | REV 02  | 710-011481  | JF6820        | FPC CPU-Enhanced         |
| PIC 0            | REV 11  | 750-017405  | DP7259        | 4x 10GE (LAN/WAN) XFP    |
| Xcvr 0           | REV 01  | 740-014279  | AA0741N1C8T   | XFP-10G-LR               |
| Xcvr 1           | REV 01  | 740-014279  | AA0746N1GAM   | XFP-10G-LR               |
| Xcvr 2           | REV 01  | 740-014279  | AA0747N1H0B   | XFP-10G-LR               |
| Xcvr 3           | REV 01  | 740-014279  | AA0748N1HZ5   | XFP-10G-LR               |
| MMB 0            | REV 03  | 710-010842  | HY7601        | ST-MMB                   |
| FPC 1            | REV 16  | 710-013037  | BBAA7398      | FPC Type 4-ES            |
| CPU              | REV 09  | 710-016744  | BBAA2329      | ST-PMB2                  |
| PIC 0            | REV 03  | 711-029996  | EB1575        | 100GE                    |
| PIC 1            | REV 06  | 750-034781  | EB9980        | 100GE CFP                |
| MMB 0            | REV 04  | 710-025563  | BBAA5325      | ST-MMB2                  |
| MMB 1            | REV 04  | 710-025563  | BBAA5444      | ST-MMB2                  |
| FPC 2            | REV 16  | 710-013037  | BBAA7185      | FPC Type 4-ES            |
| CPU              | REV 09  | 710-016744  | BBAA3522      | ST-PMB2                  |
| PIC 0            | REV 03  | 711-029996  | EB1557        | 100GE                    |
| PIC 1            | REV 05  | 750-034781  | EB4660        | 100GE CFP                |
| Xcvr 0           | REV 0   | 740-032210  | J10F73666     | CFP-100G-LR4             |
| BRIDGE 0         | REV 02  | 711-029995  | EB2237        | 100GE Bridge Board       |
| MMB 0            | REV 04  | 710-025563  | BBAA5347      | ST-MMB2                  |
| MMB 1            | REV 04  | 710-025563  | BBAA5401      | ST-MMB2                  |
| FPC 3            | REV 10  | 710-021534  | DZ0941        | FPC Type 1-ES            |
| CPU              | REV 09  | 710-016744  | DY6364        | ST-PMB2                  |
| PIC 0            | REV 13  | 750-012266  | DK9192        | 4x 1GE(LAN), IQ2         |
| Xcvr 0           | REV 01  | 740-011613  | AM0812S8WVD   | SFP-SX                   |
| Xcvr 1           |         | NON-JNPR    | PDD63Q4       | SFP-SX                   |
| Xcvr 2           |         | NON-JNPR    | PDE4G54       | SFP-SX                   |
| Xcvr 3           |         | NON-JNPR    | PD40MAG       | SFP-SX                   |
| PIC 1            | REV 01  | 750-007641  | HJ2003        | 1x G/E IQ, 1000 BASE     |
| Xcvr 0           | REV 01  | 740-011613  | AM0812S8WVG   | SFP-SX                   |
| PIC 3            | REV 17  | 750-007444  | JB6873        | 1x CHSTM1 IQ SDH, SMIR   |
| MMB 0            | REV 04  | 710-025563  | DZ0281        | ST-MMB2                  |
| FPC 4            | REV 06  | 710-013035  | DK0614        | FPC Type 3-ES            |
| CPU              | REV 07  | 710-016744  | DK1616        | ST-PMB2                  |
| PIC 0            | REV 22  | 750-007141  | DM1870        | 10x 1GE(LAN), 1000 BASE  |
| Xcvr 0           | REV 01  | 740-011782  | PCL3UKW       | SFP-SX                   |
| Xcvr 1           | REV 01  | 740-011782  | P7E0T73       | SFP-SX                   |
| Xcvr 2           | REV 01  | 740-007326  | P4TOWLRL      | SFP-SX                   |
| Xcvr 3           | REV 01  | 740-011782  | PAR1LRL       | SFP-SX                   |
| Xcvr 4           | REV 01  | 740-011782  | P9MOU3Z       | SFP-SX                   |

|            |        |            |              |                         |
|------------|--------|------------|--------------|-------------------------|
| Xcvr 5     | REV 01 | 740-011782 | P9M0U0C      | SFP-SX                  |
| Xcvr 6     | REV 01 | 740-011782 | P9M0TLG      | SFP-SX                  |
| Xcvr 7     | REV 01 | 740-011782 | P9M0U0F      | SFP-SX                  |
| Xcvr 8     | REV 01 | 740-011613 | PFA6LAP      | SFP-SX                  |
| Xcvr 9     | REV 01 | 740-011782 | PCH2P0U      | SFP-SX                  |
| PIC 1      | REV 16 | 750-009450 | CV2565       | 1x OC-192 SM SR2        |
| PIC 2      | REV 05 | 750-004424 | HH3057       | 1x 10GE(LAN),10GBASE-LR |
| PIC 3      | REV 12 | 750-013423 | DP0403       | MultiServices 500       |
| MMB 0      | REV 04 | 710-016036 | DK1988       | ST-MMB2                 |
| FPC 5      | REV 07 | 710-013560 | DR0004       | E2-FPC Type 3           |
| CPU        | REV 05 | 710-013563 | DR0089       | FPC CPU-Enhanced        |
| PIC 0      | REV 11 | 750-012793 | DR6107       | 1x 10GE(LAN/WAN) IQ2    |
| Xcvr 0     | REV 01 | 740-014289 | C743XU074    | XFP-10G-SR              |
| PIC 1      | REV 01 | 750-004695 | HD5980       | 1x Tunnel               |
| PIC 2      | REV 32 | 750-003700 | DL3770       | 1x OC-192 12xMM VSR     |
| PIC 3      | REV 12 | 750-009553 | WB8901       | 4x OC-48 SONET          |
| Xcvr 0     | REV 01 | 740-011785 | P9D1GTQ      | SFP-SR                  |
| Xcvr 1     | REV 01 | 740-011785 | PDSOMMB      | SFP-SR                  |
| Xcvr 3     | REV 01 | 740-011785 | PDE1KXP      | SFP-SR                  |
| MMB 0      | REV 07 | 710-010171 | DP7374       | MMB-5M3-288mbit         |
| MMB 1      | REV 07 | 710-010171 | DP7404       | MMB-5M3-288mbit         |
| FPC 6      | REV 07 | 710-013035 | DM0994       | FPC Type 3-ES           |
| CPU        | REV 07 | 710-016744 | DM3651       | ST-PMB2                 |
| PIC 0      | REV 07 | 750-015217 | DN4743       | 8x 1GE(TYPE3), IQ2      |
| Xcvr 3     | REV 01 | 740-011613 | AM0812S8XB0  | SFP-SX                  |
| Xcvr 4     | REV 01 | 740-011782 | PB829RB      | SFP-SX                  |
| Xcvr 5     | REV 01 | 740-011782 | P8J1SYX      | SFP-SX                  |
| PIC 1      | REV 03 | 750-003336 | HJ9954       | 4x OC-48 SONET, SMSR    |
| PIC 3      | REV 02 | 750-012793 | JM7665       | 1x 10GE(LAN/WAN) IQ2    |
| MMB 0      | REV 04 | 710-016036 | DN6913       | ST-MMB2                 |
| FPC 7      | REV 08 | 710-010845 | JM3958       | FPC Type 4              |
| CPU        | REV 04 | 710-011481 | JK3669       | FPC CPU-Enhanced        |
| PIC 0      | REV 11 | 750-017405 | DP8837       | 4x 10GE (LAN/WAN) XFP   |
| Xcvr 1     | REV 01 | 740-014279 | 753019A00277 | XFP-10G-LR              |
| Xcvr 2     | REV 02 | 740-011571 | C850XJ00P    | XFP-10G-SR              |
| Xcvr 3     | REV 01 | 740-014279 | AA0813N1RTG  | XFP-10G-LR              |
| MMB 0      | REV 04 | 710-010842 | JN1971       | ST-MMB                  |
| SPMB 0     | REV 04 | 710-023321 | DW3629       | LCC Switch CPU          |
| SPMB 1     | REV 04 | 710-023321 | DW3621       | LCC Switch CPU          |
| SIB 0      | REV 07 | 710-022594 | DW4200       | LCC SIB                 |
| B Board    | REV 07 | 710-023185 | DW3932       | LCC SIB Mezz            |
| SIB 1      | REV 07 | 710-022594 | DW4193       | LCC SIB                 |
| B Board    | REV 07 | 710-023185 | DW3904       | LCC SIB Mezz            |
| SIB 2      |        |            |              |                         |
| SIB 3      | REV 07 | 710-022594 | DW4210       | LCC SIB                 |
| B Board    | REV 06 | 710-023185 | DT5780       | LCC SIB Mezz            |
| SIB 4      | REV 08 | 710-022594 | DW8019       | LCC SIB                 |
| B Board    | REV 06 | 710-023185 | DT5795       | LCC SIB Mezz            |
| Fan Tray 0 |        |            |              | Front Top Fan Tray      |
| Fan Tray 1 |        |            |              | Front Bottom Fan Tray   |
| Fan Tray 2 |        |            |              | Rear Fan Tray -- Rev 3  |

### show chassis hardware sfc (TX Matrix Plus Router)

```
user@host> show chassis hardware sfc 0
sfc0-re0:
```

#### Hardware inventory:

| Item    | Version | Part number | Serial number | Description |
|---------|---------|-------------|---------------|-------------|
| Chassis |         |             | JN112F007AHB  | TXP         |

|                  |        |            |           |                    |
|------------------|--------|------------|-----------|--------------------|
| Midplane         | REV 05 | 710-022574 | TS4027    | SFC Midplane       |
| FPM Display      | REV 03 | 710-024027 | DX0282    | TXP FPM Display    |
| CIP 0            | REV 04 | 710-023792 | DW4889    | TXP CIP            |
| CIP 1            | REV 04 | 710-023792 | DW4887    | TXP CIP            |
| PEM 0            | Rev 07 | 740-027463 | UM26368   | Power Entry Module |
| Routing Engine 0 | REV 01 | 740-026942 | 737A-1064 | SFC RE             |
| Routing Engine 1 | REV 01 | 740-026942 | 737A-1082 | SFC RE             |
| CB 0             | REV 09 | 710-022606 | DW6099    | SFC Control Board  |
| CB 1             | REV 09 | 710-022606 | DW6096    | SFC Control Board  |
| SPMB 0           |        | BUILTIN    |           | SFC Switch CPU     |
| SPMB 1           |        | BUILTIN    |           | SFC Switch CPU     |
| SIB F13 0        | REV 04 | 710-022600 | DX0841    | F13 SIB            |
| B Board          | REV 03 | 710-023431 | DX0966    | F13 SIB Mezz       |
| SIB F13 1        | REV 04 | 750-024564 | DW5776    | F13 SIB            |
| B Board          | REV 03 | 710-023431 | DW9028    | F13 SIB            |
| SIB F13 3        | REV 04 | 750-024564 | DW5762    | F13 SIB            |
| B Board          | REV 03 | 710-023431 | DW9059    | F13 SIB            |
| SIB F13 4        | REV 04 | 750-024564 | DW5797    | F13 SIB            |
| B Board          | REV 03 | 710-023431 | DW9041    | F13 SIB            |
| SIB F13 6        | REV 04 | 750-024564 | DW5770    | F13 SIB            |
| B Board          | REV 03 | 710-023431 | DW9079    | F13 SIB Mezz       |
| SIB F13 7        | REV 04 | 750-024564 | DW5758    | F13 SIB            |
| B Board          | REV 03 | 710-023431 | DW9047    | F13 SIB            |
| SIB F13 8        | REV 04 | 750-024564 | DW5761    | F13 SIB            |
| B Board          | REV 03 | 710-023431 | DW9043    | F13 SIB Mezz       |
| SIB F13 9        | REV 04 | 750-024564 | DW5754    | F13 SIB            |
| B Board          | REV 03 | 710-023431 | DW9078    | F13 SIB Mezz       |
| SIB F13 11       | REV 04 | 710-022600 | DX0826    | F13 SIB            |
| B Board          | REV 03 | 710-023431 | DX0967    | F13 SIB Mezz       |
| SIB F13 12       | REV 04 | 750-024564 | DW5794    | F13 SIB            |
| B Board          | REV 03 | 710-023431 | DW9044    | F13 SIB Mezz       |
| SIB F2S 0/0      | REV 05 | 710-022603 | DW7897    | F2S SIB            |
| B Board          | REV 05 | 710-023787 | DW7657    | NEO PMB            |
| SIB F2S 0/2      | REV 05 | 710-022603 | DW7833    | F2S SIB            |
| B Board          | REV 05 | 710-023787 | DW7526    | NEO PMB            |
| SIB F2S 0/4      | REV 05 | 710-022603 | DW7875    | F2S SIB            |
| B Board          | REV 05 | 710-023787 | DW7588    | NEO PMB            |
| SIB F2S 0/6      | REV 05 | 710-022603 | DW7860    | F2S SIB            |
| B Board          | REV 05 | 710-023787 | DW7589    | NEO PMB            |
| SIB F2S 1/0      | REV 04 | 710-022603 | DW4820    | F2S SIB            |
| B Board          | REV 05 | 710-023787 | DW8510    | NEO PMB            |
| SIB F2S 1/2      | REV 05 | 710-022603 | DW7849    | F2S SIB            |
| B Board          | REV 05 | 710-023787 | DW7525    | NEO PMB            |
| SIB F2S 1/4      | REV 05 | 710-022603 | DW7927    | F2S SIB            |
| B Board          | REV 05 | 710-023787 | DW7556    | F2S SIB Mezz       |
| SIB F2S 1/6      | REV 05 | 710-022603 | DW7866    | F2S SIB            |
| B Board          | REV 05 | 710-023787 | DW7651    | NEO PMB            |
| SIB F2S 2/0      | REV 05 | 710-022603 | DW7880    | F2S SIB            |
| B Board          | REV 05 | 710-023787 | DW7523    | NEO PMB            |
| SIB F2S 2/2      | REV 05 | 710-022603 | DW7895    | F2S SIB            |
| B Board          | REV 05 | 710-023787 | DW7591    | NEO PMB            |
| SIB F2S 2/4      | REV 05 | 710-022603 | DW7907    | F2S SIB            |
| B Board          | REV 05 | 710-023787 | DW7590    | NEO PMB            |
| SIB F2S 2/6      | REV 05 | 710-022603 | DW7785    | F2S SIB            |
| B Board          | REV 05 | 710-023787 | DW7524    | NEO PMB            |
| SIB F2S 3/0      | REV 05 | 710-022603 | DW7782    | F2S SIB            |
| B Board          | REV 05 | 710-023787 | DW7634    | NEO PMB            |
| SIB F2S 3/2      | REV 05 | 710-022603 | DW7793    | F2S SIB            |
| B Board          | REV 05 | 710-023787 | DW7548    | NEO PMB            |
| SIB F2S 3/4      | REV 05 | 710-022603 | DW7779    | F2S SIB            |
| B Board          | REV 05 | 710-023787 | DW7587    | NEO PMB            |

|             |        |            |        |                |
|-------------|--------|------------|--------|----------------|
| SIB F2S 3/6 | REV 05 | 710-022603 | DW7930 | F2S SIB        |
| B Board     | REV 05 | 710-023787 | DW7505 | NEO PMB        |
| SIB F2S 4/0 | REV 05 | 710-022603 | DW7867 | F2S SIB        |
| B Board     | REV 05 | 710-023787 | DW7656 | NEO PMB        |
| SIB F2S 4/2 | REV 05 | 710-022603 | DW7917 | F2S SIB        |
| B Board     | REV 05 | 710-023787 | DW7640 | NEO PMB        |
| SIB F2S 4/4 | REV 05 | 710-022603 | DW7929 | F2S SIB        |
| B Board     | REV 05 | 710-023787 | DW7643 | NEO PMB        |
| SIB F2S 4/6 | REV 05 | 710-022603 | DW7870 | F2S SIB        |
| B Board     | REV 05 | 710-023787 | DW7635 | NEO PMB        |
| Fan Tray 0  | REV 06 | 760-024497 | DV7831 | Front Fan Tray |
| Fan Tray 1  | REV 06 | 760-024497 | DV9614 | Front Fan Tray |
| Fan Tray 2  | REV 06 | 760-024502 | DV9618 | Rear Fan Tray  |
| Fan Tray 3  | REV 06 | 760-024502 | DV9616 | Rear Fan Tray  |
| Fan Tray 4  | REV 06 | 760-024502 | DV7807 | Rear Fan Tray  |
| Fan Tray 5  | REV 06 | 760-024502 | DV7828 | Rear Fan Tray  |

### show chassis hardware extensive (TX Matrix Plus Router)

```
user@host> show chassis hardware extensive
```

```
sfc0-re0:
```

#### Hardware inventory:

| Item    | Version | Part number | Serial number | Description |
|---------|---------|-------------|---------------|-------------|
| Chassis |         |             | JN112F007AHB  | TXP         |

|              |            |                   |              |
|--------------|------------|-------------------|--------------|
| JeDEC Code:  | 0x7fb0     | EEPROM Version:   | 0x02         |
|              |            | S/N:              | JN112F007AHB |
| Assembly ID: | 0x052c     | Assembly Version: | 00.00        |
| Date:        | 00-00-0000 | Assembly Flags:   | 0x00         |
| ID:          | TXP        |                   |              |

#### Board Information Record:

```
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```

#### I2C Hex Data:

```
Address 0x00: 7f b0 02 ff 05 2c 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: 4a 4e 31 31 32 46 30 30 37 41 48 42 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```

|          |        |            |        |              |
|----------|--------|------------|--------|--------------|
| Midplane | REV 05 | 710-022574 | TS4027 | SFC Midplane |
|----------|--------|------------|--------|--------------|

|              |            |                   |            |
|--------------|------------|-------------------|------------|
| JeDEC Code:  | 0x7fb0     | EEPROM Version:   | 0x01       |
| P/N:         | 710-022574 | S/N:              | S/N TS4027 |
| Assembly ID: | 0x0962     | Assembly Version: | 01.05      |
| Date:        | 03-23-2009 | Assembly Flags:   | 0x00       |
| Version:     | REV 05     |                   |            |

#### ID: SFC Midplane

#### Board Information Record:

```
Address 0x00: ad 01 ff ff 00 1d b5 14 00 00 ff ff ff ff ff ff
```

#### I2C Hex Data:

```
Address 0x00: 7f b0 01 ff 09 62 01 05 52 45 56 20 30 35 00 00
Address 0x10: 00 00 00 00 37 31 30 2d 30 32 32 35 37 34 00 00
Address 0x20: 53 2f 4e 20 54 53 34 30 32 37 00 00 00 17 03 07
Address 0x30: d9 ff ff ff ad 01 ff ff 00 1d b5 14 00 00 ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
```

|             |        |                 |        |                 |
|-------------|--------|-----------------|--------|-----------------|
| FPM Display | REV 03 | 710-024027      | DX0282 | TXP FPM Display |
| JeDEC Code: | 0x7fb0 | EEPROM Version: | 0x01   |                 |

```

P/N: 710-024027 S/N: S/N DX0282
Assembly ID: 0x096c Assembly Version: 01.03
Date: 02-10-2009 Assembly Flags: 0x00
Version: REV 03
ID: TXP FPM Display FRU Model Number: CRAFT-TXP
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 09 6c 01 03 52 45 56 20 30 33 00 00
Address 0x10: 00 00 00 00 37 31 30 2d 30 32 34 30 32 37 00 00
Address 0x20: 53 2f 4e 20 44 58 30 32 38 32 00 00 00 0a 02 0f
Address 0x30: d9 ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff ff 01 00 00 00 00 00 00 00 00 00 43
Address 0x50: 52 41 46 54 2d 54 58 50 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
CIP 0 REV 04 710-023792 DW4889 TXP CIP
Jedec Code: 0x7fb0 EEPROM Version: 0x01
P/N: 710-023792 S/N: S/N DW4889
Assembly ID: 0x0969 Assembly Version: 01.04
Date: 01-26-2009 Assembly Flags: 0x00
Version: REV 04
ID: TXP CIP FRU Model Number: CIP-TXP
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

```

#### show chassis hardware clei-models (TX Matrix Plus Router)

```

user@host> show chassis hardware clei-models
sfc0-re0:

Hardware inventory:

```

| Item             | Version | Part number | CLEI code  | FRU model number   |
|------------------|---------|-------------|------------|--------------------|
| Midplane         | REV 05  | 710-022574  |            | CHAS-BP-TXP-S      |
| FPM Display      | REV 03  | 710-024027  |            | CRAFT-TXP-S        |
| CIP 0            | REV 05  | 710-023792  |            | CIP-TXP-S          |
| CIP 1            | REV 05  | 710-023792  |            | CIP-TXP-S          |
| PEM 0            | Rev 04  | 740-027463  | IPUPAFGKTA | PWR-TXP-7-60-DC    |
| PEM 1            | Rev 04  | 740-027463  | IPUPAFGKTA | PWR-TXP-7-60-DC    |
| Routing Engine 0 | REV 06  | 740-026942  |            | RE-DUO-C2600-16G-S |
| Routing Engine 1 | REV 06  | 740-026942  |            | RE-DUO-C2600-16G-S |
| CB 0             | REV 05  | 710-022606  |            | CB-TXP-S           |
| CB 1             | REV 09  | 710-022606  |            | CB-TXP-S           |
| SIB F13 0        | REV 04  | 750-024564  |            | SIB-TXP-F13        |
| SIB F13 3        | REV 04  | 750-024564  |            | SIB-TXP-F13        |
| SIB F13 8        | REV 04  | 750-024564  |            | SIB-TXP-F13        |
| SIB F13 11       | REV 04  | 750-024564  |            | SIB-TXP-F13        |
| SIB F13 12       | REV 03  | 750-024564  |            | SIB-TXP-F13        |
| SIB F2S 0/0      | REV 05  | 710-022603  |            | SIB-TXP-F2S-S      |
| SIB F2S 0/2      | REV 05  | 710-022603  |            | SIB-TXP-F2S-S      |
| SIB F2S 0/4      | REV 04  | 710-022603  |            | SIB-TXP-F2S-S      |
| SIB F2S 0/6      | REV 04  | 710-022603  |            | SIB-TXP-F2S-S      |
| SIB F2S 1/0      | REV 04  | 710-022603  |            | SIB-TXP-F2S-S      |
| SIB F2S 1/2      | REV 05  | 710-022603  |            | SIB-TXP-F2S-S      |
| SIB F2S 1/4      | REV 04  | 710-022603  |            | SIB-TXP-F2S-S      |
| SIB F2S 1/6      | REV 05  | 710-022603  |            | SIB-TXP-F2S-S      |
| SIB F2S 2/0      | REV 04  | 710-022603  |            | SIB-TXP-F2S-S      |
| SIB F2S 2/2      | REV 04  | 710-022603  |            | SIB-TXP-F2S-S      |
| SIB F2S 2/4      | REV 05  | 710-022603  |            | SIB-TXP-F2S-S      |
| SIB F2S 2/6      | REV 04  | 710-022603  |            | SIB-TXP-F2S-S      |
| SIB F2S 3/0      | REV 05  | 710-022603  |            | SIB-TXP-F2S-S      |

|             |        |            |                 |
|-------------|--------|------------|-----------------|
| SIB F2S 3/2 | REV 03 | 710-022603 | SIB-TXP-F2S-S   |
| SIB F2S 3/4 | REV 05 | 710-022603 | SIB-TXP-F2S-S   |
| SIB F2S 3/6 | REV 03 | 710-022603 | SIB-TXP-F2S-S   |
| SIB F2S 4/0 | REV 03 | 710-022603 | SIB-TXP-F2S-S   |
| SIB F2S 4/2 | REV 05 | 710-022603 | SIB-TXP-F2S-S   |
| SIB F2S 4/4 | REV 04 | 710-022603 | SIB-TXP-F2S-S   |
| SIB F2S 4/6 | REV 03 | 710-022603 | SIB-TXP-F2S-S   |
| Fan Tray 0  | REV 02 | 760-024497 | FANTRAY-TXP-H-S |
| Fan Tray 1  | REV 02 | 760-024497 | FANTRAY-TXP-H-S |
| Fan Tray 2  | REV 05 | 760-024502 | FANTRAY-TXP-V-S |
| Fan Tray 3  |        |            |                 |
| Fan Tray 4  | REV 05 | 760-024502 | FANTRAY-TXP-V-S |
| Fan Tray 5  | REV 02 | 760-024502 | FANTRAY-TXP-V-S |

```
lcc0-re0:
```

```

Hardware inventory:
```

| Item             | Version | Part number | CLEI code  | FRU model number        |
|------------------|---------|-------------|------------|-------------------------|
| Midplane         | REV 03  | 710-017247  |            | CHAS-BP-T1600-S         |
| FPM Display      | REV 01  | 710-021387  |            | CRAFT-T1600-S           |
| CIP              | REV 06  | 710-002895  |            | CIP-L-T640-S            |
| PEM 1            | Rev 02  | 740-023211  | IPUPAC8KTA | PWR-T1600-4-60-DC-S     |
| SCG 0            | REV 15  | 710-003423  |            | SCG-T-S                 |
| SCG 1            | REV 15  | 710-003423  |            | SCG-T-S                 |
| Routing Engine 0 | REV 01  | 740-026941  |            | RE-DUO-C1800-8G-S       |
| Routing Engine 1 | REV 01  | 740-026941  |            | RE-DUO-C1800-8G-S       |
| CB 0             | REV 06  | 710-022597  |            | CB-LCC-S                |
| CB 1             | REV 06  | 710-022597  |            | CB-LCC-S                |
| FPC 1            | REV 07  | 710-013035  |            | T640-FPC3-ES            |
| PIC 0            | REV 05  | 750-015217  |            | PC-8GE-TYPE3-SFP-IQ2    |
| PIC 1            | REV 03  | 750-004424  |            | PC-1XGE-LR              |
| PIC 2            | REV 01  | 750-003336  |            | PC-40C48-SON-SMSR       |
| FPC 3            | REV 12  | 710-013037  |            | T1600-FPC4-ES           |
| PIC 0            | REV 02  | 750-010850  |            | PD-10C768-SON-SR        |
| FPC 4            | REV 05  | 710-021534  |            | T640-FPC1-ES            |
| PIC 0            | REV 04  | 750-014627  |            | PB-40C3-10C12-SON-SFP   |
| PIC 1            | REV 22  | 750-005634  |            | PB-1CHOC12SMIR-QPP      |
| PIC 2            | REV 09  | 750-002911  |            | PB-4FE-TX               |
| PIC 3            | REV 08  | 750-021652  |            | PB-1CHOC12-STM4-IQE-SFP |
| FPC 5            | REV 07  | 710-007529  |            | T640-FPC3               |
| PIC 0            | REV 14  | 750-009567  |            | PC-1XGE-XENPAK          |
| PIC 1            | REV 16  | 750-007141  |            | PC-10GE-SFP             |
| PIC 2            | REV 12  | 750-009567  |            | PC-1XGE-XENPAK          |
| FPC 6            | REV 07  | 710-013035  |            | T640-FPC3-ES            |
| PIC 0            | REV 09  | 750-009567  |            | PC-1XGE-XENPAK          |
| PIC 1            | REV 06  | 750-015217  |            | PC-8GE-TYPE3-SFP-IQ2    |
| PIC 2            | REV 06  | 750-015217  |            | PC-8GE-TYPE3-SFP-IQ2    |
| FPC 7            | REV 03  | 710-021540  |            | T640-FPC2-ES            |
| PIC 0            | REV 13  | 750-001901  |            | PB-40C12-SON-SMIR       |
| PIC 1            | REV 05  | 750-001900  |            | PB-10C48-SON-SMSR       |
| PIC 2            | REV 10  | 750-008155  |            | PB-2GE-SFP-QPP          |
| PIC 3            | REV 03  | 750-014638  |            | PB-10C48-SON-B-SFP      |
| SIB 0            | REV 07  | 710-022594  |            | SIB-TXP-T1600-S         |
| SIB 1            | REV 07  | 710-022594  |            | SIB-TXP-T1600-S         |
| SIB 3            | REV 06  | 710-022594  |            | SIB-TXP-T1600-S         |
| SIB 4            | REV 08  | 710-022594  |            | SIB-TXP-T1600-S         |
| Fan Tray 0       |         |             |            | FANTRAY-T-S             |
| Fan Tray 1       |         |             |            | FANTRAY-T-S             |
| Fan Tray 2       |         |             |            | FANTRAY-TXP-R-S         |

```
lcc1-re0:
```

```

Hardware inventory:
Item Version Part number CLEI code FRU model number
Midplane REV 04 710-017247
FPM Display REV 01 710-021387
CIP REV 06 710-002895
PEM 0 Rev 02 740-023211 IPUPAC8KTA PWR-T1600-4-60-DC-S
SCG 0 REV 15 710-003423 SCG-T-S
SCG 1 REV 15 710-003423 SCG-T-S
Routing Engine 0 REV 01 740-026941 RE-DUO-C1800-8G-S
Routing Engine 1 REV 01 740-026941 RE-DUO-C1800-8G-S
CB 0 REV 06 710-022597 CB-LCC-S
CB 1 REV 06 710-022597 CB-LCC-S
FPC 0 REV 02 710-010845 T640-FPC4-ES
 PIC 0 REV 11 750-017405 PD-4XGE-XFP
FPC 1 REV 16 710-013037 T1600-FPC4-ES
 PIC 1 REV 06 750-034781 PD-1CE-CFP
FPC 2 REV 16 710-013037 T1600-FPC4-ES
 PIC 1 REV 05 750-034781 PD-1CE-CFP
FPC 3 REV 10 710-021534 T640-FPC1-ES
 PIC 0 REV 13 750-012266 PB-4GE-TYPE1-SFP-IQ2
 PIC 1 REV 01 750-007641 PE-1GE-SFP-QPP
 PIC 3 REV 17 750-007444 PB-1CHSTM1-SMIR-QPP
FPC 4 REV 06 710-013035 T640-FPC3-ES
 PIC 0 REV 22 750-007141 PC-10GE-SFP
 PIC 1 REV 16 750-009450 PC-10C192-SON-SR2
 PIC 2 REV 05 750-004424 PC-1XGE-LR
 PIC 3 REV 12 750-013423 PC-MS-500-3
FPC 5 REV 07 710-013560 T640-FPC3-E2
 PIC 0 REV 11 750-012793 PC-1XGE-TYPE3-XFP-IQ2
 PIC 1 REV 01 750-004695 PC-TUNNEL
 PIC 2 REV 32 750-003700 PC-10C192-SON-VSR
 PIC 3 REV 12 750-009553 PC-40C48-SON-SFP
FPC 6 REV 07 710-013035 T640-FPC3-ES
 PIC 0 REV 07 750-015217 PC-8GE-TYPE3-SFP-IQ2
 PIC 1 REV 03 750-003336 PC-40C48-SON-SMSR
 PIC 3 REV 02 750-012793 PC-1XGE-TYPE3-XFP-IQ2
FPC 7 REV 08 710-010845 T640-FPC4-ES
 PIC 0 REV 11 750-017405 PD-4XGE-XFP
SIB 0 REV 07 710-022594 SIB-TXP-T1600-S
SIB 1 REV 07 710-022594 SIB-TXP-T1600-S
SIB 3 REV 07 710-022594 SIB-TXP-T1600-S
SIB 4 REV 08 710-022594 SIB-TXP-T1600-S
Fan Tray 0 FANTRAY-T-S
Fan Tray 1 FANTRAY-T-S
Fan Tray 2 FANTRAY-TXP-R-S

```

#### show chassis hardware detail (TX Matrix Plus Router)

```

user@host> show chassis hardware detail
sfc0-re0:

```

```

Hardware inventory:
Item Version Part number Serial number Description
Chassis
Midplane REV 01 710-022574 JN111B023AHB TXP
FPM Display REV 03 710-024027 DW4699 SFC Midplane
CIP 0 REV 01 710-023792 DR1437 TXP FPM Display
CIP 1 REV 02 710-023792 DS4564 TXP CIP
PEM 0 Rev 07 740-027463 UM26360 TXP CIP
Routing Engine 0 REV 01 740-026942 737A-1024 Power Entry Module
SFC RE

```

```

ad0 3887 MB SMART CF 200811050193CEB1CEB1 Compact Flash
ad1 30533 MB SAMSUNG MCBQE32G8MPP-0V SY814A0762 Disk 1
Routing Engine 1 REV 01 740-026942 737A-1024 SFC RE
ad0 3887 MB SMART CF 20081105004C19A019A0 Compact Flash
ad1 30533 MB SAMSUNG MCBQE32G8MPP-0V SY814A0794 Disk 1
CB 0 REV 03 710-022606 DR7134 SFC Control Board
CB 1 REV 01 710-022606 DP8890 SFC Control Board
SPMB 0 BUILTIN SFC Switch CPU
SPMB 1 BUILTIN SFC Switch CPU
SIB F13 0 REV 03 750-024564 DT9478 F13 SIB
 B Board REV 02 710-023431 DT6554 F13 SIB
SIB F13 1 REV 03 750-024564 DT9454 F13 SIB
 B Board REV 02 710-023431 DT6551 F13 SIB
SIB F2S 0/0 REV 02 710-022603 DT2838 F2S SIB
 B Board REV 02 710-023787 DT1725 NEO PMB
SIB F2S 0/2 REV 02 710-022603 DT2824 F2S SIB
 B Board REV 02 710-023787 DT1706 NEO PMB
SIB F2S 0/4 REV 02 710-022603 DT2822 F2S SIB
 B Board REV 02 710-023787 DT1696 NEO PMB
SIB F2S 0/6 REV 02 710-022603 DT2823 F2S SIB
 B Board REV 02 710-023787 DT1717 NEO PMB
SIB F2S 1/0 REV 03 710-022603 DV0059 F2S SIB
 B Board REV 03 710-023787 DT9942 NEO PMB
SIB F2S 1/2 REV 02 710-022603 DT2826 F2S SIB
 B Board REV 02 710-023787 DT1713 NEO PMB
SIB F2S 1/4 REV 03 710-022603 DV0092 F2S SIB
 B Board REV 03 710-023787 DV0000 NEO PMB
SIB F2S 1/6 REV 03 710-022603 DV0079 F2S SIB
 B Board REV 03 710-023787 DT9972 NEO PMB
SIB F2S 2/0 REV 03 710-022603 DV0100 F2S SIB
 B Board REV 03 710-023787 DT9925 NEO PMB
SIB F2S 2/2 REV 03 710-022603 DV0050 F2S SIB
 B Board REV 03 710-023787 DV0005 NEO PMB
SIB F2S 2/4 REV 03 710-022603 DV0097 F2S SIB
 B Board REV 03 710-023787 DT9936 NEO PMB
Fan Tray 0 REV 02 760-024497 DR8286 Front Fan Tray
Fan Tray 1 REV 06 760-024497 DV9624 Front Fan Tray
Fan Tray 2 REV 02 760-024502 DR8259 Rear Fan Tray
Fan Tray 3 REV 02 760-024502 DR8270 Rear Fan Tray
Fan Tray 4 REV 02 760-024502 DR8284 Rear Fan Tray
Fan Tray 5 REV 06 760-024502 DV7813 Rear Fan Tray

```

lcc0-re0:

-----  
Hardware inventory:

| Item             | Version                                     | Part number | Serial number        | Description             |
|------------------|---------------------------------------------|-------------|----------------------|-------------------------|
| Chassis          |                                             |             | JN1101F27AHA         | T1600                   |
| Midplane         | REV 04                                      | 710-017247  | RC5317               | T Series Backplane      |
| FPM GBUS         | REV 10                                      | 710-002901  | DS8197               | T640 FPM Board          |
| FPM Display      | REV 01                                      | 710-021387  | DS6433               | T1600 FPM Display       |
| CIP              | REV 06                                      | 710-002895  | DS1493               | T Series CIP            |
| PEM 0            | Rev 08                                      | 740-017906  | UD26601              | Power Entry Module 3x80 |
| SCG 0            | REV 15                                      | 710-003423  | DP5847               | T640 Sonet Clock Gen.   |
| SCG 1            | REV 15                                      | 710-003423  | DR0924               | T640 Sonet Clock Gen.   |
| Routing Engine 0 | REV 01                                      | 740-026942  | 737F-1024            | LCC RE                  |
| ad0              | 3887 MB SMART CF                            |             | 2008110502B63E513E51 | Compact Flash           |
| ad1              | 30533 MB SAMSUNG MCBQE32G8MPP-0V SY814A1208 |             |                      | Disk 1                  |
| Routing Engine 1 | REV 01                                      | 740-026942  | 737F-1024            | LCC RE                  |
| ad0              | 3887 MB SMART CF                            |             | 2008110500F9A8A8A8A8 | Compact Flash           |
| ad1              | 30533 MB SAMSUNG MCBQE32G8MPP-0V SY814A1076 |             |                      | Disk 1                  |
| CB 0             | REV 05                                      | 710-022597  | DV4264               | LCC Control Board       |



|            |        |            |        |                       |
|------------|--------|------------|--------|-----------------------|
| CB 1       | REV 03 | 710-022597 | DP8558 | LCC Control Board     |
| FPC 0      | REV 14 | 710-013037 | DS9967 | FPC Type 4-ES         |
| CPU        | REV 08 | 710-016744 | DS3989 | ST-PMB2               |
| PIC 0      | REV 12 | 750-013198 | DL7506 | 1x Tunnel             |
| PIC 1      | REV 12 | 750-013198 | DL7505 | 1x Tunnel             |
| MMB 0      | REV 01 | 710-025563 | DS8524 | ST-MMB2               |
| MMB 1      | REV 01 | 710-025563 | DS8373 | ST-MMB2               |
| FPC 1      | REV 14 | 710-013037 | DT0027 | FPC Type 4-ES         |
| CPU        | REV 09 | 710-016744 | DS7684 | ST-PMB2               |
| PIC 0      | REV 12 | 750-013198 | DL7512 | 1x Tunnel             |
| PIC 1      | REV 12 | 750-013198 | DL7498 | 1x Tunnel             |
| MMB 0      | REV 01 | 710-025563 | DS8494 | ST-MMB2               |
| MMB 1      | REV 01 | 710-025563 | DS8436 | ST-MMB2               |
| SPMB 0     | REV 04 | 710-023321 | DV3867 | LCC Switch CPU        |
| SPMB 1     | REV 02 | 710-023321 | DP0238 | LCC Switch CPU        |
| SIB 0      | REV 06 | 710-022594 | DT8268 | LCC SIB               |
| B Board    | REV 06 | 710-023185 | DT5791 | LCC SIB Mezz          |
| SIB 1      | REV 06 | 710-022594 | DT8261 | LCC SIB               |
| B Board    | REV 06 | 710-023185 | DT5769 | LCC SIB Mezz          |
| SIB 2      | REV 04 | 710-022594 | DS2315 | LCC SIB               |
| B Board    | REV 06 | 710-023185 | DT5788 | LCC SIB Mezz          |
| SIB 3      | REV 06 | 710-022594 | DT8253 | LCC SIB               |
| B Board    | REV 06 | 710-023185 | DT5811 | LCC SIB Mezz          |
| SIB 4      | REV 06 | 710-022594 | DT8248 | LCC SIB               |
| B Board    | REV 06 | 710-023185 | DT5812 | LCC SIB Mezz          |
| Fan Tray 0 |        |            |        | Front Top Fan Tray    |
| Fan Tray 1 |        |            |        | Front Bottom Fan Tray |
| Fan Tray 2 |        |            |        | Rear Fan Tray         |

### show chassis hardware models (TX Matrix Plus Router)

```
user@host> show chassis hardware models
sfc0-re0:
```

```

```

| Hardware inventory: |         |             |               |                                |
|---------------------|---------|-------------|---------------|--------------------------------|
| Item                | Version | Part number | Serial number | FRU model number               |
| FPM Display         | REV 03  | 710-024027  | DX0282        | CRAFT-TXP                      |
| CIP 0               | REV 04  | 710-023792  | DW4889        | CIP-TXP                        |
| CIP 1               | REV 04  | 710-023792  | DW4887        | CIP-TXP                        |
| PEM 0               | Rev 07  | 740-027463  | UM26368       | yyyyyyyyyyyyyyyyyyyyyyyyyyyyyy |
| Routing Engine 0    | REV 01  | 740-026942  | 737A-1064     | RE-TXP-SFC-DU0-2600-16G        |
| Routing Engine 1    | REV 01  | 740-026942  | 737A-1082     | RE-TXP-SFC-DU0-2600-16G        |
| CB 0                | REV 09  | 710-022606  | DW6099        | CB-TXP                         |
| CB 1                | REV 09  | 710-022606  | DW6096        | CB-TXP                         |
| SIB F13 1           | REV 04  | 750-024564  | DW5776        | SIB-TXP-F13                    |
| SIB F13 3           | REV 04  | 750-024564  | DW5762        | SIB-TXP-F13                    |
| SIB F13 4           | REV 04  | 750-024564  | DW5797        | SIB-TXP-F13                    |
| SIB F13 6           | REV 04  | 750-024564  | DW5770        | SIB-TXP-F13                    |
| SIB F13 7           | REV 04  | 750-024564  | DW5758        | SIB-TXP-F13                    |
| SIB F13 8           | REV 04  | 750-024564  | DW5761        | SIB-TXP-F13                    |
| SIB F13 9           | REV 04  | 750-024564  | DW5754        | SIB-TXP-F13                    |
| SIB F13 12          | REV 04  | 750-024564  | DW5794        | SIB-TXP-F13                    |
| SIB F2S 0/0         | REV 05  | 710-022603  | DW7897        |                                |
| SIB F2S 0/2         | REV 05  | 710-022603  | DW7833        |                                |
| SIB F2S 0/4         | REV 05  | 710-022603  | DW7875        |                                |
| SIB F2S 0/6         | REV 05  | 710-022603  | DW7860        |                                |
| SIB F2S 1/0         | REV 04  | 710-022603  | DW4820        |                                |
| SIB F2S 1/2         | REV 05  | 710-022603  | DW7849        |                                |
| SIB F2S 1/4         | REV 05  | 710-022603  | DW7927        | SIB-TXP-F2S                    |
| SIB F2S 1/6         | REV 05  | 710-022603  | DW7866        |                                |
| SIB F2S 2/0         | REV 05  | 710-022603  | DW7880        |                                |

|             |        |            |        |               |
|-------------|--------|------------|--------|---------------|
| SIB F2S 2/2 | REV 05 | 710-022603 | DW7895 |               |
| SIB F2S 2/4 | REV 05 | 710-022603 | DW7907 |               |
| SIB F2S 2/6 | REV 05 | 710-022603 | DW7785 |               |
| SIB F2S 3/0 | REV 05 | 710-022603 | DW7782 |               |
| SIB F2S 3/2 | REV 05 | 710-022603 | DW7793 |               |
| SIB F2S 3/4 | REV 05 | 710-022603 | DW7779 |               |
| SIB F2S 3/6 | REV 05 | 710-022603 | DW7930 |               |
| SIB F2S 4/0 | REV 05 | 710-022603 | DW7867 |               |
| SIB F2S 4/2 | REV 05 | 710-022603 | DW7917 |               |
| SIB F2S 4/4 | REV 05 | 710-022603 | DW7929 |               |
| SIB F2S 4/6 | REV 05 | 710-022603 | DW7870 |               |
| Fan Tray 0  | REV 06 | 760-024497 | DV7831 | FANTRAY-TXP-F |
| Fan Tray 1  | REV 06 | 760-024497 | DV9614 | FANTRAY-TXP-F |
| Fan Tray 2  | REV 06 | 760-024502 | DV9618 | FANTRAY-TXP-R |
| Fan Tray 3  | REV 06 | 760-024502 | DV9616 | FANTRAY-TXP-R |
| Fan Tray 4  | REV 06 | 760-024502 | DV7807 | FANTRAY-TXP-R |
| Fan Tray 5  | REV 06 | 760-024502 | DV7828 | FANTRAY-TXP-R |

lcc0-re0:

-----  
Hardware inventory:

| Item        | Version | Part number | Serial number | FRU model number    |
|-------------|---------|-------------|---------------|---------------------|
| Midplane    | REV 03  | 710-017247  | RC3765        | CHAS-BP-T1600-S     |
| FPM Display | REV 01  | 710-021387  | DN5441        | CRAFT-T1600-S       |
| CIP         | REV 06  | 710-002895  | DP6021        | CIP-L-T640-S        |
| PEM 0       | Rev 07  | 740-017906  | UA26384       | PWR-T1600-3-80-DC-S |
| PEM 1       | Rev 07  | 740-017906  | UA26296       | PWR-T1600-3-80-DC-S |
| SCG 0       | REV 15  | 710-003423  | DR0875        | SCG-T-S             |
| CB 0        | REV 06  | 710-022597  | DW8534        | CB-LCC              |
| CB 1        | REV 06  | 710-022597  | DW8527        | CB-LCC              |
| FPC 4       | REV 12  | 710-013037  | DJ8717        | T1600-FPC4-ES       |
| PIC 0       | REV 11  | 750-017405  | DP8795        | PD-4XGE-XFP         |
| PIC 1       | REV 11  | 750-017405  | DP8794        | PD-4XGE-XFP         |
| FPC 6       | REV 14  | 710-013037  | DS5335        | T1600-FPC4-ES       |
| PIC 0       | REV 13  | 750-017405  | DS7634        | PD-4XGE-XFP         |
| PIC 1       | REV 13  | 750-017405  | DS7637        | PD-4XGE-XFP         |
| FPC 7       | REV 07  | 710-013035  | DM0990        | T1600-FPC3-ES       |
| PIC 0       | REV 16  | 750-007141  | JJ8067        | PC-10GE-SFP         |
| PIC 1       | REV 08  | 750-015749  | WE9598        | PC-10C192-SON-XFP   |
| PIC 2       | REV 10  | 750-009450  | HX6466        | PC-10C192-SON-SR2   |
| SIB 0       | REV 08  | 710-022594  | DW8033        | SIB-TXP-T1600-S     |
| SIB 1       | REV 08  | 710-022594  | DW8044        | SIB-TXP-T1600-S     |
| SIB 2       | REV 08  | 710-022594  | DW8020        | SIB-TXP-T1600-S     |
| SIB 3       | REV 08  | 710-022594  | DW8063        | SIB-TXP-T1600-S     |
| SIB 4       | REV 08  | 710-022594  | DW8064        | SIB-TXP-T1600-S     |
| Fan Tray 0  |         |             |               | FANTRAY-T-S         |
| Fan Tray 1  |         |             |               | FANTRAY-T-S         |
| Fan Tray 2  |         |             |               | FANTRAY-TXP-R-S     |

lcc1-re0:

-----  
Hardware inventory:

| Item        | Version | Part number | Serial number | FRU model number    |
|-------------|---------|-------------|---------------|---------------------|
| Midplane    | REV 04  | 710-017247  | RC5361        | CHAS-BP-T1600-S     |
| FPM Display | REV 01  | 710-021387  | DS6430        | CRAFT-T1600-S       |
| CIP         | REV 06  | 710-002895  | DS4239        | CIP-L-T640-S        |
| PEM 0       | Rev 08  | 740-017906  | UD26649       | PWR-T1600-3-80-DC-S |
| SCG 0       | REV 15  | 710-003423  | DP5820        | SCG-T-S             |
| CB 0        | REV 06  | 710-022597  | DW8523        | CB-LCC              |
| CB 1        | REV 06  | 710-022597  | DW8528        | CB-LCC              |
| FPC 4       | REV 12  | 710-013037  | DP8509        | T1600-FPC4-ES       |

|            |        |            |        |                 |
|------------|--------|------------|--------|-----------------|
| PIC 0      | REV 11 | 750-017405 | DP8808 | PD-4XGE-XFP     |
| PIC 1      | REV 11 | 750-017405 | DP7263 | PD-4XGE-XFP     |
| FPC 6      | REV 14 | 710-013037 | DS9961 | T1600-FPC4-ES   |
| PIC 0      | REV 13 | 750-017405 | DS5532 | PD-4XGE-XFP     |
| PIC 1      | REV 13 | 750-017405 | DS7639 | PD-4XGE-XFP     |
| FPC 7      | REV 03 | 710-013035 | DF5564 | T1600-FPC3-ES   |
| PIC 0      | REV 16 | 750-007141 | JJ8063 | PC-10GE-SFP     |
| SIB 0      | REV 08 | 710-022594 | DW8035 | SIB-TXP-T1600-S |
| SIB 1      | REV 10 | 710-022594 | DX7672 | SIB-TXP-T1600-S |
| SIB 2      | REV 08 | 710-022594 | DW8060 | SIB-TXP-T1600-S |
| SIB 3      | REV 08 | 710-022594 | DW8072 | SIB-TXP-T1600-S |
| SIB 4      | REV 08 | 710-022594 | DW8043 | SIB-TXP-T1600-S |
| Fan Tray 0 |        |            |        | FANTRAY-T-S     |
| Fan Tray 1 |        |            |        | FANTRAY-T-S     |
| Fan Tray 2 |        |            |        | FANTRAY-TXP-R-S |

lcc2-re0:

-----  
Hardware inventory:

| Item        | Version | Part number | Serial number | FRU model number    |
|-------------|---------|-------------|---------------|---------------------|
| Midplane    | REV 03  | 710-017247  | RC3956        | CHAS-BP-T1600-S     |
| FPM Display | REV 01  | 710-021387  | DN7030        | CRAFT-T1600-S       |
| CIP         | REV 06  | 710-002895  | DM3962        | CIP-L-T640-S        |
| PEM 0       | Rev 08  | 740-017906  | UD26519       | PWR-T1600-3-80-DC-S |
| PEM 1       | Rev 07  | 740-017906  | UC26601       | PWR-T1600-3-80-DC-S |
| SCG 0       | REV 15  | 710-003423  | DP0277        | SCG-T-S             |
| CB 0        | REV 06  | 710-022597  | DW8524        | CB-LCC              |
| CB 1        | REV 06  | 710-022597  | DW8536        | CB-LCC              |
| FPC 4       | REV 12  | 710-013037  | DR1194        | T1600-FPC4-ES       |
| PIC 0       | REV 11  | 750-017405  | DP8811        | PD-4XGE-XFP         |
| PIC 1       | REV 11  | 750-017405  | DP8823        | PD-4XGE-XFP         |
| FPC 5       | REV 12  | 710-013037  | DR1184        | T1600-FPC4-ES       |
| PIC 1       | REV 11  | 750-017405  | DP4744        | PD-4XGE-XFP         |
| FPC 6       | REV 12  | 710-013037  | DN8622        | T1600-FPC4-ES       |
| PIC 0       | REV 14  | 750-012518  | JY9924        | PD-40C192-SON-XFP   |
| PIC 1       | REV 11  | 750-017405  | DP8776        | PD-4XGE-XFP         |
| FPC 7       | REV 04  | 710-013560  | JR3968        | T640-FPC3-E2        |
| PIC 0       | REV 16  | 750-007141  | NC9330        | PC-10GE-SFP         |
| SIB 0       | REV 07  | 710-022594  | DW4217        | SIB-TXP-T1600-S     |
| SIB 1       | REV 07  | 710-022594  | DW4213        | SIB-TXP-T1600-S     |
| SIB 2       | REV 07  | 710-022594  | DW4189        | SIB-TXP-T1600-S     |
| SIB 3       | REV 07  | 710-022594  | DW4173        | SIB-TXP-T1600-S     |
| SIB 4       | REV 07  | 710-022594  | DW4201        | SIB-TXP-T1600-S     |
| Fan Tray 0  |         |             |               | FANTRAY-T-S         |
| Fan Tray 1  |         |             |               | FANTRAY-T-S         |
| Fan Tray 2  |         |             |               | FANTRAY-TXP-R-S     |

lcc3-re0:

-----  
Hardware inventory:

| Item        | Version | Part number | Serial number | FRU model number    |
|-------------|---------|-------------|---------------|---------------------|
| Midplane    | REV 04  | 710-017247  | RC5319        | CHAS-BP-T1600-S     |
| FPM Display | REV 01  | 710-021387  | DS6402        | CRAFT-T1600-S       |
| CIP         | REV 06  | 710-002895  | DR9973        | CIP-L-T640-S        |
| PEM 0       | Rev 07  | 740-017906  | UC26496       | PWR-T1600-3-80-DC-S |
| PEM 1       | Rev 07  | 740-017906  | UC26599       | PWR-T1600-3-80-DC-S |
| SCG 0       | REV 15  | 710-003423  | DP5831        | SCG-T-S             |
| CB 0        | REV 06  | 710-022597  | DW8533        | CB-LCC              |
| CB 1        | REV 06  | 710-022597  | DW8538        | CB-LCC              |
| FPC 0       | REV 14  | 710-013037  | DS5345        | T1600-FPC4-ES       |
| PIC 0       | REV 13  | 750-017405  | DS7641        | PD-4XGE-XFP         |

|            |        |            |        |                   |
|------------|--------|------------|--------|-------------------|
| PIC 1      | REV 13 | 750-017405 | DS5479 | PD-4XGE-XFP       |
| FPC 1      | REV 14 | 710-013037 | DS7338 | T1600-FPC4-ES     |
| PIC 0      | REV 13 | 750-017405 | DS7631 | PD-4XGE-XFP       |
| PIC 1      | REV 13 | 750-017405 | DS7632 | PD-4XGE-XFP       |
| FPC 2      | REV 14 | 710-013037 | DS9962 | T1600-FPC4-ES     |
| PIC 0      | REV 13 | 750-017405 | DS7581 | PD-4XGE-XFP       |
| PIC 1      | REV 13 | 750-017405 | DS7627 | PD-4XGE-XFP       |
| FPC 4      | REV 10 | 710-010845 | JZ6573 | T640-FPC4-ES      |
| PIC 0      | REV 14 | 750-012518 | JT5124 | PD-40C192-SON-XFP |
| FPC 5      | REV 14 | 710-013037 | DT0016 | T1600-FPC4-ES     |
| PIC 0      | REV 14 | 750-012518 | JY9918 | PD-40C192-SON-XFP |
| FPC 7      | REV 07 | 710-013035 | DM0967 | T1600-FPC3-ES     |
| PIC 0      | REV 16 | 750-007141 | JJ8059 | PC-10GE-SFP       |
| PIC 1      | REV 13 | 750-004695 | DM5712 | PC-TUNNEL         |
| SIB 0      | REV 07 | 710-022594 | DW4174 | SIB-TXP-T1600-S   |
| SIB 1      | REV 07 | 710-022594 | DW4207 | SIB-TXP-T1600-S   |
| SIB 2      | REV 06 | 710-022594 | DT8231 | SIB-TXP-T1600-S   |
| SIB 3      | REV 07 | 710-022594 | DW4175 | SIB-TXP-T1600-S   |
| SIB 4      | REV 07 | 710-022594 | DW4209 | SIB-TXP-T1600-S   |
| Fan Tray 0 |        |            |        | FANTRAY-T-S       |
| Fan Tray 1 |        |            |        | FANTRAY-T-S       |
| Fan Tray 2 |        |            |        | FANTRAY-TXP-R-S   |

#### show chassis hardware (16-Port 10-Gigabit Ethernet MPC with SFP+ Optics [MX Series Routers])

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

| Item             | Version | Part number | Serial number | Description          |
|------------------|---------|-------------|---------------|----------------------|
| Chassis          |         |             | JN112D865AFA  | MX960                |
| Midplane         | REV 03  | 710-013698  | TS3339        | MX960 Backplane      |
| FPM Board        | REV 03  | 710-014974  | WW6267        | Front Panel Display  |
| PDM              | Rev 03  | 740-013110  | QCS12485026   | Power Distribution   |
| Module           |         |             |               |                      |
| PEM 0            | Rev 04  | 740-013682  | QCS12434086   | PS 1.7kW; 200-240VAC |
| in               |         |             |               |                      |
| PEM 1            | Rev 04  | 740-013682  | QCS1243408Z   | PS 1.7kW; 200-240VAC |
| in               |         |             |               |                      |
| PEM 2            | Rev 04  | 740-013682  | QCS1243407X   | PS 1.7kW; 200-240VAC |
| in               |         |             |               |                      |
| Routing Engine 0 | REV 07  | 740-015113  | 9009009677    | RE-S-1300            |
| Routing Engine 1 | REV 07  | 740-015113  | 9009011510    | RE-S-1300            |
| CB 0             | REV 03  | 710-021523  | XF0394        | MX SCB               |
| CB 1             | REV 03  | 710-021523  | XF0550        | MX SCB               |
| CB 2             | REV 03  | 710-021523  | XD7455        | MX SCB               |
| FPC 4            | REV 02  | 750-028467  | JR6127        | MPC M 16x 10GE       |
| CPU              | REV 02  | 711-029089  | JX0129        | AS PMB               |
| PIC 0            |         | BUILTIN     | BUILTIN       | 4x 10GE(LAN) SFP+    |
| PIC 1            |         | BUILTIN     | BUILTIN       | 4x 10GE(LAN) SFP+    |
| PIC 2            |         | BUILTIN     | BUILTIN       | 4x 10GE(LAN) SFP+    |
| PIC 3            |         | BUILTIN     | BUILTIN       | 4x 10GE(LAN) SFP+    |
| Fan Tray 0       | REV 05  | 740-014971  | TP9990        | Fan Tray             |
| Fan Tray 1       | REV 05  | 740-014971  | VS1709        | Fan Tray             |

#### show chassis hardware (MPC3E [MX Series Routers])

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

| Item     | Version | Part number | Serial number | Description    |
|----------|---------|-------------|---------------|----------------|
| Chassis  |         |             | JN1101AFEAFB  | MX480          |
| Midplane | REV 05  | 710-017414  | TR4444        | MX480 Midplane |

|                  |             |            |             |                        |
|------------------|-------------|------------|-------------|------------------------|
| FPM Board        | REV 02      | 710-017254 | KG6056      | Front Panel Display    |
| PEM 0            | Rev 03      | 740-017330 | QCS082090FC | PS 1.2-1.7kW; 100-240V |
| PEM 1            | Rev 03      | 740-017330 | QCS082090FD | PS 1.2-1.7kW; 100-240V |
| Routing Engine 0 | REV 07      | 740-013063 | 9009004124  | RE-S-2000              |
| Routing Engine 1 | REV 07      | 740-013063 | 9009005569  | RE-S-2000              |
| CB 0             | REV 07      | 710-021523 | XZ3587      | MX SCB                 |
| CB 1             | REV 03      | 710-021523 | KH8306      | MX SCB                 |
| FPC 1            | REV 04.1.07 | 750-033205 | P1240       | MPCE Type 3D           |
| CPU              | REV 01      | 711-035209 | YL0504      | HMPD PMB 2G            |
| MIC 1            | REV 10      | 750-033199 | YX4495      | 1X100GE CFP            |
| PIC 2            |             | BUILTIN    | BUILTIN     | 1X100GE CFP            |
| Xcvr 0           | REV 01      | 740-032210 | C22CQNE     | CFP-100G-LR4           |
| FPC 2            | REV 26      | 750-016670 | KH0045      | DPCE 40x 1GE R EQ      |
| CPU              | REV 07      | 710-013713 | KF5448      | DPC PMB                |
| PIC 0            |             | BUILTIN    | BUILTIN     | 10x 1GE(LAN) EQ        |
| Xcvr 0           | REV 01      | 740-011613 | PF21JHU     | SFP-SX                 |
| PIC 1            |             | BUILTIN    | BUILTIN     | 10x 1GE(LAN) EQ        |
| Xcvr 9           | REV 01      | 740-011613 | AM0813S8ZL6 | SFP-SX                 |
| PIC 2            |             | BUILTIN    | BUILTIN     | 10x 1GE(LAN) EQ        |
| Xcvr 0           | REV 02      | 740-011613 | PGL2KYF     | SFP-SX                 |
| Xcvr 2           | REV 01      | 740-011613 | AM0806S8N4P | SFP-SX                 |
| PIC 3            |             | BUILTIN    | BUILTIN     | 10x 1GE(LAN) EQ        |
| Xcvr 5           | REV 01      | 740-011613 | AM0815S967N | SFP-SX                 |
| Xcvr 7           | REV 01      | 740-011613 | AM0806S8N1X | SFP-SX                 |
| Xcvr 8           | REV 01      | 740-011613 | AM0815S967J | SFP-SX                 |
| Xcvr 9           | REV 01      | 740-011613 | AM0815S967M | SFP-SX                 |
| FPC 3            | REV 12.2.09 | 750-033205 | YR9443      | MPCE Type 3D           |
| CPU              | REV 03      | 711-035209 | YL6931      | HMPD PMB 2G            |
| MIC 0            | REV 05      | 750-033199 | YR3269      | 1X100GE CFP            |
| PIC 0            |             | BUILTIN    | BUILTIN     | 1X100GE CFP            |
| Xcvr 0           | REV 01      | 740-032210 | ULHOKG3     | CFP-100G-LR4           |
| MIC 1            | REV 02      | 750-033199 | YG3245      | 1X100GE CFP            |
| PIC 2            |             | BUILTIN    | BUILTIN     | 1X100GE CFP            |
| Xcvr 0           | REV 01      | 740-032210 | ULHOKGF     | CFP-100G-LR4           |
| FPC 4            | REV 12.3.09 | 750-033205 | YR9437      | MPCE Type 3D           |
| CPU              | REV 03      | 711-035209 | YT5857      | HMPD PMB 2G            |
| MIC 0            | REV 05      | 750-033199 | YR3295      | 1X100GE CFP            |
| PIC 0            |             | BUILTIN    | BUILTIN     | 1X100GE CFP            |
| Xcvr 0           |             | NON-JNPR   | X12000187   | CFP-100G-SR10          |
| MIC 1            | REV 10      | 750-033199 | YX4518      | 1X100GE CFP            |
| PIC 2            |             | BUILTIN    | BUILTIN     | 1X100GE CFP            |
| Xcvr 0           | REV 01      | 740-035329 | X12J00008   | CFP-100G-SR10          |
| FPC 5            | REV 06      | 750-024884 | JW9769      | MPC Type 2 3D EQ       |
| CPU              | REV 02      | 711-028401 | JR6158      | MPC PMB 2G Proto       |
| MIC 0            | REV 05      | 750-028387 | JR6197      | 3D 4x 10GE XFP         |
| PIC 0            |             | BUILTIN    | BUILTIN     | 2x 10GE XFP            |
| Xcvr 0           | REV 01      | 740-014289 | T07M71112   | XFP-10G-SR             |
| Xcvr 1           | REV 02      | 740-014289 | T08L85610   | XFP-10G-SR             |
| PIC 1            |             | BUILTIN    | BUILTIN     | 2x 10GE XFP            |
| MIC 1            | REV 22      | 750-028392 | YM0053      | 3D 20x 1GE(LAN) SFP    |
| PIC 2            |             | BUILTIN    | BUILTIN     | 10x 1GE(LAN) SFP       |
| Xcvr 0           | REV 01      | 740-011613 | AM0703S005B | SFP-SX                 |
| Xcvr 1           | REV 01      | 740-011613 | E07L01352   | SFP-SX                 |
| PIC 3            |             | BUILTIN    | BUILTIN     | 10x 1GE(LAN) SFP       |
| Xcvr 5           | REV 01      | 740-013111 | 6500217     | SFP-T                  |

|          |        |            |         |               |
|----------|--------|------------|---------|---------------|
| Xcvr 9   | REV 02 | 740-013111 | 8499527 | SFP-T         |
| Fan Tray |        |            |         | Left Fan Tray |

The PIC number for MIC 1 always starts from 2 (even if the first MIC is a 1X100GE CFP or a legacy MIC).

### show chassis hardware (QFX3500 Switches)

```
user@switch> show chassis hardware
Hardware inventory:
Item Version Part number Serial number Description
Chassis
Routing Engine 0
FPC 0 REV 04 750-044071 BUILTIN QFX3500
CPU BUILTIN BUILTIN QFX Routing Engine
PIC 0 BUILTIN BUILTIN QFX3500-48S4Q-AFI
PIC 1 BUILTIN BUILTIN FPC CPU
MGMT BRD REV 02 750-044063 BBAR0398 48x 10G-SFP+
Xcvr 0 REV 01 740-011614 AC0946S0BD1 15x 10G-SFP+
Xcvr 1 REV 02 740-013111 A281922 QFX3500-MGMT-SFP-AFO
Power Supply 0 Rev 04 740-032091 UI00677 SFP-LX10
Power Supply 1 REV 00 740-041741 VJ00162 SFP-T
Fan Tray 0
Front Airflow
Fan Tray 1
Front Airflow
Fan Tray 2
Front Airflow
QFX Fan Tray, Back to
QFX Fan Tray, Back to
QFX Fan Tray, Back to
```

### show chassis hardware detail (QFX3500 Switches)

```
user@switch> show chassis hardware detail
Hardware inventory:
Item Version Part number Serial number Description
Chassis
Routing Engine 0
FPC 0 REV 05 750-036931 EE0823 QFX3500
CPU BUILTIN BUILTIN QFX Routing Engine
PIC 0 BUILTIN BUILTIN QFX3500-48S4Q-AFI
Xcvr 0 REV 01 740-030589 S99E270079 SFP+-10G-LPBK
Xcvr 1 REV 01 740-030589 S9AK450099 SFP+-10G-LPBK
Xcvr 2 REV 01 740-030589 S99E270078 SFP+-10G-LPBK
Xcvr 3 REV 01 740-030589 S9AK450098 SFP+-10G-LPBK
Xcvr 4 REV 01 740-030589 S99E270075 SFP+-10G-LPBK
Xcvr 5 REV 01 740-030589 S9AK450093 SFP+-10G-LPBK
Xcvr 6 REV 01 740-030589 S9AK450097 SFP+-10G-LPBK
Xcvr 7 REV 01 740-030589 S9AK450095 SFP+-10G-LPBK
Xcvr 8 REV 01 740-030589 S99E270072 SFP+-10G-LPBK
Xcvr 9 REV 01 740-030589 S99E270073 SFP+-10G-LPBK
Xcvr 10 REV 01 740-030589 S99E270080 SFP+-10G-LPBK
Xcvr 11 REV 01 740-030589 S9AK450169 SFP+-10G-LPBK
Xcvr 12 REV 01 740-030589 S99E270076 SFP+-10G-LPBK
Xcvr 13 REV 01 740-030589 S9AK450167 SFP+-10G-LPBK
Xcvr 14 REV 01 740-030589 S9AK450170 SFP+-10G-LPBK
Xcvr 15 REV 01 740-030589 S9AK450166 SFP+-10G-LPBK
Xcvr 16 REV 01 740-030589 S9AK450092 SFP+-10G-LPBK
Xcvr 17 REV 01 740-030589 S9AK450163 SFP+-10G-LPBK
Xcvr 18 REV 01 740-030589 S9AK450094 SFP+-10G-LPBK
Xcvr 19 REV 01 740-030589 S9AK450100 SFP+-10G-LPBK
```

|                |        |            |            |                |
|----------------|--------|------------|------------|----------------|
| Xcvr 20        | REV 01 | 740-030589 | S9AK450168 | SFP+-10G-LPBK  |
| Xcvr 21        | REV 01 | 740-030589 | S9AK450165 | SFP+-10G-LPBK  |
| Xcvr 22        | REV 01 | 740-030589 | S9AK450073 | SFP+-10G-LPBK  |
| Xcvr 23        | REV 01 | 740-030589 | S9AK450164 | SFP+-10G-LPBK  |
| Xcvr 24        | REV 01 | 740-030589 | S9AK450074 | SFP+-10G-LPBK  |
| Xcvr 25        | REV 01 | 740-030589 | SA62270195 | SFP+-10G-LPBK  |
| Xcvr 26        | REV 01 | 740-030589 | S9AK450078 | SFP+-10G-LPBK  |
| Xcvr 27        | REV 01 | 740-030589 | S9AK450024 | SFP+-10G-LPBK  |
| Xcvr 28        | REV 01 | 740-030589 | S9AK450027 | SFP+-10G-LPBK  |
| Xcvr 29        | REV 01 | 740-030589 | S9AK450080 | SFP+-10G-LPBK  |
| Xcvr 30        | REV 01 | 740-030589 | S9AK450030 | SFP+-10G-LPBK  |
| Xcvr 31        | REV 01 | 740-030589 | S9AK450025 | SFP+-10G-LPBK  |
| Xcvr 32        | REV 01 | 740-030589 | S9AK450023 | SFP+-10G-LPBK  |
| Xcvr 33        | REV 01 | 740-030589 | S9AK450075 | SFP+-10G-LPBK  |
| Xcvr 34        | REV 01 | 740-030589 | S9AK450161 | SFP+-10G-LPBK  |
| Xcvr 35        | REV 01 | 740-030589 | S9AK450071 | SFP+-10G-LPBK  |
| Xcvr 36        | REV 01 | 740-030589 | S9AK450072 | SFP+-10G-LPBK  |
| Xcvr 37        | REV 01 | 740-030589 | S9AK450022 | SFP+-10G-LPBK  |
| Xcvr 38        | REV 01 | 740-030589 | S9AK450021 | SFP+-10G-LPBK  |
| Xcvr 39        | REV 01 | 740-030589 | S9AK450175 | SFP+-10G-LPBK  |
| Xcvr 40        | REV 01 | 740-030589 | S9AK450162 | SFP+-10G-LPBK  |
| Xcvr 41        | REV 01 | 740-030589 | S99E270074 | SFP+-10G-LPBK  |
| Xcvr 42        | REV 01 | 740-030589 | S9AK450174 | SFP+-10G-LPBK  |
| Xcvr 43        | REV 01 | 740-030589 | S9AK450077 | SFP+-10G-LPBK  |
| Xcvr 44        | REV 01 | 740-030589 | S9AK450076 | SFP+-10G-LPBK  |
| Xcvr 45        | REV 01 | 740-030589 | S9AK450026 | SFP+-10G-LPBK  |
| Xcvr 46        | REV 01 | 740-030589 | S9AK450079 | SFP+-10G-LPBK  |
| Xcvr 47        | REV 01 | 740-030589 | S9AK450029 | SFP+-10G-LPBK  |
| PIC 1          |        | BUILTIN    | BUILTIN    | 15x 10G-SFP+   |
| Xcvr 1         | REV 01 | 740-032986 | QA170087   | QSFP+-40G-SR4  |
| Xcvr 4         | REV 01 | 740-032986 | QA360442   | QSFP+-40G-SR4  |
| Xcvr 8         | REV 01 | 740-032986 | QA170091   | QSFP+-40G-SR4  |
| Xcvr 12        | REV 01 | 740-032986 | QA170042   | QSFP+-40G-SR4  |
| MGMT BRD       | REV 08 | 750-036946 | EE0731     | QFX3500-MB     |
| Power Supply 0 | Rev 04 | 740-032091 | UI00690    | QFX PS 650W AC |
| Power Supply 1 | Rev 04 | 740-032091 | UI00679    | QFX PS 650W AC |
| Fan Tray 0     |        |            |            | QFX Fan Tray   |
| Fan Tray 1     |        |            |            | QFX Fan Tray   |

### show chassis hardware models (QFX3500 Switches)

```

user@switch> show chassis hardware models
Hardware inventory:
Item Version Part number Serial number FRU model number
Routing Engine 0 BUILTIN BUILTIN
FPC 0 REV 02 711-032234 EC4074
Power Supply 0 PSMI 2C 11-d65800 --

```

### show chassis hardware clei-models (QFX3500 Switches)

```

user@switch> show chassis hardware clei-models
Hardware inventory:
Item Version Part number CLEI code FRU model number
Routing Engine 0 BUILTIN
FPC 0 REV 02 711-032234
Power Supply 0 PSMI 2C 11-d65800

```

### show chassis hardware interconnect-device (QFabric Systems)

```

user@switch> show chassis hardware interconnect-device interconnect1
Hardware inventory:
Item Version Part number Serial number Description

```

|          |        |            |              |              |
|----------|--------|------------|--------------|--------------|
| Chassis  | REV 07 |            |              | QFX_olive    |
| Midplane | REV 07 | 750-021261 | BH0208188289 | QFX Midplane |
| CB 0     | REV 07 | 750-021261 | BH0208188289 | QFXIC08-CB4S |

### show chassis hardware node-device (QFabric Systems)

```

user@switch> show chassis hardware node-device node1
Routing Engine 0 BUILTIN BUILTIN QFX Routing Engine
node1 REV 05 711-032234 ED3694 QFX3500-48S4Q-AFI

CPU BUILTIN BUILTIN
PIC 0 BUILTIN BUILTIN
Xcvr 8 REV 01 740-030658 AD0946A028B FPC CPU
 48x 10G-SFP+
 SFP+-10G-USR
...

```

### show chassis hardware (PTX5000 Packet Transport Switch)

```

user@switch> show chassis hardware
Hardware inventory:
Item Version Part number Serial number Description
Chassis JN11D1FD7AJA PTX5000
Midplane REV 03 711-031896 ABAC5589 Midplane-8S
FPM REV 08 760-030647 EG1679 Front Panel Display
PDU 0 Rev 05 740-032019 ZE00006 DC Power Dist Unit
PSM 0 Rev 05 740-032022 ZJ00018 DC 12V Power Supply
PSM 1 Rev 04 740-032022 ZC00052 DC 12V Power Supply
PSM 2 Rev 04 740-032022 ZD00051 DC 12V Power Supply
PSM 3 Rev 05 740-032022 ZJ00060 DC 12V Power Supply
CCG 0 REV 04 750-030653 EG3703 Clock Generator
CCG 1 REV 04 750-030653 EG3698 Clock Generator
Routing Engine 0 REV 05 740-026942 P737A-002231 RE-DUO-2600
Routing Engine 1 REV 06 740-026942 P737A-002438 RE-DUO-2600
CB 0 REV 08 750-030625 EG5519 Control Board
CB 1 REV 08 750-030625 EG5516 Control Board
FPC 0 REV 18 750-036844 EJ3080 FPC
CPU REV 12 711-030686 EJ3260 SNG PMB
FPC 2 REV 13 750-036844 EG5065 FPC
CPU REV 09 711-030686 EG4082 SNG PMB
PIC 0 REV 14 750-031913 EG5127 24x 10GE(LAN) SFP+
Xcvr 0 REV 01 740-031980 143363A00240 SFP+-10G-SR
Xcvr 1 REV 01 740-031981 UK90PZ1 SFP+-10G-LR
Xcvr 2 REV 01 740-031980 AD1141A04XH SFP+-10G-SR
Xcvr 3 REV 01 740-031981 UK90Q46 SFP+-10G-LR
Xcvr 4 REV 01 740-031980 AD1141A04X4 SFP+-10G-SR
Xcvr 6 REV 01 740-031980 B11H02560 SFP+-10G-SR
Xcvr 7 REV 01 740-031980 B11C01589 SFP+-10G-SR
Xcvr 8 REV 01 740-031980 AD1141A04XF SFP+-10G-SR
Xcvr 10 REV 01 740-031980 123363A01094 SFP+-10G-SR
Xcvr 11 REV 01 740-031980 AK80LKF SFP+-10G-SR
Xcvr 12 REV 01 740-031980 183363A01528 SFP+-10G-SR
Xcvr 14 REV 01 740-031980 193363A01079 SFP+-10G-SR
Xcvr 15 REV 01 740-031980 AK80MC8 SFP+-10G-SR
Xcvr 16 REV 01 740-031980 AJCOBHC SFP+-10G-SR
Xcvr 19 REV 01 740-021309 J08D26856 SFP+-10G-LR
Xcvr 21 REV 01 740-031980 AK80KCT SFP+-10G-SR
Xcvr 22 REV 01 740-031981 UK90PZL SFP+-10G-LR
Xcvr 23 REV 01 740-031980 AK80N1V SFP+-10G-SR
FPC 3 REV 13 750-036844 EG5074 FPC
CPU REV 09 711-030686 EG4064 SNG PMB
PIC 1 REV 10 750-031903 EG0325 SNG Load

```



|            |        |            |              |                     |
|------------|--------|------------|--------------|---------------------|
| FPC 5      | REV 06 | 750-036844 | EH3198       | FPC                 |
| CPU        |        |            |              |                     |
| PIC 0      | REV 14 | 750-031913 | EG5134       | 24x 10GE(LAN) SFP+  |
| Xcvr 0     | REV 01 | 740-031980 | AK80LBH      | SFP+-10G-SR         |
| Xcvr 1     | REV 01 | 740-031980 | B11B03724    | SFP+-10G-SR         |
| Xcvr 2     | REV 01 | 740-031980 | AK80FMH      | SFP+-10G-SR         |
| Xcvr 5     | REV 01 | 740-031980 | B11J00818    | SFP+-10G-SR         |
| Xcvr 6     | REV 01 | 740-031980 | 193363A00743 | SFP+-10G-SR         |
| Xcvr 7     | REV 01 | 740-031980 | B11B06125    | SFP+-10G-SR         |
| Xcvr 10    | REV 01 | 740-031980 | B11H02529    | SFP+-10G-SR         |
| Xcvr 11    | REV 01 | 740-031980 | AK80LFB      | SFP+-10G-SR         |
| Xcvr 12    | REV 01 | 740-031980 | 193363A01061 | SFP+-10G-SR         |
| Xcvr 15    | REV 01 | 740-031980 | B11J00687    | SFP+-10G-SR         |
| Xcvr 16    | REV 01 | 740-031980 | 193363A00738 | SFP+-10G-SR         |
| Xcvr 18    | REV 01 | 740-031980 | AK80MQX      | SFP+-10G-SR         |
| Xcvr 19    | REV 01 | 740-021309 | J08C17257    | SFP+-10G-LR         |
| Xcvr 22    | REV 01 | 740-031980 | B11J00730    | SFP+-10G-SR         |
| Xcvr 23    | REV 01 | 740-031980 | AK80KEE      | SFP+-10G-SR         |
| PIC 1      | REV 08 | 750-036710 | EG3105       | 2x 40GE CFP         |
| Xcvr 0     | REV 01 | 740-034554 | B260HLT      | CFP-40G-LR4         |
| Xcvr 1     | REV 01 | 740-034554 | B11C02847    | CFP-40G-LR4         |
| FPC 6      | REV 18 | 750-036844 | EJ4391       | FPC                 |
| CPU        | REV 12 | 711-030686 | EJ3257       | SNG PMB             |
| FPC 7      | REV 18 | 750-036844 | EJ4382       | FPC                 |
| CPU        | REV 12 | 711-030686 | EJ3238       | SNG PMB             |
| SPMB 0     | REV 10 | 711-030686 | EG5418       | SNG PMB             |
| SPMB 1     | REV 09 | 711-030686 | EG5373       | SNG PMB             |
| SIB 0      | REV 07 | 750-030631 | EG4858       | SIB-I-8S            |
| SIB 1      | REV 07 | 750-030631 | EG4872       | SIB-I-8S            |
| SIB 2      | REV 07 | 750-030631 | EG4866       | SIB-I-8S            |
| SIB 3      | REV 07 | 750-030631 | EG6011       | SIB-I-8S            |
| SIB 4      | REV 07 | 750-030631 | EG4907       | SIB-I-8S            |
| SIB 5      | REV 07 | 750-030631 | EG4879       | SIB-I-8S            |
| SIB 6      | REV 07 | 750-030631 | EG4864       | SIB-I-8S            |
| SIB 7      | REV 07 | 750-030631 | EG4899       | SIB-I-8S            |
| SIB 8      | REV 07 | 750-030631 | EG4880       | SIB-I-8S            |
| Fan Tray 0 | REV 04 | 760-032784 | EG1496       | Vertical Fan Tray   |
| Fan Tray 1 | REV 04 | 760-030642 | EG1335       | Horizontal Fan Tray |
| Fan Tray 2 | REV 02 | 760-030642 | ED4952       | Horizontal Fan Tray |

### show chassis hardware clei-models (PTX5000 Packet Transport Switch)

```
user@switch> show chassis hardware clei-models
```

| Hardware inventory: |         |             |            |                     |
|---------------------|---------|-------------|------------|---------------------|
| Item                | Version | Part number | CLEI code  | FRU model number    |
| FPM                 | REV 08  | 760-030647  | PROTOXCLEI | CRAFT-PTX5000-S     |
| PDU 0               | Rev 05  | 740-032019  | IPUPAHLKAA | PWR-SAN-PDU-DC      |
| PSM 0               | Rev 05  | 740-032022  | IPUPAHNKAA | PSM-PTX-DC-120-S    |
| PSM 1               | Rev 04  | 740-032022  | 032022XXXX | PWR-SAN-12-DC       |
| PSM 2               | Rev 04  | 740-032022  | 032022XXXX | PWR-SAN-12-DC       |
| PSM 3               | Rev 05  | 740-032022  | IPUPAHNKAA | PSM-PTX-DC-120-S    |
| CCG 0               | REV 04  | 750-030653  | PROTOXCLEI | CCG-PTX-S           |
| CCG 1               | REV 04  | 750-030653  | PROTOXCLEI | CCG-PTX-S           |
| Routing Engine 0    | REV 05  | 740-026942  |            | RE-DUO-C2600-16G-S  |
| Routing Engine 1    | REV 06  | 740-026942  |            | RE-DUO-C2600-16G-S  |
| CB 0                | REV 08  | 750-030625  | PROTOXCLEI | CB-PTX-S            |
| CB 1                | REV 08  | 750-030625  | PROTOXCLEI | CB-PTX-S            |
| FPC 0               | REV 18  | 750-036844  | PROTOXCLEI | FPC-PTX-P1-A        |
| FPC 2               | REV 13  | 750-036844  | PROTOXCLEI | FPC-PTX-P1-A        |
| PIC 0               | REV 14  | 750-031913  | PROTOXCLEI | P1-PTX-24-10GE-SFPP |
| FPC 3               | REV 13  | 750-036844  | PROTOXCLEI | FPC-PTX-P1-A        |

|            |        |            |            |                     |
|------------|--------|------------|------------|---------------------|
| FPC 5      |        |            |            |                     |
| PIC 0      | REV 14 | 750-031913 | PROTOXCLEI | P1-PTX-24-10GE-SFPP |
| FPC 6      | REV 18 | 750-036844 | PROTOXCLEI | FPC-PTX-P1-A        |
| FPC 7      | REV 18 | 750-036844 | PROTOXCLEI | FPC-PTX-P1-A        |
| SIB 0      | REV 07 | 750-030631 | PROTOXCLEI | SIB-I-PTX5008       |
| SIB 1      | REV 07 | 750-030631 | PROTOXCLEI | SIB-I-PTX5008       |
| SIB 2      | REV 07 | 750-030631 | PROTOXCLEI | SIB-I-PTX5008       |
| SIB 3      | REV 07 | 750-030631 | PROTOXCLEI | SIB-I-PTX5008       |
| SIB 4      | REV 07 | 750-030631 | PROTOXCLEI | SIB-I-PTX5008       |
| SIB 5      | REV 07 | 750-030631 | PROTOXCLEI | SIB-I-PTX5008       |
| SIB 6      | REV 07 | 750-030631 | PROTOXCLEI | SIB-I-PTX5008       |
| SIB 7      | REV 07 | 750-030631 | PROTOXCLEI | SIB-I-PTX5008       |
| SIB 8      | REV 07 | 750-030631 | PROTOXCLEI | SIB-I-PTX5008       |
| Fan Tray 1 | REV 04 | 760-030642 | PROTOXCLEI | FAN-PTX-H-S         |

### show chassis hardware detail (PTX5000 Packet Transport Switch)

```

user@switch> show chassis hardware detail
Hardware inventory:
Item Version Part number Serial number Description
Chassis REV 03 711-031896 JN1D1FD7AJA PTX5000
Midplane REV 08 760-030647 EG1679 Midplane-8S
FPM REV 05 740-032019 ZE00006 Front Panel Display
PDU 0 Rev 05 740-032022 ZJ00018 DC Power Dist Unit
 PSM 0 Rev 04 740-032022 ZC00052 DC 12V Power Supply
 PSM 1 Rev 04 740-032022 ZD00051 DC 12V Power Supply
 PSM 2 Rev 05 740-032022 ZJ00060 DC 12V Power Supply
CCG 0 REV 04 750-030653 EG3703 Clock Generator
CCG 1 REV 04 750-030653 EG3698 Clock Generator
Routing Engine 0 REV 05 740-026942 P737A-002231 RE-DUO-2600
 ad0 3823 MB SMART CF 201006190039C02DC02D Compact Flash
 ad1 62720 MB SMART Lite SATA Drive 2011042300CF4C6B4C6B Disk 1
Routing Engine 1 REV 06 740-026942 P737A-002438 RE-DUO-2600
 ad0 3823 MB SMART CF 20100619053455F055F0 Compact Flash
 ad1 62720 MB SMART Lite SATA Drive 20110423000AE8E7E8E7 Disk 1
CB 0 REV 08 750-030625 EG5519 Control Board
CB 1 REV 08 750-030625 EG5516 Control Board
FPC 0 REV 18 750-036844 EJ3080 FPC
 CPU REV 12 711-030686 EJ3260 SNG PMB
FPC 2 REV 13 750-036844 EG5065 FPC
 CPU REV 09 711-030686 EG4082 SNG PMB
 PIC 0 REV 14 750-031913 EG5127 24x 10GE(LAN) SFP+
 Xcvr 0 REV 01 740-031980 143363A00240 SFP+-10G-SR
 Xcvr 1 REV 01 740-031981 UK90PZ1 SFP+-10G-LR
 Xcvr 2 REV 01 740-031980 AD1141A04XH SFP+-10G-SR
 Xcvr 3 REV 01 740-031981 UK90Q46 SFP+-10G-LR
 Xcvr 4 REV 01 740-031980 AD1141A04X4 SFP+-10G-SR
 Xcvr 6 REV 01 740-031980 B11H02560 SFP+-10G-SR
 Xcvr 7 REV 01 740-031980 B11C01589 SFP+-10G-SR
 Xcvr 8 REV 01 740-031980 AD1141A04XF SFP+-10G-SR
 Xcvr 10 REV 01 740-031980 123363A01094 SFP+-10G-SR
 Xcvr 11 REV 01 740-031980 AK80LKF SFP+-10G-SR
 Xcvr 12 REV 01 740-031980 183363A01528 SFP+-10G-SR
 Xcvr 14 REV 01 740-031980 193363A01079 SFP+-10G-SR
 Xcvr 15 REV 01 740-031980 AK80MC8 SFP+-10G-SR
 Xcvr 16 REV 01 740-031980 AJC0BHC SFP+-10G-SR
 Xcvr 19 REV 01 740-021309 J08D26856 SFP+-10G-LR
 Xcvr 21 REV 01 740-031980 AK80KCT SFP+-10G-SR
 Xcvr 22 REV 01 740-031981 UK90PZL SFP+-10G-LR
 Xcvr 23 REV 01 740-031980 AK80N1V SFP+-10G-SR

```

|            |        |            |              |                     |
|------------|--------|------------|--------------|---------------------|
| FPC 3      | REV 13 | 750-036844 | EG5074       | FPC                 |
| CPU        | REV 09 | 711-030686 | EG4064       | SNG PMB             |
| PIC 1      | REV 10 | 750-031903 | EG0325       | SNG Load            |
| FPC 5      | REV 06 | 750-036844 | EH3198       | FPC                 |
| CPU        |        |            |              |                     |
| PIC 0      | REV 14 | 750-031913 | EG5134       | 24x 10GE(LAN) SFP+  |
| Xcvr 0     | REV 01 | 740-031980 | AK80LBH      | SFP+-10G-SR         |
| Xcvr 1     | REV 01 | 740-031980 | B11B03724    | SFP+-10G-SR         |
| Xcvr 2     | REV 01 | 740-031980 | AK80FMH      | SFP+-10G-SR         |
| Xcvr 5     | REV 01 | 740-031980 | B11J00818    | SFP+-10G-SR         |
| Xcvr 6     | REV 01 | 740-031980 | 193363A00743 | SFP+-10G-SR         |
| Xcvr 7     | REV 01 | 740-031980 | B11B06125    | SFP+-10G-SR         |
| Xcvr 10    | REV 01 | 740-031980 | B11H02529    | SFP+-10G-SR         |
| Xcvr 11    | REV 01 | 740-031980 | AK80LFB      | SFP+-10G-SR         |
| Xcvr 12    | REV 01 | 740-031980 | 193363A01061 | SFP+-10G-SR         |
| Xcvr 15    | REV 01 | 740-031980 | B11J00687    | SFP+-10G-SR         |
| Xcvr 16    | REV 01 | 740-031980 | 193363A00738 | SFP+-10G-SR         |
| Xcvr 18    | REV 01 | 740-031980 | AK80MQX      | SFP+-10G-SR         |
| Xcvr 19    | REV 01 | 740-021309 | J08C17257    | SFP+-10G-LR         |
| Xcvr 22    | REV 01 | 740-031980 | B11J00730    | SFP+-10G-SR         |
| Xcvr 23    | REV 01 | 740-031980 | AK80KEE      | SFP+-10G-SR         |
| PIC 1      | REV 08 | 750-036710 | EG3105       | 2x 40GE CFP         |
| Xcvr 0     | REV 01 | 740-034554 | B260HLT      | CFP-40G-LR4         |
| Xcvr 1     | REV 01 | 740-034554 | B11C02847    | CFP-40G-LR4         |
| FPC 6      | REV 18 | 750-036844 | EJ4391       | FPC                 |
| CPU        | REV 12 | 711-030686 | EJ3257       | SNG PMB             |
| FPC 7      | REV 18 | 750-036844 | EJ4382       | FPC                 |
| CPU        | REV 12 | 711-030686 | EJ3238       | SNG PMB             |
| SPMB 0     | REV 10 | 711-030686 | EG5418       | SNG PMB             |
| SPMB 1     | REV 09 | 711-030686 | EG5373       | SNG PMB             |
| SIB 0      | REV 07 | 750-030631 | EG4858       | SIB-I-8S            |
| SIB 1      | REV 07 | 750-030631 | EG4872       | SIB-I-8S            |
| SIB 2      | REV 07 | 750-030631 | EG4866       | SIB-I-8S            |
| SIB 3      | REV 07 | 750-030631 | EG6011       | SIB-I-8S            |
| SIB 4      | REV 07 | 750-030631 | EG4907       | SIB-I-8S            |
| SIB 5      | REV 07 | 750-030631 | EG4879       | SIB-I-8S            |
| SIB 6      | REV 07 | 750-030631 | EG4864       | SIB-I-8S            |
| SIB 7      | REV 07 | 750-030631 | EG4899       | SIB-I-8S            |
| SIB 8      | REV 07 | 750-030631 | EG4880       | SIB-I-8S            |
| Fan Tray 0 | REV 04 | 760-032784 | EG1496       | Vertical Fan Tray   |
| Fan Tray 1 | REV 04 | 760-030642 | EG1335       | Horizontal Fan Tray |
| Fan Tray 2 | REV 02 | 760-030642 | ED4952       | Horizontal Fan Tray |

### show chassis hardware models (PTX5000 Packet Transport Switch)

user@switch> show chassis hardware models

Hardware inventory:

| Item             | Version | Part number | Serial number | FRU model number   |
|------------------|---------|-------------|---------------|--------------------|
| FPM              | REV 08  | 760-030647  | EG1679        | CRAFT-PTX5000-S    |
| PDU 0            | Rev 05  | 740-032019  | ZE00006       | PWR-SAN-PDU-DC     |
| PSM 0            | Rev 05  | 740-032022  | ZJ00018       | PSM-PTX-DC-120-S   |
| PSM 1            | Rev 04  | 740-032022  | ZC00052       | PWR-SAN-12-DC      |
| PSM 2            | Rev 04  | 740-032022  | ZD00051       | PWR-SAN-12-DC      |
| PSM 3            | Rev 05  | 740-032022  | ZJ00060       | PSM-PTX-DC-120-S   |
| CCG 0            | REV 04  | 750-030653  | EG3703        | CCG-PTX-S          |
| CCG 1            | REV 04  | 750-030653  | EG3698        | CCG-PTX-S          |
| Routing Engine 0 | REV 05  | 740-026942  | P737A-002231  | RE-DUO-C2600-16G-S |
| Routing Engine 1 | REV 06  | 740-026942  | P737A-002438  | RE-DUO-C2600-16G-S |
| CB 0             | REV 08  | 750-030625  | EG5519        | CB-PTX-S           |
| CB 1             | REV 08  | 750-030625  | EG5516        | CB-PTX-S           |
| FPC 0            | REV 18  | 750-036844  | EJ3080        | FPC-PTX-P1-A       |

|            |        |            |        |                     |
|------------|--------|------------|--------|---------------------|
| FPC 2      | REV 13 | 750-036844 | EG5065 | FPC-PTX-P1-A        |
| PIC 0      | REV 14 | 750-031913 | EG5127 | P1-PTX-24-10GE-SFPP |
| FPC 3      | REV 13 | 750-036844 | EG5074 | FPC-PTX-P1-A        |
| FPC 5      |        |            |        |                     |
| PIC 0      | REV 14 | 750-031913 | EG5134 | P1-PTX-24-10GE-SFPP |
| FPC 6      | REV 18 | 750-036844 | EJ4391 | FPC-PTX-P1-A        |
| FPC 7      | REV 18 | 750-036844 | EJ4382 | FPC-PTX-P1-A        |
| SIB 0      | REV 07 | 750-030631 | EG4858 | SIB-I-PTX5008       |
| SIB 1      | REV 07 | 750-030631 | EG4872 | SIB-I-PTX5008       |
| SIB 2      | REV 07 | 750-030631 | EG4866 | SIB-I-PTX5008       |
| SIB 3      | REV 07 | 750-030631 | EG6011 | SIB-I-PTX5008       |
| SIB 4      | REV 07 | 750-030631 | EG4907 | SIB-I-PTX5008       |
| SIB 5      | REV 07 | 750-030631 | EG4879 | SIB-I-PTX5008       |
| SIB 6      | REV 07 | 750-030631 | EG4864 | SIB-I-PTX5008       |
| SIB 7      | REV 07 | 750-030631 | EG4899 | SIB-I-PTX5008       |
| SIB 8      | REV 07 | 750-030631 | EG4880 | SIB-I-PTX5008       |
| Fan Tray 1 | REV 04 | 760-030642 | EG1335 | FAN-PTX-H-S         |

### show chassis hardware extensive (PTX5000 Packet Transport Switch)

```

user@switch> show chassis hardware extensive
Hardware inventory:
Item Version Part number Serial number Description
.....
PDU 0 Rev 04 740-032019 UE0003 DC Power Dist Unit
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 740-032019 S/N: S/N UE0003
Assembly ID: 0x043d Assembly Version: 04.00
Date: 11-29-2010 Assembly Flags: 0x00
Version: Rev 04 CLEI Code: 032022XXXX
ID: DC Power Dist Unit FRU Model Number: PWR-SAN-PDU-DC
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 3d 04 00 52 65 76 20 30 34 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 33 32 30 31 39 00 00
Address 0x20: 53 2f 4e 20 55 45 30 30 30 33 00 00 00 1d 0b 07
Address 0x30: da ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 30 33 32 30 32 32 58 58 58 58 50
Address 0x50: 57 52 2d 53 41 4e 2d 50 44 55 2d 44 43 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 a3 ff ff ff ff ff ff ff ff ff ff ff ff
PSM 0 Rev 04 740-032022 YG00065 DC 12V Power Supply
Module
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 740-032022 S/N: S/N YG00065
Assembly ID: 0x0440 Assembly Version: 04.00
Date: 07-30-2010 Assembly Flags: 0x00
Version: Rev 04 CLEI Code: 032022XXXX
ID: DC 12V Power Supply Module FRU Model Number: PWR-SAN-12-DC
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 40 04 00 52 65 76 20 30 34 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 33 32 30 32 32 00 00
Address 0x20: 53 2f 4e 20 59 47 30 30 30 36 35 00 00 1e 07 07
Address 0x30: da ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 30 33 32 30 32 32 58 58 58 58 50
Address 0x50: 57 52 2d 53 41 4e 2d 31 32 2d 44 43 20 20 20 20

```

```
Address 0x60: 20 20 20 20 20 20 01 00 ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff 0c ff ff ff ff ff ff ff ff ff ff ff
```

### show chassis hardware (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis hardware
Hardware inventory:
Item Version Part number Serial number Description
Chassis JN1100FB1AFB MX480
Midplane REV 05 710-017414 TR3310 MX480 Midplane
FPM Board REV 02 710-017254 KG1872 Front Panel Display
PEM 2 Rev 02 740-017343 QCS0812A00N DC Power Entry Module
PEM 3 Rev 02 740-017343 QCS0812A00U DC Power Entry Module
Routing Engine 0 REV 07 740-015113 1000740938 RE-S-1300
CB 0 REV 03 710-021523 KF4630 MX SCB
FPC 1 REV 11 750-037207 ZW9726 AS-MCC
 CPU REV 04 711-038173 ZW4819 AS-MCC PMB
 MIC 0 REV 06 750-037214 ZW3574 AS-MSC
 PIC 0 BUILTIN BUILTIN AS-MSC
 MIC 1 REV 00 750-037211 AS-MXC
 PIC 2 BUILTIN BUILTIN AS-MXC
```

### show chassis hardware extensive (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis hardware extensive
FPC 1 REV 11 750-037207 ZW9726 AS-MCC
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 750-037207 S/N: S/N ZW9726
Assembly ID: 0x0b37 Assembly Version: 01.11
Date: 02-17-2012 Assembly Flags: 0x00
Version: REV 11 CLEI Code: PROTOXCLEI
ID: AS-MCC FRU Model Number: 750-037207
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 37 01 0b 52 45 56 20 31 31 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 30 37 00 00
Address 0x20: 53 2f 4e 20 5a 57 39 37 32 36 00 00 00 11 02 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 35 30 2d 30 33 37 32 30 37 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 31 31 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 5e ff ff ff ff ff ff ff ff ff ff ff ff
CPU REV 04 711-038173 ZW4819 AS-MCC-PMB
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 711-038173 S/N: S/N ZW4819
Assembly ID: 0x0b38 Assembly Version: 01.04
Date: 12-30-2011 Assembly Flags: 0x00
Version: REV 04
ID: AS-MCC PMB
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 38 01 04 52 45 56 20 30 34 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 38 31 37 33 00 00
Address 0x20: 53 2f 4e 20 5a 57 34 38 31 39 00 00 00 1e 0c 07
Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 31 31 2d 30 33 38 31 37 33 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 30 34 00 ff ff ff ff ff ff ff
```

```

Address 0x70: ff ff ff 60 00 00 00 00 00 00 00 00 00 00 00
MIC 0 REV 06 750-037214 ZW3574 AS-MSC
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 750-037214 S/N: S/N ZW3574
Assembly ID: 0x0a44 Assembly Version: 01.06
Date: 02-19-2012 Assembly Flags: 0x00
Version: REV 06 CLEI Code: PROTOXCLEI
ID: AS-MSC FRU Model Number: 750-037214
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0a 44 01 06 52 45 56 20 30 36 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 31 34 00 00
Address 0x20: 53 2f 4e 20 5a 57 33 35 37 34 00 00 00 13 02 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 35 30 2d 30 33 37 32 31 34 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 30 36 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 60 c0 03 e5 f4 00 00 00 00 00 00 00 00
PIC 0 BUILTIN BUILTIN AS-MSC
MIC 1 REV 00 750-037211 AS-MXC
Jedec Code: 0x7fb0 EEPROM Version: 0x01
P/N: 750-037211
Assembly ID: 0x0a43 Assembly Version: 01.00
Date: 255-255-65535 Assembly Flags: 0x00
Version: REV 00
ID: AS-MXC
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 0a 43 01 00 52 45 56 20 30 30 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 31 31 00 00
Address 0x20: 00 00 00 00 00 00 00 00 00 00 00 00 00 ff ff ff
Address 0x30: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff c0 02 e6 6c 7f b0 02 ff 0a 44 01 06
PIC 2 BUILTIN BUILTIN AS-MXC

```

## show chassis pic

|                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                                               | show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Syntax (TX Matrix and TX Matrix Plus Routers)        | show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i> <lcc <i>number</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Syntax (MX Series Routers and EX Series Switches)    | show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i> <all-members> <local> <member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Syntax (MX2010 and MX2010 3D Universal Edge Routers) | show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Syntax (QFX Series)                                  | show chassis pic <interconnect-device <i>name</i> (fpc-slot <i>slot-number</i>   pic-slot <i>slot-number</i> )> <node-device <i>name</i> pic-slot <i>slot-number</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Syntax (ACX Series Universal Access Routers)         | show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Release Information                                  | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Description                                          | Display status information about the PIC installed in the specified Flexible PIC Concentrator (FPC) and PIC slot.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Options                                              | <p><b>fpc-slot <i>slot-number</i></b>—Display information about the PIC in this particular FPC slot:</p> <ul style="list-style-type: none"> <li>On a TX Matrix router, if you specify the number of the T640 router by using the <b>lcc <i>number</i></b> option (the recommended method), replace <b><i>slot-number</i></b> with a value from 0 through 7. Otherwise, replace <b><i>slot-number</i></b> with a value from 0 through 31.</li> </ul> <p>Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 router by using the <b>lcc <i>number</i></b> option (the recommended method), replace <b><i>slot-number</i></b> with a value from 0 through 7. Otherwise, replace <b><i>slot-number</i></b> with a value from 0 through 31. For example, the following commands have the same result:</p> <pre> user@host&gt; show chassis pic fpc-slot 1 lcc 1 pic-slot 1 user@host&gt; show chassis pic fpc-slot 9 pic-slot 1 </pre> <ul style="list-style-type: none"> <li>M120 routers only—Replace <b><i>slot-number</i></b> with a value from 0 through 5.</li> <li>MX80 routers only—Replace <b><i>slot-number</i></b> with a value from 0 through 1.</li> <li>MX240 routers only—Replace <b><i>slot-number</i></b> with a value from 0 through 2.</li> </ul> |

- MX480 routers only—Replace **slot-number** with a value from 0 through 5.
- MX960 routers only—Replace **slot-number** with a value from 0 through 11.
- MX2020 routers only—Replace **slot-number** with a value from 0 through 19.
- Other routers—Replace **slot-number** with a value from 0 through 7.
- EX Series switches:
  - EX3200 switches and EX4200 standalone switches—Replace **slot-number** with 0.
  - EX4200 switches in a Virtual Chassis configuration—Replace **slot-number** with a value from 0 through 9 (switch's member ID).
  - EX8208 switches—Replace **slot-number** with a value from 0 through 7 (line card).
  - EX8216 switches—Replace **slot-number** with a value from 0 through 15 (line card).
- QFX Series:
  - QFX3500 switches—Replace **slot-number** with 0. In the command output, FPC refers to a line card. The FPC number equals the slot number for the line card.
  - QFabric systems—Replace **slot-number** with any number between 0 and 15. In the command output, FPC refers to a line card. The FPC number equals the slot number for the line card.

**all-members**—(MX Series routers and EX Series switches only) (Optional) Display PIC information for all member routers in the Virtual Chassis configuration.

**interconnect-device name**—(QFabric systems only) (Optional) Display PIC information for a specified Interconnect device.

**lcc-number**—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display PIC information for a specified T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display PIC information for a specified T1600 router (or line-card chassis) that is connected to the TX Matrix Plus router. Replace **number** with a value from 0 through 3.

**local**—(MX Series routers and EX Series switches only) (Optional) Display PIC information for the local Virtual Chassis member.

**member member-id**—(MX Series routers and EX Series switches only) (Optional) Display PIC information for the specified member of the Virtual Chassis configuration. Replace **member-id** with a value of 0 or 1.

**node-device name**—(QFabric systems only) (Optional) Display PIC information for a specified Node device.

**pic-slot slot-number**—Display information about the PIC in this particular PIC slot. For routers, replace **slot-number** with a value from 0 through 3. For EX3200 and EX4200



switches, replace *slot-number* with 0 for built-in network interfaces and 1 for interfaces on uplink modules. For EX8208 and EX8216 switches, replace *slot-number* with 0. For the QFX3500 standalone switch and the QFabric system, replace *slot-number* with 0 or 1.

**Required Privilege Level** view

**Related Documentation**

- *request chassis pic*
- [show chassis hardware on page 644](#)
- *Configuring the PIC Type*
- *100-Gigabit Ethernet PIC Overview*

**List of Sample Output**

- [show chassis pic fpc-slot pic-slot on page 4437](#)
- [show chassis pic fpc-slot pic-slot \(PIC Offline\) on page 4438](#)
- [show chassis pic fpc-slot pic-slot \(FPC Offline\) on page 4438](#)
- [show chassis pic fpc-slot pic-slot \(FPC Not Present\) on page 4438](#)
- [show chassis pic fpc-slot pic-slot \(PIC Not Present\) on page 4438](#)
- [show chassis pic fpc-slot 3 pic-slot 0 \(M120 Router\) on page 4438](#)
- [show chassis pic fpc-slot pic-slot \(MX960 Router with Bidirectional Optics\) on page 4438](#)
- [show chassis pic fpc-slot pic-slot \(MX480 Router with 100-Gigabit Ethernet MIC\) on page 4439](#)
- [show chassis pic fpc-slot pic-slot \(MX240, MX480, MX960 Routers with Application Services Modular Line Card\) on page 4439](#)
- [show chassis pic fpc-slot pic-slot \(MX480 Router with MPC4E\) on page 4439](#)
- [show chassis pic fpc-slot pic-slot \(MX2010 Router\) on page 4439](#)
- [show chassis pic fpc-slot pic-slot \(MX2020 Router\) on page 4439](#)
- [show chassis pic fpc-slot pic-slot \(MX2020 Router with MPC4E\) on page 4440](#)
- [show chassis pic fpc-slot pic-slot \(T1600 Router with 100-Gigabit Ethernet PIC\) on page 4440](#)
- [show chassis pic fpc-slot pic-slot lcc \(TX Matrix Router\) on page 4440](#)
- [show chassis pic fpc-slot pic-slot lcc \(TX Matrix Plus Router\) on page 4441](#)
- [show chassis pic fpc-slot pic-slot \(Next-Generation SONET/SDH SFP\) on page 4441](#)
- [show chassis pic fpc-slot pic-slot \(12-Port T1/E1\) on page 4441](#)
- [show chassis pic fpc-slot 0 pic-slot 1 \(4x CHOC3 SONET CE SFP\) on page 4441](#)
- [show chassis pic fpc-slot 0 pic-slot 0 \(SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 4442](#)
- [show chassis pic fpc-slot 3 pic-slot 0 \(8-port Channelized SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 4442](#)
- [show chassis pic fpc-slot 5 pic-slot 0 \(4-port Channelized SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 4442](#)
- [show chassis pic fpc-slot 1 pic-slot 0 \(1-port OC192/STM64 MIC with XFP\) on page 4443](#)
- [show chassis pic fpc-slot 1 pic-slot 2 \(8-port DS3/E3 MIC\) on page 4443](#)
- [show chassis pic fpc-slot pic-slot \(OTN\) on page 4443](#)
- [show chassis pic fpc-slot pic-slot \(QFX3500 Switch\) on page 4443](#)
- [show chassis pic interconnect-device fpc-slot pic-slot \(QFabric Systems\) on page 4443](#)
- [show chassis pic node-device fpc-slot pic-slot \(QFabric System\) on page 4443](#)

[show chassis pic fpc-slot 0 pic-slot 1 \(ACX2000 Universal Access Router\) on page 4444](#)  
[show chassis pic FPC-slot 1 PIC-slot 0 \(MX Routers with Media Services Blade \[MSB\]\) on page 4445](#)  
[show chassis pic FPC slot 1, PIC slot 2 \(MX Routers with Media Services Blade \[MSB\]\) on page 4445](#)

**Output Fields** Table 62 on page 757 lists the output fields for the **show chassis pic** command. Output fields are listed in the approximate order in which they appear.

**Table 321: show chassis pic Output Fields**

| Field Name                                                   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Type                                                         | <p>PIC type.</p> <p><b>NOTE:</b> On the 1-port OC192/STM64 MICs with the SDH framing mode, the type is displayed as <b>MIC-3D-1STM64-XFP</b> and with the SONET framing mode, the type is displayed as <b>MIC-3D-1OC192-XFP</b>. By default, the 1-port OC192/STM64 MICs displays the type as <b>MIC-3D-1OC192-XFP</b>.</p>                                                                                                                                                                                                                                       |
| ASIC type                                                    | Type of ASIC on the PIC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| State                                                        | <p>Status of the PIC. State is displayed only when a PIC is in the slot.</p> <ul style="list-style-type: none"> <li>• <b>Online</b>— PIC is online and running.</li> <li>• <b>Offline</b>—PIC is powered down.</li> </ul>                                                                                                                                                                                                                                                                                                                                         |
| PIC version                                                  | PIC hardware version.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Uptime                                                       | How long the PIC has been online.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Package                                                      | (Multiservices PICs only) Services package supported: <b>Layer-2</b> or <b>Layer-3</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Port Number                                                  | Port number for the PIC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Cable Type                                                   | Type of cable connected to the port: <b>LH</b> , <b>LX</b> , or <b>SX</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| PIC Port Information (MX480 Router 100-Gigabit Ethernet CFP) | <p>Port-level information for the PIC.</p> <ul style="list-style-type: none"> <li>• Port—Port number</li> <li>• Cable type—Type of optical transceiver installed.</li> <li>• Fiber type—Type of fiber. SM is single-mode.</li> <li>• Xcvr vendor—Transceiver vendor name.</li> <li>• Xcvr vendor part number—Transceiver vendor part number.</li> <li>• Wavelength—Wavelength of the transmitted signal. Uplinks and downlinks are always 1550 nm. There is a separate fiber for each direction</li> <li>• Xcvr Firmware—Transceiver firmware version.</li> </ul> |

Table 321: show chassis pic Output Fields (*continued*)

| Field Name                                                               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>PIC Port Information<br/>(MX960 Router<br/>Bidirectional Optics )</b> | <p>Port-level information for the PIC.</p> <ul style="list-style-type: none"> <li>• Port—Port number</li> <li>• Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed. Uplink interfaces display -U. Down link interfaces display -D.</li> <li>• Fiber type—Type of fiber. SM is single-mode.</li> <li>• Xcvr vendor—Transceiver vendor name.</li> <li>• Xcvr vendor part number—Transceiver vendor part number. <ul style="list-style-type: none"> <li>• BX10-10-km bidirectional optics.</li> <li>• BX40-40-km bidirectional optics.</li> <li>• SFP-LX-40-km SFP optics.</li> </ul> </li> <li>• Wavelength—Wavelength of the transmitted signal. Uplinks are always 1310 nm. Downlinks are either 1490 nm or 1550 nm.</li> </ul> |
| <b>PIC Port Information<br/>(Next-Generation<br/>SONET/SDH SFP)</b>      | <p>Port-level information for the next-generation SONET/SDH SFP PIC.</p> <ul style="list-style-type: none"> <li>• Port—Port number.</li> <li>• Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed.</li> <li>• Fiber type—Type of fiber: <b>SM</b> (single-mode) or <b>MM</b> (multimode).</li> <li>• Xcvr vendor—Transceiver vendor name.</li> <li>• Xcvr vendor part number—Transceiver vendor part number.</li> <li>• Wavelength—Wavelength of the transmitted signal. Next-generation SONET/SDH SFPs use 1310 nm.</li> </ul>                                                                                                                                                                                                 |
| <b>Multirate Mode</b>                                                    | Rate-selectability status for the MIC: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Channelization</b>                                                    | Indicates whether channelization is enabled or disabled on the DS3/E3 MIC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## Sample Output

### show chassis pic fpc-slot pic-slot

```

user@host> show chassis pic fpc-slot 2 pic-slot 0
PIC fpc slot 2 pic slot 0 information:
 Type 10x 1GE(LAN), 1000 BASE
 ASIC type H chip
 State Online
 PIC version 1.1
 Uptime 1 day, 50 minutes, 58 seconds
PIC Port Information:
 Port Cable Xcvr Xcvr Vendor
 Number Type Vendor Name Part Number
 0 GIGE 1000EX FINISAR CORP. FTRJ8519P1BNL-J3
 1 GIGE 1000EX FINISAR CORP. FTRJ-8519-7D-JUN

```

**show chassis pic fpc-slot pic-slot (PIC Offline)**

```

user@host> show chassis pic fpc-slot 1 pic-slot 0
PIC fpc slot 1 pic slot 0 information:
 State Offline

```

**show chassis pic fpc-slot pic-slot (FPC Offline)**

```

user@host> show chassis pic fpc-slot 1 pic-slot 0
FPC 1 is not online

```

**show chassis pic fpc-slot pic-slot (FPC Not Present)**

```

user@host> show chassis pic fpc-slot 4 pic-slot 0
FPC slot 4 is empty

```

**show chassis pic fpc-slot pic-slot (PIC Not Present)**

```

user@host> show chassis pic fpc-slot 5 pic-slot 2
FPC 5, PIC 2 is empty

```

**show chassis pic fpc-slot 3 pic-slot 0 (M120 Router)**

```

user@host> show chassis pic fpc-slot 3 pic-slot 0
PIC slot 3, PIC slot 0 information:
 Type 2x G/E IQ, 1000 BASE
 ASIC type IQ GE 2 VLAN-TAG FPGA
 State Online
 PIC version 1.16
 Uptime 3 hours, 3 minutes

PIC Port Information:
 Port Cable Xcvr Xcvr Vendor
 Number Type Vendor Name Part Number
 0 GIGE 1000SX FINISAR CORP. FTRJ8519P1BNL-J3
 1 GIGE 1000SX FINISAR CORP. FTRJ-8519-7D-JUN

```

**show chassis pic fpc-slot pic-slot (MX960 Router with Bidirectional Optics)**

```

user@host> show chassis pic fpc-slot 4 pic-slot 1
FPC slot 4, PIC slot 1 information:
 Type 10x 1GE(LAN)
 State Online
 PIC version 0.0
 Uptime 18 days, 5 hours, 41 minutes, 54 seconds

PIC port information:

```

| Port | Cable type          | Fiber type | Xcvr vendor      | Xcvr vendor part number | Wavelength |
|------|---------------------|------------|------------------|-------------------------|------------|
| 0    | SFP-1000BASE-BX10-D | SM         | SumitomoElectric | SBP6H44-J3-BW-49        | 1490 nm    |
| 1    | SFP-1000BASE-BX10-D | SM         | SumitomoElectric | SBP6H44-J3-BW-49        | 1490 nm    |
| 2    | SFP-1000BASE-BX10-D | SM         | SumitomoElectric | SBP6H44-J3-BW-49        | 1490 nm    |
| 3    | SFP-1000BASE-BX10-D | SM         | OCP              | TRXBG1LXDBVM2-JW        | 1490 nm    |
| 4    | SFP-1000BASE-BX10-D | SM         | OCP              | TRXBG1LXDBVM2-JW        | 1490 nm    |
| 5    | SFP-1000BASE-BX10-U | SM         | SumitomoElectric | SBP6H44-J3-BW-31        | 1310 nm    |
| 6    | SFP-1000BASE-BX10-U | SM         | SumitomoElectric | SBP6H44-J3-BW-31        | 1310 nm    |
| 7    | SFP-1000BASE-BX10-U | SM         | OCP              | TRXBG1LXDBBMH-J1        | 1310 nm    |
| 8    | SFP-1000BASE-BX10-U | SM         | OCP              | TRXBG1LXDBBMH-J1        | 1310 nm    |
| 9    | SFP-1000BASE-BX10-U | SM         | SumitomoElectric | SBP6H44-J3-BW-31        | 1310 nm    |

**show chassis pic fpc-slot pic-slot (MX480 Router with 100-Gigabit Ethernet MIC)**

```

user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
 Type 1X100GE CFP
 State Online
 PIC version 2.10
 Uptime 4 minutes, 48 seconds

PIC port information:
 Fiber
 Port Cable type type Xcvr vendor part number Wavelength
 0 100GBASE LR4 SM FINISAR CORP. FTLC1181RDN5-J3 1310 nm

 Xcvr vendor
 firmware version
 1.8

```

**show chassis pic fpc-slot pic-slot (MX240, MX480, MX960 Routers with Application Services Modular Line Card)**

```

user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
 Type AS-MXC
 State Online
 PIC version 1.0
 Uptime 11 hours, 18 minutes, 3 seconds

```

**show chassis pic fpc-slot pic-slot (MX480 Router with MPC4E)**

```

user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 3, PIC slot 0 information:
 Type 4x10GE SFPP
 State Online
 PIC version 0.0
 Uptime 41 seconds

PIC port information:
 Fiber
 Port Cable type type Xcvr vendor part number Wave- Xcvr
 0 10GBASE SR MM OPNEXT, INC. TRS2001EM-0014 850 nm 0.0
 1 10GBASE SR MM OPNEXT, INC. TRS2001EM-0014 850 nm 0.0

```

**show chassis pic fpc-slot pic-slot (MX2010 Router)**

```

user@host> show chassis pic fpc-slot 9 pic-slot 3
FPC slot 9, PIC slot 3 information:
 Type 1X100GE CFP
 State Online
 PIC version 0.0
 Uptime 14 hours, 51 seconds

```

**show chassis pic fpc-slot pic-slot (MX2020 Router)**

```

user@host> show chassis pic fpc-slot 19 pic-slot 3
FPC slot 19, PIC slot 3 information:
 Type 4x 10GE(LAN) SFP+

```

```

State Online
PIC version 0.0
Uptime 1 day, 11 hours, 26 minutes, 36 seconds

PIC port information:

```

| Port | Cable type | Fiber type | Xcvr vendor      | part number    | Wave-length | Xcvr |
|------|------------|------------|------------------|----------------|-------------|------|
| 0    | 10GBASE SR | MM         | SumitomoElectric | SPP5200SR-J6-M | 850 nm      | 0.0  |
| 1    | 10GBASE SR | MM         | SumitomoElectric | SPP5200SR-J6-M | 850 nm      | 0.0  |
| 2    | 10GBASE SR | MM         | SumitomoElectric | SPP5200SR-J6-M | 850 nm      | 0.0  |
| 3    | 10GBASE SR | MM         | SumitomoElectric | SPP5200SR-J6-M | 850 nm      | 0.0  |

#### show chassis pic fpc-slot pic-slot (MX2020 Router with MPC4E)

```

user@host> show chassis pic fpc-slot 14 pic-slot 0
FPC slot 14, PIC slot 1 information:
Type 1X100GE CFP
State Online
PIC version 0.0
Uptime 1 day, 2 hours, 19 minutes, 18 seconds

PIC port information:

```

| Port | Cable type    | Fiber type | Xcvr vendor      | part number      | Wave-length | Xcvr |
|------|---------------|------------|------------------|------------------|-------------|------|
| 0    | 100GBASE SR10 | MM         | Reflex Photonics | CF-X12-C11801-50 | 860 nm      | 4.7  |

#### show chassis pic fpc-slot pic-slot (T1600 Router with 100-Gigabit Ethernet PIC)

```

user@host> run show chassis pic fpc-slot 3 pic-slot 1
FPC slot 3, PIC slot 1 information:
Type 100GE SLOT1
ASIC type Brooklyn 100GE FPGA
State Online
PIC version 1.3
Uptime 10 minutes, 44 seconds

PIC port information:

```

| Port | Cable type   | Fiber type | Xcvr vendor | part number      | Wavelength |
|------|--------------|------------|-------------|------------------|------------|
| 0    | 100GBASE LR4 | SM         | Opnext Inc. | TRC5E20ENFSF000F | 1310 nm    |

#### show chassis pic fpc-slot pic-slot lcc (TX Matrix Router)

```

user@host> show chassis pic fpc-slot 1 pic-slot 1 lcc 0
lcc0-re0:

PIC fpc slot 1 pic slot 1 information:
Type 4x OC-3 SONET, SMIR
ASIC type D chip
State Online

```

```

PIC version 1.2
Uptime 5 days, 2 hours, 12 minutes, 8 seconds

```

#### show chassis pic fpc-slot pic-slot lcc (TX Matrix Plus Router)

```

user@host> show chassis pic pic-slot 0 fpc-slot 8
lcc0-re0:

FPC slot 8, PIC slot 0 information:
Type 1x 10GE(LAN/WAN)
State Online
Uptime 2 hours, 46 minutes, 23 seconds

PIC port information:

Port Cable type Fiber
 type Xcvr vendor part number Wavelength
0 10GBASE ZR SM Opnext Inc. TRF7061BN-LF150 1550 nm
0 10GBASE ZR SM FINISAR CORP. FTRX-1811-3-J2 1550 nm

```

#### show chassis pic fpc-slot pic-slot (Next-Generation SONET/SDH SFP)

```

user@host> show chassis pic fpc-slot 4 pic-slot 0
FPC slot 4, PIC slot 0 information:
Type 4x OC-3 1x OC-12 SFP
ASIC type D FPGA
State Online
PIC version 1.3
Uptime 1 day, 50 minutes, 4 seconds

PIC port information:

Port Cable type Fiber
 type Xcvr vendor part number Wavelength
0 OC48 short reach SM FINISAR CORP. FTRJ1321P1BTL-J2 1310 nm
1 OC3 short reach MM OCP TRPA03MM3BAS-JE 1310 nm
2 OC3 short reach MM OCP TRXA03MM3BAS-JW 1310 nm
3 OC12 inter reach SM FINISAR CORP. FTLF1322P1BTR 1310 nm

```

#### show chassis pic fpc-slot pic-slot (12-Port T1/E1)

```

user@host> show chassis pic fpc-slot 0 pic-slot 3
FPC slot 0, PIC slot 3 information:
Type 12x T1/E1 CE
State Online
PIC version 1.1
CPU load average 1 percent
Interrupt load average 0 percent
Total DRAM size 128 MB
Memory buffer utilization 100 percent
Memory heap utilization 4 percent
Uptime 1 day, 22 hours, 28 minutes, 12 seconds
Internal Clock Synchronization Normal

```

#### show chassis pic fpc-slot 0 pic-slot 1 (4x CHOC3 SONET CE SFP)

```

user@host> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
Type 4x CHOC3 SONET CE SFP
State Online
PIC version 1.3
CPU load average 1 percent

```

```

Interrupt load average 0 percent
Total DRAM size 128 MB
Memory buffer utilization 99 percent
Memory heap utilization 4 percent
Uptime 1 day, 22 hours, 55 minutes, 37 seconds
Internal Clock Synchronization Normal

```

PIC port information:

| Port | Cable type      | Fiber type | Xcvr vendor | Xcvr vendor part number | Wavelength |
|------|-----------------|------------|-------------|-------------------------|------------|
| 0    | OC3 short reach | MM         | AVAGO       | HFBR-57E0P-JU2          | n/a        |
| 1    | OC3 short reach | MM         | AVAGO       | HFBR-57E0P-JU2          | n/a        |
| 3    | OC3 long reach  | SM         | OPNEXT INC  | TRF5456AVLB314          | 1310 nm    |

**show chassis pic fpc-slot 0 pic-slot 0 (SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)**

```
user@host> show chassis pic fpc-slot 0 pic-slot 0
```

FPC slot 0, PIC slot 0 information:

```

Type MIC-3D-80C30C12-40C48
State Online
PIC version 1.8
Uptime 3 days, 22 hours, 3 minutes, 50 seconds

```

PIC port information:

| Port | Cable type       | Fiber type | Xcvr vendor  | Xcvr vendor part number | Wavelength |
|------|------------------|------------|--------------|-------------------------|------------|
| 1    | OC12 inter reach | SM         | FINISAR CORP | FTRJ1322P1BTR-J3        | 1310 nm    |
| 7    | OC12 inter reach | SM         | FINISAR CORP | FTRJ1322P1BTR-J3        | 1310 nm    |

Multirate Mode Enabled

**show chassis pic fpc-slot 3 pic-slot 0 (8-port Channelized SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)**

```
user@host> show chassis pic fpc-slot 3 pic-slot 0
```

FPC slot 3, PIC slot 0 information:

```

Type MIC-3D-8CHOC3-4CHOC12
State Online
PIC version 1.9
Uptime 1 hour, 21 minutes, 24 seconds

```

PIC port information:

| Port | Cable type       | Fiber type | Xcvr vendor   | Xcvr vendor part number | Wavelength |
|------|------------------|------------|---------------|-------------------------|------------|
| 0    | OC12 short reach | SM         | FINISAR CORP. | FTRJ1322P1BTR-J3        | 1310 nm    |
| 1    | OC12 short reach | SM         | FINISAR CORP. | FTRJ1322P1BTR-J3        | 1310 nm    |
| 2    | OC12 inter reach | SM         | FINISAR CORP. | FTRJ1322P1BTR-J2        | 1310 nm    |
| 4    | OC12 short reach | SM         | FINISAR CORP. | FTRJ1322P1BTR-J3        | 1310 nm    |
| 5    | OC12 short reach | SM         | FINISAR CORP. | FTRJ1322P1BTR-J3        | 1310 nm    |
| 6    | OC12 short reach | SM         | FINISAR CORP. | FTRJ1322P1BTR-J3        | 1310 nm    |
| 7    | OC12 short reach | SM         | FINISAR CORP. | FTRJ1322P1BTR-J3        | 1310 nm    |

**show chassis pic fpc-slot 5 pic-slot 0 (4-port Channelized SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)**

```
user@host> show chassis pic fpc-slot 5 pic-slot 0
```

FPC slot 5, PIC slot 0 information:

```

Type MIC-3D-4CHOC3-2CHOC12
State Online
PIC version 1.9
Uptime 1 hour, 21 minutes

```

PIC port information:

| Port | Cable type | Fiber type | Xcvr vendor | Xcvr vendor part number | Wavelength |
|------|------------|------------|-------------|-------------------------|------------|
|------|------------|------------|-------------|-------------------------|------------|



|   |                  |    |               |                  |         |
|---|------------------|----|---------------|------------------|---------|
| 1 | OC12 inter reach | SM | FINISAR CORP. | FTRJ1322P1BTR-J3 | 1310 nm |
| 2 | OC12 inter reach | SM | FINISAR CORP. | FTRJ1322P1BTR-J3 | 1310 nm |
| 3 | OC12 short reach | SM | FINISAR CORP. | FTRJ1322P1BTR-J3 | 1310 nm |

#### show chassis pic fpc-slot 1 pic-slot 0 (1-port OC192/STM64 MIC with XFP)

```

user@host> show chassis pic fpc-slot 1 pic-slot 0
FPC slot 1, PIC slot 0 information:
 Type MIC-3D-10C192-XFP
 State Online
 PIC version 1.2
 Uptime 1 day, 11 hours, 4 minutes, 6 seconds

PIC port information:
 Port Cable type Fiber type Xcvr vendor Xcvr vendor part number Wavelength
 0 OC192 short reach n/a FINISAR CORP. FTLX1412M3BCL-J3 1310 nm

```

#### show chassis pic fpc-slot 1 pic-slot 2 (8-port DS3/E3 MIC)

```

user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
 Type MIC-3D-8DS3-E3
 State Online
 PIC version 1.10
 Uptime 4 days, 1 hour, 29 minutes, 19 seconds
 Channelization Mode Disabled

```

#### show chassis pic fpc-slot pic-slot (OTN)

```

user@host> show chassis pic fpc-slot 5 pic-slot 0
PIC fpc slot 5 pic slot 0 information:
 Type 1x10GE(LAN),OTN
 ASIC type H chip
 State Online
 PIC version 1.0
 Uptime 5 minutes, 50 seconds

```

#### show chassis pic fpc-slot pic-slot (QFX3500 Switch)

```

user@switch> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
 Type 48x 10G-SFP+ Builtin
 State Online
 Uptime 3 days, 3 hours, 5 minutes, 20 seconds

```

#### show chassis pic interconnect-device fpc-slot pic-slot (QFabric Systems)

```

user@switch> show chassis pic interconnect-device interconnect1 fpc-slot 9 pic-slot 0
FPC slot 9, PIC slot 0 information:
 Type 16x 40G-GE Builtin
 State Online
 Uptime 2 hours, 47 minutes, 40 seconds

```

#### show chassis pic node-device fpc-slot pic-slot (QFabric System)

```

user@switch> show chassis pic node-device node1 pic-slot 0
FPC slot node1, PIC slot 0 information:
 Type 48x 10G-SFP+ Builtin
 State Online
 Uptime 2 hours, 52 minutes, 37 seconds

```

## PIC port information:

| Port | Cable type | Fiber type | Xcvr vendor      | Xcvr vendor part number | Wavelength |
|------|------------|------------|------------------|-------------------------|------------|
| 0    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 1    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 2    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 3    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 4    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 5    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 6    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 7    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 8    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 9    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 10   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 11   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 12   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 13   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 14   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 15   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 16   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 17   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 18   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 19   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 20   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 21   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 22   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 23   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 24   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 25   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 26   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 27   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 28   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 29   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 30   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 31   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 32   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 33   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 34   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 35   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 36   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 37   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 38   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 39   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 40   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 41   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 42   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 43   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 44   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 45   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 46   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 47   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |

## show chassis pic fpc-slot 0 pic-slot 1 (ACX2000 Universal Access Router)

```

user@host> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
 Type 8x 1GE(LAN) RJ45 Builtin
 State Online
 Uptime 6 days, 2 hours, 51 minutes, 11 seconds

```

**show chassis pic FPC-slot 1 PIC-slot 0 (MX Routers with Media Services Blade [MSB])**

```
user@switch> show chassis pic fpc-slot 1 pic-slot 0
FPC slot 1, PIC slot 0 information:
 Type AS-MS
 State Online
 PIC version 1.6
 Uptime 11 hours, 17 minutes, 56 seconds
```

**show chassis pic FPC slot 1, PIC slot 2 (MX Routers with Media Services Blade [MSB])**

```
user@switch> show chassis pic fpc-slot 1 pic-slot 2
 Type AS-MXC
 State Online
 PIC version 1.0
 Uptime 11 hours, 18 minutes, 3 seconds
```

**Operational Mode Commands for Layer 2 Port-Mirroring Instances**

## show forwarding-options port-mirroring

|                                 |                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show forwarding-options port-mirroring</b><br><b>&lt;terse   detail&gt;</b><br><b>&lt;instance-name&gt;</b>                                                                                        |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.                                                                                |
| <b>Description</b>              | Display current state of port-mirroring instances.                                                                                                                                                    |
| <b>Options</b>                  | <b>terse   detail</b> —(Optional) Display the specified level of output.<br><br><b>instance-name</b> —(Optional) Display a single port-mirroring instance.                                            |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show forwarding-options next-hop-group on page 4460</a></li> </ul>                                                                               |
| <b>List of Sample Output</b>    | <a href="#">show forwarding-options port-mirroring terse on page 4447</a><br><a href="#">show forwarding-options port-mirroring detail on page 4447</a>                                               |
| <b>Output Fields</b>            | <a href="#">Table 322 on page 4446</a> lists the output fields for the <b>show forwarding-options port-mirroring</b> command. Output fields are listed in the approximate order in which they appear. |

**Table 322: show forwarding-options port-mirroring Output Fields**

| Field Name               | Field Description                                   | Level of Output |
|--------------------------|-----------------------------------------------------|-----------------|
| Instance Name            | Name of port-mirroring instance.                    | All levels      |
| Instance Id              | Instance identification number.                     | All levels      |
| State                    | Instance state, either <b>up</b> or <b>down</b> .   | All levels      |
| <b>Input parameters</b>  |                                                     |                 |
| Rate                     | Rate (ratio of packets sampled).                    | <b>detail</b>   |
| Run-length               | Run length (number of consecutive packets sampled). | <b>detail</b>   |
| Maximum-packet-length    | Maximum packet length.                              | <b>detail</b>   |
| <b>Output parameters</b> |                                                     |                 |
| Family                   | Protocol family.                                    | <b>detail</b>   |
| State                    | Instance state, either <b>up</b> or <b>down</b> .   | <b>detail</b>   |
| Destination              | Destination (next-hop group name).                  | <b>detail</b>   |

## Sample Output

### show forwarding-options port-mirroring terse

```
user@host> show forwarding-options port-mirroring terse
Instance Name Instance Id State
&global_instance 1 up
inst1 2 up
```

### show forwarding-options port-mirroring detail

```
user@host> show forwarding-options port-mirroring detail
Instance Name: &global_instance
Instance Id: 1 State: up
 Input parameters:
 Rate: 10
 Run-length: 4
 Maximum-packet-length: 0
 Output parameters:
 Family: inet State: up Destination: inet_nhg
 Family: vpls/eth-switch State: up Destination: vpls_nhg

Instance Name: inst1
Instance Id: 2 State: up
 Input parameters:
 Rate: 1
 Run-length: 0
 Maximum-packet-length: 200
 Output parameters:
 Family: inet State: up Destination: inet_nhg
 Family: vpls/eth-switch State: down Destination: vpls_nhg_2
```

## Operational Mode Commands for Firewall Filter Statistics and Logs

## clear firewall

---

|                                    |                                                                                                                                                                                                                                                         |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | clear firewall (all   counter <i>counter-name</i>   filter <i>filter-name</i>   log (all   <i>logical-system-name</i> )   logical-system <i>logical-system-name</i> )                                                                                   |
| <b>Syntax (EX Series Switches)</b> | clear firewall (all   counter <i>counter-name</i>   filter <i>filter-name</i>   log (all   <i>logical-system-name</i> )   policer counter (all   counter-id <i>counter-index</i> ))                                                                     |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>logical-system</b> option introduced in Junos OS Release 9.3.<br><b>log</b> option introduced before Junos OS Release 11.4. |
| <b>Description</b>                 | Clear statistics about configured firewall filters.<br><br>When you clear the counters of a filter, this impacts not only the counters shown by the CLI, but also the ones tracked by SNMP2.                                                            |



**NOTE:** The clear firewall command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for graceful Routing Engine switchover (GRES).

If you clear statistics for firewall filters that are applied to Trio-based DPCs and that also use the **prefix-action** action on matched packets, wait at least 5 seconds before you enter the **show firewall prefix-action-stats** command. A 5-second pause between issuing the **clear firewall** and **show firewall prefix-action-stats** commands avoids a possible timeout of the **show firewall prefix-action-stats** command.

- Options**
- all**—Clear the packet and byte counts for all filters. On EX Series switches, this option also clears the packet counts for all policer counters.
  - counter *counter-name***—Clear the packet and byte counts for a filter counter that has been configured with the counter firewall filter action.
  - filter *filter-name***—Clear the packet and byte counts for the specified firewall filter.
  - log (all | *logical-system-name*)**—Clear log entries for IPv4 firewall filters that have **then log** as an action. Use **log all** to clear all log entries or **log *logical-system-name*** to clear log entries for the specified logical system.
  - logical-system *logical-system-name***—Clear the packet and byte counts for the specified logical system.
  - policer counter (all | counter-id *counter-index*)**—(EX8200 switches only) Clear all policer counters using the **policer counter all** command, or clear a specific policer counter using the **policer counter counter-id *counter-index*** command. The value of *counter-index* can be 0, 1, or 2.

|                          |                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Required Privilege Level | clear                                                                                                                                                                                                                                                                                                                                                                                    |
| Related Documentation    | <ul style="list-style-type: none"> <li>• <a href="#">show firewall on page 4450</a></li> </ul>                                                                                                                                                                                                                                                                                           |
| List of Sample Output    | <a href="#">clear firewall all on page 4449</a><br><a href="#">clear firewall (counter counter-name) on page 4449</a><br><a href="#">clear firewall (filter filter-name) on page 4449</a><br><a href="#">clear firewall (policer counter all) (EX8200 Switch) on page 4449</a><br><a href="#">clear firewall (policer counter counter-id counter-index) (EX8200 Switch) on page 4449</a> |

## Sample Output

clear firewall all

```
user@host> clear firewall all
```

clear firewall (counter counter-name)

```
user@host> clear firewall counter port-filter-counter
```

clear firewall (filter filter-name)

```
user@host> clear firewall filter ingress-port-filter
```

clear firewall (policer counter all) (EX8200 Switch)

```
user@switch> clear firewall policer counter all
```

clear firewall (policer counter counter-id counter-index) (EX8200 Switch)

```
user@switch> clear firewall policer counter counter-id 0
```

## show firewall

---

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | <pre>show firewall   &lt;counter <i>counter-name</i>&gt;   &lt;filter <i>filter-name</i>&gt;   &lt;log&gt;   &lt;logical-system (all   <i>logical-system-name</i>)&gt;   &lt;terse&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax (EX Series Switches)</b> | <pre>show firewall   &lt;counter <i>counter-name</i>&gt;   &lt;detail&gt;   &lt;filter <i>filter-name</i>&gt;   &lt;log &lt;(detail   interface <i>interface-name</i>)&gt;&gt;   &lt;policer counters &lt;(detail   counter-id <i>counter-index</i> &lt;detail&gt;)&gt;&gt;   &lt;terse&gt;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>logical-system</b> option introduced in Junos OS Release 9.3.</p> <p><b>terse</b> option introduced in Junos OS Release 9.4.</p> <p><b>policer counters</b> option introduced in Junos OS Release 12.2 for EX Series switches.</p> <p><b>detail</b> option introduced in Junos OS Release 12.3.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                 | Display statistics about configured firewall filters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                     | <p><b>none</b>—(Optional) Display statistics about all configured firewall filters and counters. For EX Series switches, this command also displays statistics about all configured policers.</p> <p><b>counter <i>counter-name</i></b>—(Optional) Name of a filter counter.</p> <p><b>detail</b>—(EX Series switches only) (Optional) Display firewall filter statistics with enhanced policer.</p> <p><b>filter <i>filter-name</i></b>—(Optional) Name of a configured filter.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>log</b>—(Optional) Display log entries for firewall filters.</p> <p><b>log &lt;(detail   interface <i>interface-name</i>)&gt;</b>—(EX Series switches only) (Optional) Display detailed log entries of firewall activity or log information about a specific interface.</p> <p><b>policer counters &lt;(detail   counter-id <i>counter-index</i> &lt;detail&gt;)&gt;</b>—(EX8200 switches only) (Optional) Display policer counter statistics in brief or in detail.</p> <p><b>terse</b>—(Optional) Display firewall filter names only.</p> |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |



- Related Documentation**
- [clear firewall on page 4448](#)
  - [show firewall log on page 4457](#)
  - *Verifying That Firewall Filters Are Operational*
  - *Verifying That Policers Are Operational*

- List of Sample Output**
- [show firewall filter \(MX Series Router and EX Series Switch\) on page 4453](#)
  - [show firewall filter \(non MX Series Router and EX Series Switch\) on page 4453](#)
  - [show firewall filter \(Hierarchical Policers, MX Series with MPC\) on page 4453](#)
  - [show firewall filter \(Dynamic Input Filter\) on page 4453](#)
  - [show firewall \(Logical Systems\) on page 4453](#)
  - [show firewall \(counter counter-name\) on page 4454](#)
  - [show firewall log on page 4454](#)
  - [show firewall policer counters \(EX8200 Switch\) on page 4454](#)
  - [show firewall policer counters \(detail\) \(EX8200 Switch\) on page 4455](#)
  - [show firewall policer counters \(counter-id counter-index\) \(EX8200 Switch\) on page 4455](#)
  - [show firewall policer counters \(counter-id counter-index detail\) \(EX8200 Switch\) on page 4455](#)
  - [show firewall detail on page 4456](#)

- Output Fields** [Table 323 on page 4451](#) lists the output fields for the **show firewall** command. Output fields are listed in the approximate order in which they appear.

Table 323: show firewall Output Fields

| Field Name      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Filter</b>   | <p>Name of a filter that has been configured with the <b>filter</b> statement at the <b>[edit firewall]</b> hierarchy level.</p> <p>Except on EX Series switches:</p> <ul style="list-style-type: none"> <li>• When an interface-specific filter is displayed, the name of the filter is followed by the full interface name and by either <b>-i</b> for an input filter or <b>-o</b> for an output filter.</li> <li>• When dynamic filters are displayed, the name of the filter is followed by the full interface name and by either <b>-in</b> for an input filter or <b>-out</b> for an output filter. When a logical system-specific filter is displayed, the name of the filter is prefixed with two underscore (__) characters and the name of the logical system (for example, <b>__ls1/filter1</b>).</li> </ul>                        |
| <b>Counters</b> | <p>Display filter counter information:</p> <ul style="list-style-type: none"> <li>• <b>Name</b>—Name of a filter counter that has been configured with the <b>counter</b> firewall filter action.</li> <li>• <b>Bytes</b>—Number of bytes that match the filter term under which the <b>counter</b> action is specified.</li> <li>• <b>Packets</b>—Number of packets that matched the filter term under which the <b>counter</b> action is specified.</li> </ul> <p><b>NOTE:</b> On M and T series routers, firewall filters cannot count <b>ip-options</b> packets on a per option type and per interface basis. A limited work around is to use the <b>show pfe statistics ip options</b> command to see <b>ip-options</b> statistics on a per Packet Forwarding Engine (PFE) basis. See <i>show pfe statistics ip</i> for sample output.</p> |

Table 323: show firewall Output Fields (*continued*)

| Field Name                   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Policers</b>              | <p>Display policer information:</p> <ul style="list-style-type: none"> <li>• <b>Name</b>—Name of policer.</li> <li>• <b>Bytes</b>—(For two-color policers on MX Series routers and EX Series switches, and for hierarchical policers on interfaces hosted on MICs and MPCs in MX Series routers) Number of bytes that match the filter term under which the policer action is specified. This is only the number out-of-specification (out-of-spec) byte counts, not all the bytes in all packets policed by the policer.<br/>For other platforms, this field is blank.</li> <li>• <b>Packets</b>—Number of packets that matched the filter term under which the policer action is specified. This is only the number of out-of-specification (out-of-spec) packet counts, not all packets policed by the policer.</li> </ul> |
| <b>Policer Counter Index</b> | (EX8200 switch only) Global management counter ID. The counter ID value ( <i>counter-index</i> ) can be 0, 1, or 2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Green</b>                 | (EX8200 switch only) Number of packets within the limits. The number of packets is smaller than the committed information rate (CIR).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Yellow</b>                | (EX8200 switch only) Number of packets partially within the limits. The number of packets is greater than the CIR, but the burst size is within the excess burst size (EBS) limit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Discard</b>               | (EX8200 switch only) Number of discarded packets.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Bytes</b>                 | (EX8200 switch only) Number of green, yellow, red, or discarded packets in bytes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Packets</b>               | (EX8200 switch only) Number of green, yellow, red, or discarded packets.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Filter name</b>           | (EX8200 switch only) Name of the filter with a term associated to a policer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Term name</b>             | (EX8200 switch only) Name of the term associated with a policer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Policer name</b>          | (EX8200 switch only) Name of the policer that is associated with a global management counter.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## Sample Output

### show firewall filter (MX Series Router and EX Series Switch)

```
user@host> show firewall filter test
Filter: test
Counters:
Name Bytes Packets
Counter-1 0 0
Counter-2 0 0
Policers:
Name Bytes Packets
Policer-1 2770 70
```

### show firewall filter (non MX Series Router and EX Series Switch)

```
user@host> show firewall filter test
Filter: test
Counters:
Name Bytes Packets
Counter-1 0 0
Counter-2 0 0
Policers:
Name Bytes Packets
Policer-1 70
```

### show firewall filter (Hierarchical Policier, MX Series with MPC)

```
user@host> show firewall filter
FL_V4_PHY-HP-EF-AWARE-Gold=400k-MCAST=200k-Total=1M-ds-10/0/0:2:1-i

Filter: FL_V4_PHY-HP-EF-AWARE-Gold=400k-MCAST=200k-Total=1M-ds-10/0/0:2:1-i
Counters:
Name Bytes Packets
AF1x_counter-ds-10/0/0:2:1-i 0 0
AF2x_counter-ds-10/0/0:2:1-i 25529445976 24500428
AF3x_counter-ds-10/0/0:2:1-i 2182022 39482
AF4x_counter-ds-10/0/0:2:1-i 0 0
BE_counter-ds-10/0/0:2:1-i 0 0
EF_counter-ds-10/0/0:2:1-i 14817044120 12265765
STD_counter-ds-10/0/0:2:1-i 0 0
Policers:
Name Bytes Packets
POL_CE-PE_M=200k-filter-ds-10/0/0:2:1-i 5948099658 5708349
POL_CE-PE_G=400K_R=1M-filter-ds-10/0/0:2:1-i ?????????? 3572794
????????????? ?????????? ????????
```

### show firewall filter (Dynamic Input Filter)

```
user@host> show firewall filter dfwd-ge-5/0/0.1-in
Filter: dfwd-ge-5/0/0.1-in
Counters:
Name Bytes Packets
c1-ge-5/0/0.1-in 0 0
```

### show firewall (Logical Systems)

```
user@host> show firewall
```

```

Filter: __lr1/test
Counters:
Name Bytes Packets
icmp 420 5
Filter: __default_bpdu_filter__
Filter: __lr1/inet_filter1
Counters:
Name Bytes Packets
inet_tcp_count 0 0
inet_udp_count 0 0
Filter: __lr1/inet_filter2
Counters:
Name Bytes Packets
inet_icmp_count 0 0
inet_pim_count 0 0
Filter: __lr2/inet_filter1
Counters:
Name Bytes Packets
inet_tcp_count 0 0
inet_udp_count 0 0

```

#### show firewall (counter counter-name)

```

user@host> show firewall counter icmp-counter
Filter: ingress-port-voip-class-filter
Counters:
Name Bytes Packets
icmp-counter 0 0

```

#### show firewall log

```

user@host> show firewall log
Log :

Time Filter Action Interface Protocol Src Addr
 Dest Addr
08:00:53 pfe R ge-1/0/1.0 ICMP 192.168.3.5
 192.168.3.4
08:00:52 pfe R ge-1/0/1.0 ICMP 192.168.3.5
 192.168.3.4
08:00:51 pfe R ge-1/0/1.0 ICMP 192.168.3.5
 192.168.3.4
08:00:50 pfe R ge-1/0/1.0 ICMP 192.168.3.5
 192.168.3.4
08:00:49 pfe R ge-1/0/1.0 ICMP 192.168.3.5
 192.168.3.4
08:00:48 pfe R ge-1/0/1.0 ICMP 192.168.3.5
 192.168.3.4
08:00:47 pfe R ge-1/0/1.0 ICMP 192.168.3.5
 192.168.3.4

```

#### show firewall policer counters (EX8200 Switch)

```

user@switch> show firewall policer counters
Policer Counter Index 0:
Green: Bytes 73 Packets 15914
Yellow: 9 1962
Discard: 119 25942

```

```

Policer Counter Index 1:
 Bytes Packets
Green: 0 0
Yellow: 0 0
Discard: 0 0

Policer Counter Index 2:
 Bytes Packets
Green: 0 0
Yellow: 0 0
Discard: 0 0

```

### show firewall policer counters (detail) (EX8200 Switch)

```

user@switch> show firewall policer counters detail
Policer Counter Index 0:
 Bytes Packets
Green: 73 15914
Yellow: 9 1962
Discard: 119 25942

Filter name Term name Policer name
myfilter polcr-term-1 myfilter-polcr-1
inet-filter-ae ae-snmp policer-1
inet-filter-ae ae-ssh policer-2

Policer Counter Index 1:
 Bytes Packets
Green: 0 0
Yellow: 0 0
Discard: 0 0

Filter name Term name Policer name

Policer Counter Index 2:
 Bytes Packets
Green: 0 0
Yellow: 0 0
Discard: 0 0

Filter name Term name Policer name

```

### show firewall policer counters (counter-id counter-index) (EX8200 Switch)

```

user@switch> show firewall policer counters counter-id 0
Policer Counter Index 0:
 Bytes Packets
Green: 73 15914
Yellow: 9 1962
Discard: 119 25942

```

### show firewall policer counters (counter-id counter-index detail) (EX8200 Switch)

```

user@switch> show firewall policer counters counter-id 0 detail
Policer Counter Index 0:
 Bytes Packets
Green: 73 15914
Yellow: 9 1962
Discard: 119 25942

Filter name Term name Policer name

```

|                |              |                  |
|----------------|--------------|------------------|
| myfilter       | polcr-term-1 | myfilter-polcr-1 |
| inet-filter-ae | ae-snmp      | policer-1        |
| inet-filter-ae | ae-ssh       | policer-2        |

#### show firewall detail

```
user@host> show firewall detail
Filter: __default_bpdu_filter__
```

Filter: foo

Counters:

Name

c1

| Bytes    | Packets |
|----------|---------|
| 17652140 | 160474  |

Policers:

Name

P1-t1

| Bytes                  | Packets |
|------------------------|---------|
| 0                      | 18286   |
| 0 18446744073709376546 |         |
| 0 18446744073709358260 |         |

OOS

Offered

Transmitted

## show firewall log

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                      | show firewall log<br><detail><br><interface <i>interface-name</i> ><br><logical-system ( <i>logical-system-name</i>   all)>                                                                                                                                                                                                                                                                    |
| <b>Syntax (EX Series Switches)</b> | show firewall log<br><detail><br><interface <i>interface-name</i> >                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>logical-system</b> option introduced in Junos OS Release 9.3.                                                                                                                                                                                                      |
| <b>Description</b>                 | Display log information about firewall filters.                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                     | <b>none</b> —Display log information about firewall filters.<br><br><b>detail</b> —(Optional) Display detailed information.<br><br><b>interface <i>interface-name</i></b> —(Optional) Display log information about a specific interface.<br><br><b>logical-system (<i>logical-system-name</i>   all)</b> —(Optional) Perform this operation on all logical systems or on a particular system. |
| <b>Required Privilege Level</b>    | view                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>List of Sample Output</b>       | <a href="#">show firewall log on page 4458</a><br><a href="#">show firewall log detail on page 4458</a>                                                                                                                                                                                                                                                                                        |
| <b>Output Fields</b>               | <a href="#">Table 324 on page 4457</a> lists the output fields for the <b>show firewall log</b> command. Output fields are listed in the approximate order in which they appear.                                                                                                                                                                                                               |

**Table 324: show firewall log Output Fields**

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Time of Log</b> | Time that the event occurred.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Filter</b>      | <p>Name of a filter that has been configured with the <b>filter</b> statement at the <b>[edit firewall]</b> hierarchy level.</p> <ul style="list-style-type: none"> <li>A hyphen (-) indicates that the packet was handled by the Packet Forwarding Engine.</li> <li>A space (no hyphen) indicates the packet was handled by the Routing Engine.</li> <li>The notation <b>pfe</b> indicates packets logged by the Packet Forwarding Engine hardware filters.</li> </ul> |

Table 324: show firewall log Output Fields (*continued*)

| Field Name          | Field Description                                                                                                    |
|---------------------|----------------------------------------------------------------------------------------------------------------------|
| Filter Action       | Filter action: <ul style="list-style-type: none"> <li>• A—Accept</li> <li>• D—Discard</li> <li>• R—Reject</li> </ul> |
| Name of Interface   | Ingress interface for the packet.                                                                                    |
| Name of protocol    | Packet's protocol name: <b>egp, gre, icmp, ipip, ospf, pim, rsvp, tcp, or udp</b> .                                  |
| Packet length       | Length of the packet.                                                                                                |
| Source address      | Packet's source address.                                                                                             |
| Destination address | Packet's destination address and port.                                                                               |

## Sample Output

### show firewall log

```

user@host>show firewall log
Time Filter Action Interface Protocol Src Addr Dest Addr
13:10:12 pfe D rlsq0.902 ICMP 180.1.177.2 180.1.177.1
13:10:11 pfe D rlsq0.902 ICMP 180.1.177.2 180.1.177.1

```

### show firewall log detail

```

user@host> show firewall log detail
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0Name of protocol: TCP, Packet Length: 50824, Source address:
172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 1020, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0

```



Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,  
Destination address: 192.168.70.66:513  
....

## Operational Mode Commands for Next-Hop Groups for Layer 2 Port Mirroring

## show forwarding-options next-hop-group

|                                 |                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <b>show forwarding-options next-hop-group</b><br><b>&lt;terse   brief   detail&gt;</b><br><b>&lt;group-name&gt;</b>                                                                                                                  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.                                                                                                               |
| <b>Description</b>              | Display current state of next-hop groups.                                                                                                                                                                                            |
| <b>Options</b>                  | <b>terse   brief   detail</b> —(Optional) Display the specified level of output.<br><br><b>group-name</b> —(Optional) Display a single next-hop group.                                                                               |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show forwarding-options port-mirroring on page 4446</a></li> </ul>                                                                                                              |
| <b>List of Sample Output</b>    | <a href="#">show forwarding-options next-hop-group terse on page 4461</a><br><a href="#">show forwarding-options next-hop-group brief on page 4461</a><br><a href="#">show forwarding-options next-hop-group detail on page 4461</a> |
| <b>Output Fields</b>            | <a href="#">Table 325 on page 4460</a> lists the output fields for the <b>show forwarding-options next-hop-group</b> command. Output fields are listed in the approximate order in which they appear.                                |

**Table 325: show forwarding-options next-hop-group Output Fields**

| Field Name                           | Field Description                                            | Level of Output     |
|--------------------------------------|--------------------------------------------------------------|---------------------|
| <b>Next-hop-group</b>                | Name of next-hop group.                                      | All levels          |
| <b>Type</b>                          | Next-hop group type, such as <b>inet</b> or <b>layer-2</b> . | All levels          |
| <b>State</b>                         | Next-hop group state, either <b>up</b> or <b>down</b> .      | All levels          |
| <b>Members Interfaces</b>            | Names of interfaces to which next-hop group members belong.  | <b>brief detail</b> |
| <b>Members Subgroup</b>              | Names of subgroups to which next-hop group members belong.   | <b>brief detail</b> |
| <b>Number of members configured</b>  | Number of next-hop group members configured.                 | <b>detail</b>       |
| <b>Number of members that are up</b> | Number of next-hop group members that are up.                | <b>detail</b>       |

Table 325: show forwarding-options next-hop-group Output Fields (*continued*)

| Field Name                      | Field Description                | Level of Output |
|---------------------------------|----------------------------------|-----------------|
| Number of subgroups configured  | Number of subgroups configured.  | detail          |
| Number of subgroups that are up | Number of subgroups that are up. | detail          |

## Sample Output

### show forwarding-options next-hop-group terse

```

user@host> show forwarding-options next-hop-group terse
Next-hop-group Type State
inet_nhg inet up
vpls_nhg layer-2 up
vpls_nhg_2 layer-2 down

```

### show forwarding-options next-hop-group brief

```

user@host> show forwarding-options next-hop-group brief
Next-hop-group: inet_nhg
Type: inet State: up
Members Interfaces:
 ge-2/0/2.101 next-hop 101.2.0.2

Next-hop-group: vpls_nhg
Type: layer-2 State: up
Members Interfaces:
 ge-2/0/1.100
 ge-2/2/9.0
Members Subgroup: vpls_subg
Members Interfaces:
 ge-2/0/1.101
 ge-2/2/9.1

Next-hop-group: vpls_nhg_2
Type: layer-2 State: down

```

### show forwarding-options next-hop-group detail

```

user@host> show forwarding-options next-hop-group detail
Next-hop-group: inet_nhg
Type: inet State: up
Number of members configured : 2
Number of members that are up : 1
Number of subgroups configured : 0
Number of subgroups that are up : 0
Members Interfaces: State
 ge-2/0/2.101 next-hop 101.2.0.2 up
 ge-2/2/8.2 next-hop 2.8.0.2 down

Next-hop-group: vpls_nhg
Type: layer-2 State: up
Number of members configured : 2

```

```
Number of members that are up : 2
Number of subgroups configured : 1
Number of subgroups that are up : 1
Members Interfaces: State
 ge-2/0/1.100 up
 ge-2/2/9.0 up
Members Subgroup: vpls_subg up
 Number of members configured : 2
 Number of members that are up : 2
 Members Interfaces:
 ge-2/0/1.101 up
 ge-2/2/9.1 up

Next-hop-group: vpls_nhg_2
Number of members configured : 2
Number of members that are up : 0
Number of subgroups configured : 0
Number of subgroups that are up : 0
Type: layer-2 State: down
Members Interfaces: State
 ge-2/2/1.100 down
 ge-2/3/9.0 down
```

## CHAPTER 18

# Routing Policy and Packet Filtering

- [Firewall Filters on page 4463](#)
- [Traffic Policers on page 4802](#)

## Firewall Filters

---

- [Overview on page 4463](#)
- [Configuration on page 4547](#)
- [Administration on page 4703](#)

## Overview

- [Introduction to Stateless Firewall Filters on page 4463](#)
- [Standard Firewall Filters Overview on page 4474](#)
- [Standard Firewall Filter Match Conditions Overview on page 4488](#)
- [Introduction to Standard Firewall Filters for Fragment Handling on page 4503](#)
- [Introduction to Standard Firewall Filters Configuration on page 4503](#)
- [Introduction to Service Filters Configuration on page 4521](#)
- [Introduction to Simple Filters Configuration on page 4529](#)
- [Introduction to Firewall Filters Configuration in Logical Systems on page 4535](#)

## Introduction to Stateless Firewall Filters

---

- [Stateless Firewall Filter Overview on page 4463](#)
- [Stateless Firewall Filter Types on page 4465](#)
- [Stateless Firewall Filter Components on page 4466](#)
- [Stateless Firewall Filter Application Points on page 4472](#)

### ***Stateless Firewall Filter Overview***

This topic covers the following information:

- [Packet Flow Control on page 4464](#)
- [Stateless and Stateful Firewall Filters on page 4464](#)
- [Purpose of Stateless Firewall Filters on page 4464](#)

### ***Packet Flow Control***

To influence which packets are allowed to transit the system and to apply special actions to packets as necessary, you can configure *stateless firewall filters*. A stateless firewall specifies a sequence of one or more packet-filtering rules, called *filter terms*. A filter term specifies *match conditions* to use to determine a match and *actions* to take on a matched packet. A stateless firewall filter enables you to manipulate any packet of a particular protocol family, including fragmented packets, based on evaluation of Layer 3 and Layer 4 header fields. You typically apply a stateless firewall filter to one or more interfaces that have been configured with protocol family features. You can apply a stateless firewall filter to an ingress interface, an egress interface, or both.

### ***Data Packet Flow Control***

To control the flow of data packets transiting the device as the packets are being forwarded from a source to a destination, you can apply stateless firewall filters to the input or output of the router's or switch's physical interfaces.

To enforce a specified bandwidth and maximum burst size for traffic sent or received on an interface, you can configure *policers*. Policers are a specialized type of stateless firewall filter and a primary component of the Junos OS *class-of-service* (CoS).

### ***Local Packet Flow Control***

To control the flow of local packets between the physical interfaces and the Routing Engine, you can apply stateless firewall filters to the input or output of the *loopback interface*. The loopback interface (**lo0**) is the interface to the Routing Engine and carries no data packets.

### ***Stateless and Stateful Firewall Filters***

A stateless firewall filter, also known as an *access control list* (ACL), does not statefully inspect traffic. Instead, it evaluates packet contents statically and does not keep track of the state of network connections. In contrast, a *stateful firewall filter* uses connection state information derived from other applications and past communications in the data flow to make dynamic control decisions.

The *Junos OS Firewall Filters and Traffic Policers Configuration Guide* describes *stateless firewall filters* supported on T Series, M Series, MX Series routers and EX Series switches.

### ***Purpose of Stateless Firewall Filters***

The basic purpose of a stateless firewall filter is to enhance security through the use of packet filtering. Packet filtering enables you to inspect the components of incoming or outgoing packets and then perform the actions you specify on packets that match the criteria you specify. The typical use of a stateless firewall filter is to protect the Routing Engine processes and resources from malicious or untrusted packets.

#### **Related Documentation**

- [Router Data Flow Overview](#)
- [Stateless Firewall Filter Types on page 4465](#)
- [“Traffic Policing Overview on page 1440” in the Junos OS Firewall Filters and Traffic Policers Configuration Guide](#)

- [“Packet Flow Through the CoS Process Overview on page 1273”](#) in the *Junos OS Class of Service Configuration Guide*

### **Stateless Firewall Filter Types**

This topic covers the following information:

- [Standard Stateless Firewall Filters on page 4465](#)
- [Service Filters on page 4465](#)
- [Simple Filters on page 4465](#)

### **Standard Stateless Firewall Filters**

The Junos OS standard stateless firewall filters support a rich set of packet-matching criteria that you can use to match on specific traffic and perform specific actions, such as forwarding or dropping packets that match the criteria you specify. You can configure firewall filters to protect the local router (or switch) or to protect another device that is either directly or indirectly connected to the local router (or switch). For example, you can use the filters to restrict the local packets that pass from the router's (or switch's) physical interfaces to the Routing Engine. Such filters are useful in protecting the IP services that run on the Routing Engine, such as Telnet, SSH, and BGP, from denial-of-service attacks.



**NOTE:** If you configured targeted broadcast for virtual routing and forwarding (VRF) by including the `forward-and-send-to-re` statement, any firewall filter that is configured on the Routing Engine loopback interface (lo0) cannot be applied to the targeted broadcast packets that are forwarded to the Routing Engine. This is because broadcast packets are forwarded as flood next hop traffic and not as local next hop traffic, and you can only apply a firewall filter to local next hop routes for traffic directed toward the Routing Engine.

### **Service Filters**

A service filter defines packet-filtering (a set of match conditions and a set of actions) for IPv4 or IPv6 traffic. You can apply a service filter to the inbound or outbound traffic at an adaptive services interface to perform packet filtering on traffic before it is accepted for service processing. You can also apply a service filter to the traffic that is returning to the services interface after service processing to perform postservice processing.

Service filters filter IPv4 and IPv6 traffic only and can be applied to logical interfaces on Adaptive Services PICs, MultiServices PICs, and MultiServices DPCs only. Service filters are not supported on J Series devices and Branch SRX devices.

### **Simple Filters**

Simple filters are supported on Gigabit Ethernet intelligent queuing (IQ2) and Enhanced Queuing Dense Port Concentrator (EQ DPC) interfaces only. Unlike standard filters, simple filters support IPv4 traffic only and have a number of restrictions. For example, you cannot configure a terminating action for a simple filter. Simple filters always accept packets. Also, simple filters can be applied only as input filters. They are not supported

on outbound traffic. Simple filters are recommended for metropolitan Ethernet applications.

**Related Documentation**

- [Stateless Firewall Filter Overview on page 4463](#)
- [Stateless Firewall Filter Components on page 4466](#)

**Stateless Firewall Filter Components**

This topic covers the following information:

- [Protocol Family on page 4466](#)
- [Filter Type on page 4467](#)
- [Terms on page 4468](#)
- [Match Conditions on page 4469](#)
- [Actions on page 4470](#)

**Protocol Family**

Under the **firewall** statement, you can specify the protocol family for which you want to filter traffic.

[Table 326 on page 4466](#) describes the firewall filter protocol families.

**Table 326: Firewall Filter Protocol Families**

| Type of Traffic to Be Filtered     | Protocol Family Configuration Statement | Comments                                                                     |
|------------------------------------|-----------------------------------------|------------------------------------------------------------------------------|
| Protocol Independent               | <b>family any</b>                       | All protocol families configured on a logical interface.                     |
| Internet Protocol version 4 (IPv4) | <b>family inet</b>                      | The <b>family inet</b> statement is optional for IPv4.                       |
| Internet Protocol version 6 (IPv6) | <b>family inet6</b>                     |                                                                              |
| MPLS                               | <b>family mpls</b>                      |                                                                              |
| MPLS-tagged IPv4                   | <b>family mpls</b>                      | Supports matching on IP addresses and ports, up to five MPLS stacked labels. |
| MPLS-tagged IPv6                   | <b>family mpls</b>                      | Supports matching on IP addresses and ports, up to five MPLS stacked labels. |
| Virtual private LAN service (VPLS) | <b>family vpls</b>                      |                                                                              |
| Layer 2 Circuit Cross-Connection   | <b>family ccc</b>                       |                                                                              |



Table 326: Firewall Filter Protocol Families (*continued*)

| Type of Traffic to Be Filtered | Protocol Family Configuration Statement                                                                    | Comments                                       |
|--------------------------------|------------------------------------------------------------------------------------------------------------|------------------------------------------------|
| Layer 2 Bridging               | <b>family bridge</b> (for MX Series routers) and <b>family ethernet-switching</b> (for EX Series switches) | MX Series routers and EX Series switches only. |

**Filter Type**

Under the **family *family-name*** statement, you can specify the type and name of the filter you want to configure.

[Table 327 on page 4467](#) describes the firewall filter types.

Table 327: Filter Types

| Filter Type               | Filter Configuration Statement   | Description                                                                                                                                                                                                                                                                                                                                     |
|---------------------------|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Stateless Firewall Filter | <b>filter <i>filter-name</i></b> | <p>Filters the following traffic types:</p> <ul style="list-style-type: none"> <li>• Protocol independent</li> <li>• IPv4</li> <li>• IPv6</li> <li>• MPLS</li> <li>• MPLS-tagged IPv4</li> <li>• MPLS-tagged IPv6</li> <li>• VPLS</li> <li>• Layer 2 CCC</li> <li>• Layer 2 bridging (MX Series routers and EX Series switches only)</li> </ul> |

Table 327: Filter Types (*continued*)

| Filter Type    | Filter Configuration Statement                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Service Filter | <b>service-filter</b><br><i>service-filter-name</i> | <p>Defines packet-filtering to be applied to ingress or egress before it is accepted for service processing or applied to returning service traffic after service processing has completed.</p> <p>Filters the following traffic types:</p> <ul style="list-style-type: none"> <li>• IPv4</li> <li>• IPv6</li> </ul> <p>Supported at logical interfaces configured on the following hardware only:</p> <ul style="list-style-type: none"> <li>• Adaptive Services (AS) PICs on M Series and T Series routers</li> <li>• Multiservices (MS) PICs on M Series and T Series routers</li> <li>• Multiservices (MS) DPCs on MX Series routers (and EX Series switches)</li> </ul> |
| Simple Filter  | <b>simple-filter</b><br><i>simple-filter-name</i>   | <p>Defines packet filtering to be applied to ingress traffic only.</p> <p>Filters the following traffic type:</p> <ul style="list-style-type: none"> <li>• IPv4</li> </ul> <p>Supported at logical interfaces configured on the following hardware only:</p> <ul style="list-style-type: none"> <li>• Gigabit Ethernet Intelligent Queuing (IQ2) PICs installed on M120, M320, or T Series routers</li> <li>• Enhanced Queuing Dense Port Concentrators (EQ DPCs) installed on MX Series routers (and EX Series switches)</li> </ul>                                                                                                                                         |

**Terms**

Under the **filter**, **service-filter**, or **simple-filter** statement, you must configure at least one firewall filter *term*. A term is a named structure in which match conditions and actions are defined. Within a firewall filter, you must configure a unique name for each term.



**TIP:** You cannot apply a different filter on each direction of traffic on the same interface. As a result, it is common to create firewall filters with multiple terms.

All stateless firewall filters contain one or more terms, and each term consists of two components—match conditions and actions. The match conditions define the values or fields that the packet must contain to be considered a match. If a packet is a match, the

corresponding action is taken. By default, a packet that does not match a firewall filter is discarded.

If a packet arrives on an interface for which no firewall filter is applied for the incoming traffic on that interface, the packet is accepted by default.



**NOTE:** A firewall filter with a large number of terms can adversely affect both the configuration commit time and the performance of the Routing Engine.

Additionally, you can configure a stateless firewall filter within the term of another filter. This method enables you to add common terms to multiple filters without having to modify all filter definitions. You can configure one filter with the desired common terms, and configure this filter as a term in other filters. Consequently, to make a change in these common terms, you need to modify only one filter that contains the common terms, instead of multiple filters.

### **Match Conditions**

A firewall filter term must contain at least one packet-filtering criteria, called a *match condition*, to specify the field or value that a packet must contain in order to be considered a match for the firewall filter term. For a match to occur, the packet must match all the conditions in the term. If a packet matches a firewall filter term, the router (or switch) takes the configured action on the packet.

If a firewall filter term contains multiple match conditions, a packet must meet *all* match conditions to be considered a match for the firewall filter term.

If a single match condition is configured with multiple values, such as a range of values, a packet must match only *one* of the values to be considered a match for the firewall filter term.

The scope of match conditions you can specify in a firewall filter term depends on the protocol family under which the firewall filter is configured. You can define various match conditions, including the IP source address field, IP destination address field, TCP or UDP source port field, IP protocol field, Internet Control Message Protocol (ICMP) packet type, IP options, TCP flags, incoming logical or physical interface, and outgoing logical or physical interface.

Each protocol family supports a different set of match conditions, and some match conditions are supported only on certain routing devices. For example, a number of match conditions for VPLS traffic are supported only on the MX Series 3D Universal Edge Routers.

In the **from** statement in a firewall filter term, you specify characteristics that the packet must have for the action in the subsequent **then** statement to be performed. The characteristics are referred to as *match conditions*. The packet must match all conditions in the **from** statement for the action to be performed, which also means that the order of the conditions in the **from** statement is not important.

If an individual match condition can specify a list of values (such as multiple source and destination addresses) or a range of numeric values, a match occurs if any of the values matches the packet.

If a filter term does not specify match conditions, the term accepts all packets and the actions specified in the term's **then** statement are optional.



**NOTE:**

Some of the numeric range and bit-field match conditions allow you to specify a text synonym. For a complete list of synonyms:

- If you are using the J-Web interface, select the synonym from the appropriate list.
- If you are using the CLI, type a question mark (?) after the **from** statement.

---

### **Actions**

The actions specified in a firewall filter term define the actions to take for any packet that matches the conditions specified in the term.

Actions that are configured within a single term are all taken on traffic that matches the conditions configured.



**BEST PRACTICE:** We strongly recommend that you explicitly configure one or more actions per firewall filter term. Any packet that matches all the conditions of the term is automatically accepted unless the term specifies other or additional actions.

---

Firewall filter actions fall into the following categories:

#### ***Filter-Terminating Actions***

A filter-terminating action halts all evaluation of a firewall filter for a specific packet. The router (or switch) performs the specified action, and no additional terms are examined.

#### ***Nonterminating Actions***

Nonterminating actions are used to perform other functions on a packet, such as incrementing a counter, logging information about the packet header, sampling the packet data, or sending information to a remote host using the system log functionality.

The presence of a nonterminating action, such as **count**, **log**, or **syslog**, without an explicit terminating action, such as **accept**, **discard**, or **reject**, results in a default terminating action of **accept**. If you do not want the firewall filter action to terminate, use the **next term** action after the nonterminating action.

In this example, term 2 is never evaluated, because term 1 has the implicit default **accept** terminating action.

[edit firewall filter test]

```

term 1 {
 from {
 source-address {
 0.0.0.0/0;
 }
 }
 then {
 log;
 <accept> #By default if not specified
 }
}
term 2 {
 then {
 reject;
 }
}

```

In this example, term 2 is evaluated, because term 1 has the explicit **next term** flow control action.

```

[edit firewall filter test]
term 1 {
 from {
 source-address {
 0.0.0.0/0;
 }
 }
 then {
 log;
 next term;
 }
}
term 2 {
 then {
 reject;
 }
}

```

### **Flow Control Action**

For standard stateless firewall filters only, the action **next term** enables the router (or switch) to perform configured actions on the packet and then evaluate the following term in the filter, rather than terminating the filter.

A maximum of 1024 **next term** actions are supported per standard stateless firewall filter configuration. If you configure a standard filter that exceeds this limit, your candidate configuration results in a commit error.

### **Related Documentation**

- [Stateless Firewall Filter Types on page 4465](#)
- “Inserting a New Identifier in a Junos Configuration” in the *CLI User Guide*

### Stateless Firewall Filter Application Points

After you define the firewall filter, you must apply it to an application point. These application points include logical interfaces, physical interfaces, routing interfaces, and routing instances.

In most cases, you can apply a firewall filter as an *input* filter or an *output* filter, or both at the same time. Input filters take action on packets being received on the specified interface, whereas output filters take action on packets that are transmitted through the specified interface.

You typically apply one filter with multiple terms to a single logical interface, to incoming traffic, outbound traffic, or both. However, there are times when you might want to chain together multiple firewall filters (with single or multiple terms) and apply them to an interface. You use an *input list* to apply multiple firewall filters to the incoming traffic on an interface. You use an *output list* to apply multiple firewall filters to the outbound traffic on an interface. You can include up to 16 filters in an input list or an output list.

There is no limit to the number of filters and counters you can set, but there are some practical considerations. More counters require more terms, and a large number of terms can take a long time to process during a commit operation. However, filters with more than 4000 terms and counters have been implemented successfully.

Table 328 on page 4472 describes each point to which you can apply a firewall filter. For each application point, the table describes the types of firewall filters supported at that point, the router (or switch) hierarchy level at which the filter can be applied, and any platform-specific limitations.

**Table 328: Stateless Firewall Filter Configuration and Application Summary**

| Filter Type                                                                                                                                                                                                                                                                                                            | Application Point                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Restrictions                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Stateless firewall filter</b></p> <p>Configure by including the <b>filter filter-name</b> statement the <b>[edit firewall]</b> hierarchy level:</p> <pre>filter filter-name;</pre> <p><b>NOTE:</b> If you do not include the <b>family</b> statement, the firewall filter processes IPv4 traffic by default.</p> | <p><b>Logical interface</b></p> <p>Apply at the <b>[edit interfaces interface-name unit unit-number family inet]</b> hierarchy level by including the <b>input filter-name</b> or <b>output filter-name</b> statements:</p> <pre>filter {   input filter-name;   output filter-name; }</pre> <p><b>NOTE:</b> A filter configured with the implicit <b>inet</b> protocol family cannot be included in an input filter list or an output filter list.</p> <p><b>NOTE:</b> On T4000 Type 5 FPCs, a filter attached at the Layer 2 application point (that is, at the logical interface level) is unable to match with the forwarding class of a packet that is set by a Layer 3 classifier such as DSCP, DSCP V6, <b>inet-precedence</b>, and <b>mpls-exp</b>.</p> | <p>Supported on the following routers:</p> <ul style="list-style-type: none"> <li>• T Series routers</li> <li>• M320 routers</li> <li>• M7i routers with the enhanced CFEB (CFEB-e)</li> <li>• M10i routers with the enhanced CFEB-e</li> </ul> <p>Also supported on the following Modular Port Concentrators (MPCs) on MX Series routers:</p> <ul style="list-style-type: none"> <li>• 10-Gigabit Ethernet MPC</li> <li>• 60-Gigabit Ethernet Queuing MPC</li> <li>• 60-Gigabit Ethernet Enhanced Queuing MPC</li> <li>• 100-Gigabit Ethernet MPC</li> </ul> <p>• Also supported on EX Series switches</p> |

Table 328: Stateless Firewall Filter Configuration and Application Summary (*continued*)

| Filter Type                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Application Point                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Restrictions                                                                     |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
| <p><b>Stateless firewall filter</b></p> <p>Configure at the <code>[edit firewall family <i>family-name</i>]</code> hierarchy level by including the following statement:</p> <pre>filter <i>filter-name</i>;</pre> <p>The <i>family-name</i> can be any of the following protocol families:</p> <ul style="list-style-type: none"> <li>• any</li> <li>• bridge</li> <li>• ethernet-switching</li> <li>• ccc</li> <li>• inet</li> <li>• inet6</li> <li>• mpls</li> <li>• vpls</li> </ul> | <p><b>Protocol family on a logical interface</b></p> <p>Apply at the <code>[edit interfaces <i>interface-name</i> unit <i>unit-number</i> family <i>family-name</i>]</code> hierarchy level by, including the <code>input</code>, <code>input-list</code>, <code>output</code>, or <code>output-list</code> statements:</p> <pre>filter {   input <i>filter-name</i>;   input-list [ <i>filter-names</i> ];   output <i>filter-name</i>;   output-list [ <i>filter-names</i> ]; }</pre>                                                                                                                                                                                                                              | <p>The protocol family <b>bridge</b> is supported only on MX Series routers.</p> |
| Stateless firewall filter                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>Routing Engine loopback interface</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                                                                                  |
| <p><b>Service filter</b></p> <p>Configure at the <code>[edit firewall family (inet   inet6)]</code> hierarchy level by including the following statement:</p> <pre>service-filter   <i>service-filter-name</i>;</pre>                                                                                                                                                                                                                                                                   | <p><b>Family inet or inet6 on a logical interface</b></p> <p>Apply at the <code>[edit interfaces <i>interface-name</i> unit <i>unit-number</i> family (inet   inet6)]</code> hierarchy level by using the <code>service-set</code> statement to apply a service filter as an input or output filter to a service set:</p> <pre>service {   input {     service-set <i>service-set-name</i>     service-filter <i>filter-name</i>;   }   output {     service-set <i>service-set-name</i>     service-filter <i>filter-name</i>;   } }</pre> <p>Configure a service set at the <code>[edit services]</code> hierarchy level by including the following statement:</p> <pre>service-set <i>service-set-name</i>;</pre> | <p>Supported only on Adaptive Services (AS) and Multiservices (MS) PICs.</p>     |

Table 328: Stateless Firewall Filter Configuration and Application Summary (*continued*)

| Filter Type                                                                                                                                                                                                           | Application Point                                                                                                                                                                                                                                                                                                                                                                                       | Restrictions                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Postservice filter</b><br><br>Configure at the <b>[edit firewall family (inet   inet6)]</b> hierarchy level by including the following statement:<br><br><pre>service-filter   service-filter-name;</pre>          | <b>Family inet or inet6 on a logical interface</b><br><br>Apply at the <b>[edit interfaces interface-name unit unit-number family (inet   inet6)]</b> hierarchy level by including the <b>post-service-filter</b> statement to apply a service filter as an input filter:<br><br><pre>service {   input {     post-service-filter filter-name;   } }</pre>                                              | A postservice filter is applied to traffic returning to the services interface after service processing. The filter is applied only if a service set is configured and selected.                                                                                                                                                                             |
| <b>Simple filter</b><br><br>Configure at the <b>[edit firewall family inet]</b> hierarchy level by including the following statement:<br><br><pre>simple-filter filter-name</pre>                                     | <b>Family inet on a logical interface</b><br><br>Apply at the <b>[edit interfaces interface-name unit unit-number family inet]</b> hierarchy level by including the following statement:<br><br><pre>simple-filter simple-filter-name;</pre>                                                                                                                                                            | Simple filters can only be applied as input filters.<br><br>Supported on the following platforms only: <ul style="list-style-type: none"> <li>Gigabit Ethernet intelligent queuing (IQ2) PICs on the M120, M320, and T Series routers.</li> <li>Enhanced Queuing Dense Port Concentrators (EQ DPC) on MX Series routers (and EX Series switches).</li> </ul> |
| <b>Reverse packet forwarding (RPF) check filter</b><br><br>Configured at the <b>[edit firewall family (inet   inet6)]</b> hierarchy level by including the following statement:<br><br><pre>filter filter-name;</pre> | <b>Family inet or inet6 on a logical interface</b><br><br>Apply at the <b>[edit interfaces interface-name unit unit-number family (inet   inet6)]</b> hierarchy level by including the following statement:<br><br><pre>rpf-check fail-filter filter-name</pre> to apply the stateless firewall filter as an RPF check filter.<br><br><pre>rpf-check {   fail-filter filter-name;   mode loose; }</pre> | Supported on MX Series routers and EX Series switches only.                                                                                                                                                                                                                                                                                                  |

- Related Documentation**
- [Stateless Firewall Filter Components on page 4466](#)
  - [Supported Standards for Filtering on page 4703](#)

### Standard Firewall Filters Overview

- [Standard Stateless Firewall Filter Overview on page 4475](#)
- [How Standard Firewall Filters Evaluate Packets on page 4475](#)
- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Guidelines for Applying Standard Firewall Filters on page 4483](#)
- [Understanding How to Use Standard Firewall Filters on page 4486](#)



### ***Standard Stateless Firewall Filter Overview***

Firewall filters provide a means of protecting your router (and switch) from excessive traffic transiting the router (and switch) to a network destination or destined for the Routing Engine. Firewall filters that control local packets can also protect your router (and switch) from external incidents.

You can configure a firewall filter to do the following:

- Restrict traffic destined for the Routing Engine based on its source, protocol, and application.
- Limit the traffic rate of packets destined for the Routing Engine to protect against flood, or denial-of-service (DoS) attacks.
- Address special circumstances associated with fragmented packets destined for the Routing Engine. Because the device evaluates every packet against a firewall filter (including fragments), you must configure the filter to accommodate fragments that do not contain packet header information. Otherwise, the filter discards all but the first fragment of a fragmented packet.

#### **Related Documentation**

- [Stateless Firewall Filter Types on page 4465](#)
- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Guidelines for Applying Standard Firewall Filters on page 4483](#)
- [Understanding How to Use Standard Firewall Filters on page 3293](#)

### ***How Standard Firewall Filters Evaluate Packets***

This topic covers the following information:

- [Firewall Filter Packet Evaluation Overview on page 4475](#)
- [Packet Evaluation at a Single Firewall Filter on page 4476](#)
- [Best Practice: Explicitly Accept Any Traffic That Is Not Specifically Discarded on page 4477](#)
- [Best Practice: Explicitly Reject Any Traffic That Is Not Specifically Accepted on page 4477](#)
- [Multiple Firewall Filters Attached to a Single Interface on page 4477](#)
- [Single Firewall Filter Attached to Multiple Interfaces on page 4478](#)

### ***Firewall Filter Packet Evaluation Overview***

The following sequence describes how the device evaluates a packet entering or exiting an interface if the input or output traffic at a device interface is associated with a firewall filter. Packet evaluation proceeds as follows:

1. The device evaluates the packet against the terms in the firewall filter sequentially, beginning with the first term in the filter.
  - If the packet matches all the conditions specified in a term, the device performs all the actions specified in that term.

- If the packet does not match all the conditions specified in a term, the device proceeds to the next term in the filter (if a subsequent term exists) and evaluates the packet against that term.
  - If the packet does not match any term in the firewall filter, the device implicitly discards the packet.
2. Unlike service filters and simple filters, firewall filters support the **next term** action, which is neither a terminating action nor a nonterminating action but a flow control action.
- If the matched term includes the **next term** action, the device continues evaluation of the packet at the next term within the firewall filter.
  - If the matched term does not include the **next term** action, evaluation of the packet against the given firewall filter ends at this term. The device does not evaluate the packet against any subsequent terms in this filter.
- A maximum of 1024 **next term** actions are supported per firewall filter configuration. If you configure a firewall filter that exceeds this limit, your candidate configuration results in a commit error.
3. The device stops evaluating a packet against a given firewall filter when either the packet matches a term without the **next term** action or the packet fails to match the last term in the firewall filter.

#### ***Packet Evaluation at a Single Firewall Filter***

Table 329 on page 4476 describes packet-filtering behaviors at a device interface associated with a single firewall filter.

**Table 329: Packet Evaluation at a Single Firewall Filter**

| Firewall Filter Event                                                                                               | Action                                                                                                  | Subsequent Action                                                                                                                                                                      |
|---------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The firewall filter term does not specify any match conditions.                                                     | The term matches all packets by default, and so the device performs the actions specified by that term. | If the term actions include the <b>next term</b> action, the device continues evaluation of the packet against the next term within the firewall filter (if a subsequent term exists). |
| The packet matches all conditions specified by the firewall filter term.                                            | The device performs the actions specified by that term.                                                 | If the term actions include the <b>next term</b> action, the device continues evaluation of the packet against the next term within the firewall filter (if a subsequent term exists). |
| The packet matches all conditions specified by the firewall filter term, but the term does not specify any actions. | The device implicitly accepts the packet.                                                               | If the term actions include the <b>next term</b> action, the device continues evaluation of the packet against the next term                                                           |

Table 329: Packet Evaluation at a Single Firewall Filter (*continued*)

| Firewall Filter Event                                                           | Action                                                                                                                                                                                                                                                                                                                                                                             | Subsequent Action                                                                                                    |
|---------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
|                                                                                 |                                                                                                                                                                                                                                                                                                                                                                                    | within the firewall filter (if a subsequent term exists).                                                            |
| The packet does not match all conditions specified by the firewall filter term. | The device does not perform the actions specified by that term.                                                                                                                                                                                                                                                                                                                    | The device continues evaluation of the packet against the next term within the filter (if a subsequent term exists). |
| The packet does not match any term in the filter                                | <p>The device implicitly discards the packet</p> <p>Every firewall filter configuration includes an implicit <b>discard</b> action at the end of the filter. This implicit terminating action is equivalent to including the following example term <b>t_explicit_discard</b> as the final term in the firewall filter:</p> <pre>term t_explicit_discard {   then discard; }</pre> |                                                                                                                      |

**Best Practice: Explicitly Accept Any Traffic That Is Not Specifically Discarded**

You might want a firewall filter to accept any traffic that the filter does not specifically discard. In this case, we recommend that you configure the firewall filter with a final term that specifies the **accept** terminating action.

In the following example snippet, configuring the **t\_allow\_all\_else** term as the final term in the firewall filter explicitly configures the firewall filter to accept any traffic that the filter did not specifically discard :

```
term t_allow_all_else {
 then accept;
}
```

Following this best practice can simplify troubleshooting of the firewall filter.

**Best Practice: Explicitly Reject Any Traffic That Is Not Specifically Accepted**

On the other hand, you might want a firewall filter to reject any traffic that the firewall filter does not specifically accept. In this case, we recommend that you configure the firewall filter with a final term that specifies the **reject** terminating action.

In the following example snippet, configuring the **t\_deny\_all\_else** term as the final term in the firewall filter explicitly configures the firewall filter to reject any traffic that the filter did not specifically accept:

```
term t_deny_all_else {
 then reject;
}
```

Following this best practice can simplify troubleshooting of the firewall filter.

**Multiple Firewall Filters Attached to a Single Interface**

On supported device interfaces, you can attach multiple firewall filters to a single interface. For more information, see [“Multiple Standard Firewall Filters Applied as a List Overview” on page 4505](#).



**NOTE:** On supported interfaces, you can attach a protocol-independent (family any) firewall filter and a protocol-specific (family inet or family inet6) firewall filter to the same interface. The protocol-independent firewall filter executes first. For more information, see [“Guidelines for Applying Standard Firewall Filters” on page 4483](#).

---

### ***Single Firewall Filter Attached to Multiple Interfaces***

On supported interfaces, you can associate a single firewall filter with multiple interfaces, and Junos OS creates an *interface-specific instance* of that firewall filter for each associated interface.

- Junos OS associates each interface-specific instantiation of a firewall filter with a system-generated, interface-specific name.
- For any **count** actions in the filter terms, the Packet Forwarding Engine maintains separate, interface-specific counters, and Junos OS associates each counter with a system-generated, interface-specific name.
- For any **policer** actions in the filter terms, Junos OS creates separate, interface-specific instances of the policer actions.

For more information, see *Interface-Specific Firewall Filter Instances Overview*.

#### **Related Documentation**

- [Standard Firewall Filter Match Conditions for Protocol-Independent Traffic on page 4706](#)
- [Standard Firewall Filter Terminating Actions on page 4742](#)
- [How Service Filters Evaluate Packets on page 4523](#)
- [How Simple Filters Evaluate Packets on page 4530](#)
- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Understanding How to Use Standard Firewall Filters on page 3293](#)

### ***Guidelines for Configuring Standard Firewall Filters***

This topic covers the following information:

- [Statement Hierarchy for Configuring Standard Firewall Filters on page 4478](#)
- [Standard Firewall Filter Protocol Families on page 4479](#)
- [Standard Firewall Filter Names and Options on page 4480](#)
- [Standard Firewall Filter Terms on page 4480](#)
- [Standard Firewall Filter Match Conditions on page 4480](#)
- [Standard Firewall Filter Actions on page 4482](#)

### ***Statement Hierarchy for Configuring Standard Firewall Filters***

To configure a standard firewall filter, you can include the following statements. For an IPv4 standard firewall filter, the **family inet** statement is optional.

```
firewall {
```

```

family family-name {
 filter filter-name {
 accounting-profile name;
 interface-specific;
 physical-interface-filter;
 term term-name {
 filter filter-name;
 }
 term term-name {
 from {
 match-conditions;
 ip-version ipv4 {
 match-conditions;
 protocol (tcp | udp) {
 match conditions;
 }
 }
 }
 then {
 actions;
 }
 }
 }
}

```

You can include the firewall configuration at one of the following hierarchy levels:

- [edit]
- [edit logical-systems *logical-system-name*]



**NOTE:** For stateless firewall filtering, you must allow the output tunnel traffic through the firewall filter applied to input traffic on the interface that is the next-hop interface toward the tunnel destination. The firewall filter affects only the packets exiting the router (or switch) by way of the tunnel.

### Standard Firewall Filter Protocol Families

A standard firewall filter configuration is specific to a particular protocol family. Under the **firewall** statement, include one of the following statements to specify the protocol family for which you want to filter traffic:

- **family any**—To filter protocol-independent traffic.
- **family inet**—To filter Internet Protocol version 4 (IPv4) traffic.
- **family inet6**—To filter Internet Protocol version 6 (IPv6) traffic.
- **family mpls**—To filter MPLS traffic.
- **family vpls**—To filter virtual private LAN service (VPLS) traffic.
- **family ccc**—To filter Layer 2 circuit cross-connection (CCC) traffic.

- **family bridge**—To filter Layer 2 bridging traffic for MX Series 3D Universal Edge Routers only.
- **family ethernet-switching**—To filter Layer 2 (Ethernet) traffic.

The **family *family-name*** statement is required only to specify a protocol family other than IPv4. To configure an IPv4 firewall filter, you can configure the filter at the **[edit firewall]** hierarchy level without including the **family inet** statement, because the **[edit firewall]** and **[edit firewall family inet]** hierarchy levels are equivalent.

### ***Standard Firewall Filter Names and Options***

Under the **family *family-name*** statement, you can include **filter *filter-name*** statements to create and name standard firewall filters. The filter name can contain letters, numbers, and hyphens (-) and be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" ").

At the **[edit firewall family *family-name* filter *filter-name*]** hierarchy level, the following statements are optional:

- **accounting-profile**
- **interface-specific**
- **physical-interface-filter**

### ***Standard Firewall Filter Terms***

Under the **filter *filter-name*** statement, you can include **term *term-name*** statements to create and name filter terms.

- You must configure at least one term in a firewall filter.
- You must specify a unique name for each term within a firewall filter. The term name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" ").
- The order in which you specify terms within a firewall filter configuration is important. Firewall filter terms are evaluated in the order in which they are configured. By default, new terms are always added to the end of the existing filter. You can use the **insert** configuration mode command to reorder the terms of a firewall filter.

At the **[edit firewall family *family-name* filter *filter-name* term *term-name*]** hierarchy level, the **filter *filter-name*** statement is not valid in the same term as **from** or **then** statements. When included at this hierarchy level, the **filter *filter-name*** statement is used to *nest* firewall filters.

### ***Standard Firewall Filter Match Conditions***

Standard firewall filter match conditions are specific to the type of traffic being filtered.

With the exception of MPLS-tagged IPv4 or IPv6 traffic, you specify the term's match conditions under the **from** statement. For MPLS-tagged IPv4 traffic, you specify the term's IPv4 address-specific match conditions under the **ip-version ipv4** statement and the term's IPv4 port-specific match conditions under the **protocol (tcp | udp)** statement.

For MPLS-tagged IPv6 traffic, you specify the term's IPv6 address-specific match conditions under the **ip-version ipv6** statement and the term's IPv6 port-specific match conditions under the **protocol (tcp | udp)** statement.

[Table 330 on page 4481](#) describes the types of traffic for which you can configure standard stateless firewall filters.

**Table 330: Standard Firewall Filter Match Conditions by Protocol Family**

| Traffic Type                 | Hierarchy Level at Which Match Conditions Are Specified                                                                                                                                                                                                                                             |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Protocol-independent         | <p><b>[edit firewall family any filter <i>filter-name</i> term <i>term-name</i>]</b></p> <p>For the complete list of match conditions, see <a href="#">“Standard Firewall Filter Match Conditions for Protocol-Independent Traffic” on page 4706</a>.</p>                                           |
| IPv4                         | <p><b>[edit firewall family inet filter <i>filter-name</i> term <i>term-name</i>]</b></p> <p>For the complete list of match conditions, see <a href="#">“Standard Firewall Filter Match Conditions for IPv4 Traffic” on page 4707</a>.</p>                                                          |
| IPv6                         | <p><b>[edit firewall family inet6 filter <i>filter-name</i> term <i>term-name</i>]</b></p> <p>For the complete list of match conditions, see <a href="#">“Standard Firewall Filter Match Conditions for IPv6 Traffic” on page 4716</a>.</p>                                                         |
| MPLS                         | <p><b>[edit firewall family mpls filter <i>filter-name</i> term <i>term-name</i>]</b></p> <p>For the complete list of match conditions, see <a href="#">“Standard Firewall Filter Match Conditions for MPLS Traffic” on page 3565</a>.</p>                                                          |
| IPv4 addresses in MPLS flows | <p><b>[edit firewall family mpls filter <i>filter-name</i> term <i>term-name</i> ip-version ipv4 ]</b></p> <p>For the complete list of match conditions, see <a href="#">“Standard Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic” on page 4725</a>.</p>                     |
| IPv4 ports in MPLS flows     | <p><b>[edit firewall family mpls filter <i>filter-name</i> term <i>term-name</i> ip-version ipv4 protocol (tcp   udp)]</b></p> <p>For the complete list of match conditions, see <a href="#">“Standard Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic” on page 4725</a>.</p> |
| IPv6 addresses in MPLS flows | <p><b>[edit firewall family mpls filter <i>filter-name</i> term <i>term-name</i> ip-version ipv6 ]</b></p> <p>For the complete list of match conditions, see <a href="#">“Standard Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic” on page 4725</a>.</p>                     |
| IPv6 ports in MPLS flows     | <p><b>[edit firewall family mpls filter <i>filter-name</i> term <i>term-name</i> ip-version ipv6 protocol (tcp   udp)]</b></p> <p>For the complete list of match conditions, see <a href="#">“Standard Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic” on page 4725</a>.</p> |
| VPLS                         | <p><b>[edit firewall family vpls filter <i>filter-name</i> term <i>term-name</i>]</b></p> <p>For the complete list of match conditions, see <a href="#">“Standard Firewall Filter Match Conditions for VPLS Traffic” on page 4727</a>.</p>                                                          |

**Table 330: Standard Firewall Filter Match Conditions by Protocol Family (*continued*)**

| Traffic Type                                    | Hierarchy Level at Which Match Conditions Are Specified                                                                                                                                                                                                                         |
|-------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Layer 2 CCC                                     | <p>[edit firewall family ccc filter <i>filter-name</i> term <i>term-name</i>]</p> <p>For the complete list of match conditions, see “Standard Firewall Filter Match Conditions for Layer 2 CCC Traffic” on page 4734.</p>                                                       |
| Layer 2 Bridging                                | <p>[edit firewall family bridge filter <i>filter-name</i> term <i>term-name</i>]</p>                                                                                                                                                                                            |
| (MX Series routers and EX Series switches only) | <p>[edit firewall family ethernet-switching filter <i>filter-name</i> term <i>term-name</i>]<br/>(for EX Series switches only)</p> <p>For the complete list of match conditions, see “Standard Firewall Filter Match Conditions for Layer 2 Bridging Traffic” on page 4737.</p> |

If you specify an IPv6 address in a match condition (the **address**, **destination-address**, or **source-address** match conditions), use the syntax for text representations described in RFC 2373, *IP Version 6 Addressing Architecture*. For more information about IPv6 addresses, see “IPv6 Overview” and “IPv6 Standards” in the *Junos OS Routing Protocols Configuration Guide*.

### Standard Firewall Filter Actions

Under the **then** statement for a standard stateless firewall filter term, you can specify the actions to be taken on a packet that matches the term.

Table 331 on page 4482 summarizes the types of actions you can specify in a standard stateless firewall filter term.

**Table 331: Standard Firewall Filter Action Categories**

| Type of Action | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Comment                                                          |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| Terminating    | <p>Halts all evaluation of a firewall filter for a specific packet. The router (or switch) performs the specified action, and no additional terms are used to examine the packet.</p> <p>You can specify only one <i>terminating action</i> in a standard firewall filter. You can, however, specify one terminating action with one or more <i>nonterminating actions</i> in a single term. For example, within a term, you can specify <b>accept</b> with <b>count</b> and <b>syslog</b>.</p> | See “Standard Firewall Filter Terminating Actions” on page 4742. |



Table 331: Standard Firewall Filter Action Categories (*continued*)

| Type of Action | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Comment                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Nonterminating | Performs other functions on a packet (such as incrementing a counter, logging information about the packet header, sampling the packet data, or sending information to a remote host using the system log functionality), but any additional terms are used to examine the packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | See <a href="#">"Standard Firewall Filter Nonterminating Actions"</a> on page 4744.                                                                                                                                                                                                                                                                                                                                                                              |
| Flow control   | <p>For standard stateless firewall filters only, the <b>next term</b> action directs the router (or switch) to perform configured actions on the packet and then, rather than terminate the filter, use the next term in the filter to evaluate the packet. If the <b>next term</b> action is included, the matching packet is evaluated against the next term in the firewall filter. Otherwise, the matching packet is not evaluated against subsequent terms in the firewall filter.</p> <p>For example, when you configure a term with the nonterminating action <b>count</b>, the term's action changes from an implicit <b>discard</b> to an implicit <b>accept</b>. The <b>next term</b> action forces the continued evaluation of the firewall filter.</p> | <p>You cannot configure the <b>next term</b> action with a terminating action in the same filter term. However, you can configure the next term action with another nonterminating action in the same filter term.</p> <p>A maximum of 1024 <b>next term</b> actions are supported per standard stateless firewall filter configuration. If you configure a standard filter that exceeds this limit, your candidate configuration results in a commit error.</p> |

**Related Documentation**

- [Guidelines for Applying Standard Firewall Filters on page 4483](#)
- [Understanding How to Use Standard Firewall Filters on page 3293](#)

**Guidelines for Applying Standard Firewall Filters**

This topic covers the following information:

- [Applying Firewall Filters Overview on page 4483](#)
- [Statement Hierarchy for Applying Firewall Filters on page 4484](#)
- [Restrictions on Applying Firewall Filters on page 4485](#)

**Applying Firewall Filters Overview**

You can apply a standard firewall filter to a loopback interface on the router or to a physical or logical interface on the router. [Table 332 on page 4484](#) summarizes the behavior of firewall filters based on the point to which you attach the filter.

Table 332: Firewall Filter Behavior by Filter Attachment Point

| Filter Attachment Point                                                                    | Filter Behavior                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Loopback interface                                                                         | The router's loopback interface, <b>lo0</b> , is the interface to the Routing Engine and carries no data packets. When you apply a firewall filter to the loopback interface, the filter evaluates the local packets received or transmitted by the Routing Engine.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Physical interface or logical interface                                                    | When you apply a filter to a physical interface on the router or to a logical interface (or member of an aggregated Ethernet bundle defined on the interface), the filter evaluates all data packet that pass through that interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Multiple interfaces                                                                        | <p>You can use the same firewall filter one or more times.</p> <p>On M Series routers, except the M120 and M320 routers, if you apply a firewall filter to multiple interfaces, the filter acts on the sum of traffic entering or exiting those interfaces.</p> <p>On T Series, M120, M320, and MX Series routers, interfaces are distributed among multiple packet-forwarding components. On these routers, you can configure firewall filters and service filters that, when applied to multiple interfaces, act on the individual traffic streams entering or exiting each interface, regardless of the sum of traffic on the multiple interfaces.</p> <p>For more information, see <i>Interface-Specific Firewall Filter Instances Overview</i>.</p>                                                                                                                                                                                                                                                                                                             |
| Single interface with protocol-independent and protocol-specific firewall filters attached | <p>For interfaces hosted on the following hardware only, you can attach a protocol-independent (<b>family any</b>) firewall filter and a protocol-specific (<b>family inet</b> or <b>family inet6</b>) firewall filter simultaneously. The protocol-independent firewall executes first.</p> <ul style="list-style-type: none"> <li>ACX Series Universal Access Routers</li> <li>Flexible PIC Concentrators (FPCs) in M7i and M10i Multiservice Edge Routers</li> <li>Modular Interface Cards (MICs) and Modular Port Concentrators (MPCs) in MX Series 3D Universal Edge Routers</li> <li>T Series Core Routers</li> </ul> <p><b>NOTE:</b></p> <p>Interfaces hosted on the following hardware do not support protocol-independent firewall filters:</p> <ul style="list-style-type: none"> <li>Forwarding Engine Boards (FEBs) in M120 routers</li> <li>Enhanced III FPCs in M320 routers</li> <li>FPC2 and FPC3 modules in MX Series routers</li> <li>Dense Port Concentrators (DPCs) in MX Series routers</li> <li>PTX Series Packet Transport Routers</li> </ul> |

#### Statement Hierarchy for Applying Firewall Filters

To apply a standard firewall filter to a logical interface, configure the **filter** statement for the logical interface defined under either the **[edit]** or **[edit logical-systems logical-system-name]** hierarchy level. Under the **filter** statement, you can include one or more of the following statements: **group group-number**, **input filter-name**, **input-list filter-name**, **output filter-name**, or **output-list filter-name**. The hierarchy level at which you attach the **filter** statement depends on the filter type and device type you are configuring.

### Protocol-Independent Firewall Filters on MX Series Routers

To apply a protocol-independent firewall filter to a logical interface on an MX Series router, configure the **filter** statement *directly* under the logical unit:

```
interfaces {
 interface-name {
 unit logical-unit-number {
 filter {
 group group-number;
 input filter-name;
 input-list [filter-names];
 output filter-name;
 output-list [filter-names];
 }
 }
 }
}
```

### All Other Firewall Filters on Logical Interfaces

To apply a standard firewall filter to a logical interface for all cases *other than* a protocol-independent filter on an MX Series router, configure the **filter** statement under the protocol family:

```
interfaces {
 interface-name {
 unit logical-unit-number {
 family family-name {
 ...
 filter {
 group group-number;
 input filter-name;
 input-list [filter-names];
 output filter-name;
 output-list [filter-names];
 }
 }
 }
 }
}
```

### Restrictions on Applying Firewall Filters

- [Number of Input and Output Filters Per Logical Interface on page 4485](#)
- [MPLS and Layer 2 CCC Firewall Filters in Lists on page 4486](#)
- [Layer 2 CCC Firewall Filters on MX Series Routers and EX Series Switches on page 4486](#)

### Number of Input and Output Filters Per Logical Interface

**Input filters**—Although you can use the same filter multiple times, you can apply only one input filter or one input filter list to an interface.

- To specify a single firewall filter to be used to evaluate packets received on the interface, include the **input filter-name** statement in the **filter** stanza.

- To specify an ordered list of firewall filters to be used to evaluate packets received on the interface, include the **input-list [ *filter-names* ]** statement in the **filter** stanza. You can specify up to 16 firewall filters for the filter input list.

**Output filters**—Although you can use the same filter multiple times, you can apply only one output filter or one output filter list to an interface.

- To specify a single firewall filter to be used to evaluate packets transmitted on the interface, include the **output *filter-name*** statement in the **filter** stanza.
- To specify an ordered list of firewall filters to be used to evaluate packets transmitted on the interface, include the **output-list [ *filter-names* ]** statement in the **filter** stanza. You can specify up to 16 firewall filters in a filter output list.

### ***MPLS and Layer 2 CCC Firewall Filters in Lists***

The **input-list *filter-names*** and **output-list *filter-names*** statements for firewall filters for the **ccc** and **mpls** protocol families are supported on all interfaces with the exception of the following:

- Management interfaces and internal Ethernet interfaces (**fxp** or **em0**)
- Loopback interfaces (**lo0**)
- USB modem interfaces (**umd**)

### ***Layer 2 CCC Firewall Filters on MX Series Routers and EX Series Switches***

Only on MX Series routers and EX Series switches, you cannot apply a Layer 2 CCC stateless firewall filter (a firewall filter configured at the **[edit firewall filter family ccc]** hierarchy level) as an output filter. On MX Series routers and EX Series switches, firewall filters configured for the **family ccc** statement can be applied only as input filters.

#### **Related Documentation**

- [family \(Firewall\) on page 4788](#)
- [family \(Interfaces\) on page 1883](#)
- [filter \(Applying to a Logical Interface\) on page 4790](#)
- [filter \(Configuring\) on page 1422](#)
- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Understanding How to Use Standard Firewall Filters on page 3293](#)

### ***Understanding How to Use Standard Firewall Filters***

This topic covers the following information:

- [Using Standard Firewall Filters to Affect Local Packets on page 4486](#)
- [Using Standard Firewall Filters to Affect Data Packets on page 4487](#)

#### ***Using Standard Firewall Filters to Affect Local Packets***

On a router (or switch), you can configure one physical loopback interface, **lo0**, and one or more addresses on the interface. The loopback interface is the interface to the Routing

Engine, which runs and monitors all the control protocols. The loopback interface carries local packets only. Standard firewall filters applied to the loopback interface affect the local packets destined for or transmitted from the Routing Engine.



**NOTE:** When you create an additional loopback interface, it is important to apply a filter to it so the Routing Engine is protected. We recommend that when you apply a filter to the loopback interface, you include the `apply-groups` statement. Doing so ensures that the filter is automatically inherited on every loopback interface, including `lo0` and other loopback interfaces.

### ***Trusted Sources***

The typical use of a standard stateless firewall filter is to protect the Routing Engine processes and resources from malicious or untrusted packets. To protect the processes and resources owned by the Routing Engine, you can use a standard stateless firewall filter that specifies which protocols and services, or applications, are allowed to reach the Routing Engine. Applying this type of filter to the loopback interface ensures that the local packets are from a trusted source and protects the processes running on the Routing Engine from an external attack.

### ***Flood Prevention***

You can create standard stateless firewall filters that limit certain TCP and ICMP traffic destined for the Routing Engine. A router (or switch) without this kind of protection is vulnerable to TCP and ICMP flood attacks, which are also called denial-of-service (DoS) attacks. For example:

- A TCP flood attack of SYN packets initiating connection requests can overwhelm the device until it can no longer process legitimate connection requests, resulting in denial of service.
- An ICMP flood can overload the device with so many echo requests (ping requests) that it expends all its resources responding and can no longer process valid network traffic, also resulting in denial of service.

Applying the appropriate firewall filters to the Routing Engine protects against these types of attacks.

### ***Using Standard Firewall Filters to Affect Data Packets***

Standard firewall filters that you apply to your router's (or switch's) transit interfaces evaluate only the user data packets that transit the router (or switch) from one interface directly to another as they are being forwarded from a source to a destination. To protect the network as a whole from unauthorized access and other threats at specific interfaces, you can apply firewall filters router (or switch) transit interfaces.

#### **Related Documentation**

- [How Standard Firewall Filters Evaluate Packets on page 4475](#)
- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Guidelines for Applying Standard Firewall Filters on page 4483](#)

## Standard Firewall Filter Match Conditions Overview

---

- [Firewall Filter Match Conditions Based on Numbers or Text Aliases on page 4488](#)
- [Firewall Filter Match Conditions Based on Bit-Field Values on page 4489](#)
- [Firewall Filter Match Conditions Based on Address Fields on page 4493](#)
- [Firewall Filter Match Conditions Based on Address Classes on page 4502](#)

### ***Firewall Filter Match Conditions Based on Numbers or Text Aliases***

This topic covers the following information:

- [Matching on a Single Numeric Value on page 4488](#)
- [Matching on a Range of Numeric Values on page 4488](#)
- [Matching on a Text Alias for a Numeric Value on page 4488](#)
- [Matching on a List of Numeric Values or Text Aliases on page 4488](#)

#### ***Matching on a Single Numeric Value***

You can specify a firewall filter match condition based on whether a particular packet field value is a specified numeric value. In the following example, a match occurs if the packet source port number is 25:

```
[edit firewall family inet filter filter1 term term1 from]
user@host# set source-port 25
```

#### ***Matching on a Range of Numeric Values***

You can specify a firewall filter match condition based on whether a particular packet field value falls within a specified range of numeric values. In the following example, a match occurs for source ports values from 1024 through 65,535, inclusive:

```
[edit firewall family inet filter filter2 term term1 from]
user@host# set source-port 1024-65535
```

#### ***Matching on a Text Alias for a Numeric Value***

You can specify a firewall filter match condition based on whether a particular packet field value is a numeric value that you specify by using a text string as an *alias* for the numeric value. In the following example, a match occurs if the packet source port number is 25. For the **source-port** and **destination-port** match conditions, the text alias **smtp** corresponds to the numeric value 25.

```
[edit firewall family inet filter filter3 term term1 from]
user@host# set source-port smtp
```

#### ***Matching on a List of Numeric Values or Text Aliases***

You can specify a firewall filter match condition based on whether a particular packet field value matches any one of multiple numeric values or text aliases that you specify within square brackets and delimited by spaces. In the following example, a match occurs if the packet source port number is any of the following values: 20 (which corresponds to the text aliases **ftp-data**), 25, or any value from 1024 through 65535.

```
[edit firewall family inet filter filter3 term term1 from]
user@host# set source-port [smtp ftp-data 25 1024-65535]
```

**Related Documentation**

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Firewall Filter Match Conditions Based on Bit-Field Values on page 4489](#)
- [Firewall Filter Match Conditions Based on Address Fields on page 4493](#)
- [Firewall Filter Match Conditions Based on Address Classes on page 4502](#)

**Firewall Filter Match Conditions Based on Bit-Field Values**

- [Match Conditions for Bit-Field Values on page 4489](#)
- [Match Conditions for Common Bit-Field Values or Combinations on page 4490](#)
- [Logical Operators for Bit-Field Values on page 4491](#)
- [Matching on a Single Bit-Field Value or Text Alias on page 4492](#)
- [Matching on Multiple Bit-Field Values or Text Aliases on page 4492](#)
- [Matching on a Negated Bit-Field Value on page 4492](#)
- [Matching on the Logical OR of Two Bit-Field Values on page 4493](#)
- [Matching on the Logical AND of Two Bit-Field Values on page 4493](#)
- [Grouping Bit-Field Match Conditions on page 4493](#)

**Match Conditions for Bit-Field Values**

Table 333 on page 4489 lists the firewall filter match conditions that are based on whether certain bit fields in a packet are set or not set. The second and third columns list the types of traffic for which the match condition is supported.

**Table 333: Binary and Bit-Field Match Conditions for Firewall Filters**

| Bit-Field Match Condition           | Match Values                                                                                        | Protocol Families for Standard Stateless Firewall Filters | Protocol Families for Service Filters |
|-------------------------------------|-----------------------------------------------------------------------------------------------------|-----------------------------------------------------------|---------------------------------------|
| <b>fragment-flags <i>flags</i></b>  | Hexadecimal values or text aliases for the three-bit IP fragmentation flags field in the IP header. | <b>family inet</b>                                        | <b>family inet</b>                    |
| <b>fragment-offset <i>value</i></b> | Hexadecimal values or text aliases for the 13-bit fragment offset field in the IP header.           | <b>family inet</b>                                        | <b>family inet</b>                    |

<sup>†</sup> The Junos OS does not automatically check the first fragment bit when matching TCP flags for IPv4 traffic. To check the first fragment bit for IPv4 traffic only, use the **first-fragment** match condition.

**Table 333: Binary and Bit-Field Match Conditions for Firewall Filters** (*continued*)

| Bit-Field Match Condition           | Match Values                                                                                                | Protocol Families for Standard Stateless Firewall Filters                                                                                                   | Protocol Families for Service Filters     |
|-------------------------------------|-------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|
| <b>tcp-flags value</b> <sup>†</sup> | Hexadecimal values or text aliases for the low-order 6 bits of the 8-bit TCP flags field in the TCP header. | <b>family inet</b><br><b>family inet6</b><br><b>family vpls</b><br><b>family bridgefamily</b><br><b>ethernet-switching</b><br>(only for EX Series switches) | <b>family inet</b><br><b>family inet6</b> |

<sup>†</sup> The Junos OS does not automatically check the first fragment bit when matching TCP flags for IPv4 traffic. To check the first fragment bit for IPv4 traffic only, use the **first-fragment** match condition.

#### **Match Conditions for Common Bit-Field Values or Combinations**

Table 334 on page 4490 describes firewall filter match conditions that are based on whether certain commonly used values or *combinations* of bit fields in a packet are set or not set.

You can use text synonyms to specify some common bit-field matches. In the previous example, you can specify **tcp-initial** as the same match condition.



#### **NOTE:**

Some of the numeric range and bit-field match conditions allow you to specify a text synonym. For a complete list of synonyms:

- If you are using the J-Web interface, select the synonym from the appropriate list.
- If you are using the CLI, type a question mark (?) after the from statement.

**Table 334: Bit-Field Match Conditions for Common Combinations**

| Match Condition       | Description                                                                                                                        | Protocol Families for Standard Stateless Firewall Filters | Protocol Families for Service Filters |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------|---------------------------------------|
| <b>first-fragment</b> | Text alias for the bit-field match condition <b>fragment-offset 0</b> , which indicates the first fragment of a fragmented packet. | <b>family inet</b>                                        | <b>family inet</b>                    |



**Table 334: Bit-Field Match Conditions for Common Combinations** (*continued*)

| Match Condition        | Description                                                                                                                                                             | Protocol Families for Standard Stateless Firewall Filters | Protocol Families for Service Filters |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------|---------------------------------------|
| <b>is-fragment</b>     | Text alias for the bit-field match condition <b>fragment-offset 0 except</b> , which indicates a trailing fragment of a fragmented packet.                              | <b>family inet</b>                                        | <b>family inet</b>                    |
| <b>tcp-established</b> | Alias for the bit-field match condition <b>tcp-flags "(ack   rst)"</b> , which indicates an established TCP session, but not the first packet of a TCP connection.      | <b>family inet</b><br><b>family inet6</b>                 | —                                     |
| <b>tcp-initial</b>     | Alias for the bit-field match condition <b>tcp-flags "(!ack &amp; syn)"</b> , which indicates the first packet of a TCP connection, but not an established TCP session. | <b>family inet</b><br><b>family inet6</b>                 | —                                     |

#### Logical Operators for Bit-Field Values

Table 335 on page 4491 lists the logical operators you can apply to *single* bit-field values when specifying stateless firewall filter match conditions. The operators are listed in order, from highest precedence to lowest precedence. Operations are left-associative, meaning that the operations are processed from left to right.

**Table 335: Bit-Field Logical Operators**

| Precedence Order | Bit-Field Logical Operator                                                                                                         | Description                                                                                                 |
|------------------|------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------|
| 1                | <b>(<i>complex-match-condition</i>)</b>                                                                                            | Grouping—The complex match condition is evaluated before any operators outside the parentheses are applied. |
| 2                | <b>! <i>match-condition</i></b>                                                                                                    | Negation—A match occurs if the match condition is false.                                                    |
| 3                | <b><i>match-condition-1</i> &amp; <i>match-condition-2</i></b><br>or<br><b><i>match-condition-1</i> + <i>match-condition-2</i></b> | Logical AND—A match occurs if both match conditions are true.                                               |
| 4                | <b><i>match-condition-1</i>   <i>match-condition-2</i></b><br>or<br><b><i>match-condition-1</i> , <i>match-condition-2</i></b>     | Logical OR—A match occurs if either match condition is true.                                                |

### ***Matching on a Single Bit-Field Value or Text Alias***

For the **fragment-flags** and **tcp-flags** bit-match conditions, you can specify firewall filter match conditions based on whether a particular bit in the packet field is set or not set.

- Numeric value to specify a single bit—You can specify a single bit-field match condition by using a numeric value that has one bit set. Depending on the match condition, you can specify a decimal value, a binary value, or a hexadecimal value. To specify a binary value, specify the number with the prefix **b**. To specify a hexadecimal value, specify the number with the prefix **0x**.

In the following example, a match occurs if the **RST** bit in the TCP flags field is set:

```
[edit firewall family inet filter filter_tcp_rst_number term term1 from]
user@host# set tcp-flags 0x04
```

- Text alias to specify a single bit—You generally specify a single bit-field match condition by using a text alias enclosed in double-quotation marks (" ").

In the following example, a match occurs if the **RST** bit in the TCP flags field is set:

```
[edit firewall family inet filter filter_tcp_rst_alias term term1 from]
user@host# set tcp-flags "rst"
```

### ***Matching on Multiple Bit-Field Values or Text Aliases***

You can specify a firewall filter match condition based on whether a particular set of bits in a packet field are set.

- Numeric values to specify multiple set bits—When you specify a numeric value whose binary representation has more than one set bit, the value is treated as a logical AND of the set bits.

In the following example, the two match conditions are the same. A match occurs if either bit **0x01** or **0x02** is not set:

```
[edit firewall family inet filter reset_or_not_initial_packet term term5 from]
user@host# set tcp-flags "!0x3"
user@host# set tcp-flags "!(0x01 & 0x02)"
```

- Text aliases that specify common bit-field matches—You can use text aliases to specify some common bit-field matches. You specify these matches as a single keyword.

In the following example, the **tcp-established** condition, which is an alias for **"(ack | rst)"**, specifies that a match occurs on TCP packets other than the first packet of a connection:

```
[edit firewall family inet filter reset_or_not_initial_packet term term6 from]
user@host# set tcp-established
```

### ***Matching on a Negated Bit-Field Value***

To negate a match, precede the value with an exclamation point.

In the following example, a match occurs if the **RST** bit in the TCP flags field is *not* set:

```
[edit firewall family inet filter filter_tcp_rst term term1 from]
user@host# set tcp-flags "!rst"
```

### ***Matching on the Logical OR of Two Bit-Field Values***

You can use the *logical OR operator* (`|` or `,`) to specify that a match occurs if a bit field matches either of two bit-field values specified.

In the following example, a match occurs if the packet is *not* the initial packet in a TCP session:

```
[edit firewall family inet filter not_initial_packet term term3 from]
user@host# set tcp-flags "!syn | ack"
```

In a TCP session, the SYN flag is set only in the initial packet sent, while the ACK flag is set in all packets sent after the initial packet. In a packet that is not the initial packet in a TCP session, either the SYN flag is not set or the ACK flag is set.

### ***Matching on the Logical AND of Two Bit-Field Values***

You can use the *logical AND operator* (`&` or `+`) to specify that a match occurs if a bit field matches both of two bit-field values specified.

In the following example, a match occurs if the packet is the initial packet in a TCP session:

```
[edit firewall family inet filter initial_packet term term2 from]
user@host# set tcp-flags "syn & !ack"
```

In a TCP session, the SYN flag is set only in the initial packet sent, while the ACK flag is set in all packets sent after the initial packet. In a packet that is an initial packet in a TCP session, the SYN flag is set and the ACK flag is not set.

### ***Grouping Bit-Field Match Conditions***

You can use the *logical grouping notation* to specify that the complex match condition inside the parentheses is evaluated before any operators outside the parentheses are applied.

In the following example, a match occurs if the packet is a TCP reset or if the packet is not the initial packet in the TCP session:

```
[edit firewall family inet filter reset_or_not_initial_packet term term4 from]
user@host# set tcp-flags "!(syn & !ack) | rst"
```

In a TCP session, the SYN flag is set only in the initial packet sent, while the ACK flag is set in all packets sent after the initial packet. In a packet that is *not* the initial packet in a TCP session, the SYN flag is not set and the ACK field is set.

#### **Related Documentation**

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Firewall Filter Match Conditions Based on Numbers or Text Aliases on page 4488](#)
- [Firewall Filter Match Conditions Based on Address Fields on page 4493](#)
- [Firewall Filter Match Conditions Based on Address Classes on page 4502](#)

### ***Firewall Filter Match Conditions Based on Address Fields***

You can configure firewall filter match conditions that evaluate packet address fields—IPv4 source and destination addresses, IPv6 source and destination addresses,

or media access control (MAC) source and destination addresses—against specified addresses or prefix values.

- [Implied Match on the '0/0 except' Address for Firewall Filter Match Conditions Based on Address Fields on page 4494](#)
- [Matching an Address Field to a Subnet Mask or Prefix on page 4494](#)
- [Matching an Address Field to an Excluded Value on page 4495](#)
- [Matching Either IP Address Field to a Single Value on page 4499](#)
- [Matching an Address Field to Noncontiguous Prefixes on page 4499](#)
- [Matching an Address Field to a Prefix List on page 4501](#)

#### ***Implied Match on the '0/0 except' Address for Firewall Filter Match Conditions Based on Address Fields***

Every firewall filter match condition based on a set of addresses or address prefixes is associated with an implicit match on the address **0.0.0.0/0 except** (for IPv4 or VPLS traffic) or **0:0:0:0:0:0:0:0/0 except** (for IPv6 traffic). As a result, any packet whose specified address field does not match any of the specified addresses or address prefixes fails to match the entire term.

#### ***Matching an Address Field to a Subnet Mask or Prefix***

You can specify a single match condition to match a source address or destination address that falls within a specified address prefix.

#### ***IPv4 Subnet Mask Notation***

For an IPv4 address, you can specify a subnet mask value rather than a prefix length. For example:

```
[edit firewall family inet filter filter_on_dst_addr term term3 from]
user@host# set address 10.0.0.10/255.0.0.255
```

#### ***Prefix Notation***

To specify the address prefix, use the notation **prefix/prefix-length**. In the following example, a match occurs if a destination address matches the prefix **10.0.0.0/8**:

```
[edit firewall family inet filter filter_on_dst_addr term term1 from]
user@host# set destination-address 10.0.0.0/8
```

#### ***Default Prefix Length for IPv4 Addresses***

If you do not specify **/prefix-length** for an IPv4 address, the prefix length defaults to **/32**. The following example illustrates the default prefix value:

```
[edit firewall family inet filter filter_on_dst_addr term term2 from]
user@host# set destination-address 10
user@host# show
destination-address {
 10.0.0.0/32;
}
```

### Default Prefix Length for IPv6 Addresses

If you do not specify */prefix-length* for an IPv6 address, the prefix length defaults to */128*. The following example illustrates the default prefix value:

```
[edit firewall family inet6 filter filter_on_dst_addr term term1 from]
user@host# set destination-address ::10
user@host# show
destination-address {
 ::10/128;
}
```

### Default Prefix Length for MAC Addresses

If you do not specify */prefix-length* for a media access control (MAC) address of a VPLS, Layer 2 CCC, or Layer 2 bridging packet, the prefix length defaults to */48*. The following example illustrates the default prefix value:

```
[edit firewall family vpls filter filter_on_dst_mac_addr term term1 from]
user@host# set destination-mac-address 01:00:0c:cc:cc:cd
user@host# show
destination-address {
 01:00:0c:cc:cc:cd/48;
}
```

### Matching an Address Field to an Excluded Value

For the address-field match conditions, you can include the **except** keyword to specify that a match occurs for an address field that does not match the specified address or prefix.

### Excluding IP Addresses in IPv4 or IPv6 Traffic

For the following IPv4 and IPv6 match conditions, you can include the **except** keyword to specify that a match occurs for an IP address field that does not match the specified IP address or prefix:

- **address address except**—A match occurs if either the source IP address or the destination IP address does not match the specified address or prefix.
- **source-address address except**—A match occurs if the source IP address does not match the specified address or prefix.
- **destination-address address except**—A match occurs if the destination IP address does not match the specified address or prefix.

In the following example, a match occurs for any IPv4 destination addresses that fall under the **192.168.10.0/8** prefix, except for addresses that fall under **192.168.0.0/16**. All other addresses implicitly do not match this condition.

```
[edit firewall family inet filter filter_on_dst_addr term term1 from]
user@host# set 192.168.0.0/16 except
user@host# set 192.168.10.0/8
user@host# show
destination-address {
 192.168.0.0/16 except;
 192.168.10.0/8;
}
```

```
}
```

In the following example, a match occurs for any IPv4 destination address that does not fall within the prefix 10.1.1.0/24:

```
[edit firewall family inet filter filter_on_dst_addr term term24 from]
user@host# set destination-address 0.0.0.0/0
user@host# set destination-address 10.1.1.0/24 except
user@host# show
destination-address {
 0.0.0.0/0;
 10.1.1.0/24 except;
}
```

#### *Excluding IP Addresses in VPLS or Layer 2 Bridging Traffic*

For the following VPLS and Layer 2 bridging match conditions on MX Series routers and EX Series switches only, you can include the **except** keyword to specify that a match occurs for an IP address field that does not match the specified IP address or prefix:

- **ip-address address except**—A match occurs if either the source IP address or the destination IP address does not match the specified address or prefix.
- **source-ip-address address except**—A match occurs if the source IP address does not match the specified address or prefix.
- **destination-ip-address address except**—A match occurs if the destination IP address does not match the specified address or prefix.

In the following example for filtering VPLS traffic on an MX Series router and on an EX Series switch, a match occurs if the source IP address falls within the exception range of **55.0.1.0/255.0.255.0** and the destination IP address matches **55.0.0.0/8**:

```
[edit]
firewall {
 family vpls {
 filter fvpls {
 term 1 {
 from {
 ip-address {
 55.0.0.0/8;
 55.0.1.0/255.0.255.0 except;
 }
 }
 then {
 count from-55/8;
 discard;
 }
 }
 }
 }
}
```

#### ***Excluding MAC Addresses in VPLS or Layer 2 Bridging Traffic***

For the following VPLS or Layer 2 bridging traffic match conditions, you can include the **except** keyword to specify that a match occurs for a MAC address field that does not match the specified MAC address or prefix:

- **source-mac-address address except**—A match occurs if the source MAC address does not match the specified address or prefix.
- **destination-mac-address address except**—A match occurs if either the destination MAC address does not match the specified address or prefix.

#### ***Excluding All Addresses Requires an Explicit Match on the '0/0' Address***

If you specify a firewall filter match condition that consists of one or more address-*exception* match conditions (address match conditions that use the **except** keyword) but no *matchable* address match conditions, packets that do not match any of the configured prefixes fails the overall match operation. To configure a firewall filter term of address-exception match conditions to match any address that is not in the prefix list, include an explicit match of **0/0** so that the term contain a matchable address.

For the following example firewall filter for IPv4 traffic, the **from-trusted-addresses** term fails to discard matching traffic, and the **INTRUDERS-COUNT** counter is missing from the output of the **show firewall** operational mode command:

```
[edit]
user@host# show policy-options
prefix-list TRUSTED-ADDRESSES {
 10.2.1.0/24;
 192.168.122.0/24;
}
```

```
[edit firewall family inet filter protect-RE]
```

```
user@host# show
```

```
term from-trusted-addresses {
 from {
 source-prefix-list {
 TRUSTED-ADDRESSES except;
 }
 protocol icmp;
 }
 then {
 count INTRUDERS-COUNT;
 discard;
 }
}
term other-icmp {
 from {
 protocol icmp;
 }
 then {
 count VALID-COUNT;
 accept;
 }
}
term all {
 then accept;
}
```

```
[edit]
```

```
user@host# run show firewall
```

```
Filter: protect-RE
```

```
Counters:
```

| Name        | Bytes | Packets |
|-------------|-------|---------|
| VALID-COUNT | 2770  | 70      |

```
Filter: __default_bpdu_filter__
```

To cause a filter term of address-exception match conditions to match any address that is not in the prefix list, include an explicit match of **0/0** in the set of match conditions:

```
[edit firewall family inet filter protect-RE]
```

```
user@host# show term from-trusted-addresses
```

```
from {
 source-address {
 0.0.0.0/0;
 }
 source-prefix-list {
 TRUSTED-ADDRESSES except;
 }
 protocol icmp;
}
```

With the addition of the **0.0.0.0/0** source prefix address to the match condition, the **from-trusted-addresses** term discards matching traffic, and the INTRUDERS-COUNT counter displays in the output of the **show firewall** operational mode command:

```
[edit]
```

```
user@host# run show firewall
```

```
Filter: protect-RE
```

```
Counters:
```



| Name                            | Bytes | Packets |
|---------------------------------|-------|---------|
| VALID-COUNT                     | 2770  | 70      |
| INTRUDERS-COUNT                 | 420   | 5       |
| Filter: __default_bpdu_filter__ |       |         |

### ***Matching Either IP Address Field to a Single Value***

For IPv4 and IPv6 traffic and for VPLS and Layer 2 bridging traffic only on MX Series routers and on EX Series switches, you can use a single match condition to match a single address or prefix value to either the source or destination IP address field.

### ***Matching Either IP Address Field in IPv4 or IPv6 Traffic***

For IPv4 or IPv6 traffic, you can use a single match condition to specify the same address or prefix value as the match for either the source or destination IP address field. Instead of creating separate filter terms that specify the same address for the **source-address** and **destination-address** match conditions, you use only the **address** match condition. A match occurs if *either* the source IP address *or* the destination IP address matches the specified address or prefix.

If you use the **except** keyword with the **address** match condition, a match occurs if *both* the source IP address and the destination IP address match the specified value *before* the exception applies.

In a firewall filter term that specifies either the **source-address** or the **destination-address** match condition, you cannot also specify the **address** match condition.

### ***Matching Either IP Address Field in VPLS or Layer 2 Bridging Traffic***

For VPLS or Layer 2 bridging traffic on MX Series routers and EX Series switches only, you can use a single match condition to specify the same address or prefix value as the match for either the source or destination IP address field. Instead of creating separate filter terms that specify the same address for the **source-ip-address** and **destination-ip-address** match conditions, you use only the **ip-address** match condition. A match occurs if *either* the source IP address *or* the destination IP address matches the specified address or prefix.

If you use the **except** keyword with the **ip-address** match condition, a match occurs if *both* the source IP address and the destination IP address match the specified value *before* the exception applies.

In a firewall filter term that specifies either the **source-ip-address** or the **destination-ip-address** match condition, you cannot also specify the **ip-address** match condition.

### ***Matching an Address Field to Noncontiguous Prefixes***

For IPv4 traffic only, specify a single match condition to match the IP source or destination address field to any prefix specified. The prefixes do not need to be contiguous. That is, the prefixes under the **source-address** or **destination-address** match condition do not need to be adjacent or neighboring to one another.

In the following example, a match occurs if a destination address matches either the **10.0.0.0/8** prefix or the **192.168.0.0/32** prefix:

```
[edit firewall family inet filter filter_on_dst_addr term term5 from]
user@host# set destination-address 10.0.0.0/8
user@host# set destination-address 192.168.0.0/32
user@host# show
destination-address {
 destination-address 10.0.0.0/8;
 destination-address 192.168.0.0/32;
}
```

The order in which you specify the prefixes within the match condition is not significant. Packets are evaluated against all the prefixes in the match condition to determine whether a match occurs. If prefixes overlap, longest-match rules are used to determine whether a match occurs. A match condition of noncontiguous prefixes includes an implicit **0/0 except** statement, which means that any prefix that does not match any prefix included in the match condition is explicitly considered not to match.

Because the prefixes are order-independent and use longest-match rules, longer prefixes subsume shorter ones as long as they are the same type (whether you specify **except** or not). This is because anything that would match the longer prefix would also match the shorter one.

Consider the following example:

```
[edit firewall family inet filter filter_on_src_addr term term1 from]
source-address {
 172.16.0.0/10;
 172.16.2.0/24 except;
 192.168.1.0;
 192.168.1.192/26 except;
 192.168.1.254;
 172.16.3.0/24; # ignored
 10.2.2.2 except; # ignored
}
```

Within the **source-address** match condition, two addresses are ignored. The **172.16.3.0/16** value is ignored because it falls under the address **172.16.0.0/10**, which is the same type. The **10.2.2.2 except** value is ignored because it is subsumed by the implicit **0.0.0.0/0 except** match value.

Suppose the following source IP address are evaluated by this firewall filter:

- Source IP address **172.16.1.2**—This address matches the **172.16.0.0/10** prefix, and thus the action in the **then** statement is taken.
- Source IP address **172.16.2.2**—This address matches the **172.16.2.0/24** prefix. Because this prefix is negated (that is, includes the **except** keyword), an explicit *mismatch* occurs. The next term in the filter is evaluated, if there is one. If there are no more terms, the packet is discarded.
- Source IP address **10.1.2.3**—This address does not match any of the prefixes included in the **source-address** condition. Instead, it matches the implicit **0.0.0.0/0 except** at

the end of the list of prefixes configured under the **source-address** match condition, and is considered to be a mismatch.

The **172.16.3.0/24** statement is ignored because it falls under the address **172.16.0.0/10**—both are the same type.

The **10.2.2.2 except** statement is ignored because it is subsumed by the implicit **0.0.0.0/0 except** statement at the end of the list of prefixes configured under the **source-address** match condition.



**BEST PRACTICE:** When a firewall filter term includes the **from address address** match condition and a subsequent term includes the **from source-address address** match condition for the same address, packets might be processed by the latter term before they are evaluated by any intervening terms. As a result, packets that should be rejected by the intervening terms might be accepted instead, or packets that should be accepted might be rejected instead.

To prevent this from occurring, we recommend that you do the following. For every firewall filter term that contains the **from address address** match condition, replace that term with two separate terms: one that contains the **from source-address address** match condition, and another that contains the **from destination-address address** match condition.

### *Matching an Address Field to a Prefix List*

You can define a list of IPv4 or IPv6 address prefixes for use in a routing policy statement or in a stateless firewall filter match condition that evaluates packet address fields.

To define a list of IPv4 or IPv6 address prefixes, include the **prefix-list prefix-list** statement.

```
prefix-list name {
 ip-addresses;
 apply-path path;
}
```

You can include the statement at the following hierarchy levels:

- **[edit policy-options]**
- **[edit logical-systems logical-system-name policy-options]**

After you have defined a prefix list, you can use it when specifying a firewall filter match condition based on an IPv4 or IPv6 address prefix.

```
[edit firewall family family-name filter filter-name term term-name]
from {
 source-prefix-list {
 prefix-lists;
 }
 destination-prefix-list {
 prefix-lists;
 }
}
```

```
}
```

**Related Documentation**

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Firewall Filter Match Conditions Based on Numbers or Text Aliases on page 4488](#)
- [Firewall Filter Match Conditions Based on Bit-Field Values on page 4489](#)
- [Firewall Filter Match Conditions Based on Address Classes on page 4502](#)

***Firewall Filter Match Conditions Based on Address Classes***

For IPv4 and IPv6 traffic only, you can use class-based firewall filter conditions to match packet fields based on source class or destination class.

- [Source-Class Usage on page 4502](#)
- [Destination-Class Usage on page 4502](#)
- [Guidelines for Applying SCU or DCU Firewall Filters to Output Interfaces on page 4502](#)

***Source-Class Usage***

A *source class* is a set of source prefixes grouped together and given a class name. To configure a firewall filter term that matches an IP source address field to one or more source classes, use the **source-class class-name** match condition under the **[edit firewall family (inet | inet6) filter filter-name term term-name from]** hierarchy level.

*Source-class usage* (SCU) enables you to monitor the amount of traffic originating from a specific prefix. With this feature, usage can be tracked and customers can be billed for the traffic they receive.

***Destination-Class Usage***

A *destination class* is a set of destination prefixes grouped together and given a class name. To configure a firewall filter term that matches an IP destination address field to one or more destination classes, use the **destination-class class-name** match condition at the **[edit firewall family (inet | inet6) filter filter-name term term-name from]** hierarchy level.

*Destination-class usage* (DCU) enables you can track how much traffic is sent to a specific prefix in the core of the network originating from one of the specified interfaces.

Note, however, that DCU limits your ability to keep track of traffic moving in the reverse direction. It can account for all traffic that arrives on a core interface and heads toward a specific customer, but it cannot count traffic that arrives on a core interface from a specific prefix.

***Guidelines for Applying SCU or DCU Firewall Filters to Output Interfaces***

When applying a SCU or DCU firewall filter to an interface, keep the following guidelines in mind:

- Output interfaces—Class-based firewall filter match conditions work only for firewall filters that you apply to output interfaces. This is because the SCU and DCU are determined after route lookup occurs.

- Input interfaces—Although you can specify a source class and destination class for an input firewall filter, the counters are incremented only if the firewall filter is applied on the output interface.
- Output interfaces for tunnel traffic—SCU and DCU are not supported on the interfaces you configure as the output interface for tunnel traffic for transit packets exiting the router (or switch) through the tunnel.

**Related Documentation**

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Standard Firewall Filter Match Conditions for IPv4 Traffic on page 4707](#)
- [Standard Firewall Filter Match Conditions for IPv6 Traffic on page 4716](#)
- [Source Class Usage](#)
- [Firewall Filter Match Conditions Based on Numbers or Text Aliases on page 4488](#)
- [Firewall Filter Match Conditions Based on Bit-Field Values on page 4489](#)
- [Firewall Filter Match Conditions Based on Address Fields on page 4493](#)

---

**Introduction to Standard Firewall Filters for Fragment Handling**

---

- [Firewall Filters That Handle Fragmented Packets Overview on page 4503](#)

***Firewall Filters That Handle Fragmented Packets Overview***

You can create stateless firewall filters that handle fragmented packets destined for the Routing Engine. By applying these policies to the Routing Engine, you protect against the use of IP fragmentation as a means to disguise TCP packets from a firewall filter.

For example, consider an IP packet that is fragmented into the smallest allowable fragment size of 8 bytes (a 20-byte IP header plus an 8-byte payload). If this IP packet carries a TCP packet, the first fragment (fragment offset of 0) that arrives at the device contains only the TCP source and destination ports (first 4 bytes), and the sequence number (next 4 bytes). The TCP flags, which are contained in the next 8 bytes of the TCP header, arrive in the second fragment (fragment offset of 1).

On all SRX Series and J Series devices, fragmented packets are not sampled correctly by the firewall filter. When file sampling, port-mirroring and CFLOW is applied on an interface in output direction, packets are sampled before fragmenting the packet and packet-capture captures packet after fragmentation.

See RFC 1858, *Security Considerations for IP Fragment Filtering*.

**Related Documentation**

- [Understanding How to Use Standard Firewall Filters on page 3293](#)
- [Example: Configuring a Stateless Firewall Filter to Handle Fragments on page 4613](#)

---

**Introduction to Standard Firewall Filters Configuration**

---

- [Stateless Firewall Filters That Reference Policers Overview on page 4504](#)
- [Multiple Standard Firewall Filters Applied as a List Overview on page 4505](#)

- [Guidelines for Applying Multiple Standard Firewall Filters as a List on page 4509](#)
- [Multiple Standard Firewall Filters in a Nested Configuration Overview on page 4510](#)
- [Guidelines for Nesting References to Multiple Standard Firewall Filters on page 4511](#)
- [Interface-Specific Firewall Filter Instances Overview on page 4513](#)
- [Filtering Packets Received on a Set of Interface Groups Overview on page 4514](#)
- [Filtering Packets Received on an Interface Set Overview on page 4515](#)
- [Filter-Based Forwarding Overview on page 4515](#)
- [Accounting for Standard Firewall Filters Overview on page 4517](#)
- [System Logging Overview on page 4517](#)
- [System Logging of Events Generated for the Firewall Facility on page 4518](#)
- [Logging of Packet Headers Evaluated by a Firewall Filter Term on page 4520](#)

### ***Stateless Firewall Filters That Reference Policers Overview***

Policing, or rate limiting, is an important component of firewall filters that lets you limit the amount of traffic that passes into or out of an interface.

A stateless firewall filter that references a policer can provide protection from denial-of-service (DOS) attacks. Traffic that exceeds the rate limits configured for the policer is either discarded or marked as lower priority than traffic that conforms to the configured rate limits. Packets can be marked for a lower priority by being set to a specific output queue, set to a specific packet loss priority (PLP) level, or both. When necessary, low-priority traffic can be discarded to prevent congestion.

A policer specifies two types of rate limits on traffic:

- **Bandwidth limit**—The average traffic rate permitted, specified as a number of bits per second.
- **Maximum burst size**—The packet size permitted for bursts of data that exceed the bandwidth limit.

Policing uses an algorithm to enforce a limit on average bandwidth while allowing bursts up to a specified maximum value. You can use policing to define specific classes of traffic on an interface and apply a set of rate limits to each class. After you name and configure a policer, it is stored as a template. You can then apply the policer in an interface configuration or, to rate-limit packet-filtered traffic only, in a firewall filter configuration.

For an IPv4 firewall filter term only, you can also specify a *prefix-specific action* as a nonterminating action that applies a policer to the matched packets. A prefix-specific action applies additional matching criteria on the filter-matched packets based on specified address prefix bits and then associates the matched packets with a counter and policer instance for that filter term or for all terms in the firewall filter.

To apply a policer or a prefix action to packet-filtered traffic, you can use the following firewall filter nonterminating actions:

- ***policer* *policer-name***
- ***three-color-policer* (*single-rate* | *two-rate*) *policer-name***

- **prefix-action** *action-name*



**NOTE:** The packet lengths that a policer considers depends on the address family of the firewall filter. See [“Understanding the Frame Length for Policing Packets” on page 4811](#).

#### Related Documentation

- [Standard Firewall Filter Nonterminating Actions on page 4744](#)
- [Traffic Policing Overview on page 1440](#)
- [Prefix-Specific Counting and Policing Overview on page 4874](#)

#### ***Multiple Standard Firewall Filters Applied as a List Overview***

This topic covers the following information:

- [The Challenge: Simplify Large-Scale Firewall Filter Administration on page 4505](#)
- [A Solution: Apply Lists of Firewall Filters on page 4505](#)
- [Configuration of Multiple Filters for Filter Lists on page 4506](#)
- [Application of Filter Lists to a Router or Switch Interface on page 4506](#)
- [Interface-Specific Names for Filter Lists on page 4507](#)
- [How Filter Lists Evaluate Packets When the Matched Term Includes Terminating or Next Term Actions on page 4507](#)
- [How Filter Lists Evaluate Packets When the List Includes Protocol-Independent and IP Firewall Filters on page 4508](#)

#### ***The Challenge: Simplify Large-Scale Firewall Filter Administration***

Typically, you apply a single stateless firewall filter to an interface in the input or output direction or both. However, this approach might not be practical, when you have a device configured with many interfaces. In large environments, you want the flexibility of being able to modify filtering terms common to multiple interfaces without having to reconfigure the filter of every affected interface.

In general, the solution is to apply an effectively “chained” structure of multiple stateless firewall filters to a single interface. You partition your filtering terms into multiple firewall filters that each perform a filtering task. You can then choose which filtering tasks you want to perform for a given interface and apply the filtering tasks to that interface. In this way, you only manage the configuration for a filtering task in a single firewall filter.

The Junos OS policy framework provides two options for managing the application of multiple separate firewall filters to individual router (or switch) interfaces. One option is to apply multiple filters as a single input list or output list. The other option is to reference a stateless firewall filter from within the term of another stateless firewall filter.

#### ***A Solution: Apply Lists of Firewall Filters***

The most straightforward way to avoid configuring duplicate filtering terms common to multiple stateless firewall filters is to configure multiple firewall filters and then apply a

customized *list* of filters to each interface. The Junos OS uses the filters—in the order in which they appear in the list—to evaluate packets that transit the interface. If you need to modify filtering terms shared across multiple interfaces, you only need to modify one firewall filter that contains those terms.



**NOTE:** In contrast with the alternative approach (configuring nested firewall filters) applying firewall filter lists combines multiple firewall filters at each interface application point.

---

### ***Configuration of Multiple Filters for Filter Lists***

Configuring firewall filters to be applied in unique lists for each router (or switch) interface involves separating shared packet-filtering rules from interface-specific packet-filtering rules as follows:

- **Unique filters**—For each set of packet-filtering rules unique to a specific interface, configure a separate firewall filter that contains only the filtering terms for that interface.
- **Shared filters**—For each set of packet-filtering rules common across two or more interfaces, consider configuring a separate firewall filter that contains the shared filtering terms.



**TIP:** When planning for a large number firewall filters to be applied using filter lists, administrators often organize the shared filters by filtering criteria, by the services to which customers subscribe, or by the purposes of the interfaces.

---

### ***Application of Filter Lists to a Router or Switch Interface***

Applying a list of firewall filters to an interface is a matter of selecting the filters that meet the packet-filtering requirements of that interface. For each interface, you can include an **input-list** or **output-list** statement (or both) within the **filter** stanza to specify the relevant filters in the order in which they are to be used:

- Include any filters that contain common filtering terms relevant to the interface.
- Include the filter that contain only the filtering terms unique to the interface.



### Interface-Specific Names for Filter Lists

Because a filter list is configured under an interface, the resulting concatenated filter is *interface-specific*.



**NOTE:** When a filter list is configured under an interface, the resulting concatenated filter is interface-specific, regardless whether the firewall filters in the filter list are configured as interface-specific or not. Furthermore, the instantiation of interface-specific firewall filters not only create separate instances of any firewall filter counters, but also separate instances of any policer actions. Any policers applied through an action specified in the firewall filter configuration are applied separately to each interface in the interface group.

The system-generated name of an interface-specific filter consists of the full interface name followed by either '-i' for an input filter list or '-o' for an output filter list.

- **Input filter list name**—For example, if you use the **input-list** statement to apply a chain of filters to logical interface **ge-1/3/0.0**, the Junos OS uses the following name for the filter:

**ge-1/3/0.0-i**

- **Output filter list name**—For example, if you use the **output-list** statement to apply a chain of filters to logical interface **fe-0/1/2.0**, the Junos OS uses the following name for the filter:

**fe-0/1/2.0-o**

You can use the interface-specific name of a filter list when you enter a Junos OS operational mode command that specifies a stateless firewall filter name.

### How Filter Lists Evaluate Packets When the Matched Term Includes Terminating or Next Term Actions

The device evaluates a packet against the filters in a list sequentially, beginning with the first filter in the list until either a terminating action occurs or the packet is implicitly discarded.

[Table 336 on page 4508](#) describes how a firewall filter list evaluates a packet based on whether the matched term specifies a terminating action and the **next term** action. The **next term** action is neither a terminating action nor a nonterminating action but a *flow control* action.

Table 336: Firewall Filter List Behavior

| Firewall Filter Actions Included in the Matched Term |           | Term Description                                                                                             | Packet-Filtering Behavior                                                                                                                                                                                                                                         |
|------------------------------------------------------|-----------|--------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Terminating                                          | next term |                                                                                                              |                                                                                                                                                                                                                                                                   |
| Yes                                                  | —         | The matched term includes a terminating action (such as <b>discard</b> ) but not the <b>next term</b> action | The device executes the terminating action. No subsequent terms in the filter and no subsequent filters in the list are used to evaluate the packet.                                                                                                              |
| —                                                    | Yes       | The matched term includes the <b>next term</b> action, but it does not include any terminating actions.      | The device executes any nonterminating actions, then the device evaluates the packet against the next term in the filter or the next filter in the list.                                                                                                          |
| —                                                    | —         | The matched term includes neither the <b>next term</b> action nor any terminating actions.                   | The device executes any nonterminating actions, then the device implicitly accepts the packet. Because the <b>accept</b> action is a terminating action, no subsequent terms in the filter and no subsequent filters in the list are used to evaluate the packet. |

For information about terminating actions, see “[Standard Firewall Filter Terminating Actions](#)” on page 4742.



**NOTE:** You cannot configure the **next term** action with a terminating action in the same firewall filter term.

#### *How Filter Lists Evaluate Packets When the List Includes Protocol-Independent and IP Firewall Filters*

On a single interface associated with a protocol-independent (**family any**) firewall filter and a protocol-specific (**family inet** or **family inet6**) firewall filter simultaneously, the protocol-independent firewall filter executes first.

The terminating action of the first filter determines whether the second filter also evaluates the packet:

- If the first filter terminates by executing the **accept** action, the second filter also evaluates the packet.
- If the first filter terminates without any terms matching the packet (an *implicit discard* action), the second filter also evaluates the packet.
- If the first filter terminates by executing an *explicit discard* action, the second filter does not evaluate the packet.

#### **Related Documentation**

- [How Standard Firewall Filters Evaluate Packets on page 4475](#)
- [Guidelines for Applying Multiple Standard Firewall Filters as a List on page 4509](#)
- [Example: Applying Lists of Multiple Standard Firewall Filters on page 4621](#)

### ***Guidelines for Applying Multiple Standard Firewall Filters as a List***

This topic covers the following information:

- [Statement Hierarchy for Applying Lists of Multiple Firewall Filters on page 4509](#)
- [Filter Input Lists and Output Lists for Router or Switch Interfaces on page 4509](#)
- [Types of Filters Supported in Lists on page 4509](#)
- [Restrictions on Applying Filter Lists for MPLS or Layer 2 CCC Traffic on page 4510](#)

### ***Statement Hierarchy for Applying Lists of Multiple Firewall Filters***

To apply a single filter to the input or output direction of a router (or switch) logical interface, you include the **input** *filter-name* or **output** *filter-name* statement under the **filter** stanza for a protocol family.

To apply a list of multiple filters to the input or output direction of a router (or switch) logical interface, include the **input-list** [ *filter-names* ] or **output-list** [ *filter-names* ] statement under the **filter** stanza for a protocol family:

```

interfaces {
 interface-name {
 unit logical-unit-number {
 family family-name {
 filter {
 ...filter-options...
 input-list [filter-names];
 output-list [filter-names];
 }
 }
 }
 }
}

```

You can include the interface configuration at one of the following hierarchy levels:

- [edit]
- [edit logical-systems *logical-system-name*]

### ***Filter Input Lists and Output Lists for Router or Switch Interfaces***

When applying a list of firewall filters as a list, the following limitations apply:

- You can specify up to 16 firewall filters for a filter input list.
- You can specify up to 16 firewall filters for a filter output list.

### ***Types of Filters Supported in Lists***

Lists of multiple firewall filters applied to a router (or switch) interface support standard stateless firewall filters only. You cannot apply lists containing service filters or simple filters to a router (or switch) interface.

***Restrictions on Applying Filter Lists for MPLS or Layer 2 CCC Traffic***

When applying stateless firewall filters that evaluate MPLS traffic (**family mpls**) or Layer 2 circuit cross-connection traffic (**family ccc**), you can use the **input-list [ *filter-names* ]** and **output-list [ *filter-names* ]** statements for all interfaces except the following:

- Management and internal Ethernet (**fxp**) interfaces
- Loopback (**lo0**) interfaces
- USB modem (**umd**) interfaces

**Related Documentation**

- [Multiple Standard Firewall Filters Applied as a List Overview on page 4505](#)
- [Example: Applying Lists of Multiple Standard Firewall Filters on page 4621](#)

***Multiple Standard Firewall Filters in a Nested Configuration Overview***

This topic covers the following information:

- [The Challenge: Simplify Large-Scale Firewall Filter Administration on page 4510](#)
- [A Solution: Configure Nested References to Firewall Filters on page 4510](#)
- [Configuration of Nested Firewall Filters on page 4511](#)
- [Application of Nested Firewall Filters to a Router or Switch Interface on page 4511](#)

***The Challenge: Simplify Large-Scale Firewall Filter Administration***

Typically, you apply a single stateless firewall filter to an interface in the input or output direction or both. This approach might not be practical, however, when you have a router (or switch) configured with many, even hundreds of interfaces. In an environment of this scale, you want the flexibility of being able to modify filtering terms common to multiple interfaces without having to reconfigure the filter of every affected interface.

In general, the solution is to apply an effectively “chained” structure of multiple stateless firewall filters to a single interface. You partition your filtering terms into multiple firewall filters configured so that you can apply a unique filter to each router (or switch) interface but also apply common filters to multiple router (or switch) interfaces as required. The Junos OS policy framework provides two options for managing the application of multiple separate firewall filters to individual router (or switch) interfaces. One option is to apply multiple filters as a single input list or output list. The other option is to reference a stateless firewall filter from within the term of another stateless firewall filter.

***A Solution: Configure Nested References to Firewall Filters***

The most structured way to avoid configuring duplicate filtering terms common to multiple stateless firewall filters is to configure multiple stateless firewall filters so that each filter includes the shared filtering terms by *referencing* a separate filter that contains the common filtering terms. The Junos OS uses the filter terms—in the order in which they appear in the filter definition—to evaluate packets that transit the interface. If you need to modify filtering terms shared across multiple interfaces, you only need to modify one firewall filter.



**NOTE:** Similar to the alternative approach (applying a list of firewall filters), configuring a nested firewall filter combines multiple firewall filters into a new firewall filter definition.

### ***Configuration of Nested Firewall Filters***

Configuring a nested firewall filter for each router (or switch) interface involves separating shared packet-filtering rules from interface-specific packet-filtering rules as follows:

- For each set of packet-filtering rules common across multiple interfaces, configure a separate firewall filter that contains the shared filtering terms.
- For each router (or switch) interface, configure a separate firewall filter that contains:
  - All the filtering terms unique to that interface.
  - An additional filtering term that includes a **filter** reference to the firewall filter that contains the common filtering terms.

### ***Application of Nested Firewall Filters to a Router or Switch Interface***

Applying nested firewall filters is no different from applying an unnested firewall filter. For each interface, you can include an **input** or **output** statement (or both) within the **filter** stanza to specify the appropriate nested firewall filter.

Applying nested firewall filters to an interface, the shared filtering terms and the interface-specific firewall filters are applied through a *single nested firewall filter* that includes other filters through the **filter** statement within a separate filtering term.

#### **Related Documentation**

- [Guidelines for Nesting References to Multiple Standard Firewall Filters on page 4511](#)
- [Example: Nesting References to Multiple Standard Firewall Filters on page 4626](#)

### ***Guidelines for Nesting References to Multiple Standard Firewall Filters***

This topic covers the following information:

- [Statement Hierarchy for Configuring Nested Firewall Filters on page 4511](#)
- [Filter-Defining Terms and Filter-Referencing Terms on page 4512](#)
- [Types of Filters Supported in Nested Configurations on page 4512](#)
- [Number of Filter References in a Single Filter on page 4512](#)
- [Depth of Filter Nesting on page 4512](#)

### ***Statement Hierarchy for Configuring Nested Firewall Filters***

To reference a filter from within a filter, include the **filter filter-name** statement as a separate filter term:

```
firewall firewall-name {
 family family-name {
 filter filter-name {
 term term-name {
```

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

You can include the firewall configuration at one of the following hierarchy levels:

- [edit]
- [edit logical-systems *logical-system-name*]

### ***Filter-Defining Terms and Filter-Referencing Terms***

You cannot configure a firewall filter term that both references another firewall filter and defines a match condition or action. If a firewall filter term includes the **filter** statement, then it cannot also include the **from** or **then** statement.

For example, the firewall filter term **term term1** in the configuration is *not* valid:

```
[edit]
firewall {
 family inet {
 filter filter_1 {
 term term1 {
 filter filter_2;
 from {
 source-address 1.1.1.1/32;
 }
 then {
 accept;
 }
 }
 }
 }
}
```

In order for **term term1** to be a valid filter term, you must either remove the **filter filter\_2** statement or remove both the **from** and **then** stanzas.

### ***Types of Filters Supported in Nested Configurations***

Nested configurations of firewall filters support standard stateless firewall filters only. You cannot use service filters or simple filters in a nested firewall filter configuration.

### ***Number of Filter References in a Single Filter***

The total number of filters referenced from within a filter cannot exceed 256.

### ***Depth of Filter Nesting***

The Junos OS supports a single level of firewall filter nesting. If **filter\_1** references **filter\_2**, you cannot configure a filter that references a filter that references **filter\_1**.

#### **Related Documentation**

- [Multiple Standard Firewall Filters in a Nested Configuration Overview on page 4510](#)
- [Example: Nesting References to Multiple Standard Firewall Filters on page 4626](#)

### *Interface-Specific Firewall Filter Instances Overview*

This topic covers the following information:

- [Instantiation of Interface-Specific Firewall Filters on page 4513](#)
- [Interface-Specific Names for Firewall Filter Instances on page 4513](#)
- [Interface-Specific Firewall Filter Counters on page 4514](#)
- [Interface-Specific Firewall Filter Policers on page 4514](#)

### *Instantiation of Interface-Specific Firewall Filters*

On T Series, M120, M320, MX Series routers, and EX Series switches, you can enable the Junos OS to automatically create an interface-specific instance of a firewall filter for each interface to which you apply the filter. If you enable interface-specific instantiation of a firewall filter and then apply that filter to multiple interfaces, any **count** actions or **policer** actions configured in the filter terms act on the traffic stream entering or exiting each individual interface, regardless of the sum of traffic on the multiple interfaces.

You can enable this option per firewall filter by including the **interface-specific** statement in the filter configuration.



**NOTE:** On T Series, M120, M320, MX Series routers, and EX Series switches, interfaces are distributed among multiple packet-forwarding components.

Interface-specific firewall filtering is not supported on M Series routers other than the M120 and M320 routers. If you apply a firewall filter to multiple interfaces on an M Series router other than the M120 or M320 routers, the filter acts on the sum of traffic entering or exiting those interfaces.

Interface-specific firewall filtering is supported for standard stateless firewall filters and for service filters. Interface-specific instances are not supported for simple filters.

### *Interface-Specific Names for Firewall Filter Instances*

When the Junos OS creates a separate instance of a firewall filter for a logical interface, the instance is associated with an interface-specific name. The system-generated name of a firewall filter instance consists of the name of the configured filter followed by a hyphen ('-'), the full interface name, and either '-i' for an input filter instance or '-o' for an output filter instance.

- **Input filter instance name**—For example, if you apply the interface-specific firewall filter **filter\_s\_tcp** to the input at logical interface **at-1/1/1.0**, the Junos OS instantiates an interface-specific filter instance with the following system-generated name:

**filter\_s\_tcp-at-1/1/1.0-i**

- **Output filter instance name**—For example, if you apply the interface-specific firewall filter **filter\_s\_tcp** to the output at logical interface **ge-2/2/2.2**, the Junos OS instantiates an interface-specific filter instance with the following system-generated name:

**count\_s\_tcp-ge-2/2/2.2-o**

You can use the interface-specific name of a filter instance when you enter a Junos OS operational mode command that specifies a stateless firewall filter name.



**TIP:** When you configure a firewall filter with interface-specific instances enabled, we recommend you limit the filter name to *52 bytes* in length. This is because firewall filter names are restricted to *64 bytes* in length. If a system-generated filter instance name exceeds this maximum length, the policy framework software might reject the instance name.

---

### ***Interface-Specific Firewall Filter Counters***

Instantiation of interface-specific firewall filters causes the Packet Forwarding Engine to maintain any counters for the firewall filter separately for each interface. You specify interface-specific counters per firewall filter term by specifying the **count counter-name** non-terminating action.

The system-generated name of an interface-specific firewall filter counter consists of the name of the configured counter followed by a hyphen ('-'), the full interface name, and either '-i' for an input filter instance or '-o' for an output filter instance.

- **Interface-specific input filter counter name**—For example, suppose you configure the filter counter **count\_tcp** for an interface-specific firewall filter. If the filter is applied to the input at logical interface **at-1/1/1.0**, the Junos OS creates the following system-generated counter name:

**count\_tcp-at-1/1/1.0-i**

- **Interface-specific output filter counter name**—For example, suppose you configure the filter counter **count\_udp** for an interface-specific firewall filter. If the filter is applied to the output at logical interface **ge-2/2/2.2**, the Junos OS creates the following system-generated counter name:

**count\_udp-ge-2/2/2.2-o**

### ***Interface-Specific Firewall Filter Policers***

Instantiation of interface-specific firewall filters not only creates separate instances of any firewall filter counters but also creates separate instances of any policer actions. Any policers applied through an action specified in the firewall filter configuration are applied separately to each interface in the interface group. You specify interface-specific policers per firewall filter term by specifying the **policer policer-name** non-terminating action.

#### **Related Documentation**

- [Statement Hierarchy for Configuring Interface-Specific Firewall Filters on page 4773](#)
- [Statement Hierarchy for Applying Interface-Specific Firewall Filters on page 4774](#)
- [Example: Configuring Interface-Specific Firewall Filter Counters](#)

### ***Filtering Packets Received on a Set of Interface Groups Overview***

You can configure a firewall filter term that matches packets tagged for a specified *interface group* or set of interface groups. An interface group consists of one or more



logical interfaces with the same group number. Packets received on an interface in an interface group are tagged as being part of that group.

For standard stateless firewall filters, you can filter packets received on an interface group for IPv4, IPv6, virtual private LAN service ( VPLS), Layer 2 circuit cross-connection (CCC), and Layer 2 bridging traffic. For service filters, you can filter packets received on an interface group for either IPv4 or IPv6 traffic.



**NOTE:** You can also configure a firewall filter term that matches on packets tagged for a specified *interface set*. For more information, see [“Filtering Packets Received on an Interface Set Overview”](#) on page 4515.

**Related  
Documentation**

- [Statement Hierarchy for Assigning Interfaces to Interface Groups](#) on page 4774
- [Statement Hierarchy for Configuring a Filter to Match on a Set of Interface Groups](#) on page 4775
- [Example: Filtering Packets Received on an Interface Group](#) on page 4634

***Filtering Packets Received on an Interface Set Overview***

You can configure a standard stateless firewall filter term that matches packets tagged for a specified *interface set*. An interface set groups two or more physical or logical interfaces into a single interface-set name. You can filter packets received on an interface set for protocol-independent, IPv4, IPv6, MPLS, VPLS, or bridging traffic.



**NOTE:** You can also configure a standard stateless firewall filter term or a service filter term that matches on packets tagged for a specified *interface group*. For more information, see [“Filtering Packets Received on a Set of Interface Groups Overview”](#) on page 4514.

**Related  
Documentation**

- [Statement Hierarchy for Defining an Interface Set](#) on page 4777
- [Statement Hierarchy for Configuring a Filter to Match on an Interface Set](#) on page 4777
- [Example: Configuring a Rate-Limiting Filter Based on Destination Class](#) on page 4617
- [Example: Filtering Packets Received on an Interface Set](#) on page 4638

***Filter-Based Forwarding Overview***

- [Filters That Classify Packets or Direct Them to Routing Instances](#) on page 4516
- [Input Filtering to Classify and Forward Packets Within the Router or Switch](#) on page 4516
- [Output Filtering to Forward Packets to Another Routing Table](#) on page 4517
- [Restrictions for Applying Filter-Based Forwarding](#) on page 4517

### ***Filters That Classify Packets or Direct Them to Routing Instances***

For IPv4 or IPv6 traffic only, you can use stateless firewall filters in conjunction with forwarding classes and routing instances to control how packets travel in a network. This is called *filter-based forwarding* (FBF).

You can define a filtering term that matches incoming packets based on source address and then classifies matching packets to a specified forwarding class. This type of filtering can be configured to grant certain types of traffic preferential treatment or to improve load balancing. To configure a stateless firewall filter to classify packets to a forwarding class, configure a term with the *nonterminating action* **forwarding-class class-name**.

You can also define a filtering term that directs matching packets to a specified routing instance. This type of filtering can be configured to route specific types of traffic through a firewall or other security device before the traffic continues on its path. To configure a stateless firewall filter to direct traffic to a routing instance, configure a term with the *terminating action* **routing-instance routing-instance-name <topology topology-name>** to specify the routing instance to which matching packets will be forwarded.

To forward traffic to the master routing instance, reference **routing-instance default** in the firewall configuration, as shown here:

```
[edit firewall]
family inet {
 filter test {
 term 1 {
 then {
 routing-instance default;
 }
 }
 }
}
```



**NOTE:** Do not reference **routing-instance master**. This does not work.

---

### ***Input Filtering to Classify and Forward Packets Within the Router or Switch***

You can configure filters to classify packets based on source address and specify the forwarding path the packets take within the router or switch by configuring a filter on the ingress interface.

For example, you can use this filter for applications to differentiate traffic from two clients that have a common access layer (for example, a Layer 2 switch) but are connected to different Internet service providers (ISPs). When the filter is applied, the router or switch can differentiate the two traffic streams and direct each to the appropriate network. Depending on the media type the client is using, the filter can use the source IP address to forward the traffic to the corresponding network through a tunnel. You can also configure filters to classify packets based on IP protocol type or IP precedence bits.

### ***Output Filtering to Forward Packets to Another Routing Table***

You can also forward packets based on output filters by configuring a filter on the egress interfaces. In the case of port mirroring, it is useful for port-mirrored packets to be distributed to multiple monitoring PICs and collection PICs based on patterns in packet headers. FBF on the port-mirroring egress interface must be configured.

Packets forwarded to the output filter have been through at least one route lookup when an FBF filter is configured on the egress interface. After the packet is classified at the egress interface by the FBF filter, it is redirected to another routing table for further route lookup.

### ***Restrictions for Applying Filter-Based Forwarding***

An interface configured with filter-based forwarding does not support source-class usage (SCU) filter matching and unicast reverse-path forwarding (RPF) check filters.

#### **Related Documentation**

- [Example: Configuring Filter-Based Forwarding on the Source Address on page 4644](#)
- [Example: Configuring Filter-Based Forwarding on Logical Systems on page 3298](#)

### ***Accounting for Standard Firewall Filters Overview***

Juniper Networks devices can collect various kinds of data about traffic passing through the device. You can set up one or more accounting profiles that specify some common characteristics of this data, including the following:

- Fields used in the accounting records.
- Number of files that the routing platform retains before discarding, and the number of bytes per file.
- Polling period that the system uses to record the data

There are several types of accounting profiles: interface, firewall filter, source class and destination class usage, and Routing Engine. If you apply the same profile name to both a firewall filter and an interface, it causes an error.

#### **Related Documentation**

- [Statement Hierarchy for Configuring Firewall Filter Accounting Profiles on page 4783](#)
- [Statement Hierarchy for Applying Firewall Filter Accounting Profiles on page 4784](#)
- [Example: Configuring Statistics Collection for a Standard Firewall Filter on page 4669](#)

### ***System Logging Overview***

The Junos OS generates system log messages (also called *syslog messages*) to record *system events* that occur on the device. Events consist of routine operations, failure and error conditions, and critical conditions that might require urgent resolution. This system logging utility is similar to the UNIX **syslogd** utility.

Each Junos OS system log message belongs to a message category, called a *facility*, that reflects the hardware- or software-based source of the triggering event. A group of messages belonging to the same facility are either generated by the same software

process or concern a similar hardware condition or user activity (such as authentication attempts). Each system log message is also preassigned a *severity*, which indicates how seriously the triggering event affects router (or switch) functions. Together, the facility and severity of an event are known as the message *priority*. The content of a syslog message identifies the Junos OS *process* that generates the message and briefly describes the operation or error that occurred.

By default, syslog messages that have a severity of **info** or more serious are written to the main system log file **messages** in the **/var/log** directory of the local Routing Engine. To configure global settings and facility-specific settings that override these default values, you can include statements at the **[edit system syslog]** hierarchy level.

For all syslog facilities or for a specified facility, you can configure the syslog message utility to redirect messages of a specified severity to a specified file instead of to the main system log file. You can also configure the syslog message utility to write syslog messages of a specified severity, for all syslog facilities or for a specified facility, to additional destinations. In addition to writing syslog messages to a log file, you can write syslog messages to the terminal sessions of any logged-in users, to the router (or switch) console, or to a remote host or the other Routing Engine.

At the global level—for all system logging messages, regardless of facility, severity, or destination—you can override the default values for file-archiving properties and the default timestamp format.

**Related  
Documentation**

- [System Logging of Events Generated for the Firewall Facility on page 4518](#)
- [Logging of Packet Headers Evaluated by a Firewall Filter Term on page 4520](#)
- [Example: Configuring Logging for a Stateless Firewall Filter Term](#)

### ***System Logging of Events Generated for the Firewall Facility***

System log messages generated for firewall filter actions belong to the **firewall** facility. Just as you can for any other Junos OS system logging facility, you can direct **firewall** facility syslog messages to one or more specific destinations: to a specified file, to the terminal session of one or more logged in users (or to all users), to the router (or switch) console, or to a remote host or the other Routing Engine on the router (or switch).

When you configure a syslog message destination for **firewall** facility syslog messages, you include a statement at the **[edit system syslog]** hierarchy level, and you specify the **firewall** facility name together with a severity level. Messages from the **firewall** that are rated at the specified level or more severe are logged to the destination.

System log messages with the **DFWD\_** prefix are generated by the firewall process (**dfwd**), which manages compilation and downloading of Junos OS firewall filters. System log messages with the **PFE\_FW\_** prefix are messages about firewall filters, generated by the Packet Forwarding Engine controller, which manages packet forwarding functions. For more information, see the *Junos OS System Log Messages Reference*.

[Table 337 on page 4519](#) lists the system log destinations you can configure for the **firewall** facility.

Table 337: Syslog Message Destinations for the Firewall Facility

| Destination                             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Configuration Statements                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| File                                    | <p>Configuring this option keeps the <b>firewall</b> syslog messages out of the main system log file.</p> <p>To include priority and facility with messages written to the file, include the <b>explicit-priority</b> statement.</p> <p>To override the default standard message format, which is based on a UNIX system log format, include the <b>structured-data</b> statement.<sup>†</sup></p>                                                                                         | <pre>[edit] system {   syslog {     file <i>filename</i> {       firewall <i>severity</i>;       allow-duplicates; # File option       archive <i>archive-options</i>; # File option       explicit-priority; # File option       structured-data; # File option     }     allow-duplicates; # All destinations     archive <i>archive-options</i>; # All files     time-format (<i>option</i>); # Local destinations   } }</pre>                                     |
| Terminal session                        | <p>Configuring this option causes a copy of the <b>firewall</b> syslog messages to be written to the specified terminal sessions. Specify one or more user names, or specify <b>*</b> for all logged in users.</p>                                                                                                                                                                                                                                                                         | <pre>[edit] system {   syslog {     user (<i>username</i>   *) {       firewall <i>severity</i>;     }     time-format (<i>option</i>); # Local destinations   } }</pre>                                                                                                                                                                                                                                                                                              |
| Router (or switch) console              | <p>Configuring this option causes a copy of the <b>firewall</b> syslog messages to be written to the router (or switch) console.</p>                                                                                                                                                                                                                                                                                                                                                       | <pre>[edit] system {   syslog {     console {       firewall <i>severity</i>;     }     time-format (<i>option</i>); # Local destinations   } }</pre>                                                                                                                                                                                                                                                                                                                 |
| Remote host or the other Routing Engine | <p>Configuring this option causes a copy of the <b>firewall</b> syslog messages to be written to the specified remote host or to the other Routing Engine.</p> <p>To override the default alternative facility for forwarding <b>firewall</b> syslog messages to a remote machine (<b>local3</b>), include the <b>facility-override firewall</b> statement.</p> <p>To include priority and facility with messages written to the file, include the <b>explicit-priority</b> statement.</p> | <pre>[edit] system {   syslog {     host (<i>hostname</i>   other-routing-engine) {       firewall <i>severity</i>;       allow-duplicates; # Host option       archive <i>archive-options</i>; # File option       facility-override firewall; # Host option       explicit-priority; # Host option     }     allow-duplicates; # All destinations     archive <i>archive-options</i>; # All files     time-format (<i>option</i>); # Local destinations   } }</pre> |

<sup>†</sup> When the **structured-data** statement is included, other statements that specify the format for messages written to the file are ignored (the **explicit-priority** statement at the [edit system syslog file *filename*] hierarchy level and the **time-format** statement at the [edit system syslog] hierarchy level).

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

**Sep 07 08:00:10**

To include the year, the millisecond, or both in the timestamp for all system logging messages, regardless of the facility, include one of the following statement at the **[edit system syslog]** hierarchy level:

- **time-format year;**
- **time-format millisecond;**
- **time-format year millisecond;**

The following example illustrates the format for a timestamp that includes both the millisecond (401) and the year (2010):

**Sep 07 08:00:10.401.2010**

#### Related Documentation

- [System Logging Overview on page 4517](#)
- [Logging of Packet Headers Evaluated by a Firewall Filter Term on page 4520](#)
- *Example: Configuring Logging for a Stateless Firewall Filter Term*
- “Junos OS System Logging Facilities and Message Severity Levels” in the *Junos OS System Basics Configuration Guide*
- “Junos OS System Log Configuration Hierarchy” in the *Junos OS System Basics Configuration Guide*
- “Junos OS Default System Log Settings” in the *Junos OS System Basics Configuration Guide*
- “Logging Messages in Structured-Data Format” in the *Junos OS System Basics Configuration Guide*
- “Including the Year or Millisecond in Timestamps” in the *Junos OS System Basics Configuration Guide*
- “Changing the Alternative Facility Name for Remote System Log Messages” in the *Junos OS System Basics Configuration Guide*
- “Junos OS System Log Alternate Facilities for Remote Logging” in the *Junos OS System Basics Configuration Guide*

#### **Logging of Packet Headers Evaluated by a Firewall Filter Term**

Built in to the stateless firewall filtering software is the capability to log packet-header information for the packets evaluated by a stateless firewall filter term. You can write the packet header information to the system log file on the local Routing Engine or to a firewall filter buffer in the Packet Forwarding Engine. Logging of packet headers evaluated by firewall filters is supported for standard stateless firewall filters for IPv4 or IPv6 traffic only. Service filters and simple filters do not support logging of packet headers.

Table 338 on page 4521 lists the packet-header logs you can configure for a firewall filter action.

**Table 338: Packet-Header Logs for Stateless Firewall Filter Terms**

| Log                                                              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Configuration Statements                                                                                                                                                                          |
|------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syslog message destinations configured for the firewall facility | <p>Configure this option by using the <b>syslog</b> nonterminating action.</p> <p><b>NOTE:</b> Packet header information is interspersed with event messages.</p> <p>To list log files, enter the <b>show log</b> operational mode command without command options.</p> <p>To display log file contents for a specific file in the <b>/var/log</b> directory on the local Routing Engine, enter the <b>show log filename</b> operational mode command or the <b>file show /var/log/filename</b> operational mode command.</p> <p>To clear log file contents, enter the <b>clear log filename &lt;all&gt;</b> operational mode command. If you include the <b>all</b> option, the specified log file is truncated, all archived versions of the log file are deleted.</p> | <pre>firewall {   family {     filter filter-name {       from {         match-conditions;       }       then {         ...         syslog;         terminating-action;       }     }   } }</pre> |
| Buffer in the Packet Forwarding Engine                           | <p>Configure this option by using the <b>log</b> nonterminating action.</p> <p><b>NOTE:</b> Restarting the router (or switch) causes the contents of this buffer to be cleared.</p> <p>To display the local log entries for firewall filters, enter the <b>show firewall log</b> operational mode command.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <pre>firewall {   family {     filter filter-name {       from {         match-conditions;       }       then {         ...         log;         terminating-action;       }     }   } }</pre>    |

- Related Documentation**
- [System Logging Overview on page 4517](#)
  - [System Logging of Events Generated for the Firewall Facility on page 4518](#)
  - [Example: Configuring Logging for a Stateless Firewall Filter Term](#)

### Introduction to Service Filters Configuration

- [Service Filter Overview on page 4522](#)
- [How Service Filters Evaluate Packets on page 4523](#)
- [Guidelines for Configuring Service Filters on page 4524](#)
- [Guidelines for Applying Service Filters on page 4526](#)

### Service Filter Overview

This topic covers the following information:

- [Services on page 4522](#)
- [Service Rules on page 4522](#)
- [Service Rule Refinement on page 4522](#)
- [Service Filter Counters on page 4522](#)

### Services

The Adaptive Services Physical Interface Cards (PICs), Multiservices PICs, and Multiservices Dense Port Concentrators (DPCs) provide *adaptive services interfaces*. Adaptive services interfaces enable you to coordinate a special range of services on a single PIC or DPC by configuring a set of services and applications.



**NOTE:** Service filters are not supported on T4000 routers, J Series devices and Branch SRX devices.

### Service Rules

A *service set* is an optional definition you can apply to the traffic at an adaptive services interface. A service set enables you to configure combinations of directional rules and default settings that control the behavior of each service in the service set.

### Service Rule Refinement

When you apply a service set to the traffic at an adaptive services interface, you can optionally use *service filters* to refine the target of the set of services and also to process traffic. Service filters enable you to manipulate traffic by performing packet filtering to a defined set of services on an adaptive services interface before the traffic is delivered to its destination. You can apply a service filter to traffic before packets are accepted for input or output service processing or after packets return from input service processing.

### Service Filter Counters

Like standard firewall filters, service filters support counting of matched packets. When you display counters for a service filter, however, the syntax for specifying the filter name includes the name of the *service set* to which the service filter is applied.

- To enable counting of the packets matched by a service filter term, specify the **count counter-name** nonterminating action in that term.
- To display counters for service filters, use the **show firewall filter filter-name <counter counter-name>** operational mode command, and specify the **filter-name** as follows:

**`__service-service-set-name:service-filter-name`**



For example, suppose you configure a service filter named **out\_filter** with a counter named **out\_counter** and apply that service filter to a logical interface to direct certain packets for processing by the output services associated with the service set **nat\_set**. In this scenario, the syntax for using the **show firewall** operational mode command to display the counter is as follows:

```
[edit]
user@host> show firewall filter __service-nat_set:out_filter counter out_counter
```

#### Related Documentation

- [Stateless Firewall Filter Types on page 4465](#)
- [How Service Filters Evaluate Packets on page 4523](#)
- [Guidelines for Configuring Service Filters on page 4524](#)
- [Guidelines for Applying Service Filters on page 4526](#)
- [Example: Configuring and Applying Service Filters on page 4678](#)
- “Adaptive Services Overview” in the *Junos Services Interfaces Configuration Release 11.2*
- “Configuring Service Sets to be Applied to Services Interfaces” in the *Junos Services Interfaces Configuration Release 11.2*
- “Configuring Service Rules” in the *Junos Services Interfaces Configuration Release 11.2*

#### **How Service Filters Evaluate Packets**

This topic covers the following information:

- [Service Filters That Contain a Single Term on page 4523](#)
- [Service Filters That Contain Multiple Terms on page 4524](#)
- [Service Filter Terms That Do Not Contain Any Match Conditions on page 4524](#)
- [Service Filter Terms That Do Not Contain Any Actions on page 4524](#)
- [Service Filter Default Action on page 4524](#)

#### **Service Filters That Contain a Single Term**

For a service filter that consists of a single term, the policy framework software evaluates a packet as follows:

- If the packet matches all the conditions, the actions are taken.
- If the packet matches all the conditions and no actions are specified, the packet is accepted.
- If the packet does not match all the conditions, it is discarded.

### ***Service Filters That Contain Multiple Terms***

For a service filter that consists of multiple terms, the policy framework software evaluates a packet against the terms in the filter sequentially, beginning with the first term in the filter, until either the packet matches all the conditions in one of the terms or there are no more terms in the filter.

- If the packet matches all the conditions in a term, the actions in that term are performed and evaluation of the packet ends at that term. Any subsequent terms in the filter are not used.
- If the packet does not match all the conditions in the term, evaluation of the packet proceeds to the next term in the filter.

### ***Service Filter Terms That Do Not Contain Any Match Conditions***

For service filters with a single term and for filters with multiple terms, if a term does not contain any match conditions, the actions are taken on any packet evaluated.

### ***Service Filter Terms That Do Not Contain Any Actions***

If a term does not contain any actions, and if the packet matches the conditions in the term, the packet is accepted.

### ***Service Filter Default Action***

Each service filter has an *implicit skip* action at the end of the filter, which is equivalent to including the following example term **explicit\_skip** as the final term in the service filter:

```
term explicit_skip {
 then skip;
}
```

By default, if a packet matches none of the terms in a service filter, the packet bypasses service processing.

#### **Related Documentation**

- [Service Filter Overview on page 4522](#)
- [Guidelines for Configuring Service Filters on page 4524](#)
- [Guidelines for Applying Service Filters on page 4526](#)
- [Example: Configuring and Applying Service Filters on page 4678](#)

### ***Guidelines for Configuring Service Filters***

This topic covers the following information:

- [Statement Hierarchy for Configuring Service Filters on page 4525](#)
- [Service Filter Protocol Families on page 4525](#)
- [Service Filter Names on page 4525](#)
- [Service Filter Terms on page 4525](#)
- [Service Filter Match Conditions on page 4526](#)
- [Service Filter Terminating Actions on page 4526](#)

### Statement Hierarchy for Configuring Service Filters

To configure a service filter, include the **service-filter service-filter-name** statement at the **[edit firewall family (inet | inet6)]** hierarchy level:

```
[edit]
firewall {
 family (inet | inet6) {
 service-filter service-filter-name {
 term term-name {
 from {
 match-conditions;
 }
 then {
 actions;
 }
 }
 }
 }
}
```

Individual statements supported under the **service-filter service-filter-name** statement are described separately in this topic and are illustrated in the example of configuring and applying a service filter.

### Service Filter Protocol Families

You can configure service filters to filter IPv4 traffic (**family inet**) and IPv6 traffic (**family inet6**) only. No other protocol families are supported for service filters.

### Service Filter Names

Under the **family inet** or **family inet6** statement, you can include **service-filter service-filter-name** statements to create and name service filters. The filter name can contain letters, numbers, and hyphens (-) and be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" ").

### Service Filter Terms

Under the **service-filter service-filter-name** statement, you can include **term term-name** statements to create and name filter terms.

- You must configure at least one term in a firewall filter.
- You must specify a unique name for each term within a firewall filter. The term name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" ").
- The order in which you specify terms within a firewall filter configuration is important. Firewall filter terms are evaluated in the order in which they are configured. By default, new terms are always added to the end of the existing filter. You can use the **insert** configuration mode command to reorder the terms of a firewall filter.

### ***Service Filter Match Conditions***

Service filter terms support only a subset of the IPv4 and IPv6 match conditions that are supported for standard stateless firewall filters.

If you specify an IPv6 address in a match condition (the **address**, **destination-address**, or **source-address** match conditions), use the syntax for text representations described in RFC 2373, *IP Version 6 Addressing Architecture*. For more information about IPv6 addresses, see “IPv6 Overview” and “IPv6 Standards” in the *Junos OS Routing Protocols Configuration Guide*.

### ***Service Filter Terminating Actions***

When configuring a service filter term, you must specify one of the following filter-terminating actions:

- **service**
- **skip**



**NOTE:** These actions are unique to service filters.

---

Service filter terms support only a subset of the IPv4 and IPv6 nonterminating actions that are supported for standard stateless firewall filters:

- **count** *counter-name*
- **log**
- **port-mirror**
- **sample**

Service filters do not support the **next** action.

#### **Related Documentation**

- [Service Filter Overview on page 4522](#)
- [How Service Filters Evaluate Packets on page 4523](#)
- [Guidelines for Applying Service Filters on page 4526](#)
- [Service Filter Match Conditions for IPv4 or IPv6 Traffic on page 4760](#)
- [Service Filter Terminating Actions on page 4766](#)
- [Service Filter Nonterminating Actions on page 4766](#)
- [Example: Configuring and Applying Service Filters on page 4678](#)

### ***Guidelines for Applying Service Filters***

This topic covers the following information:

- [Restrictions for Adaptive Services Interfaces on page 4527](#)
- [Statement Hierarchy for Applying Service Filters on page 4527](#)

- [Associating Service Rules with Adaptive Services Interfaces on page 4528](#)
- [Filtering Traffic Before Accepting Packets for Service Processing on page 4528](#)
- [Postservice Filtering of Returning Service Traffic on page 4529](#)

### ***Restrictions for Adaptive Services Interfaces***

The following restrictions apply to adaptive services interfaces and service filters.

#### ***Adaptive Services Interfaces***

You can apply a service filter to IPv4 or IPv6 traffic associated with a service set at an *adaptive services interface* only. Adaptive services interfaces are supported for the following hardware only:

- Adaptive Services (AS) PICs on M Series and T Series routers
- Multiservices (MS) PICs on M Series and T Series routers
- Multiservices (MS) DPCs on MX Series routers (and EX Series switches)

#### ***System Logging to a Remote Host from M Series Routers***

Logging of adaptive services interfaces messages to an external server by means of the **fxp0** or **em0** port is not supported on M Series routers. The architecture does not support system logging traffic out of a management interface. Instead, access to an external server is supported on a Packet Forwarding Engine interface.

#### ***Statement Hierarchy for Applying Service Filters***

You can enable packet filtering of IPv4 or IPv6 traffic before a packet is accepted for input or output service processing. To do this, apply a service filter to the adaptive services interface input or output in conjunction with an interface service set.

You can also enable packet filtering of IPv4 or IPv6 traffic that is returning to the Packet Forwarding Engine after input service processing completes. To do this, apply a post-service filter to the adaptive services interface input.

The following configuration shows the hierarchy levels at which you can apply the service filters to adaptive services interfaces:

```
[edit]
interfaces {
 interface-name {
 unit unit-number {
 family (inet | inet6) {
 service {
 input {
 service-set service-set-name service-filter service-filter-name;
 post-service-filter service-filter-name;
 }
 output {
 service-set service-set-name service-filter service-filter-name;
 }
 }
 }
 }
 }
}
```

```
}
}
```

### ***Associating Service Rules with Adaptive Services Interfaces***

To define and group the service rules to be applied to an adaptive services interface, you define an *interface service set* by including the **service-set service-set-name** statement at the **[edit services]** hierarchy level.

To apply an interface service set to the input and output of an adaptive services interface, you include the **service-set service-set-name** at the following hierarchy levels:

- **[edit interfaces interface-name unit unit-number input]**
- **[edit interfaces interface-name unit unit-number output]**

If you apply a service set to one direction of an adaptive services interface but do not apply a service set to the other direction, an error occurs when you commit the configuration.

The adaptive services PIC performs different actions depending on whether the packet is sent to the PIC for input service or for output service. For example, you can configure a single service set to perform Network Address Translation (NAT) in one direction and destination NAT (dNAT) in the other direction.

### ***Filtering Traffic Before Accepting Packets for Service Processing***

To filter IPv4 or IPv6 traffic before accepting packets for input or output service processing, include the **service-set service-set-name service-filter service-filter-name** at one of the following interfaces:

- **[edit interfaces interface-name unit unit-number family (inet | inet6) service input]**
- **[edit interfaces interface-name unit unit-number family (inet | inet6) service output]**

For the **service-set-name**, specify a service set configured at the **[edit services service-set]** hierarchy level.

The service set retains the input interface information even after services are applied, so that functions such as filter-class forwarding and destination class usage (DCU) that depend on input interface information continue to work.

The following requirements apply to filtering inbound or outbound traffic before accepting packets for service processing:

- You configure the same service set on the input and output sides of the interface.
- If you include the **service-set** statement without an optional **service-filter** definition, the Junos OS assumes the match condition is true and selects the service set for processing automatically.
- The service filter is applied only if a service set is configured and selected.

You can include more than one service set definition on each side of an interface. The following guidelines apply:

- If you include multiple service sets, the router (or switch) software evaluates them in the order in which they appear in the configuration. The system executes the first service set for which it finds a match in the service filter and ignores the subsequent definitions.
- A maximum of six service sets can be applied to an interface.
- When you apply multiple service sets to an interface, you must also configure and apply a service filter to the interface.

### ***Postservice Filtering of Returning Service Traffic***

As an option to filtering of IPv4 or IPv6 input service traffic, you can apply a service filter to IPv4 or IPv6 traffic that is returning to the services interface after the service set is executed. To apply a service filter in this manner, include the **post-service-filter service-filter-name** statement at the **[edit interfaces interface-name unit unit-number family (inet | inet6) service input]** hierarchy level.

#### **Related Documentation**

- [Service Filter Overview on page 4522](#)
- [How Service Filters Evaluate Packets on page 4523](#)
- [Guidelines for Configuring Service Filters on page 4524](#)
- [Example: Configuring and Applying Service Filters on page 4678](#)
- “Adaptive Services Overview” in the *Junos Services Interfaces Configuration Release 11.2*
- “Configuring Service Sets to be Applied to Services Interfaces” in the *Junos Services Interfaces Configuration Release 11.2*
- “Configuring Service Rules” in the *Junos Services Interfaces Configuration Release 11.2*

### ***Introduction to Simple Filters Configuration***

- [Simple Filter Overview on page 4529](#)
- [How Simple Filters Evaluate Packets on page 4530](#)
- [Guidelines for Configuring Simple Filters on page 4531](#)
- [Guidelines for Applying Simple Filters on page 4534](#)

### ***Simple Filter Overview***

Simple filters are supported on Gigabit Ethernet intelligent queuing 2 (IQ2) and Enhanced Queuing Dense Port Concentrator (DPC) interfaces only.

Simple filters are recommended for metropolitan Ethernet applications.

#### **Related Documentation**

- [How Simple Filters Evaluate Packets on page 4530](#)
- [Guidelines for Configuring Simple Filters on page 4531](#)
- [Guidelines for Applying Simple Filters on page 4534](#)

- [Example: Configuring and Applying a Simple Filter on page 4683](#)

### ***How Simple Filters Evaluate Packets***

This topic covers the following information:

- [Simple Filters That Contain a Single Term on page 4530](#)
- [Simple Filters That Contain Multiple Terms on page 4530](#)
- [Simple Filter Terms That Do Not Contain Any Match Conditions on page 4530](#)
- [Simple Filter Terms That Do Not Contain Any Actions on page 4530](#)
- [Simple Filter Default Action on page 4530](#)

### ***Simple Filters That Contain a Single Term***

For a simple filter that consists of a single term, the policy framework software evaluates a packet as follows:

- If the packet matches all the conditions, the actions are taken.
- If the packet matches all the conditions and no actions are specified, the packet is accepted.
- If the packet does not match all the conditions, it is discarded.

### ***Simple Filters That Contain Multiple Terms***

For a simple filter that consists of multiple terms, the policy framework software evaluates a packet against the terms in the filter sequentially, beginning with the first term in the filter, until either the packet matches all the conditions in one of the terms or there are no more terms in the filter.

- If the packet matches all the conditions in a term, the actions in that term are performed and evaluation of the packet ends at that term. Any subsequent terms in the filter are not used.
- If the packet does not match all the conditions in the term, evaluation of the packet proceeds to the next term in the filter.

### ***Simple Filter Terms That Do Not Contain Any Match Conditions***

For simple filters with a single term and for filters with multiple terms, if a term does not contain any match conditions, the actions are taken on any packet evaluated.

### ***Simple Filter Terms That Do Not Contain Any Actions***

If a simple filter term does not contain any actions, and if the packet matches the conditions in the term, the packet is accepted.

### ***Simple Filter Default Action***

Each simple filter has an *implicit discard* action at the end of the filter, which is equivalent to including the following example term **explicit\_discard** as the final term in the simple filter:

```
term explicit_discard {
```



```

 then discard;
}

```

By default, if a packet matches none of the terms in a simple filter, the packet is discarded.

#### Related Documentation

- [Simple Filter Overview on page 4529](#)
- [Guidelines for Configuring Simple Filters on page 4531](#)
- [Guidelines for Applying Simple Filters on page 4534](#)
- [Example: Configuring and Applying a Simple Filter on page 4683](#)

#### *Guidelines for Configuring Simple Filters*

This topic covers the following information:

- [Statement Hierarchy for Configuring Simple Filters on page 4531](#)
- [Simple Filter Protocol Families on page 4531](#)
- [Simple Filter Names on page 4532](#)
- [Simple Filter Terms on page 4532](#)
- [Simple Filter Match Conditions on page 4532](#)
- [Simple Filter Terminating Actions on page 4533](#)
- [Simple Filter Nonterminating Actions on page 4533](#)

#### *Statement Hierarchy for Configuring Simple Filters*

To configure a simple filter, include the **simple-filter** *simple-filter-name* statement at the **[edit firewall family inet]** hierarchy level.

```

[edit]
firewall {
 family inet {
 simple-filter simple-filter-name {
 term term-name {
 from {
 match-conditions;
 }
 then {
 actions;
 }
 }
 }
 }
}

```

Individual statements supported under the **simple-filter** *simple-filter-name* statement are described separately in this topic and are illustrated in the example of configuring and applying a simple filter.

#### *Simple Filter Protocol Families*

You can configure simple filters to filter IPv4 traffic (**family inet**) only. No other protocol family is supported for simple filters.

### ***Simple Filter Names***

Under the **family inet** statement, you can include **simple-filter *simple-filter-name*** statements to create and name simple filters. The filter name can contain letters, numbers, and hyphens (-) and be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" ").

### ***Simple Filter Terms***

Under the **simple-filter *simple-filter-name*** statement, you can include **term *term-name*** statements to create and name filter terms.

- You must configure at least one term in a firewall filter.
- You must specify a unique name for each term within a firewall filter. The term name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" ").
- The order in which you specify terms within a firewall filter configuration is important. Firewall filter terms are evaluated in the order in which they are configured. By default, new terms are always added to the end of the existing filter. You can use the **insert** configuration mode command to reorder the terms of a firewall filter.

Simple filters do *not* support the **next term** action.

### ***Simple Filter Match Conditions***

Simple filter terms support only a subset of the IPv4 match conditions that are supported for standard stateless firewall filters.

Unlike standard stateless firewall filters, the following restrictions apply to simple filters:

- On MX Series routers with the Enhanced Queuing DPC and on EX Series switches, simple filters do *not* support the **forwarding-class** match condition.
- Simple filters support only one **source-address** and one **destination-address** prefix for each filter term. If you configure multiple prefixes, only the last one is used.
- Simple filters do *not* support multiple source addresses and destination addresses in a single term. If you configure multiple addresses, only the last one is used.
- Simple filters do *not* support negated match conditions, such as the **protocol-except** match condition or the **exception** keyword.
- Simple filters support a range of values for **source-port** and **destination-port** match conditions only. For example, you can configure **source-port 400-500** or **destination-port 600-700**.
- Simple filters do *not* support noncontiguous mask values.

[Table 339 on page 4533](#) lists the simple filter match conditions.

Table 339: Simple Filter Match Conditions

| Match Condition                                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>destination-address</b><br><i>destination-address</i> | Match IP destination address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>destination-port</b><br><i>number</i>                 | <p>TCP or UDP destination port field.</p> <p>If you configure this match condition, we recommend that you also configure the <b>protocol</b> match statement to determine which protocol is being used on the port.</p> <p>In place of the numeric value, you can specify one of the following text aliases (the port numbers are also listed): <b>afs</b> (1483), <b>bgp</b> (179), <b>biff</b> (512), <b>bootpc</b> (68), <b>bootps</b> (67), <b>cmd</b> (514), <b>cvspserver</b> (2401), <b>dhcp</b> (67), <b>domain</b> (53), <b>eklogin</b> (2105), <b>ekshell</b> (2106), <b>exec</b> (512), <b>finger</b> (79), <b>ftp</b> (21), <b>ftp-data</b> (20), <b>http</b> (80), <b>https</b> (443), <b>ident</b> (113), <b>imap</b> (143), <b>kerberos-sec</b> (88), <b>klogin</b> (543), <b>kpasswd</b> (761), <b>krb-prop</b> (754), <b>krbupdate</b> (760), <b>kshell</b> (544), <b>ldap</b> (389), <b>login</b> (513), <b>mobileip-agent</b> (434), <b>mobilip-mn</b> (435), <b>msdp</b> (639), <b>netbios-dgm</b> (138), <b>netbios-ns</b> (137), <b>netbios-ssn</b> (139), <b>nfsd</b> (2049), <b>nnntp</b> (119), <b>ntalk</b> (518), <b>ntp</b> (123), <b>pop3</b> (110), <b>pptp</b> (1723), <b>printer</b> (515), <b>radacct</b> (1813), <b>radius</b> (1812), <b>rip</b> (520), <b>rkinit</b> (2108), <b>smtp</b> (25), <b>snmp</b> (161), <b>snmptrap</b> (162), <b>snpp</b> (444), <b>socks</b> (1080), <b>ssh</b> (22), <b>sunrpc</b> (111), <b>syslog</b> (514), <b>tacacs-ds</b> (65), <b>talk</b> (517), <b>telnet</b> (23), <b>tftp</b> (69), <b>timed</b> (525), <b>who</b> (513), or <b>xmcp</b> (177).</p> |
| <b>forwarding-class</b> <i>class</i>                     | <p>Match the forwarding class of the packet.</p> <p>Specify <b>assured-forwarding</b>, <b>best-effort</b>, <b>expedited-forwarding</b>, or <b>network-control</b>.</p> <p>For information about forwarding classes and router-internal output queues, see the <i>Junos OS Class of Service Configuration Guide</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>protocol</b> <i>number</i>                            | <p>IP protocol field. In place of the numeric value, you can specify one of the following text aliases (the field values are also listed): <b>ah</b> (51), <b>dstdopts</b> (60), <b>egp</b> (8), <b>esp</b> (50), <b>fragment</b> (44), <b>gre</b> (47), <b>hop-by-hop</b> (0), <b>icmp</b> (1), <b>icmp6</b> (58), <b>icmpv6</b> (58), <b>igmp</b> (2), <b>ipip</b> (4), <b>ipv6</b> (41), <b>ospf</b> (89), <b>pim</b> (103), <b>rsvp</b> (46), <b>sctp</b> (132), <b>tcp</b> (6), <b>udp</b> (17), or <b>vrrp</b> (112).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>source-address</b><br><i>ip-source-address</i>        | Match the IP source address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>source-port</b> <i>number</i>                         | <p>Match the UDP or TCP source port field.</p> <p>If you configure this match condition, we recommend that you also configure the <b>protocol</b> match statement to determine which protocol is being used on the port.</p> <p>In place of the numeric field, you can specify one of the text aliases listed for <b>destination-port</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

### Simple Filter Terminating Actions

Simple filters do *not* support explicitly configurable terminating actions, such as **accept**, **reject**, and **discard**. Terms configured in a simple filter always accept packets.

Simple filters do *not* support the **next** action.

### Simple Filter Nonterminating Actions

Simple filters support only the following nonterminating actions:

- **forwarding-class** (*forwarding-class* | **assured-forwarding** | **best-effort** | **expedited-forwarding** | **network-control**)



**NOTE:** On MX Series routers and EX Series switches with the Enhanced Queuing DPC, the forwarding class is not supported as a from match condition.

- **loss-priority (high | low | medium-high | medium-low)**

Simple filters do not support actions that perform other functions on a packet (such as incrementing a counter, logging information about the packet header, sampling the packet data, or sending information to a remote host using the system log functionality).

#### Related Documentation

- [Simple Filter Overview on page 4529](#)
- [How Simple Filters Evaluate Packets on page 4530](#)
- [Guidelines for Applying Simple Filters on page 4534](#)
- [Example: Configuring and Applying a Simple Filter on page 4683](#)

#### *Guidelines for Applying Simple Filters*

This topic covers the following information:

- [Statement Hierarchy for Applying Simple Filters on page 4534](#)
- [Restrictions for Applying Simple Filters on page 4534](#)

#### *Statement Hierarchy for Applying Simple Filters*

You can apply a simple filter to the IPv4 ingress traffic at a logical interface by including the **simple-filter** input *simple-filter-name* statement at the [edit interfaces *interface-name* unit *unit-number* family inet] hierarchy level.

```
[edit]
interfaces {
 interface-name {
 unit logical-unit-number {
 family inet {
 simple-filter {
 input filter-name;
 }
 }
 }
 }
}
```

#### *Restrictions for Applying Simple Filters*

You can apply a simple filter to the ingress IPv4 traffic at a logical interface configured on the following hardware only:

- Gigabit Ethernet intelligent queuing (IQ2) PICs installed on M120, M320, or T Series routers.
- Enhanced Queuing Dense Port Concentrators (EQ DPCs) installed on MX Series routers and EX Series switches.

For more information about Ethernet IQ2 PICs and EQ DPCs and related features, see the *Junos OS Class of Service Configuration Guide*. For more information about configuring the MX Series routers, on which EQ DPCs are supported, see the *Junos OS Layer 2 Configuration Guide*.

The following additional restrictions pertain to applying simple filters:

- Simple filters are not supported on Modular Port Concentrator (MPC) interfaces, including Enhanced Queuing MPC interfaces.
- Simple filters are not supported for interfaces in an aggregated-Ethernet bundle.
- You can apply simple filters to **family inet** traffic only. No other protocol family is supported.
- You can apply simple filters to ingress traffic only. Egress traffic is not supported.
- You can apply only a single simple filter to a supported logical interface. Input lists are not supported.

**Related  
Documentation**

- [Simple Filter Overview on page 4529](#)
- [How Simple Filters Evaluate Packets on page 4530](#)
- [Guidelines for Configuring Simple Filters on page 4531](#)
- [Example: Configuring and Applying a Simple Filter on page 4683](#)

---

### Introduction to Firewall Filters Configuration in Logical Systems

- [Stateless Firewall Filters in Logical Systems Overview on page 4535](#)
- [Guidelines for Configuring and Applying Firewall Filters in Logical Systems on page 4536](#)
- [References from a Firewall Filter in a Logical System to Subordinate Objects on page 4539](#)
- [References from a Firewall Filter in a Logical System to Nonfirewall Objects on page 4540](#)
- [References from a Nonfirewall Object in a Logical System to a Firewall Filter on page 4542](#)

#### ***Stateless Firewall Filters in Logical Systems Overview***

This topic covers the following information:

- [Logical Systems on page 4535](#)
- [Stateless Firewall Filters in Logical Systems on page 4536](#)
- [Identifiers for Firewall Objects in Logical Systems on page 4536](#)

#### ***Logical Systems***

With the Junos OS, you can partition a single physical router or switch into multiple logical devices that perform independent routing tasks. Because logical systems perform a subset of the tasks once handled by the physical router or switch, logical systems offer an effective way to maximize the use of a single router or switch.

### ***Stateless Firewall Filters in Logical Systems***

You can configure a separate set of stateless firewall filters for each logical system on a router or switch. To configure a filter in a logical system, you must define the filter in the **firewall** stanza at the **[edit logical-systems *logical-system-name*]** hierarchy level, and you must apply the filter to a logical interface that is also configured at the **[edit logical-systems *logical-system-name*]** hierarchy level.

### ***Identifiers for Firewall Objects in Logical Systems***

To identify firewall objects configured under logical systems, operational **show** commands and firewall-related SNMP MIB objects include a **\_\_logical-system-name/** prefix in the object name. For example, firewall objects configured under the **ls1** logical system include **\_\_ls1/** as the prefix.

#### **Related Documentation**

- [Stateless Firewall Filter Types on page 4465](#)
- [Guidelines for Configuring and Applying Firewall Filters in Logical Systems on page 4536](#)
- [Unsupported Firewall Filter Statements for Logical Systems on page 4767](#)
- [Unsupported Actions for Firewall Filters in Logical Systems on page 4769](#)
- [Example: Configuring a Stateless Firewall Filter to Protect a Logical System Against ICMP Floods on page 3295](#)
- [“Introduction to Logical Systems on page 3259” in the \*Logical Systems Configuration Guide\*](#)
- [“Logical Systems Operations and Restrictions on page 3263” in the \*Logical Systems Configuration Guide\*](#)

### ***Guidelines for Configuring and Applying Firewall Filters in Logical Systems***

This topic covers the following information:

- [Statement Hierarchy for Configuring Firewall Filters in Logical Systems on page 4536](#)
- [Filter Types in Logical Systems on page 4537](#)
- [Firewall Filter Protocol Families in Logical Systems on page 4537](#)
- [Firewall Filter Match Conditions in Logical Systems on page 4538](#)
- [Firewall Filter Actions in Logical Systems on page 4538](#)
- [Statement Hierarchy for Applying Firewall Filters in Logical Systems on page 4538](#)

### ***Statement Hierarchy for Configuring Firewall Filters in Logical Systems***

To configure a firewall filter in a logical system, include the **filter**, **service-filter**, or **simple-filter** statement at the **[edit logical-systems *logical-system-name* firewall family *family-name*]** hierarchy level.

```
[edit]
logical systems {
 logical-system-name {
 firewall {
 family family-name {
```

```
filter filter-name {
 interface-specific;
 physical-interface-filter;
 term term-name {
 filter filter-name;
 from {
 match-conditions;
 }
 then {
 actions;
 }
 }
}

service-filter filter-name { # For 'family inet' or 'family inet6' only.
 term term-name {
 from {
 match-conditions;
 }
 then {
 actions;
 }
 }
}

simple-filter filter-name { # For 'family inet' only.
 term term-name {
 from {
 match-conditions;
 }
 then {
 actions;
 }
 }
}
}
```

## Filter Types in Logical Systems

There are no special restrictions on the types of stateless firewall filter types that you can configure in logical systems.

In a logical system, you can use the same types of stateless firewall filters that are available on a physical router or switch:

- Standard stateless firewall filters
- Service filters
- Simple filters

## Firewall Filter Protocol Families in Logical Systems

There are no special restrictions on the protocol families supported with stateless firewall filters in logical systems.

In a logical system, you can filter the same protocol families as you can on a physical router or switch.

- Standard stateless firewall filters—In logical systems, you can filter the following traffic types: protocol-independent, IPv4, IPv6, MPLS, MPLS-tagged IPv4 or IPv6, VPLS, Layer 2 circuit cross-connection, and Layer 2 bridging.
- Service filters—In logical systems, you can filter IPv4 and IPv6 traffic.
- Simple filters—In logical systems, you can filter IPv4 traffic only.

### ***Firewall Filter Match Conditions in Logical Systems***

There are no special restrictions on the match conditions supported with stateless firewall filters in logical systems.

### ***Firewall Filter Actions in Logical Systems***

There are no special restrictions on the actions supported with stateless firewall filters in logical systems.

### ***Statement Hierarchy for Applying Firewall Filters in Logical Systems***

To apply a firewall filter in a logical system, include the **filter** *filter-name*, **service-filter** *service-filter-name*, or **simple-filter** *simple-filter-name* statement to a logical interface in the logical system.

The following configuration shows the hierarchy levels at which you can apply the statements:

```
[edit]
logical-systems logical-system-name {
 interfaces {
 interface-name {
 unit logical-unit-number {
 family family-name {
 filter {
 group group-name;
 input filter-name;
 input-list [filter-names];
 output filter-name;
 output-list [filter-names]
 }
 rpf-check { # For 'family inet' or 'family inet6' only.
 fail-filter filter-name;
 mode loose;
 }
 service { # For 'family inet' or 'family inet6' only.
 input {
 service-set service-set-name <service-filter service-filter-name>;
 post-service-filter service-filter-name;
 }
 output {
 service-set service-set-name <service-filter service-filter-name>;
 }
 }
 simple-filter { # For 'family inet' only.
```



```

 input simple-filter-name;
 }
 }
 }
}

```

#### Related Documentation

- [Stateless Firewall Filters in Logical Systems Overview on page 4535](#)
- [References from a Firewall Filter in a Logical System to Subordinate Objects on page 4539](#)
- [References from a Firewall Filter in a Logical System to Nonfirewall Objects on page 4540](#)
- [References from a Nonfirewall Object in a Logical System to a Firewall Filter on page 4542](#)
- [Example: Configuring a Stateless Firewall Filter to Protect a Logical System Against ICMP Floods on page 3295](#)
- [Unsupported Firewall Filter Statements for Logical Systems on page 4767](#)
- [Unsupported Actions for Firewall Filters in Logical Systems on page 4769](#)

#### *References from a Firewall Filter in a Logical System to Subordinate Objects*

This topic covers the following information:

- [Resolution of References from a Firewall Filter to Subordinate Objects on page 4539](#)
- [Valid Reference from a Firewall Filter to a Subordinate Object on page 4539](#)

#### *Resolution of References from a Firewall Filter to Subordinate Objects*

If a firewall filter defined in a logical system references a subordinate object (for example, a policer or prefix list), that subordinate object must be defined within the **firewall** stanza of the same logical system. For example, if a firewall filter configuration references a policer, the firewall filter and the policer must be configured under the same **[edit logical-systems logical-system-name firewall]** hierarchy level.

This rule applies even if the same policer is configured under the main firewall configuration or if the same policer is configured as part of a firewall in another logical system.

#### *Valid Reference from a Firewall Filter to a Subordinate Object*

In this example, the firewall filter **filter1** references the policer **pol1**. Both **filter1** and **pol1** are defined under the same firewall object. This configuration is valid. If **pol1** had been defined under another firewall object, the configuration would not be valid.

```

[edit]
logical systems {
 ls-A {
 firewall {
 policer pol1 {
 if-exceeding {
 bandwidth-limit 401k;
 burst-size-limit 50k;
 }
 }
 }
 }
}

```

```
 then discard;
 }
 filter filter1 {
 term one {
 from {
 source-address 12.1.0.0/16;
 }
 then {
 reject host-unknown;
 }
 }
 term two {
 from {
 source-address 12.2.0.0/16;
 }
 then policer pol1;
 }
 }
}
```

**Related  
Documentation**

- [Stateless Firewall Filters in Logical Systems Overview on page 4535](#)
- [Guidelines for Configuring and Applying Firewall Filters in Logical Systems on page 4536](#)
- [References from a Firewall Filter in a Logical System to Nonfirewall Objects on page 4540](#)
- [References from a Nonfirewall Object in a Logical System to a Firewall Filter on page 4542](#)

***References from a Firewall Filter in a Logical System to Nonfirewall Objects***

This topic covers the following information:

- [Resolution of References from a Firewall Filter to Nonfirewall Objects on page 4540](#)
- [Valid Reference to a Nonfirewall Object Outside of the Logical System on page 4541](#)

***Resolution of References from a Firewall Filter to Nonfirewall Objects***

In many cases, a firewall configuration references objects outside the firewall configuration. As a general rule, the referenced object must be defined under the same logical system as the referencing object. However, there are cases when the configuration of the referenced object is not supported at the `[edit logical-systems logical-system-name]` hierarchy level.

### *Valid Reference to a Nonfirewall Object Outside of the Logical System*

This example configuration illustrates an exception to the general rule that the objects referenced by a firewall filter in a logical system must be defined under the same logical system as the referencing object.

In the following scenario, the service filter **inetsf1** is applied to IPv4 traffic associated with the service set **fred** at the logical interface **fe-0/3/2.0**, which is on an adaptive services interface.

- Service filter **inetsf1** is defined in **ls-B** and references prefix list **prefix1**.
- Service set **fred** is defined at the main services hierarchy level, and the policy framework software searches the **[edit services]** hierarchy for the definition of the **fred** service set.

Because service rules cannot be configured in logical systems, firewall filter configurations in the **[edit logical-systems logical-system logical-system-name]** hierarchy are allowed to reference *service sets* outside the logical system hierarchy.

```
[edit]
logical-systems {
 ls-B {
 interfaces {
 fe-0/3/2 {
 unit 0 {
 family inet {
 service {
 input {
 service-set fred service-filter inetsf1;
 }
 }
 }
 }
 }
 }
 }
 policy-options {
 prefix-list prefix1 {
 1.1.0.0/16;
 1.2.0.0/16;
 1.3.0.0/16;
 }
 }
 firewall { # Under logical-system 'ls-B'.
 family inet {
 filter filter1 {
 term one {
 from {
 source-address {
 12.1.0.0/16;
 }
 }
 then {
 reject host-unknown;
 }
 }
 term two {
```

```
 from {
 source-address {
 12.2.0.0/16;
 }
 }
 then policer pol1;
 }
}
service-filter inetsf1 {
 term term1 {
 from {
 source-prefix-list {
 prefix1;
 }
 }
 then count prefix1;
 }
}
}
policer pol1 {
 if-exceeding {
 bandwidth-limit 401k;
 burst-size-limit 50k;
 }
 then discard;
}
}
}
} # End of logical systems configuration.
services { # Main services hierarchy level.
 service-set fred {
 max-flows 100;
 interface-service {
 service-interface sp-1/2/0.0;
 }
 }
}
```

**Related Documentation**

- [Stateless Firewall Filters in Logical Systems Overview on page 4535](#)
- [Guidelines for Configuring and Applying Firewall Filters in Logical Systems on page 4536](#)
- [References from a Firewall Filter in a Logical System to Subordinate Objects on page 4539](#)
- [References from a Nonfirewall Object in a Logical System to a Firewall Filter on page 4542](#)

***References from a Nonfirewall Object in a Logical System to a Firewall Filter***

This topic covers the following information:

- [Resolution of References from a Nonfirewall Object to a Firewall Filter on page 4543](#)
- [Invalid Reference to a Firewall Filter Outside of the Logical System on page 4543](#)
- [Valid Reference to a Firewall Filter Within the Logical System on page 4544](#)
- [Valid Reference to a Firewall Filter Outside of the Logical System on page 4546](#)

**Resolution of References from a Nonfirewall Object to a Firewall Filter**

If a nonfirewall filter object in a logical system references an object in a firewall filter configured in a logical system, the reference is resolved using the following logic:

- If the nonfirewall filter object is configured in a logical system that includes firewall filter configuration statements, the policy framework software searches the **[edit logical-systems logical-system-name firewall]** hierarchy level. Firewall filter configurations that belong to *other* logical systems or to the main **[edit firewall]** hierarchy level are not searched.
- If the nonfirewall filter object is configured in a logical system that does not include any firewall filter configuration statements, the policy framework software searches the firewall configurations defined at the **[edit firewall]** hierarchy level.

**Invalid Reference to a Firewall Filter Outside of the Logical System**

This example configuration illustrates an unresolvable reference from a nonfirewall object in a logical system to a firewall filter.

In the following scenario, the stateless firewall filters **filter1** and **fred** are applied to the logical interface **fe-0/3/2.0** in the logical system **ls-C**.

- Filter **filter1** is defined in **ls-C**.
- Filter **fred** is defined in the main firewall configuration.

Because **ls-C** contains firewall filter statements (for **filter1**), the policy framework software resolves references to and from firewall filters by searching the **[edit logical systems ls-C firewall]** hierarchy level. Consequently, the reference from **fe-0/3/2.0** in the logical system to **fred** in the main firewall configuration cannot be resolved.

```
[edit]
logical-systems {
 ls-C {
 interfaces {
 fe-0/3/2 {
 unit 0 {
 family inet {
 filter {
 input-list [filter1 fred];
 }
 }
 }
 }
 }
 }
 firewall { # Under logical system 'ls-C'.
 family inet {
 filter filter1 {
 term one {
 from {
 source-address 12.1.0.0/16;
 }
 then {
 reject host-unknown;
 }
 }
 }
 }
 }
}
```

```
 }
 term two {
 from {
 source-address 12.2.0.0/16;
 }
 then policer pol1;
 }
 }
}
policer pol1 {
 if-exceeding {
 bandwidth-limit 401k;
 burst-size-limit 50k;
 }
 then discard;
}
}
}
} # End of logical systems
firewall { # Under the main firewall hierarchy level
 family inet {
 filter fred {
 term one {
 from {
 source-address 11.1.0.0/16;
 }
 then {
 log;
 reject host-unknown;
 }
 }
 }
 }
}
} # End of main firewall configurations.
```

#### ***Valid Reference to a Firewall Filter Within the Logical System***

This example configuration illustrates resolvable references from a nonfirewall object in a logical system to two firewall filter.

In the following scenario, the stateless firewall filters **filter1** and **fred** are applied to the logical interface **fe-0/3/2.0** in the logical system **ls-C**.

- Filter **filter1** is defined in **ls-C**.
- Filter **fred** is defined in **ls-C** and also in the main firewall configuration.

Because **ls-C** contains firewall filter statements, the policy framework software resolves references to and from firewall filters by searching the **[edit logical systems ls-C firewall]** hierarchy level. Consequently, the references from **fe-0/3/2.0** in the logical system to **filter1** and **fred** use the stateless firewall filters configured in **ls-C**.

```
[edit]
logical-systems {
 ls-C {
 interfaces {
 fe-0/3/2 {
```

```

 unit 0 {
 family inet {
 filter {
 input-list [filter1 fred];
 }
 }
 }
}
firewall { # Under logical system 'ls-C'.
 family inet {
 filter filter1 {
 term one {
 from {
 source-address 12.1.0.0/16;
 }
 then {
 reject host-unknown;
 }
 }
 term two {
 from {
 source-address 12.2.0.0/16;
 }
 then policer pol1;
 }
 }
 filter fred { # This 'fred' is in 'ls-C'.
 term one {
 from {
 source-address 10.1.0.0/16;
 }
 then {
 log;
 reject host-unknown;
 }
 }
 }
 }
}
policer pol1 {
 if-exceeding {
 bandwidth-limit 401k;
 burst-size-limit 50k;
 }
 then discard;
}
}
} # End of logical systems configurations.
firewall { # Main firewall filter hierarchy level
 family inet {
 filter fred {
 term one {
 from {
 source-address 11.1.0.0/16;
 }
 }
 }
 }
}

```

```
 then {
 log;
 reject host-unknown;
 }
 }
}
}
} # End of main firewall configurations.
```

### ***Valid Reference to a Firewall Filter Outside of the Logical System***

This example configuration illustrates resolvable references from a nonfirewall object in a logical system to two firewall filter.

In the following scenario, the stateless firewall filters **filter1** and **fred** are applied to the logical interface **fe-0/3/2.0** in the logical system **ls-C**.

- Filter **filter1** is defined in the main firewall configuration.
- Filter **fred** is defined in the main firewall configuration.

Because **ls-C** does not contain any firewall filter statements, the policy framework software resolves references to and from firewall filters by searching the **[edit firewall]** hierarchy level. Consequently, the references from **fe-0/3/2.0** in the logical system to **filter1** and **fred** use the stateless firewall filters configured in the main firewall configuration.

```
[edit]
logical-systems {
 ls-C {
 interfaces {
 fe-0/3/2 {
 unit 0 {
 family inet {
 filter {
 input-list [filter1 fred];
 }
 }
 }
 }
 }
 }
}
} # End of logical systems configurations.
firewall { # Main firewall hierarchy level.
 family inet {
 filter filter1 {
 term one {
 from {
 source-address 12.1.0.0/16;
 }
 then {
 reject host-unknown;
 }
 }
 term two {
 from {
 source-address 12.2.0.0/16;
 }
 }
 }
 }
}
```



```

 then policer pol1;
 }
}
filter fred {
 term one {
 from {
 source-address 11.1.0.0/16;
 }
 then {
 log;
 reject host-unknown;
 }
 }
}
}
policer pol1 {
 if-exceeding {
 bandwidth-limit 701k;
 burst-size-limit 70k;
 }
 then discard;
}
} # End of main firewall configurations.

```

#### Related Documentation

- [Stateless Firewall Filters in Logical Systems Overview on page 4535](#)
- [Guidelines for Configuring and Applying Firewall Filters in Logical Systems on page 4536](#)
- [References from a Firewall Filter in a Logical System to Subordinate Objects on page 4539](#)
- [References from a Firewall Filter in a Logical System to Nonfirewall Objects on page 4540](#)

## Configuration

- [Standard Firewall Filter Configurations to Match Packets on page 4548](#)
- [Standard Firewall Filters to Count Packets on page 4558](#)
- [Standard Firewall Filters to Act on Packets on page 4569](#)
- [Standard Firewall Filters for Trusted Sources on page 4576](#)
- [Standard Firewall Filters for Flood Prevention on page 4600](#)
- [Standard Firewall Filters for Fragment Handling on page 4613](#)
- [Standard Firewall Filters for Setting Rate Limits on page 4617](#)
- [Examples of Standard Firewall Filters Configuration on page 4620](#)
- [Examples of Standard Firewall Filters Configuration Options on page 4669](#)
- [Service Filters Configuration on page 4678](#)
- [Simple Filters Configuration on page 4683](#)
- [Firewall Filters Configuration in Logical Systems on page 4687](#)
- [Configuration Statements on page 4691](#)

## Standard Firewall Filter Configurations to Match Packets

---

- [Example: Configuring a Filter to Match on IPv6 Flags on page 4548](#)
- [Example: Configuring a Filter to Match on Port and Protocol Fields on page 4549](#)
- [Example: Configuring a Filter to Match on Two Unrelated Criteria on page 4552](#)
- [Example: Configuring a Filter to Limit TCP Access to a Port Based On a Prefix List on page 4555](#)

### **Example: Configuring a Filter to Match on IPv6 Flags**

This example shows how to configure a filter to match on IPv6 TCP flags.

- [Requirements on page 4548](#)
- [Overview on page 4548](#)
- [Configuration on page 4548](#)
- [Verification on page 4549](#)

### **Requirements**

No special configuration beyond device initialization is required before configuring this example.

### **Overview**

In this example, you configure a filter to match on IPv6 TCP flags. You can use this example to configure IPv6 TCP flags in the SRX100, SRX210, SRX240, SRX650, and J Series security devices and in M Series, MX Series, and T Series routing devices.

### **Configuration**

#### **Step-by-Step Procedure**

To configure a filter to match on IPv6 TCP flags:

1. Include the family statement at the firewall hierarchy level, specifying **inet6** as the protocol family.  
  
[edit]  
user@host# **edit firewall family inet6**
2. Create the stateless firewall filter.  
  
[edit firewall family inet6]  
user@host# **edit filter tcpfilt**
3. Define the first term for the filter.  
  
[edit firewall family inet6 filter tcpfilt]  
user@host# **edit term 1**
4. Define the source address match conditions for the term.  
  
[edit firewall family inet6 filter tcpfilt term 1]  
user@host# **set from next-header tcp tcp-flags syn**
5. Define the actions for the term.  
  
[edit firewall family inet6 filter tcpfilt term 1]

```
user@host# set then count tcp_syn_pkt log accept
```

6. If you are done configuring the device, commit the configuration.

```
[edit firewall family inet6 filter tcpfilt term 1]
```

```
user@host top
```

```
[edit]
```

```
user@host# commit
```

### Verification

To confirm that the configuration is working properly, enter the **show firewall filter tcpfilt** command.

### Example: Configuring a Filter to Match on Port and Protocol Fields

This example shows how to configure a standard stateless firewall filter to match on destination port and protocol fields.

- [Requirements on page 4549](#)
- [Overview on page 4549](#)
- [Configuration on page 4549](#)
- [Verification on page 4552](#)

### Requirements

No special configuration beyond device initialization is required before configuring this example.

### Overview

In this example, you configure a stateless firewall filter that accepts all IPv4 packets except for TCP and UDP packets. TCP and UDP packets are accepted if destined for the SSH port or the Telnet port. All other packets are rejected.

### Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

- [Configure the Stateless Firewall Filter on page 4550](#)
- [Apply the Stateless Firewall Filter to a Logical Interface on page 4550](#)
- [Confirm and Commit Your Candidate Configuration on page 4551](#)

### CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level:

```
set firewall family inet filter filter1 term term1 from protocol-except tcp
set firewall family inet filter filter1 term term1 from protocol-except udp
set firewall family inet filter filter1 term term1 then accept
set firewall family inet filter filter1 term term2 from address 192.168.0.0/16
set firewall family inet filter filter1 term term2 then reject
set firewall family inet filter filter1 term term3 from destination-port ssh
```

```
set firewall family inet filter filter1 term term3 from destination-port telnet
set firewall family inet filter filter1 term term3 then accept
set firewall family inet filter filter1 term term4 then reject
set interfaces ge-0/0/1 unit 0 family inet address 10.1.2.3/30
set interfaces ge-0/0/1 unit 0 family inet filter input filter1
```

### *Configure the Stateless Firewall Filter*

#### **Step-by-Step Procedure**

To configure the stateless firewall filter **filter1**:

1. Create the IPv4 stateless firewall filter.

```
[edit]
user@host# edit firewall family inet filter filter1
```

2. Configure a term to accept all traffic except for TCP and UDP packets.

```
[edit firewall family inet filter filter1]
user@host# set term term1 from protocol-except tcp
user@host# set term term1 from protocol-except udp
user@host# set term term1 then accept
```

3. Configure a term to reject packets to or from the 192.168/16 prefix.

```
[edit firewall family inet filter filter1]
user@host# set term term2 from address 192.168.0.0/16
user@host# set term term2 then reject
```

4. Configure a term to accept packets destined for either the SSH port or the Telnet port.

```
[edit firewall family inet filter filter1]
user@host# set term term3 from destination-port ssh
user@host# set term term3 from destination-port telnet
user@host# set term term3 then accept
```

5. Configure the last term to reject all packets.

```
[edit firewall family inet filter filter1]
user@host# set term term4 then reject
```

### *Apply the Stateless Firewall Filter to a Logical Interface*

#### **Step-by-Step Procedure**

To apply the stateless firewall filter to a logical interface:

1. Configure the logical interface to which you will apply the stateless firewall filter.

```
[edit]
user@host# edit interfaces ge-0/0/1 unit 0 family inet
```

2. Configure the interface address for the logical interface.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
user@host# set address 10.1.2.3/30
```

3. Apply the stateless firewall filter to the logical interface.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
user@host# set filter input filter1
```

*Confirm and Commit Your Candidate Configuration***Step-by-Step  
Procedure**

To confirm and then commit your candidate configuration:

1. Confirm the configuration of the stateless firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
 filter filter1 {
 term term1 {
 from {
 protocol-except [tcp udp];
 }
 then {
 accept;
 }
 }
 term term2 {
 from {
 address 192.168/16;
 }
 then {
 reject;
 }
 }
 term term3 {
 from {
 destination-port [ssh telnet];
 }
 then {
 accept;
 }
 }
 term term4 {
 then {
 reject;
 }
 }
 }
}
```

2. Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
 unit 0 {
 family inet {
 filter {
 input filter1;
 }
 }
 }
}
```

```
 address 10.1.2.3/30;
 }
}
```

3. If you are done configuring the device, commit your candidate configuration.

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

### **Verification**

To confirm that the configuration is working properly, enter the [show firewall filter filter1](#) operational mode command.

### **Related Documentation**

- [Understanding How to Use Standard Firewall Filters on page 3293](#)
- [Example: Configuring a Filter to Match on IPv6 Flags on page 4548](#)
- [Example: Configuring a Filter to Match on Two Unrelated Criteria on page 4552](#)

### **Example: Configuring a Filter to Match on Two Unrelated Criteria**

This example shows how to configure a standard stateless firewall filter to match on two unrelated criteria.

- [Requirements on page 4552](#)
- [Overview on page 4552](#)
- [Configuration on page 4552](#)
- [Verification on page 4554](#)

### **Requirements**

No special configuration beyond device initialization is required before configuring this example.

### **Overview**

In this example, you use a standard stateless firewall filter to match IPv4 packets that are either OSPF packets or packets that come from an address in the prefix **10.108/16**, and send an **administratively-prohibited** ICMP message for all packets that do not match.

### **Configuration**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configuring the IPv4 Firewall Filter on page 4553](#)
- [Applying the IPv4 Firewall Filter to a Logical Interface on page 4554](#)

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

```
set firewall family inet filter ospf_or_131 term protocol_match from protocol ospf
set firewall family inet filter ospf_or_131 term address-match from source-address
 10.108.0.0/16
set interfaces ge-0/0/1 unit 0 family inet address 10.1.2.3/30
set interfaces ge-0/0/1 unit 0 family inet filter input ospf_or_131
```

### *Configuring the IPv4 Firewall Filter*

**Step-by-Step Procedure** To configure the IPv4 firewall filter:

1. Enable configuration of the IPv4 firewall filter.

```
[edit]
user@host# edit firewall family inet filter ospf_or_131
```

2. Configure the first term to accept OSPF packets.

```
[edit firewall family inet filter ospf_or_131]
user@host# set term protocol_match from protocol ospf
```

Packets that match the condition are accepted by default. Because another term follows this term, packets that do not match this condition are evaluated by the next term.

3. Configure the second term to accept packets from any IPv4 address in a particular prefix.

```
[edit firewall family inet filter ospf_or_131]
user@host# set term address_match from source-address 10.108.0.0/16
```

Packets that match this condition are accepted by default. Because this is the last term in the filter, packets that do not match this condition are discarded by default.

**Results** Confirm the configuration of the stateless firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
 filter ospf_or_131 {
 term protocol_match {
 from {
 protocol ospf;
 }
 }
 term address_match {
 from {
 source-address {
 10.108.0.0/16;
 }
 }
 }
 }
}
```

```
}
}
```

### *Applying the IPv4 Firewall Filter to a Logical Interface*

**Step-by-Step Procedure** To apply the stateless firewall filter to a logical interface:

1. Enable configuration of a logical interface.

```
[edit]
user@host# edit interfaces ge-0/0/1 unit 0 family inet
```

2. Configure an IP address for the logical interface.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
user@host# set address 10.1.2.3/30
```

3. Apply the IPv4 firewall filter to the logical interface.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
user@host# set filter input ospf_or_131
```

**Results** Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
 unit 0 {
 family inet {
 filter {
 input ospf_or_131;
 }
 address 10.1.2.3/30;
 }
 }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

### *Verification*

To confirm that the configuration is working properly, enter the **show firewall filter ospf\_or\_131** operational mode command.

**Related Documentation**

- [Understanding How to Use Standard Firewall Filters on page 3293](#)
- [Example: Configuring a Filter to Match on IPv6 Flags on page 4548](#)
- [Example: Configuring a Filter to Match on Port and Protocol Fields on page 4549](#)



**Example: Configuring a Filter to Limit TCP Access to a Port Based On a Prefix List**

This example shows how to configure a standard stateless firewall filter that limits certain TCP and Internet Control Message Protocol (ICMP) traffic destined for the Routing Engine by specifying a list of prefix sources that contain allowed BGP peers.

- [Requirements on page 4555](#)
- [Overview on page 4555](#)
- [Configuration on page 4555](#)
- [Verification on page 4557](#)

**Requirements**

No special configuration beyond device initialization is required before configuring this example.

**Overview**

In this example, you create a stateless firewall filter that blocks all TCP connection attempts to port 179 from all requesters except BGP peers that have a specified prefix.

A source prefix list, **plist\_bgp179**, is created that specifies the list of source prefixes that contain allowed BGP peers.

The stateless firewall filter **filter\_bgp179** matches all packets from the source prefix list **plist\_bgp179** to the destination port number 179.

**Configuration**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

- [Configure the Filter on page 4556](#)
- [Results on page 4556](#)

**CLI Quick Configuration**

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

```
set policy-options prefix-list plist_bgp179 apply-path "protocols bgp group <*> neighbor
<*>"
set firewall family inet filter filter_bgp179 term 1 from source-address 0.0.0.0/0
set firewall family inet filter filter_bgp179 term 1 from source-prefix-list plist_bgp179 except
set firewall family inet filter filter_bgp179 term 1 from destination-port bgp
set firewall family inet filter filter_bgp179 term 1 then reject
set firewall family inet filter filter_bgp179 term 2 then accept
set interfaces lo0 unit 0 family inet filter input filter_bgp179
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
```

### Configure the Filter

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure the filter:

1. Expand the prefix list **bgp179** to include all prefixes pointed to by the BGP peer group defined by **protocols bgp group <\*> neighbor <\*>**.

```
[edit policy-options prefix-list plist_bgp179]
user@host# set apply-path "protocols bgp group <*> neighbor <*>"
```

2. Define the filter term that rejects TCP connection attempts to port 179 from all requesters except the specified BGP peers.

```
[edit firewall family inet filter filter_bgp179]
user@host# set term term1 from source-address 0.0.0.0/0
user@host# set term term1 from source-prefix-list bgp179 except
user@host# set term term1 from destination-port bgp
user@host# set term term1 then reject
```

3. Define the other filter term to accept all packets.

```
[edit firewall family inet filter filter_bgp179]
user@host# set term term2 then accept
```

4. Apply the firewall filter to the loopback interface.

```
[edit interfaces lo0 unit 0 family inet]
user@host# set filter input filter_bgp179
user@host# set address 127.0.0.1/32
```

### Results

From configuration mode, confirm your configuration by entering the **show firewall**, **show interfaces**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show firewall
family inet {
 filter filter_bgp179 {
 term 1 {
 from {
 source-address {
 0.0.0.0/0;
 }
 source-prefix-list {
 plist_bgp179 except;
 }
 destination-port bgp;
 }
 then {
 reject;
 }
 }
 }
}
```

```

 term 2 {
 then {
 accept;
 }
 }
}

user@host# show interfaces
lo0 {
 unit 0 {
 family inet {
 filter {
 input filter_bgp179;
 }
 address 127.0.0.1/32;
 }
 }
}

user@host# show policy-options
prefix-list plist_bgp179 {
 apply-path "protocols bgp group <*> neighbor <*>";
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Repeat the procedure, where appropriate, for every BGP-enabled device in the network, using the appropriate interface names and addresses for each BGP-enabled device.

### Verification

Confirm that the configuration is working properly.

### Displaying the Firewall Filter Applied to the Loopback Interface

**Purpose** Verify that the firewall filter **filter\_bgp179** is applied to the IPv4 input traffic at logical interface **lo0.0**.

**Action** Use the **show interfaces statistics** operational mode command for logical interface **lo0.0**, and include the **detail** option. Under the **Protocol inet** section of the command output section, the **Input Filters** field displays the name of the stateless firewall filter applied to the logical interface in the input direction:

```

[edit]
user@host> show interfaces statistics lo0.0 detail
Logical interface lo0.0 (Index 321) (SNMP ifIndex 16) (Generation 130)
Flags: SNMP-Traps Encapsulation: Unspecified
Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0

```

```
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol inet, MTU: Unlimited, Generation: 145, Route table: 0
Flags: Sendbroadcast-pkt-to-re
Input Filters: filter_bgp179
Addresses, Flags: Primary
Destination: Unspecified, Local: 127.0.0.1, Broadcast: Unspecified,
Generation: 138
```

**Related  
Documentation**

- [Understanding How to Use Standard Firewall Filters on page 3293](#)
- [Firewall Filter Match Conditions Based on Address Fields on page 4493](#)
- [Example: Configuring a Stateless Firewall Filter to Protect Against TCP and ICMP Floods on page 4600](#)
- [Example: Configuring a Filter to Accept Packets Based on IPv6 TCP Flags on page 4610](#)
- [“prefix-list on page 4795”](#) in the *Routing Policy Configuration Guide*

---

**Standard Firewall Filters to Count Packets**

---

- [Example: Configuring a Filter to Count Accepted and Rejected Packets on page 4558](#)
- [Example: Configuring a Filter to Count and Discard IP Options Packets on page 4561](#)
- [Example: Configuring a Filter to Count IP Options Packets on page 4564](#)

**Example: Configuring a Filter to Count Accepted and Rejected Packets**

This example shows how to configure a firewall filter to count packets.

- [Requirements on page 4558](#)
- [Overview on page 4558](#)
- [Configuration on page 4559](#)
- [Verification on page 4561](#)

**Requirements**

No special configuration beyond device initialization is required before configuring this example.

**Overview**

In this example, you use a stateless firewall filter to reject all addresses except 192.168.5.0/24.

**Topology**

In the first term, the match condition **address 192.168.5.0/24 except** causes this address to be considered a mismatch, and this address is passed to the next term in the filter. The match condition **address 0.0.0.0/0** matches all other packets, and these are counted, logged, and rejected.

In the second term, all packets that passed through the first term (that is, packets whose address matches **192.168.5.0/24**) are counted, logged, and accepted.

### Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configure the Stateless Firewall Filter on page 4559](#)
- [Apply the Stateless Firewall Filter to a Logical Interface on page 4560](#)
- [Confirm and Commit Your Candidate Configuration on page 4560](#)

### CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set firewall family inet filter fire1 term 1 from address 192.168.5.0/24 except
set firewall family inet filter fire1 term 1 from address 0.0.0.0/0
set firewall family inet filter fire1 term 1 then count reject_pref1_1
set firewall family inet filter fire1 term 1 then log
set firewall family inet filter fire1 term 1 then reject
set firewall family inet filter fire1 term 2 then count reject_pref1_2
set firewall family inet filter fire1 term 2 then log
set firewall family inet filter fire1 term 2 then accept
set interfaces ge-0/0/1 unit 0 family inet filter input fire1
set interfaces ge-0/0/1 unit 0 family inet address 10.1.2.3/30
```

### Configure the Stateless Firewall Filter

### Step-by-Step Procedure

To configure the stateless firewall filter **fire1**:

1. Create the stateless firewall filter **fire1**.  
  

```
[edit]
user@host# edit firewall family inet filter fire1
```
2. Configure the first term to reject all addresses except those to or from the **192.168.5.0/24** prefix and then count, log, and reject all other packets.  
  

```
[edit firewall family inet filter fire1]
user@host# set term 1 from address 192.168.5.0/24 except
user@host# set term 1 from address 0.0.0.0/0
user@host# set term 1 then count reject_pref1_1
user@host# set term 1 then log
user@host# set term 1 then reject
```
3. Configure the next term to count, log, and accept packets in the **192.168.5.0/24** prefix.  
  

```
[edit firewall family inet filter fire1]
user@host# set term 2 then count reject_pref1_2
user@host# set term 2 then log
user@host# set term 2 then accept
```

### *Apply the Stateless Firewall Filter to a Logical Interface*

#### **Step-by-Step Procedure**

To apply the stateless firewall filter to a logical interface:

1. Configure the logical interface to which you will apply the stateless firewall filter.

```
[edit]
user@host# edit interfaces ge-0/0/1 unit 0 family inet
```

2. Configure the interface address for the logical interface.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
user@host# set address 10.1.2.3/30
```

3. Apply the stateless firewall filter to the logical interface.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
user@host# set filter input fire1
```

### *Confirm and Commit Your Candidate Configuration*

#### **Step-by-Step Procedure**

To confirm and then commit your candidate configuration:

1. Confirm the configuration of the stateless firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
 filter fire1 {
 term 1 {
 from {
 address {
 192.168.5.0/24 except;
 0.0.0.0/0;
 }
 }
 then {
 count reject_pref1_1;
 log;
 reject;
 }
 }
 term 2 {
 then {
 count reject_pref1_2;
 log;
 accept;
 }
 }
 }
}
```

2. Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
 unit 0 {
 family inet {
 filter {
 input fire1;
 }
 address 10.1.2.3/30;
 }
 }
}
```

3. If you are done configuring the device, commit your candidate configuration.

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

### Verification

To confirm that the configuration is working properly, enter the **show firewall filter fire1** operational mode command. You can also display the log and individual counters separately by using the following forms of the command:

- **show firewall counter reject\_pref1\_1**
- **show firewall counter reject\_pref1\_2**
- **show firewall log**

### Related Documentation

- [Understanding How to Use Standard Firewall Filters on page 3293](#)
- [Example: Configuring a Filter to Count IP Options Packets on page 4564](#)
- [Example: Configuring a Filter to Count and Discard IP Options Packets on page 4561](#)

### Example: Configuring a Filter to Count and Discard IP Options Packets

This example shows how to configure a standard stateless firewall to count packets.

- [Requirements on page 4561](#)
- [Overview on page 4562](#)
- [Configuration on page 4562](#)
- [Verification on page 4564](#)

### Requirements

No special configuration beyond device initialization is required before configuring this example.

Because the filter term matches on *any* IP option value, the filter term can use the **count** nonterminating action without the **discard** terminating action or (alternatively) without requiring an interface on a 10-Gigabit Ethernet Modular Port Concentrator (MPC), 60-Gigabit Ethernet MPC, 60-Gigabit Queuing Ethernet MPC, or 60-Gigabit Ethernet Enhanced Queuing MPC on an MX Series router.

### Overview

In this example, you use a standard stateless firewall filter to count and discard packets that include any IP option value but accept all other packets.

The IP option header field is an optional field in IPv4 headers only. The **ip-options** and **ip-options-except** match conditions are supported for standard stateless firewall filters and service filters only.



**NOTE:** On M and T series routers, firewall filters cannot count **ip-options** packets on a per option type and per interface basis. A limited work around is to use the `show pfe statistics ip options` command to see **ip-options** statistics on a per Packet Forwarding Engine (PFE) basis. See *show pfe statistics ip* for sample output.

### Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configure the Stateless Firewall Filter on page 4562](#)
- [Apply the Stateless Firewall Filter to a Logical Interface on page 4563](#)
- [Confirm and Commit Your Candidate Configuration on page 4563](#)

#### CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set firewall family inet filter block_ip_options term 10 from ip-options any
set firewall family inet filter block_ip_options term 10 then count option_any
set firewall family inet filter block_ip_options term 10 then discard
set firewall family inet filter block_ip_options term 999 then accept
set interfaces ge-0/0/1 unit 0 family inet filter input block_ip_options
set interfaces ge-0/0/1 unit 0 family inet address 10.1.2.3/30
```

#### Configure the Stateless Firewall Filter

#### Step-by-Step Procedure

To configure the stateless firewall filter:

1. Create the stateless firewall filter **block\_ip\_options**.

```
[edit]
user@host# edit firewall family inet filter block_ip_options
```



2. Configure the first term to count and discard packets that include any IP options header fields.

```
[edit firewall family inet filter block_ip_options]
user@host# set term 10 from ip-options any
user@host# set term 10 then count option_any
user@host# set term 10 then discard
```

3. Configure the other term to accept all other packets.

```
[edit firewall family inet filter block_ip_options]
user@host# set term 999 then accept
```

### *Apply the Stateless Firewall Filter to a Logical Interface*

#### **Step-by-Step Procedure**

To apply the stateless firewall filter to a logical interface:

1. Configure the logical interface to which you will apply the stateless firewall filter.

```
[edit]
user@host# edit interfaces ge-0/0/1 unit 0 family inet
```

2. Configure the interface address for the logical interface.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
user@host# set address 10.1.2.3/30
```

3. Apply the stateless firewall filter to the logical interface.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
user@host# set filter input block_ip_options
```

### *Confirm and Commit Your Candidate Configuration*

#### **Step-by-Step Procedure**

To confirm and then commit your candidate configuration:

1. Confirm the configuration of the stateless firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
 filter block_ip_options {
 term 10 {
 from {
 ip-options any;
 }
 then {
 count option_any;
 discard;
 }
 }
 term 999 {
 then accept;
 }
 }
}
```

2. Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
 unit 0 {
 family inet {
 filter {
 input block_ip_options;
 }
 address 10.1.2.3/30;
 }
 }
}
```

3. If you are done configuring the device, commit your candidate configuration.

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

### Verification

To confirm that the configuration is working properly, enter the **show firewall filter block\_ip\_options** operational mode command. To display the count of discarded packets separately, enter the **show firewall count option\_any** form of the command.

### Related Documentation

- [Understanding How to Use Standard Firewall Filters on page 3293](#)
- [Example: Configuring a Filter to Count Accepted and Rejected Packets on page 4558](#)
- [Example: Configuring a Filter to Count IP Options Packets on page 4564](#)

### Example: Configuring a Filter to Count IP Options Packets

This example shows how use a stateless firewall filter to count individual IP options packets:

- [Requirements on page 4564](#)
- [Overview on page 4565](#)
- [Configuration on page 4565](#)
- [Verification on page 4569](#)

### Requirements

This example uses an interface on a 10-Gigabit Ethernet Modular Port Concentrator (MPC), 60-Gigabit Ethernet MPC, 60-Gigabit Queuing Ethernet MPC, or 60-Gigabit Ethernet Enhanced Queuing MPC on an MX Series router. This interface enables you to apply an IPv4 firewall filter (standard or service filter) that can use the **count**, **log**, and **syslog** nonterminating actions on packets that match a *specific ip-option* value without having to also use the **discard** terminating action.

No special configuration beyond device initialization is required before configuring this example.

### Overview

In this example, you use a stateless firewall filter to count IP options packets but not block any traffic. Also, the filter logs packets that have loose or strict source routing.

The IP option header field is an optional field in IPv4 headers only. The **ip-options** and **ip-options-except** match conditions are supported for standard stateless firewall filters and service filters only.



**NOTE:** On M and T series routers, firewall filters cannot count **ip-options** packets on a per option type and per interface basis. A limited work around is to use the `show pfe statistics ip options` command to see **ip-options** statistics on a per Packet Forwarding Engine (PFE) basis. See *show pfe statistics ip* for sample output.

### Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configure the Stateless Firewall Filter on page 4566](#)
- [Apply the Stateless Firewall Filter to a Logical Interface on page 4567](#)
- [Confirm and Commit Your Candidate Configuration on page 4567](#)

### CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set firewall family inet filter ip_options_filter term match_strict_source from ip-options
strict-source-route
set firewall family inet filter ip_options_filter term match_strict_source then count
strict_source_route
set firewall family inet filter ip_options_filter term match_strict_source then log
set firewall family inet filter ip_options_filter term match_strict_source then accept
set firewall family inet filter ip_options_filter term match_loose_source from ip-options
loose-source-route
set firewall family inet filter ip_options_filter term match_loose_source then count
loose_source_route
set firewall family inet filter ip_options_filter term match_loose_source then log
set firewall family inet filter ip_options_filter term match_loose_source then accept
set firewall family inet filter ip_options_filter term match_record from ip-options
record-route
set firewall family inet filter ip_options_filter term match_record then count record_route
set firewall family inet filter ip_options_filter term match_record then accept
set firewall family inet filter ip_options_filter term match_timestamp from ip-options
timestamp
set firewall family inet filter ip_options_filter term match_timestamp then count timestamp
```

```
set firewall family inet filter ip_options_filter term match_timestamp then accept
set firewall family inet filter ip_options_filter term match_router_alert from ip-options
router-alert
set firewall family inet filter ip_options_filter term match_router_alert then count
router_alert
set firewall family inet filter ip_options_filter term match_router_alert then accept
set firewall family inet filter ip_options_filter term match_all then accept
set interfaces ge-0/0/1 unit 0 family inet address 10.1.2.3/30
set interfaces ge-0/0/1 unit 0 family inet filter input ip_options_filter
```

### *Configure the Stateless Firewall Filter*

#### **Step-by-Step Procedure**

To configure the stateless firewall filter `ip_option_filter`:

1. Create the stateless firewall filter `ip_option_filter`.  

```
[edit]
user@host# edit firewall family inet filter ip_options_filter
```
2. Configure the first term to count, log, and accept packets with the `strict_source_route` IP optional header field.  

```
[edit firewall family inet filter ip_option_filter]
user@host# set term match_strict_source from ip-options strict_source_route
user@host# set term match_strict_source then count strict_source_route
user@host# set term match_strict_source then log
user@host# set term match_strict_source then accept
```
3. Configure the next term to count, log, and accept packets with the `loose-source-route` IP optional header field.  

```
[edit firewall family inet filter ip_option_filter]
user@host# set term match_loose_source from ip-options loose-source-route
user@host# set term match_loose_source then count loose_source_route
user@host# set term match_loose_source then log
user@host# set term match_loose_source then accept
```
4. Configure the next term to count and accept packets with the `record-route` IP optional header field.  

```
[edit firewall family inet filter ip_option_filter]
user@host# set term match_record from ip-options record-route
user@host# set term match_record then count record_route
user@host# set term match_record then accept
```
5. Configure the next term to count and accept packets with the `timestamp` IP optional header field.  

```
[edit firewall family inet filter ip_option_filter]
user@host# set term match_timestamp from ip-options timestamp
user@host# set term match_timestamp then count timestamp
user@host# set term match_timestamp then accept
```
6. Configure the next term to count and accept packets with the `router-alert` IP optional header field.  

```
[edit firewall family inet filter ip_option_filter]
user@host# set term match_router_alert from ip-options router-alert
user@host# set term match_router_alert then count router_alert
```

```
user@host# set term match_router_alert then accept
```

7. Create the last term to accept any packet without incrementing any counters.

```
[edit firewall family inet filter ip_option_filter]
user@host# set term match_all then accept
```

### *Apply the Stateless Firewall Filter to a Logical Interface*

#### **Step-by-Step Procedure**

To apply the stateless firewall filter to a logical interface:

1. Configure the logical interface to which you will apply the stateless firewall filter.

```
[edit]
user@host# edit interfaces ge-0/0/1 unit 0 family inet
```

2. Configure the interface address for the logical interface.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
user@host# set address 10.1.2.3/30
```

3. Apply the stateless firewall filter to the logical interface.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
user@host# set filter input ip_options_filter
```

### *Confirm and Commit Your Candidate Configuration*

#### **Step-by-Step Procedure**

To confirm and then commit your candidate configuration:

1. Confirm the configuration of the stateless firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
 filter ip_options_filter {
 term match_strict_source {
 from {
 ip-options strict-source-route;
 }
 then {
 count strict_source_route;
 log;
 accept;
 }
 }
 }
 term match_loose_source {
 from {
 ip-options loose-source-route;
 }
 then {
 count loose_source_route;
 log;
 accept;
 }
 }
}
```

```
}
term match_record {
 from {
 ip-options record-route;
 }
 then {
 count record_route;
 accept;
 }
}
term match_timestamp {
 from {
 ip-options timestamp;
 }
 then {
 count timestamp;
 accept;
 }
}
term match_router_alert {
 from {
 ip-options router-alert;
 }
 then {
 count router_alert;
 accept;
 }
}
term match_all {
 then accept;
}
}
```

2. Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
 unit 0 {
 family inet {
 filter {
 input ip_option_filter;
 }
 address 10.1.2.3/30;
 }
 }
}
```

3. If you are done configuring the device, commit your candidate configuration.

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

### Verification

To confirm that the configuration is working properly, enter the [show firewall filter ip\\_option\\_filter](#) operational mode command. You can also display the log and individual counters separately by using the following forms of the command:

- `show firewall counter strict_source_route`
- `show firewall counter loose_source_route`
- `show firewall counter record_route`
- `show firewall counter timestamp`
- `show firewall counter router_alert`
- `show firewall log`

### Related Documentation

- [Understanding How to Use Standard Firewall Filters on page 3293](#)
- [Example: Configuring a Filter to Count Accepted and Rejected Packets on page 4558](#)
- [Example: Configuring a Filter to Count and Discard IP Options Packets on page 4561](#)

---

### Standard Firewall Filters to Act on Packets

- [Example: Configuring a Filter to Set the DSCP Bit to Zero on page 4569](#)
- [Example: Configuring a Filter to Count and Sample Accepted Packets on page 4572](#)

#### *Example: Configuring a Filter to Set the DSCP Bit to Zero*

This example shows how to configure a standard stateless firewall filter based on the Differentiated Services code point (DSCP).

- [Requirements on page 4569](#)
- [Overview on page 4569](#)
- [Configuration on page 4569](#)
- [Verification on page 4571](#)

### Requirements

No special configuration beyond device initialization is required before configuring this example.

### Overview

In this example, you use a stateless firewall filter to match packets on DSCP bit patterns. If the DSCP is **2**, the packet is classified to the **best-effort** forwarding class, and the DSCP is set to **0**. If the DSCP is **3**, the packet is classified to the **best-effort** forwarding class.

### Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configure the Stateless Firewall Filter on page 4570](#)
- [Apply the Stateless Firewall Filter to a Logical Interface on page 4570](#)
- [Confirm and Commit Your Candidate Configuration on page 4571](#)

#### CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set firewall filter filter1 term 1 from dscp 2
set firewall filter filter1 term 1 then forwarding-class best-effort
set firewall filter filter1 term 1 then dscp 0
set firewall filter filter1 term 2 from dscp 3
set firewall filter filter1 term 2 then forwarding-class best-effort
set interfaces ge-0/1/0 unit 0 family inet filter input filter1
```

#### *Configure the Stateless Firewall Filter*

#### Step-by-Step Procedure

To configure the stateless firewall filter **filter1**:

1. Create the stateless firewall filter.

```
[edit]
user@host# edit firewall filter filter1
```

2. Configure the first term to match a packet with a DSCP of **2**, change the DSCP to **0**, and classify the packet to the **best-effort** forwarding class.

```
[edit firewall filter filter1]
user@host# set term 1 from dscp 2
user@host# set term 1 then forwarding-class best-effort
user@host# set term 1 then dscp 0
```

3. Configure the other term to match a packet with a DSCP of **3** and classify the packet to the **best-effort** forwarding class.

```
[edit firewall filter filter1]
user@host# set term 2 from dscp 3
user@host# set term 2 then forwarding-class best-effort
```

#### *Apply the Stateless Firewall Filter to a Logical Interface*

#### Step-by-Step Procedure

To apply the stateless firewall filter to the logical interface corresponding to the VPN routing and forwarding (VRF) instance:

1. Configure the logical interface to which you will apply the stateless firewall filter.

```
[edit]
user@host# edit interfaces ge-0/1/0 unit 0 family inet
```

2. Apply the stateless firewall filter to the logical interface.

```
[input filter1]
user@host# set filter input filter1
```



**Confirm and Commit Your Candidate Configuration****Step-by-Step Procedure**

To confirm and then commit your candidate configuration:

1. Confirm the configuration of the stateless firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show firewall
filter filter1 {
 term term1 {
 from {
 dscp 2;
 }
 then {
 forwarding-class best-effort;
 dscp 0;
 }
 }
 term term2 {
 from {
 dscp 3;
 }
 then {
 forwarding-class best-effort;
 }
 }
}
```

2. Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/1/0 {
 unit 0 {
 family inet {
 filter input filter1;
 }
 }
}
```

3. If you are done configuring the device, commit your candidate configuration.

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

**Verification**

To confirm that the configuration is working properly, enter the following operational mode commands:

- **show class-of-service**—Displays the entire class-of-service (CoS) configuration, including system-chosen defaults.

- **show class-of-service classifier type dscp**—Displays only the classifiers of the DSCP for IPv4 type.

**Related  
Documentation**

- [Understanding How to Use Standard Firewall Filters on page 3293](#)
- [Example: Configuring a Filter to Count and Sample Accepted Packets on page 4572](#)

**Example: Configuring a Filter to Count and Sample Accepted Packets**

This example shows how to configure a standard stateless firewall filter to count and sample accepted packets.

- [Requirements on page 4572](#)
- [Overview on page 4572](#)
- [Configuration on page 4572](#)
- [Verification on page 4574](#)

**Requirements**

No special configuration beyond device initialization is required before configuring this example.

Before you begin, configure traffic sampling by including the **sampling** statement at the [\[edit forwarding-options\]](#) hierarchy level.

**Overview**

In this example, you use a standard stateless firewall filter to count and sample all packets received on a logical interface.



**NOTE:** When you enable reverse path forwarding (RPF) on an interface with an input filter for firewall log and count, the input firewall filter does not log the packets rejected by RPF, although the rejected packets are counted. To log the rejected packets, use an RPF check fail filter.

**Configuration**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configure the Stateless Firewall Filter on page 4573](#)
- [Apply the Stateless Firewall Filter to a Logical Interface on page 4573](#)
- [Confirm and Commit Your Candidate Configuration on page 4574](#)

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

```
set firewall family inet filter sam term all then count count_sam
set firewall family inet filter sam term all then sample
set interfaces at-2/0/0 unit 301 family inet address 10.1.2.3/30
set interfaces at-2/0/0 unit 301 family inet filter input sam
```

#### *Configure the Stateless Firewall Filter*

**Step-by-Step Procedure** To configure the stateless firewall filter **sam**:

1. Create the stateless firewall filter **sam**.  
  
[edit]  
user@host# edit firewall family inet filter sam
2. Configure the term to count and sample all packets.  
  
[edit firewall family inet filter sam]  
user@host# set term all then count count\_sam  
user@host# set term all then sample

#### *Apply the Stateless Firewall Filter to a Logical Interface*

**Step-by-Step Procedure** To apply the stateless firewall filter to a logical interface:

1. Configure the logical interface to which you will apply the stateless firewall filter.  
  
[edit]  
user@host# edit interfaces ge-0/0/1 unit 0 family inet
2. Configure the interface address for the logical interface.  
  
[edit interfaces ge-0/0/1 unit 0 family inet]  
user@host# set address 10.1.2.3/30
3. Apply the stateless firewall filter to the logical interface.  
  
[edit interfaces ge-0/0/1 unit 0 family inet]  
user@host# set filter input sam



**NOTE:** The Junos OS does not sample packets originating from the router or switch. If you configure a filter and apply it to the output side of an interface, then only the transit packets going through that interface are sampled. Packets that are sent from the Routing Engine to the Packet Forwarding Engine are not sampled.

### *Confirm and Commit Your Candidate Configuration*

#### **Step-by-Step Procedure**

To confirm and then commit your candidate configuration:

1. Confirm the configuration of the stateless firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
 filter sam {
 term all {
 then {
 count count_sam;
 sample; # default action is accept
 }
 }
 }
}
```

2. Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
interfaces {
 at-2/0/0 {
 unit 301 {
 family inet {
 filter {
 input sam;
 }
 address 10.1.2.3/30;
 }
 }
 }
}
```

3. If you are done configuring the device, commit your candidate configuration.

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

### *Verification*

Confirm that the configuration is working properly.

- [Displaying the Packet Counter on page 4574](#)
- [Displaying the Firewall Filter Log Output on page 4575](#)
- [Displaying the Sampling Output on page 4575](#)

### *Displaying the Packet Counter*

**Purpose** Verify that the firewall filter is evaluating packets.

**Action** user@host> show firewall filter sam  
 Filter:  
 Counters:

| Name  | Bytes | Packets |
|-------|-------|---------|
| sam   |       |         |
| sam-1 | 98    | 8028    |

### Displaying the Firewall Filter Log Output

**Purpose** Display the packet header information for all packets evaluated by the firewall filter.

**Action** user@host> show firewall log

| Time     | Filter | A | Interface    | Pro | Source address | Destination address |
|----------|--------|---|--------------|-----|----------------|---------------------|
| 23:09:09 | -      | A | at-2/0/0.301 | TCP | 10.2.0.25      | 10.211.211.1:80     |
| 23:09:07 | -      | A | at-2/0/0.301 | TCP | 10.2.0.25      | 10.211.211.1:56     |
| 23:09:07 | -      | A | at-2/0/0.301 | ICM | 10.2.0.25      | 10.211.211.1:49552  |
| 23:02:27 | -      | A | at-2/0/0.301 | TCP | 10.2.0.25      | 10.211.211.1:56     |
| 23:02:25 | -      | A | at-2/0/0.301 | TCP | 10.2.0.25      | 10.211.211.1:80     |
| 23:01:22 | -      | A | at-2/0/0.301 | ICM | 10.2.2.101     | 10.211.211.1:23251  |
| 23:01:21 | -      | A | at-2/0/0.301 | ICM | 10.2.2.101     | 10.211.211.1:16557  |
| 23:01:20 | -      | A | at-2/0/0.301 | ICM | 10.2.2.101     | 10.211.211.1:29471  |
| 23:01:19 | -      | A | at-2/0/0.301 | ICM | 10.2.2.101     | 10.211.211.1:26873  |

**Meaning** This output file contains the following fields:

- **Time**—Time at which the packet was received (not shown in the default).
- **Filter**—Name of a filter that has been configured with the **filter** statement at the **[edit firewall]** hierarchy level. A hyphen (-) or the abbreviation **pfe** indicates that the packet was handled by the Packet Forwarding Engine. A space (no hyphen) indicates that the packet was handled by the Routing Engine.
- **A**—Filter action:
  - **A**—Accept (or next term)
  - **D**—Discard
  - **R**—Reject
- **Interface**—Interface on which the filter is configured.



**NOTE:** We strongly recommend that you always explicitly configure an action in the **then** statement.

- **Pro**—Packet's protocol name or number.
- **Source address**—Source IP address in the packet.
- **Destination address**—Destination IP address in the packet.

### Displaying the Sampling Output

**Purpose** Verify that the sampling output contains appropriate data.

| Action | File      | Size                   | Last changed                  |
|--------|-----------|------------------------|-------------------------------|
|        | wtmp.0.gz | Size: 15017,           | Last changed: Dec 19 13:15:54 |
|        |           | Size: 493,             | Last changed: Nov 19 13:47:29 |
|        | wtmp.2.gz | Size: 57,              | Last changed: Oct 20 15:24:34 |
|        |           | Pipe through a command |                               |

```
user@host> show log /var/tmp/sam
```

# Apr 7 15:48:50

| Time           |               | Dest<br>addr  | Src<br>addr | Dest<br>port | Src<br>port | Proto | TOS | Pkt<br>len | Intf<br>num | IP<br>frag | TCP<br>flags |
|----------------|---------------|---------------|-------------|--------------|-------------|-------|-----|------------|-------------|------------|--------------|
| Apr 7 15:48:54 | 192.168.9.194 | 192.168.9.195 | 0           | 0            | 1           | 0x0   | 84  | 8          | 0x0         | 0x0        |              |
| Apr 7 15:48:55 | 192.168.9.194 | 192.168.9.195 | 0           | 0            | 1           | 0x0   | 84  | 8          | 0x0         | 0x0        |              |
| Apr 7 15:48:56 | 192.168.9.194 | 192.168.9.195 | 0           | 0            | 1           | 0x0   | 84  | 8          | 0x0         | 0x0        |              |

## Related Documentation

- Understanding How to Use Standard Firewall Filters on page 3293
- *Example: Configuring a Filter to Set the DSCP Bit to Zero*

## Standard Firewall Filters for Trusted Sources

- [Example: Configuring a Stateless Firewall Filter to Accept Traffic from Trusted Sources on page 4576](#)
- [Example: Configuring a Filter to Block Telnet and SSH Access on page 4581](#)
- [Example: Configuring a Filter to Block TFTP Access on page 4586](#)
- [Example: Configuring a Filter to Accept OSPF Packets from a Prefix on page 4589](#)
- [Example: Configuring a Filter to Accept DHCP Packets Based on Address on page 4592](#)
- [Example: Configuring a Filter to Block TCP Access to a Port Except from Specified BGP Peers on page 4595](#)

### Example: Configuring a Stateless Firewall Filter to Accept Traffic from Trusted Sources

This example shows how to create a stateless firewall filter that protects the Routing Engine from traffic originating from untrusted sources.

- Requirements on page 4576
- Overview on page 4576
- Configuration on page 4577
- Verification on page 4579

## Requirements

No special configuration beyond device initialization is required before configuring stateless firewall filters.

## Overview

In this example, you create a stateless firewall filter called protect-RE that discards all traffic destined for the Routing Engine except SSH and BGP protocol packets from specified trusted sources. This example includes the following firewall filter terms:

- **ssh-term**—Accepts TCP packets with a source address of 192.168.122.0/24 and a destination port that specifies SSH.

- **bgp-term**—Accepts TCP packets with a source address of 10.2.1.0/24 and a destination port that specifies BGP.
- **discard-rest-term**—For all packets that are not accepted by **ssh-term** or **bgp-term**, creates a firewall filter log and system logging records, then discards all packets.



**NOTE:** You can move terms within the firewall filter using the `insert` command. See *insert* in the *CLI User Guide*.

### Configuration

#### CLI Quick Configuration

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

```
set firewall family inet filter protect-RE term ssh-term from source-address
 192.168.122.0/24
set firewall family inet filter protect-RE term ssh-term from protocol tcp
set firewall family inet filter protect-RE term ssh-term from destination-port ssh
set firewall family inet filter protect-RE term ssh-term then accept
set firewall family inet filter protect-RE term bgp-term from source-address 10.2.1.0/24
set firewall family inet filter protect-RE term bgp-term from protocol tcp
set firewall family inet filter protect-RE term bgp-term from destination-port bgp
set firewall family inet filter protect-RE term bgp-term then accept
set firewall family inet filter protect-RE term discard-rest-term then log
set firewall family inet filter protect-RE term discard-rest-term then syslog
set firewall family inet filter protect-RE term discard-rest-term then discard
set interfaces lo0 unit 0 family inet filter input protect-RE
```

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure the stateless firewall filter:

1. Create the stateless firewall filter.

```
[edit]
user@host# edit firewall family inet filter protect-RE
```

2. Create the first filter term.

```
[edit firewall family inet filter protect-RE]
user@host# edit term ssh-term
```

3. Define the protocol, destination port, and source address match conditions for the term.

```
[edit firewall family inet filter protect-RE term ssh-term]
user@host# set from protocol tcp destination-port ssh source-address
 192.168.122.0/24
```

4. Define the actions for the term.

```
[edit firewall family inet filter protect-RE term ssh-term]
user@host# set then accept
```

5. Create the second filter term.

```
[edit firewall family inet filter protect-RE]
user@host# edit term bgp-term
```

6. Define the protocol, destination port, and source address match conditions for the term.

```
[edit firewall family inet filter protect-RE term bgp-term]
user@host# set from protocol tcp destination-port bgp source-address 10.2.1.0/24
```

7. Define the action for the term.

```
[edit firewall family inet filter protect-RE term bgp-term]
user@host# set then accept
```

8. Create the third filter term.

```
[edit firewall family inet filter protect-RE]
user@host# edit term discard-rest-term
```

9. Define the action for the term.

```
[edit firewall family inet filter protect-RE term discard-rest]
user@host# set then log syslog discard
```

10. Apply the filter to the input side of the Routing Engine interface.

```
[edit]
user@host# set interfaces lo0 unit 0 family inet filter input protect-RE
```

**Results** Confirm your configuration by entering the **show firewall** command and the **show interfaces lo0** command from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show firewall
family inet {
 filter protect-RE {
 term ssh-term {
 from {
 source-address {
 192.168.122.0/24;
 }
 protocol tcp;
 destination-port ssh;
 }
 then accept;
 }
 term bgp-term {
 from {
 source-address {
 10.2.1.0/24;
 }
 }
 }
 }
}
```



```

 protocol tcp;
 destination-port bgp;
 }
 then accept;
}
term discard-rest-term {
 then {
 log;
 syslog;
 discard;
 }
}
}
}

user@host# show interfaces lo0
unit 0 {
 family inet {
 filter {
 input protect-RE;
 }
 address 127.0.0.1/32;
 }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

```

[edit]
user@host# commit

```

### **Verification**

To confirm that the configuration is working properly, perform these tasks:

- [Displaying Stateless Firewall Filter Configurations on page 4579](#)
- [Verifying a Services, Protocols, and Trusted Sources Firewall Filter on page 4579](#)
- [Displaying Stateless Firewall Filter Logs on page 4580](#)

### **Displaying Stateless Firewall Filter Configurations**

|                |                                                                                                                                                                                                                                                                            |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b> | Verify the configuration of the firewall filter.                                                                                                                                                                                                                           |
| <b>Action</b>  | From configuration mode, enter the <b>show firewall</b> command and the <b>show interfaces lo0</b> command.                                                                                                                                                                |
| <b>Meaning</b> | Verify that the output shows the intended configuration of the firewall filter. In addition, verify that the terms are listed in the order in which you want the packets to be tested. You can move terms within a firewall filter by using the <b>insert</b> CLI command. |

### **Verifying a Services, Protocols, and Trusted Sources Firewall Filter**

|                |                                                                 |
|----------------|-----------------------------------------------------------------|
| <b>Purpose</b> | Verify that the actions of the firewall filter terms are taken. |
|----------------|-----------------------------------------------------------------|

**Action** Send packets to the device that match the terms. In addition, verify that the filter actions are *not* taken for packets that do not match.

- Use the **ssh *host-name*** command from a host at an IP address that matches **192.168.122.0/24** to verify that you can log in to the device using only SSH from a host with this address prefix.
- Use the **show route summary** command to verify that the routing table on the device does not contain any entries with a protocol other than **Direct**, **Local**, **BGP**, or **Static**.

## Sample Output

```
% ssh 192.168.249.71
%ssh host
user@host's password:
--- JUNOS 6.4-20040518.0 (JSERIES) #0: 2004-05-18 09:27:50 UTC

user@host>

user@host> show route summary
Router ID: 192.168.249.71

inet.0: 34 destinations, 34 routes (33 active, 0 holddown, 1 hidden)
 Direct: 10 routes, 9 active
 Local: 9 routes, 9 active
 BGP: 10 routes, 10 active
 Static: 5 routes, 5 active
...
```

**Meaning** Verify the following information:

- You can successfully log in to the device using SSH.
- The **show route summary** command does not display a protocol other than **Direct**, **Local**, **BGP**, or **Static**.

### *Displaying Stateless Firewall Filter Logs*

**Purpose** Verify that packets are being logged. If you included the **log** or **syslog** action in a term, verify that packets matching the term are recorded in the firewall log or your system logging facility.

**Action** From operational mode, enter the **show firewall log** command.

## Sample Output

```
user@host> show firewall log
Log :
Time Filter Action Interface Protocol Src Addr Dest Addr
15:11:02 pfe D ge-0/0/0.0 TCP 172.17.28.19 192.168.70.71
15:11:01 pfe D ge-0/0/0.0 TCP 172.17.28.19 192.168.70.71
15:11:01 pfe D ge-0/0/0.0 TCP 172.17.28.19 192.168.70.71
15:11:01 pfe D ge-0/0/0.0 TCP 172.17.28.19 192.168.70.71
...
```

**Meaning** Each record of the output contains information about the logged packet. Verify the following information:

- Under **Time**, the time of day the packet was filtered is shown.
- The **Filter** output is always **pfe**.
- Under **Action**, the configured action of the term matches the action taken on the packet—**A** (accept), **D** (discard), **R** (reject).
- Under **Interface**, the inbound (ingress) interface on which the packet arrived is appropriate for the filter.
- Under **Protocol**, the protocol in the IP header of the packet is appropriate for the filter.
- Under **Src Addr**, the source address in the IP header of the packet is appropriate for the filter.
- Under **Dest Addr**, the destination address in the IP header of the packet is appropriate for the filter.

**Related Documentation**

- *Junos OS Feature Support Reference for SRX Series and J Series Devices*
- *show route summary* in the *Junos OS Operational Mode Commands*
- [show firewall on page 4450](#) in the *Junos OS Operational Mode Commands*
- [show firewall log on page 4457](#) in the *Junos OS Operational Mode Commands*
- *show interfaces (Loopback)* in the *Junos OS Operational Mode Commands*

**Example: Configuring a Filter to Block Telnet and SSH Access**

- [Requirements on page 4581](#)
- [Overview on page 4581](#)
- [Configuration on page 4582](#)
- [Verification on page 4584](#)

**Requirements**

You must have access to a remote host that has network connectivity with this router or switch.

**Overview**

In this example, you create an IPv4 stateless firewall filter that logs and rejects Telnet or SSH access packets unless the packet is destined for or originates from the 192.168.1.0/24 subnet.

- To match packets destined for or originating from the **address 192.168.1.0/24** subnet, you use the **address 192.168.1.0/24** IPv4 match condition.
- To match packets destined for or originating from a TCP port, Telnet port, or SSH port, you use the **protocol tcp**, **port telnet**, and **telnet ssh** IPv4 match conditions.

### Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configure the Stateless Firewall Filter on page 4582](#)
- [Apply the Firewall Filter to the Loopback Interface on page 4583](#)
- [Confirm and Commit Your Candidate Configuration on page 4583](#)

#### CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set firewall family inet filter local_acl term terminal_access from address 192.168.1.0/24
set firewall family inet filter local_acl term terminal_access from protocol tcp
set firewall family inet filter local_acl term terminal_access from port ssh
set firewall family inet filter local_acl term terminal_access from port telnet
set firewall family inet filter local_acl term terminal_access then accept
set firewall family inet filter local_acl term terminal_access_denied from protocol tcp
set firewall family inet filter local_acl term terminal_access_denied from port ssh
set firewall family inet filter local_acl term terminal_access_denied from port telnet
set firewall family inet filter local_acl term terminal_access_denied then log
set firewall family inet filter local_acl term terminal_access_denied then reject
set firewall family inet filter local_acl term default-term then accept
set interfaces lo0 unit 0 family inet filter input local_acl
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
```

### Configure the Stateless Firewall Filter

#### Step-by-Step Procedure

To configure the stateless firewall filter that selectively blocks Telnet and SSH access:

1. Create the stateless firewall filter **local\_acl**.

```
[edit]
user@myhost# edit firewall family inet filter local_acl
```

2. Define the filter term **terminal\_access**.

```
[edit firewall family inet filter local_acl]
user@myhost# set term terminal_access from address 192.168.1.0/24
user@myhost# set term terminal_access from protocol tcp
user@myhost# set term terminal_access from port ssh
user@myhost# set term terminal_access from port telnet
user@myhost# set term terminal_access then accept
```

3. Define the filter term **terminal\_access\_denied**.

```
[edit firewall family inet filter local_acl]
user@myhost# set term terminal_access_denied from protocol tcp
user@myhost# set term terminal_access_denied from port ssh
user@myhost# set term terminal_access_denied from port telnet
user@myhost# set term terminal_access_denied then log
user@myhost# set term terminal_access_denied then reject
```

```
user@myhost# set term default-term then accept
```

### *Apply the Firewall Filter to the Loopback Interface*

#### **Step-by-Step Procedure**

- To apply the firewall filter to the loopback interface:

```
[edit]
user@myhost# set interfaces lo0 unit 0 family inet filter input local_acl
user@myhost# set interfaces lo0 unit 0 family inet address 127.0.0.1/32
```

### *Confirm and Commit Your Candidate Configuration*

#### **Step-by-Step Procedure**

To confirm and then commit your candidate configuration:

1. Confirm the configuration of the stateless firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@myhost# show firewall
family inet {
 filter local_acl {
 term terminal_access {
 from {
 address {
 192.168.1.0/24;
 }
 protocol tcp;
 port [ssh telnet];
 }
 then accept;
 }
 term terminal_access_denied {
 from {
 protocol tcp;
 port [ssh telnet];
 }
 then {
 log;
 reject;
 }
 }
 term default-term {
 then accept;
 }
 }
}
```

2. Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@myhost# show interfaces
lo0 {
 unit 0 {
```

```
family inet {
 filter {
 input local_acl;
 }
 address 127.0.0.1/32;
}
}
```

3. If you are done configuring the device, commit your candidate configuration.

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

### ***Verification***

Confirm that the configuration is working properly.

- [Verifying Accepted Packets on page 4584](#)
- [Verifying Logged and Rejected Packets on page 4585](#)

### ***Verifying Accepted Packets***

**Purpose** Verify that the actions of the firewall filter terms are taken.

- Action** 1. Clear the firewall log on your router or switch.

```
user@myhost> clear firewall log
```

2. From a host at an IP address *within* the 192.168.1.0/24 subnet, use the **ssh *hostname*** command to verify that you can log in to the device using only SSH. This packet should be accepted, and the packet header information for this packet should not be logged in the firewall filter log buffer in the Packet Forwarding Engine.

```
user@host-A> ssh myhost
user@myhosts's password:
--- JUNOS 11.1-20101102.0 built 2010-11-02 04:48:46 UTC

% cli
user@myhost>
```

3. From a host at an IP address *within* the 192.168.1.0/24 subnet, use the **telnet *hostname*** command to verify that you can log in to your router or switch using only Telnet. This packet should be accepted, and the packet header information for this packet should not be logged in the firewall filter log buffer in the Packet Forwarding Engine.

```
user@host-A> telnet myhost
Trying 192.168.249.71...
Connected to myhost-fxp0.acme.net.
Escape character is '^J'.

host (ttyp0)

login: user
Password:

--- JUNOS 11.1-20101102.0 built 2010-11-02 04:48:46 UTC

% cli
user@myhost>
```

4. Use the **show firewall log** command to verify that the routing table on the device does not contain any entries with a source address in the 192.168.1.0/24 subnet.

```
user@myhost> show firewall log
```

#### ***Verifying Logged and Rejected Packets***

- Purpose** Verify that the actions of the firewall filter terms are taken.

- Action** 1. Clear the firewall log on your router or switch.

```
user@myhost> clear firewall log
```

2. From a host at an IP address *outside of* the 192.168.1.0/24 subnet, use the **ssh hostname** command to verify that you cannot log in to the device using only SSH. This packet should be rejected, and the packet header information for this packet should be logged in the firewall filter log buffer in the Packet Forwarding Engine.

```
user@host-B ssh myhost
ssh: connect to host sugar port 22: Connection refused
--- JUNOS 11.1-20101102.0 built 2010-11-02 04:48:46 UTC
%
```

3. From a host at an IP address *outside of* the 192.168.1.0/24 subnet, use the **telnet hostname** command to verify that you can log in to the device using only Telnet. This packet should be rejected, and the packet header information for this packet should be logged in the firewall filter log buffer in the PFE.

```
user@host-B> telnet myhost
Trying 192.168.249.71...
telnet: connect to address 192.168.187.3: Connection refused
telnet: Unable to connect to remote host
%
```

4. Use the **show firewall log** command to verify that the routing table on the device does not contain any entries with a source address in the 192.168.1.0/24 subnet.

```
user@myhost> show firewall log
```

| Time     | Filter    | Action | Interface | Protocol | Src Addr      | Dest Addr     |
|----------|-----------|--------|-----------|----------|---------------|---------------|
| 18:41:25 | local_acl | R      | fxp0.0    | TCP      | 192.168.187.5 | 192.168.187.1 |
| 18:41:25 | local_acl | R      | fxp0.0    | TCP      | 192.168.187.5 | 192.168.187.1 |
| 18:41:25 | local_acl | R      | fxp0.0    | TCP      | 192.168.187.5 | 192.168.187.1 |
| ...      |           |        |           |          |               |               |
| 18:43:06 | local_acl | R      | fxp0.0    | TCP      | 192.168.187.5 | 192.168.187.1 |
| 18:43:06 | local_acl | R      | fxp0.0    | TCP      | 192.168.187.5 | 192.168.187.1 |
| 18:43:06 | local_acl | R      | fxp0.0    | TCP      | 192.168.187.5 | 192.168.187.1 |
| ...      |           |        |           |          |               |               |

#### Related Documentation

- [Logging of Packet Headers Evaluated by a Firewall Filter Term on page 4520](#)
- [Understanding How to Use Standard Firewall Filters on page 3293](#)
- [Example: Configuring a Stateless Firewall Filter to Accept Traffic from Trusted Sources on page 4576](#)
- [Example: Configuring a Filter to Block TFTP Access on page 4586](#)
- [Example: Configuring a Filter to Accept OSPF Packets from a Prefix on page 4589](#)
- [Example: Configuring a Filter to Accept DHCP Packets Based on Address on page 4592](#)

#### Example: Configuring a Filter to Block TFTP Access

- [Requirements on page 4587](#)
- [Overview on page 4587](#)



- [Configuration on page 4587](#)
- [Verification on page 4589](#)

### Requirements

No special configuration beyond device initialization is required before configuring this example.

### Overview

By default, to decrease vulnerability to denial-of-service (DoS) attacks, the Junos OS filters and discards Dynamic Host Configuration Protocol (DHCP) or Bootstrap Protocol (BOOTP) packets that have a source address of 0.0.0.0 and a destination address of 255.255.255.255. This default filter is known as a unicast RPF check. However, some vendors' equipment automatically accepts these packets.

To interoperate with other vendors' equipment, you can configure a filter that checks for both of these addresses and overrides the default RPF-check filter by accepting these packets. In this example, you block Trivial File Transfer Protocol (TFTP) access, logging any attempts to establish TFTP connections.

### Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configure the Stateless Firewall Filter on page 4587](#)
- [Apply the Firewall Filter to the Loopback Interface on page 4588](#)
- [Confirm and Commit Your Candidate Configuration on page 4588](#)

#### CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set firewall family inet filter tftp_access_control term one from protocol udp
set firewall family inet filter tftp_access_control term one from port tftp
set firewall family inet filter tftp_access_control term one then log
set firewall family inet filter tftp_access_control term one then discard
set interfaces lo0 unit 0 family inet filter input tftp_access_control
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
```

### Configure the Stateless Firewall Filter

#### Step-by-Step Procedure

To configure the stateless firewall filter that selectively blocks TFTP access:

1. Create the stateless firewall filter **tftp\_access\_control**.

```
[edit]
user@host# edit firewall family inet filter tftp_access_control
```

2. Specify a match on packets received on UDP port 69.

```
[edit firewall family inet filter tftp_access_control]
```

```
user@host# set term one from protocol udp
user@host# set term one from port tftp
```

3. Specify that matched packets be logged to the buffer on the Packet Forwarding Engine and then discarded.

```
[edit firewall family inet filter tftp_access_control]
user@host# set term one then log
user@host# set term one then discard
```

### *Apply the Firewall Filter to the Loopback Interface*

**Step-by-Step Procedure** To apply the firewall filter to the loopback interface:

- [edit]  
user@host# set interfaces lo0 unit 0 family inet filter input tftp\_access\_control  
user@host# set interfaces lo0 unit 0 family inet address 127.0.0.1/32

### *Confirm and Commit Your Candidate Configuration*

**Step-by-Step Procedure** To confirm and then commit your candidate configuration:

1. Confirm the configuration of the stateless firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
 filter tftp_access_control {
 term one {
 from {
 protocol udp;
 port tftp;
 }
 then {
 log;
 discard;
 }
 }
 }
}
```

2. Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
lo0 {
 unit 0 {
 family inet {
 filter {
 input tftp_access_control;
 }
 address 127.0.0.1/32;
 }
 }
}
```

```
}
}
```

3. If you are done configuring the device, commit your candidate configuration.

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

### **Verification**

Confirm that the configuration is operating properly:

- [Verifying Logged and Discarded Packets on page 4589](#)

### **Verifying Logged and Discarded Packets**

**Purpose** Verify that the actions of the firewall filter terms are taken.

**Action** To

1. Clear the firewall log on your router or switch.  

```
user@myhost> clear firewall log
```
2. From another host, send a packet to UDP port 69 on this router or switch.

### **Related Documentation**

- [Understanding How to Use Standard Firewall Filters on page 3293](#)
- [Example: Configuring a Stateless Firewall Filter to Accept Traffic from Trusted Sources on page 4576](#)
- [Example: Configuring a Filter to Block Telnet and SSH Access on page 4581](#)
- [Example: Configuring a Filter to Accept OSPF Packets from a Prefix on page 4589](#)
- [Example: Configuring a Filter to Accept DHCP Packets Based on Address on page 4592](#)

### **Example: Configuring a Filter to Accept OSPF Packets from a Prefix**

This example shows how to configure a standard stateless firewall filter to accept packets from a trusted source.

- [Requirements on page 4589](#)
- [Overview on page 4590](#)
- [Configuration on page 4590](#)
- [Verification on page 4592](#)

### **Requirements**

No special configuration beyond device initialization is required before configuring this example.

### Overview

In this example, you create a filter that accepts only OSPF packets from an address in the prefix 10.108.0.0/16, discarding all other packets with an **administratively-prohibited** ICMP message.

### Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configure the Stateless Firewall Filter on page 4590](#)
- [Apply the Firewall Filter to the Loopback Interface on page 4591](#)
- [Confirm and Commit Your Candidate Configuration on page 4591](#)

#### CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set firewall family inet filter ospf_filter term term1 from source-address 10.108.0.0/16
set firewall family inet filter ospf_filter term term1 from protocol ospf
set firewall family inet filter ospf_filter term term1 then accept
set firewall family inet filter ospf_filter term default-term then reject
 administratively-prohibited
set interfaces ge-0/0/1 unit 0 family inet address 10.1.2.3/30
set interfaces ge-0/0/1 unit 0 family inet filter input ospf_filter
```

### Configure the Stateless Firewall Filter

#### Step-by-Step Procedure

To configure the stateless firewall filter **ospf\_filter**:

1. Create the filter.

```
[edit]
user@host# edit firewall family inet filter ospf_filter
```

2. Configure the term that accepts packets.

```
[edit firewall family inet filter ospf_filter]
user@host# set term term1 from source-address 10.108.0.0/16
user@host# set term term1 from protocol ospf
user@host# set term term1 then accept
```

3. Configure the term that rejects all other packets.

```
[edit firewall family inet filter ospf_filter]
user@host# set term default_term then reject administratively-prohibited
```

*Apply the Firewall Filter to the Loopback Interface***Step-by-Step Procedure**

To apply the firewall filter to the loopback interface:

1. Configure the interface.

```
[edit]
user@host# edit interfaces ge-0/0/1 unit 0 family inet
```

2. Configure the logical interface IP address.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
user@host# set address 10.1.2.3/30
```

3. Apply the filter to the input.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
user@host# set filter input ospf_filter
```

*Confirm and Commit Your Candidate Configuration***Step-by-Step Procedure**

To confirm and then commit your candidate configuration:

1. Confirm the configuration of the stateless firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
 filter ospf_filter {
 term term1 {
 from {
 source-address {
 10.108.0.0/16;
 }
 protocol ospf;
 }
 then {
 accept;
 }
 }
 term default_term {
 then {
 reject administratively-prohibited; # default reject action
 }
 }
 }
}
```

2. Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
lo0 {
```

```
unit 0 {
 family inet {
 filter {
 input ospf_filter;
 }
 address 10.1.2.3/30;
 }
}
```

3. If you are done configuring the device, commit your candidate configuration.

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

### **Verification**

To confirm that the configuration is working properly, enter the [show firewall filter ospf\\_filter](#) operational mode command.

### **Related Documentation**

- [Understanding How to Use Standard Firewall Filters on page 3293](#)
- [Example: Configuring a Stateless Firewall Filter to Accept Traffic from Trusted Sources on page 4576](#)
- [Example: Configuring a Filter to Block Telnet and SSH Access on page 4581](#)
- [Example: Configuring a Filter to Block TFTP Access on page 4586](#)
- [Example: Configuring a Filter to Accept DHCP Packets Based on Address on page 4592](#)

### **Example: Configuring a Filter to Accept DHCP Packets Based on Address**

This example shows how to configure a standard stateless firewall filter to accept packets from a trusted source.

- [Requirements on page 4592](#)
- [Overview on page 4592](#)
- [Configuration on page 4593](#)
- [Verification on page 4594](#)

### **Requirements**

This example is supported only on MX Series routers and EX Series switches.

### **Overview**

In this example, you create a filter (**rpfdhcp**) that accepts DHCP packets with a source address of **0.0.0.0** and a destination address of **255.255.255.255**.

**Configuration**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

- [Configure the Stateless Firewall Filter on page 4593](#)
- [Apply the Firewall Filter to the Loopback Interface on page 4593](#)
- [Confirm and Commit Your Candidate Configuration on page 4594](#)

**CLI Quick Configuration**

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set firewall family inet filter rpf_dhcp term dhcp_term from source-address 0.0.0.0/32
set firewall family inet filter rpf_dhcp term dhcp_term from destination-address
 255.255.255.255/32
set firewall family inet filter rpf_dhcp term dhcp_term then accept
set interfaces ge-0/0/1 unit 0 family inet address 10.1.2.3/30
set interfaces ge-0/0/1 unit 0 family inet filter input sam
```

**Configure the Stateless Firewall Filter****Step-by-Step Procedure**

To configure the stateless firewall filter:

1. Create the stateless firewall filter **rpf\_dhcp**.  
  
[edit]  
user@host# edit firewall family inet filter rpf\_dhcp
2. Configure the term to match packets with a source address of **0.0.0.0** and a destination address of **255.255.255.255**.  
  
[edit firewall family inet filter rpf\_dhcp]  
user@host# set term dhcp\_term from source-address 0.0.0.0/32  
user@host# set term dhcp\_term from destination-address 255.255.255.255/32
3. Configure the term to accept packets that match the specified conditions.  
  
[edit firewall family inet filter rpf\_dhcp]  
set term dhcp\_term then accept

**Apply the Firewall Filter to the Loopback Interface****Step-by-Step Procedure**

To apply the filter to the input at the loopback interface:

1. Apply the **rpf\_dhcp** filter if packets are not arriving on an expected path.  
  
[edit]  
user@host# set interfaces lo0 unit 0 family inet **rpf-check** fail-filter rpf\_dhcp
2. Configure an address for the loopback interface.  
  
[edit]  
user@host# set interfaces lo0 unit 0 family inet address 127.0.0.1/32

### *Confirm and Commit Your Candidate Configuration*

#### **Step-by-Step Procedure**

To confirm and then commit your candidate configuration:

1. Confirm the configuration of the stateless firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
 filter rpf_dhcp {
 term dhcp_term {
 from {
 source-address {
 0.0.0.0/32;
 }
 destination-address {
 255.255.255.255/32;
 }
 }
 then accept;
 }
 }
}
```

2. Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
lo0 {
 unit 0 {
 family inet {
 filter {
 rpf-check {
 fail-filter rpf_dhcp;
 mode loose;
 }
 }
 address 127.0.0.1/32;
 }
 }
}
```

3. When you are done configuring the device, commit your candidate configuration.

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

### *Verification*

To confirm that the configuration is working properly, enter the **show firewall** operational mode command.



**Related Documentation**

- [Understanding How to Use Standard Firewall Filters on page 3293](#)
- [Example: Configuring a Stateless Firewall Filter to Accept Traffic from Trusted Sources on page 4576](#)
- [Example: Configuring a Filter to Block Telnet and SSH Access on page 4581](#)
- [Example: Configuring a Filter to Block TFTP Access on page 4586](#)
- [Example: Configuring a Filter to Accept OSPF Packets from a Prefix on page 4589](#)

**Example: Configuring a Filter to Block TCP Access to a Port Except from Specified BGP Peers**

This example shows how to configure a standard stateless firewall filter that blocks all TCP connection attempts to port 179 from all requesters except from specified BGP peers.

- [Requirements on page 4595](#)
- [Overview on page 4595](#)
- [Configuration on page 4596](#)
- [Verification on page 4599](#)

**Requirements**

No special configuration beyond device initialization is required before you configure this example.


**Overview**

In this example, you create a stateless firewall filter that blocks all TCP connection attempts to port 179 from all requesters except the specified BGP peers.

The stateless firewall filter **filter\_bgp179** matches all packets from the directly connected interfaces on Device A and Device B to the destination port number 179.

[Figure 94 on page 4595](#) shows the topology used in this example. Device C attempts to make a TCP connection to Device E. Device E blocks the connection attempt. This example shows the configuration on Device E.

**Figure 94: Typical Network with BGP Peer Sessions**



### Configuration

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

**Device C**

```
set interfaces ge-1/2/0 unit 10 description to-E
set interfaces ge-1/2/0 unit 10 family inet address 10.10.10.10/30
set protocols bgp group external-peers type external
set protocols bgp group external-peers peer-as 17
set protocols bgp group external-peers neighbor 10.10.10.9
set routing-options autonomous-system 22
```

**Device E**

```
set interfaces ge-1/2/0 unit 0 description to-A
set interfaces ge-1/2/0 unit 0 family inet address 10.10.10.1/30
set interfaces ge-1/2/1 unit 5 description to-B
set interfaces ge-1/2/1 unit 5 family inet address 10.10.10.5/30
set interfaces ge-1/0/0 unit 9 description to-C
set interfaces ge-1/0/0 unit 9 family inet address 10.10.10.9/30
set interfaces lo0 unit 2 family inet filter input filter_bgp179
set interfaces lo0 unit 2 family inet address 192.168.0.1/32
set protocols bgp group external-peers type external
set protocols bgp group external-peers peer-as 22
set protocols bgp group external-peers neighbor 10.10.10.2
set protocols bgp group external-peers neighbor 10.10.10.6
set protocols bgp group external-peers neighbor 10.10.10.10
set routing-options autonomous-system 17
set firewall family inet filter filter_bgp179 term 1 from source-address 10.10.10.2/32
set firewall family inet filter filter_bgp179 term 1 from source-address 10.10.10.6/32
set firewall family inet filter filter_bgp179 term 1 from destination-port bgp
set firewall family inet filter filter_bgp179 term 1 then accept
set firewall family inet filter filter_bgp179 term 2 then reject
```

### Configuring Device E

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure Device E with a stateless firewall filter that blocks all TCP connection attempts to port 179 from all requestors except specified BGP peers:

1. Configure the interfaces.

```
user@E# set interfaces ge-1/2/0 unit 0 description to-A
user@E# set interfaces ge-1/2/0 unit 0 family inet address 10.10.10.1/30
```

```
user@E# set interfaces ge-1/2/1 unit 5 description to-B
user@E# set interfaces ge-1/2/1 unit 5 family inet address 10.10.10.5/30
```

```
user@E# set interfaces ge-1/0/0 unit 9 description to-C
user@E# set interfaces ge-1/0/0 unit 9 family inet address 10.10.10.9/30
```

2. Configure BGP.

```
[edit protocols bgp group external-peers]
user@E# set type external
user@E# set peer-as 22
user@E# set neighbor 10.10.10.2
user@E# set neighbor 10.10.10.6
user@E# set neighbor 10.10.10.10
```

3. Configure the autonomous system number.

```
[edit routing-options]
user@E# set autonomous-system 17
```

4. Define the filter term that accepts TCP connection attempts to port 179 from the specified BGP peers.

```
[edit firewall family inet filter filter_bgp179]
user@E# set term 1 from source-address 10.10.10.2/32
user@E# set term 1 from source-address 10.10.10.6/32
user@E# set term 1 from destination-port bgp
user@E# set term 1 then accept
```

5. Define the other filter term to reject packets from other sources.

```
[edit firewall family inet filter filter_bgp179]
user@E# set term 2 then reject
```

6. Apply the firewall filter to the loopback interface.

```
[edit interfaces lo0 unit 2 family inet]
user@E# set filter input filter_bgp179
user@E# set address 192.168.0.1/32
```

**Results** From configuration mode, confirm your configuration by entering the **show firewall**, **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@E# show firewall
family inet {
 filter filter_bgp179 {
 term 1 {
 from {
 source-address {
 10.10.10.2/32;
 10.10.10.6/32;
 }
 destination-port bgp;
 }
 then accept;
 }
 term 2 {
 then {
 reject;
 }
 }
 }
}
```

```
user@E# show interfaces
lo0 {
 unit 2 {
 family inet {
 filter {
 input filter_bgp179;
 }
 address 192.168.0.1/32;
 }
 }
}
ge-1/2/0 {
 unit 0 {
 description to-A;
 family inet {
 address 10.10.10.1/30;
 }
 }
}
ge-1/2/1 {
 unit 5 {
 description to-B;
 family inet {
 address 10.10.10.5/30;
 }
 }
}
ge-1/0/0 {
 unit 9 {
 description to-C;
 family inet {
 address 10.10.10.9/30;
 }
 }
}
```

```
user@E# show protocols
bgp {
 group external-peers {
 type external;
 peer-as 22;
 neighbor 10.10.10.2;
 neighbor 10.10.10.6;
 neighbor 10.10.10.10;
 }
}
```

```
user@E# show routing-options
autonomous-system 17;
```

If you are done configuring the device, enter **commit** from configuration mode.

**Verification**

Confirm that the configuration is working properly.

- [Verifying That the Filter Is Configured on page 4599](#)
- [Verifying the TCP Connections on page 4599](#)
- [Monitoring Traffic on the Interfaces on page 4599](#)

**Verifying That the Filter Is Configured**

**Purpose** Make sure that the filter is listed in output of the `show firewall filter` command.

**Action** `user@E> show firewall filter filter_bgp179`  
 Filter: filter\_bgp179

**Verifying the TCP Connections**

**Purpose** Verify the TCP connections.

**Action** From operational mode, run the `show system connections extensive` command on Device C and Device E.

The output on Device C shows the attempt to establish a TCP connection. The output on Device E shows that connections are established with Device A and Device B only.

`user@C> show system connections extensive | match 10.10.10`

|      |   |   |                  |                 |          |
|------|---|---|------------------|-----------------|----------|
| tcp4 | 0 | 0 | 10.10.10.9.51872 | 10.10.10.10.179 | SYN_SENT |
|------|---|---|------------------|-----------------|----------|

`user@E> show system connections extensive | match 10.10.10`

|      |   |   |                  |                  |             |
|------|---|---|------------------|------------------|-------------|
| tcp4 | 0 | 0 | 10.10.10.5.179   | 10.10.10.6.62096 | ESTABLISHED |
| tcp4 | 0 | 0 | 10.10.10.6.62096 | 10.10.10.5.179   | ESTABLISHED |
| tcp4 | 0 | 0 | 10.10.10.1.179   | 10.10.10.2.61506 | ESTABLISHED |
| tcp4 | 0 | 0 | 10.10.10.2.61506 | 10.10.10.1.179   | ESTABLISHED |

**Monitoring Traffic on the Interfaces**

**Purpose** Use the `monitor traffic` command to compare the traffic on an interface that establishes a TCP connection with the traffic on an interface that does not establish a TCP connection.

**Action** From operational mode, run the `monitor traffic` command on the Device E interface to Device B and on the Device E interface to Device C. The following sample output verifies that in the first example, acknowledgment (**ack**) messages are received. In the second example, **ack** messages are not received.

```
user@E> monitor traffic size 1500 interface ge-1/2/1.5
19:02:49.700912 Out IP 10.10.10.5.bgp > 10.10.10.6.62096: P
3330573561:3330573580(19) ack 915601686 win 16384 <nop,nop,timestamp 1869518816
1869504850>: BGP, length: 19
19:02:49.801244 In IP 10.10.10.6.62096 > 10.10.10.5.bgp: . ack 19 win 16384
<nop,nop,timestamp 1869518916 1869518816>
19:03:03.323018 In IP 10.10.10.6.62096 > 10.10.10.5.bgp: P 1:20(19) ack 19 win
16384 <nop,nop,timestamp 1869532439 1869518816>: BGP, length: 19
19:03:03.422418 Out IP 10.10.10.5.bgp > 10.10.10.6.62096: . ack 20 win 16384
```

```
<nop,nop,timestamp 1869532539 1869532439>
19:03:17.220162 Out IP 10.10.10.5.bgp > 10.10.10.6.62096: P 19:38(19) ack 20 win
16384 <nop,nop,timestamp 1869546338 1869532439>: BGP, length: 19
19:03:17.320501 In IP 10.10.10.6.62096 > 10.10.10.5.bgp: . ack 38 win 16384
<nop,nop,timestamp 1869546438 1869546338>
```

```
user@E> monitor traffic size 1500 interface ge-1/0/0.9
```

```
18:54:20.175471 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,nop,wscale 0,nop,nop,timestamp 1869009240 0,sackOK,eol>
18:54:23.174422 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,nop,wscale 0,nop,nop,timestamp 1869012240 0,sackOK,eol>
18:54:26.374118 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,nop,wscale 0,nop,nop,timestamp 1869015440 0,sackOK,eol>
18:54:29.573799 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,sackOK,eol>
18:54:32.773493 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,sackOK,eol>
18:54:35.973185 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,sackOK,eol>
```

#### Related Documentation

- [Understanding How to Use Standard Firewall Filters on page 3293](#)
- [Example: Configuring a Stateless Firewall Filter to Protect Against TCP and ICMP Floods on page 4600](#)
- [Example: Configuring a Filter to Accept Packets Based on IPv6 TCP Flags on page 4610](#)

---

#### Standard Firewall Filters for Flood Prevention

- [Example: Configuring a Stateless Firewall Filter to Protect Against TCP and ICMP Floods on page 4600](#)
- [Example: Configuring a Filter to Accept Packets Based on IPv6 TCP Flags on page 4610](#)

#### *Example: Configuring a Stateless Firewall Filter to Protect Against TCP and ICMP Floods*

This example shows how to create a stateless firewall filter that protects against TCP and ICMP denial-of-service attacks.

- [Requirements on page 4600](#)
- [Overview on page 4600](#)
- [Configuration on page 4602](#)
- [Verification on page 4606](#)

#### **Requirements**

No special configuration beyond device initialization is required before configuring stateless firewall filters.

#### **Overview**

In this example, you create a stateless firewall filter called **protect-RE** that polices TCP and ICMP packets. This example includes the following policers:

- **tcp-connection-policer**—Limits the traffic rate of the TCP packets to 500,000 bps and the burst size to 15,000 bytes. Packets that exceed the traffic rate are discarded.

- **icmp-policer**—Limits the traffic rate of the ICMP packets to 1,000,000 bps and the burst size to 15,000 bytes. Packets that exceed the traffic rate are discarded.

When specifying limits, the bandwidth limit can be from 32,000 bps to 32,000,000,000 bps and the burst-size limit can be from 1,500 bytes through 100,000,000 bytes. Use the following abbreviations when specifying limits: k (1,000), m (1,000,000), and g (1,000,000,000).

Each policer is incorporated into the action of a filter term. This example includes the following terms:

- **tcp-connection-term**—Policies certain TCP packets with a source address of 192.168.0.0/24 or 10.0.0.0/24. These addresses are defined in the **trusted-addresses** prefix list.

Filtered packets include **tcp-established** packets. The **tcp-established** match condition is an alias for the bit-field match condition **tcp-flags "(ack | rst)"**, which indicates an established TCP session, but not the first packet of a TCP connection.

- **icmp-term**—Polices ICMP packets. All ICMP packets are counted in the **icmp-counter** counter.




**NOTE:** You can move terms within the firewall filter by using the `insert` command. See *insert* in the *CLI User Guide*.

You can apply a stateless firewall to the input or output sides, or both, of an interface. To filter packets transiting the device, apply the firewall filter to any non-Routing Engine interface. To filter packets originating from, or destined for, the Routing Engine, apply the firewall filter to the loopback (lo0) interface.

Figure 95 on page 4601 shows the sample network.

Figure 95: Firewall Filter to Protect Against TCP and ICMP Floods



Because this firewall filter limits Routing Engine traffic to TCP packets, routing protocols that use other transport protocols for Layer 4 cannot successfully establish sessions when this filter is active. To demonstrate, this example sets up OSPF between Device R1 and Device R2.

[“CLI Quick Configuration” on page 4602](#) shows the configuration for all of the devices in [Figure 95 on page 4601](#).

The section [“Step-by-Step Procedure” on page 4603](#) describes the steps on Device R2.

### **Configuration**

|                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>CLI Quick Configuration</b> | To quickly configure the stateless firewall filter, copy the following commands to a text file, remove any line breaks, and then paste the commands into the CLI.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Device R1</b>               | <pre>set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.1/30 set interfaces lo0 unit 0 family inet address 192.168.0.1/32 primary set interfaces lo0 unit 0 family inet address 172.16.0.1/32 set protocols bgp group ext type external set protocols bgp group ext export send-direct set protocols bgp group ext peer-as 200 set protocols bgp group ext neighbor 10.0.0.2 set protocols ospf area 0.0.0.0 interface fe-1/2/0.0 set protocols ospf area 0.0.0.0 interface lo0.0 passive set policy-options policy-statement send-direct term 1 from protocol direct set policy-options policy-statement send-direct term 1 then accept set routing-options router-id 192.168.0.1 set routing-options autonomous-system 100</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Device R2</b>               | <pre>set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.2/30 set interfaces lo0 unit 0 family inet filter input protect-RE set interfaces lo0 unit 0 family inet address 192.168.0.2/32 primary set interfaces lo0 unit 0 family inet address 172.16.0.2/32 set protocols bgp group ext type external set protocols bgp group ext export send-direct set protocols bgp group ext neighbor 10.0.0.1 peer-as 100 set protocols ospf area 0.0.0.0 interface lo0.0 passive set protocols ospf area 0.0.0.0 interface fe-1/2/0.0 set policy-options prefix-list trusted-addresses 10.0.0.0/24 set policy-options prefix-list trusted-addresses 192.168.0.0/24 set policy-options policy-statement send-direct term 1 from protocol direct set policy-options policy-statement send-direct term 1 then accept set routing-options router-id 192.168.0.2 set routing-options autonomous-system 200 set firewall family inet filter protect-RE term tcp-connection-term from source-prefix-list trusted-addresses set firewall family inet filter protect-RE term tcp-connection-term from protocol tcp set firewall family inet filter protect-RE term tcp-connection-term from tcp-established set firewall family inet filter protect-RE term tcp-connection-term then policer tcp-connection-policer set firewall family inet filter protect-RE term tcp-connection-term then accept set firewall family inet filter protect-RE term icmp-term from source-prefix-list trusted-addresses set firewall family inet filter protect-RE term icmp-term from protocol icmp set firewall family inet filter protect-RE term icmp-term then policer icmp-policer set firewall family inet filter protect-RE term icmp-term then count icmp-counter set firewall family inet filter protect-RE term icmp-term then accept set firewall policer tcp-connection-policer filter-specific set firewall policer tcp-connection-policer if-exceeding bandwidth-limit 1m set firewall policer tcp-connection-policer if-exceeding burst-size-limit 15k set firewall policer tcp-connection-policer then discard</pre> |



```

set firewall policer icmp-policer filter-specific
set firewall policer icmp-policer if-exceeding bandwidth-limit 1m
set firewall policer icmp-policer if-exceeding burst-size-limit 15k
set firewall policer icmp-policer then discard

```

### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure stateless firewall filter policers:

1. Configure the device interfaces.

```

[edit interfaces fe-1/2/0 unit 0 family inet]
user@R2# set address 10.0.0.2/30

```

```

[edit interfaces lo0 unit 0 family inet]
user@R2# set address 192.168.0.2/32 primary
user@R2# set address 172.16.0.2/32

```

2. Configure the BGP peering session.

```

[edit protocols bgp group ext]
user@R2# set type external
user@R2# set export send-direct
user@R2# set neighbor 10.0.0.1 peer-as 100

```

3. Configure the autonomous system (AS) number and router ID.

```

[edit routing-options]
user@R2# set autonomous-system 200
user@R2# set router-id 192.168.0.2

```

4. Configure OSPF.

```

[edit protocols ospf area 0.0.0.0]
user@R2# set interface lo0.0 passive
user@R2# set interface fe-1/2/0.0

```

5. Define the list of trusted addresses.

```

[edit policy-options prefix-list trusted-addresses]
user@R2# set 10.0.0.0/24
user@R2# set 192.168.0.0/24

```

6. Configure a policy to advertise direct routes.

```

[edit policy-options policy-statement send-direct term 1]
user@R2# set from protocol direct
user@R2# set then accept

```

7. Configure the TCP policer.

```

[edit firewall policer tcp-connection-policer]
user@R2# set filter-specific
user@R2# set if-exceeding bandwidth-limit 1m
user@R2# set if-exceeding burst-size-limit 15k
user@R2# set then discard

```

8. Create the ICMP policer.

```
[edit firewall policer icmp-policer]
user@R2# set filter-specific
user@R2# set if-exceeding bandwidth-limit 1m
user@R2# set if-exceeding burst-size-limit 15k
user@R2# set then discard
```

9. Configure the TCP filter rules.

```
[edit firewall family inet filter protect-RE term tcp-connection-term]
user@R2# set from source-prefix-list trusted-addresses
user@R2# set from protocol tcp
user@R2# set from tcp-established
user@R2# set then policer tcp-connection-policer
user@R2# set then accept
```

10. Configure the ICMP filter rules.

```
[edit firewall family inet filter protect-RE term icmp-term]
user@R2# set from source-prefix-list trusted-addresses
user@R2# set from protocol icmp
user@R2# set then policer icmp-policer
user@R2# set then count icmp-counter
user@R2# set then accept
```

11. Apply the filter to the loopback interface.

```
[edit interfaces lo0 unit 0]
user@R2# set family inet filter input protect-RE
```

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

```
user@R2# show interfaces
fe-1/2/0 {
 unit 0 {
 family inet {
 address 10.0.0.2/30;
 }
 }
}
lo0 {
 unit 0 {
 family inet {
 filter {
 input protect-RE;
 }
 address 192.168.0.2/32 {
 primary;
 }
 address 172.16.0.2/32;
 }
 }
}
```

```
user@R2# show protocols
bgp {
 group ext {
 type external;
 export send-direct;
 neighbor 10.0.0.1 {
 peer-as 100;
 }
 }
}
ospf {
 area 0.0.0.0 {
 interface lo0.0 {
 passive;
 }
 interface fe-1/2/0.0;
 }
}

user@R2# show policy-options
prefix-list trusted-addresses {
 10.0.0.0/24;
 192.168.0.0/24;
}
policy-statement send-direct {
 term 1 {
 from protocol direct;
 then accept;
 }
}

user@R2# show routing-options
router-id 192.168.0.2;
autonomous-system 200;

user@R2# show firewall
family inet {
 filter protect-RE {
 term tcp-connection-term {
 from {
 source-prefix-list {
 trusted-addresses;
 }
 protocol tcp;
 tcp-established;
 }
 then {
 policer tcp-connection-policer;
 accept;
 }
 }
 term icmp-term {
 from {
 source-prefix-list {
 trusted-addresses;
 }
 protocol icmp;
 }
 }
 }
}
```

```

 }
 then {
 policer icmp-policer;
 count icmp-counter;
 accept;
 }
}
}
}
policer tcp-connection-policer {
 filter-specific;
 if-exceeding {
 bandwidth-limit 1m;
 burst-size-limit 15k;
 }
 then discard;
}
policer icmp-policer {
 filter-specific;
 if-exceeding {
 bandwidth-limit 1m;
 burst-size-limit 15k;
 }
 then discard;
}
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.



**NOTE:** To verify the TCP policer, you can use a packet generation tool. This task is not shown here.

- [Displaying Stateless Firewall Filter That Are in Effect on page 4606](#)
- [Using telnet to Verify the tcp-established Condition in the TCP Firewall Filter on page 4607](#)
- [Using telnet to Verify the Trusted Prefixes Condition in the TCP Firewall Filter on page 4608](#)
- [Using OSPF to Verify the TCP Firewall Filter on page 4608](#)
- [Verifying the ICMP Firewall Filter on page 4609](#)

### Displaying Stateless Firewall Filter That Are in Effect

**Purpose** Verify the configuration of the firewall filter.

**Action** From operational mode, enter the **show firewall** command.

```
user@R2> show firewall
```

```
Filter: protect-RE
```

```
Counters:
```

| Name         | Bytes | Packets |
|--------------|-------|---------|
| icmp-counter | 0     | 0       |

|                        |       |         |  |
|------------------------|-------|---------|--|
| Policers:              |       |         |  |
| Name                   | Bytes | Packets |  |
| icmp-policer           |       | 0       |  |
| tcp-connection-policer |       | 0       |  |

**Meaning** The output shows the filter, the counter, and the policers that are in effect on Device R2.

### *Using telnet to Verify the tcp-established Condition in the TCP Firewall Filter*

**Purpose** Make sure that telnet traffic works as expected.

**Action** Verify that the device can establish only TCP sessions with hosts that meet the **from tcp-established** condition..

1. From Device R2, make sure that the BGP session with Device R1 is established.

```
user@R2> show bgp summary | match down
Groups: 1 Peers: 1 Down peers: 0
```

2. From Device R2, telnet to Device R1.

```
user@R2> telnet 192.168.0.1
Trying 192.168.0.1...
Connected to R1.acme.net.
Escape character is '^['.
```

R1 (ttyp4)

Login:

3. From Device R1, telnet to Device R2.

```
user@R1> telnet 192.168.0.2
Trying 192.168.0.2...
telnet: connect to address 192.168.0.2: Operation timed out
telnet: Unable to connect to remote host
```

4. On Device R2, deactivate the **from tcp-established** match condition.

```
[edit firewall family inet filter protect-RE term tcp-connection-term]
user@R2# deactivate from tcp-established
user@R2# commit
```

5. From Device R1, try again to telnet to Device R2.

```
user@R1> telnet 192.168.0.1
Trying 192.168.0.2...
Connected to R2.acme.net.
Escape character is '^['.
```

R2 (ttyp4)

Login:

**Meaning** Verify the following information:

- As expected, the BGP session is established. The **from tcp-established** match condition is not expected to block BGP session establishment.

- From Device R2, you can telnet to Device R1. Device R1 has no firewall filter configured, so this is the expected behavior.
- From Device R1, you cannot telnet to Device R2. Telnet uses TCP as the transport protocol, so this result might be surprising. The cause for the lack of telnet connectivity is the **from tcp-established** match condition. This match condition limits the type of TCP traffic that is accepted of Device R2. After this match condition is deactivated, the telnet session is successful.

### *Using telnet to Verify the Trusted Prefixes Condition in the TCP Firewall Filter*

**Purpose** Make sure that telnet traffic works as expected.

**Action** Verify that the device can establish only telnet sessions with a host at an IP address that matches one of the trusted source addresses. For example, log in to the device with the **telnet** command from another host with one of the trusted address prefixes. Also, verify that telnet sessions with untrusted source addresses are blocked.

1. From Device R1, telnet to Device R2 from an untrusted source address.

```
user@R1> telnet 172.16.0.2 source 172.16.0.1
Trying 172.16.0.2...
^C
```

2. From Device R2, add 172.16/16 to the list of trusted prefixes.

```
[edit policy-options prefix-list trusted-addresses]
user@R2# set 172.16.0.0/16
user@R2# commit
```

3. From Device R1, try again to telnet to Device R2.

```
user@R1> telnet 172.16.0.2 source 172.16.0.1
Trying 172.16.0.2...
Connected to R2.acme.net.
Escape character is '^J'.
```

```
R2 (ttyp4)
```

```
login:
```

**Meaning** Verify the following information:

- From Device R1, you cannot telnet to Device R2 with an untrusted source address. After the 172.16/16 prefix is added to the list of trusted prefixes, the telnet request from source address 172.16.0.1 is accepted.
- OSPF session establishment is blocked. OSPF does not use TCP as its transport protocol. After the **from protocol tcp** match condition is deactivated, OSPF session establishment is not blocked.

### *Using OSPF to Verify the TCP Firewall Filter*

**Purpose** Make sure that OSPF traffic works as expected.

\_\_\_\_\_

```
--- 192.168.0.2 ping statistics ---
600 packets transmitted, 536 packets received, 10% packet loss
pinground-trip min/avg/max/stddev = 2.976/3.405/42.380/2.293 ms
```

3. From Device R2, check the firewall statistics.

```
user@R2> show firewall
```

```
Filter: protect-RE
```

```
Counters:
```

| Name         | Bytes   | Packets |
|--------------|---------|---------|
| icmp-counter | 1180804 | 1135    |

```
Policers:
```

| Name                   | Bytes | Packets |
|------------------------|-------|---------|
| icmp-policer           |       | 66      |
| tcp-connection-policer |       | 0       |

4. From an untrusted source address on Device R1, send a ping request to Device R2's loopback interface.

```
user@R1> ping 172.16.0.2 source 172.16.0.1
```

```
PING 172.16.0.2 (172.16.0.2): 56 data bytes
```

```
^C
```

```
--- 172.16.0.2 ping statistics ---
```

```
14 packets transmitted, 0 packets received, 100% packet loss
```

**Meaning** Verify the following information:

- The ping output shows that 10% packet loss is occurring.
- The ICMP packet counter is incrementing, and the icmp-policer is incrementing.
- Device R2 does not send ICMP responses to the **ping 172.16.0.2 source 172.16.0.1** command.

**Related  
Documentation**

- [Example: Configuring a Stateless Firewall Filter to Accept Traffic from Trusted Sources on page 4576](#)
- [Two-Color Policer Configuration Overview on page 4813](#)
- [Junos OS Feature Support Reference for SRX Series and J Series Devices](#)

**Example: Configuring a Filter to Accept Packets Based on IPv6 TCP Flags**

This example shows how to configure a standard stateless firewall filter to accept packets from a trusted source.

- [Requirements on page 4610](#)
- [Overview on page 4611](#)
- [Configuration on page 4611](#)
- [Verification on page 4612](#)

**Requirements**

No special configuration beyond device initialization is required before configuring this example.



### Overview

In this example, you create a filter that accepts packets with specific IPv6 TCP flags.

### Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

- [Configure the Stateless Firewall Filter on page 4611](#)
- [Apply the Firewall Filter to the Loopback Interface on page 4611](#)
- [Confirm and Commit Your Candidate Configuration on page 4612](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```

set firewall family inet6 filter tcp_filter term 1 from next-header tcp
set firewall family inet6 filter tcp_filter term 1 from tcp-flags syn
set firewall family inet6 filter tcp_filter term 1 then count tcp_syn_pkt
set firewall family inet6 filter tcp_filter term 1 then log
set firewall family inet6 filter tcp_filter term 1 then accept
set interfaces lo0 unit 0 family inet6 filter input tcp_filter
set interfaces lo0 unit 0 family inet6 address ::10.34.1.0/120

```

### Configure the Stateless Firewall Filter

### Step-by-Step Procedure

To configure the firewall filter

1. Create the IPv6 stateless firewall filter **tcp\_filter**.
 

```

[edit]
user@host# edit firewall family inet6 filter tcp_filter

```
2. Specify that a packet matches if it is the initial packet in a TCP session and the next header after the IPv6 header is type TCP.
 

```

[edit firewall family inet6 filter tcp_filter]
user@host# set term 1 from next-header tcp
user@host# set term 1 from tcp-flags syn

```
3. Specify that matched packets are counted, logged to the buffer on the Packet Forwarding Engine, and accepted.
 

```

[edit firewall family inet6 filter tcp_filter]
user@host# set term 1 then count tcp_syn_pkt
user@host# set term 1 then log
user@host# set term 1 then accept

```

### Apply the Firewall Filter to the Loopback Interface

### Step-by-Step Procedure

To apply the firewall filter to the loopback interface:

- ```

[edit]
user@host# set interfaces lo0 unit 0 family inet6 filter input tcp_filter
user@host# set interfaces lo0 unit 0 family inet6 address ::10.34.1.0/120

```

Confirm and Commit Your Candidate Configuration

Step-by-Step Procedure

To confirm and then commit your candidate configuration:

1. Confirm the configuration of the stateless firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show firewall
family inet6 {
  filter tcp_filter {
    term 1 {
      from {
        next-header tcp;
        tcp-flags syn;
      }
      then {
        count tcp_syn_pkt;
        log;
        accept;
      }
    }
  }
}
```

2. Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
lo0 {
  unit 0 {
    family inet6 {
      filter {
        input tcp_filter;
      }
      address ::10.34.1.0/120;
    }
  }
}
```

3. When you are done configuring the device, commit your candidate configuration.

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

Verification

To confirm that the configuration is working properly, enter the **show firewall** operational mode command.

Related Documentation

- [Understanding How to Use Standard Firewall Filters on page 3293](#)

- [Example: Configuring a Stateless Firewall Filter to Protect Against TCP and ICMP Floods on page 4600](#)
- [Example: Configuring a Filter to Block TCP Access to a Port Except from Specified BGP Peers on page 4595](#)

Standard Firewall Filters for Fragment Handling

- [Example: Configuring a Stateless Firewall Filter to Handle Fragments on page 4613](#)

Example: Configuring a Stateless Firewall Filter to Handle Fragments

This example shows how to create a stateless firewall filter that handles packet fragments.

- [Requirements on page 4613](#)
- [Overview on page 4613](#)
- [Configuration on page 4614](#)
- [Verification on page 4617](#)

Requirements

No special configuration beyond device initialization is required before configuring stateless firewall filters.

Overview

In this example, you create a stateless firewall filter called **fragment-RE** that accepts fragmented packets originating from 10.2.1.0/24 and destined for the BGP port. This example includes the following firewall filter terms:

- **not-from-prefix-term**—Discards packets that are not from 10.2.1.0/24 to ensure that subsequent terms in the firewall filter are matched against packets from 10.2.1.0/24 only.
- **small-offset-term**—Discards small (1–5) offset packets to ensure that subsequent terms in the firewall filter can be matched against all the headers in the packet. In addition, the term adds a record to the system logging destinations for the firewall facility.
- **not-fragmented-term**—Accepts unfragmented TCP packets with a destination port that specifies the BGP protocol. A packet is considered unfragmented if the MF flag is not set and the fragment offset equals 0.
- **first-fragment-term**—Accepts the first fragment of a fragmented TCP packet with a destination port that specifies the BGP protocol.
- **fragment-term**—Accepts all fragments that were not discarded by **small-offset-term**. (packet fragments 6–8191). However, only those fragments that are part of a packet containing a first fragment accepted by **first-fragment-term** are reassembled by the destination device.

Packet fragments offset can be from 1 through 8191.



NOTE: You can move terms within the firewall filter by using the `insert` command. For more information, see “*insert*” in the *CLI User Guide*.

Configuration

CLI Quick Configuration

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

```
set firewall family inet filter fragment-RE term not-from-prefix-term from source-address 0.0.0.0/0
set firewall family inet filter fragment-RE term not-from-prefix-term from source-address 10.2.1.0/24 except
set firewall family inet filter fragment-RE term not-from-prefix-term then discard
set firewall family inet filter fragment-RE term small-offset-term from fragment-offset 1-5
set firewall family inet filter fragment-RE term small-offset-term then syslog
set firewall family inet filter fragment-RE term small-offset-term then discard
set firewall family inet filter fragment-RE term not-fragmented-term from fragment-offset 0
set firewall family inet filter fragment-RE term not-fragmented-term from fragment-flags "!more-fragments"
set firewall family inet filter fragment-RE term not-fragmented-term from protocol tcp
set firewall family inet filter fragment-RE term not-fragmented-term from destination-port bgp
set firewall family inet filter fragment-RE term not-fragmented-term then accept
set firewall family inet filter fragment-RE term first-fragment-term from first-fragment
set firewall family inet filter fragment-RE term first-fragment-term from protocol tcp
set firewall family inet filter fragment-RE term first-fragment-term from destination-port bgp
set firewall family inet filter fragment-RE term first-fragment-term then accept
set firewall family inet filter fragment-RE term fragment-term from fragment-offset 6-8191
set firewall family inet filter fragment-RE term fragment-term then accept
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see “[Using the CLI Editor in Configuration Mode](#)” on page 4704 in the *CLI User Guide*.

To configure the stateless firewall filter:

1. Define the stateless firewall filter.

```
[edit]
user@host# edit firewall family inet filter fragment-RE
```

2. Configure the first term for the filter.

```
[edit firewall family inet filter fragment-RE ]
user@host# set term not-from-prefix-term from source-address 0.0.0.0/0
user@host# set term not-from-prefix-term from source-address 10.2.1.0/24 except
user@host# set term not-from-prefix-term then discard
```

3. Define the second term for the filter.

```
[edit firewall family inet filter fragment-RE]
user@host# edit term small-offset-term
```
4. Define the match conditions for the term.

```
[edit firewall family inet filter fragment-RE term small-offset-term]
user@host# set from fragment-offset 1-5
```
5. Define the action for the term.

```
[edit firewall family inet filter fragment-RE term small-offset-term]
user@host# set then syslog discard
```
6. Define the third term for the filter.

```
[edit]
user@host# edit firewall family inet filter fragment-RE term not-fragmented-term
```
7. Define the match conditions for the term.

```
[edit firewall family inet filter fragment-RE term not-fragmented-term]
user@host# set from fragment-flags "!more-fragments" fragment-offset 0 protocol
tcp destination-port bgp
```
8. Define the action for the term.

```
[edit firewall family inet filter fragment-RE term not-fragmented-term]
user@host# set then accept
```
9. Define the fourth term for the filter.

```
[edit]
user@host# edit firewall family inet filter fragment-RE term first-fragment-term
```
10. Define the match conditions for the term.

```
[edit firewall family inet filter fragment-RE term first-fragment-term]
user@host# set from first-fragment protocol tcp destination-port bgp
```
11. Define the action for the term.

```
[edit firewall family inet filter fragment-RE term first-fragment-term]
user@host# set then accept
```
12. Define the last term for the filter.

```
[edit]
user@host# edit firewall family inet filter fragment-RE term fragment-term
```
13. Define the match conditions for the term.

```
[edit firewall family inet filter fragment-RE term fragment-term]
user@host# set from fragment-offset 6-8191
```
14. Define the action for the term.

```
[edit firewall family inet filter fragment-RE term fragment-term]
user@host# set then accept
```

Results Confirm your configuration by entering the **show firewall** command from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show firewall
family inet {
  filter fragment-RE {
    term not-from-prefix-term {
      from {
        source-address {
          0.0.0.0/0;
          10.2.1.0/24 except;
        }
      }
      then discard;
    }
    term small-offset-term {
      from {
        fragment-offset 1-5;
      }
      then {
        syslog;
        discard;
      }
    }
    term not-fragmented-term {
      from {
        fragment-offset 0;
        fragment-flags "!more-fragments";
        protocol tcp;
        destination-port bgp;
      }
      then accept;
    }
    term first-fragment-term {
      from {
        first-fragment;
        protocol tcp;
        destination-port bgp;
      }
      then accept;
    }
    term fragment-term {
      from {
        fragment-offset 6-8191;
      }
      then accept;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

To confirm that the configuration is working properly, perform these tasks:

- [Displaying Stateless Firewall Filter Configurations on page 4617](#)
- [Verifying a Firewall Filter that Handles Fragments on page 4617](#)

Displaying Stateless Firewall Filter Configurations

Purpose Verify the configuration of the firewall filter. You can analyze the flow of the filter terms by displaying the entire configuration.

Action From configuration mode, enter the **show firewall** command.

Meaning Verify that the output shows the intended configuration of the firewall filter. In addition, verify that the terms are listed in the order in which you want the packets to be tested. You can move terms within a firewall filter by using the **insert** CLI command.

Verifying a Firewall Filter that Handles Fragments

Purpose Verify that the actions of the firewall filter terms are taken.

Action Send packets to the device that match the terms.

Meaning Verify that packets from 10.2.1.0/24 with small fragment offsets are recorded in the device's system logging destinations for the firewall facility.

Related Documentation • *show route summary* in the *Junos OS Operational Mode Commands*.

Standard Firewall Filters for Setting Rate Limits

- [Example: Configuring a Rate-Limiting Filter Based on Destination Class on page 4617](#)

Example: Configuring a Rate-Limiting Filter Based on Destination Class

This example shows how to configure a rate-limiting stateless firewall filter.

- [Requirements on page 4617](#)
- [Overview on page 4618](#)
- [Configuration on page 4618](#)
- [Verification on page 4620](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Before you begin, configure the destination class **class1**.

Overview

In this example, you use a stateless firewall filter to set rate limits based on a destination class.

To activate a policer from within a stateless firewall filter configuration:

- Create a template for the policer by including the **policer policer-name** statement.
- Reference the policer in a filter term that specifies the policer in the **policer policer-name** nonterminating action.

You can also activate a policer by including the **policer (input | output) policer-template-name** statement at a logical interface.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

- [Configure the Stateless Firewall Filter on page 4618](#)
- [Apply the Stateless Firewall Filter to a Logical Interface on page 4619](#)
- [Confirm and Commit Your Candidate Configuration on page 4619](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set firewall filter rl_dclass1 policer police_class1 if-exceeding bandwidth-limit 25
set firewall filter rl_dclass1 policer police_class1 if-exceeding burst-size-limit 1000
set firewall filter rl_dclass1 policer police_class1 then discard
set firewall filter rl_dclass1 term term1 from destination-class class1
set firewall filter rl_dclass1 term term1 then policer police_class1
set interfaces ge-0/0/1 unit 0 family inet address 10.1.2.3/30
set interfaces ge-0/0/1 unit 0 family inet filter input ospf_or_131
```

Configure the Stateless Firewall Filter

Step-by-Step Procedure

To configure the stateless firewall filter **rl_dclass1** with policer **police_class1** for destination class **class1**:

1. Create the stateless firewall filter **rl_dclass1**.

```
[edit]
user@host# edit firewall filter rl_dclass1
```

2. Configure the policer template **police_class1**.

```
[edit firewall filter rl_dclass1]
user@host# set policer police_class1 if-exceeding bandwidth-limit 25
user@host# set policer police_class1 if-exceeding burst-size-limit 1000
user@host# set policer police_class1 then discard
```


3. Configure a filter term that uses policer **police_class1** to rate-limit traffic for destination class **class1**.

```
[edit firewall filter rl_dclass1]
user@host# set term term1 from destination-class class1
user@host# set term term1 then policer police_class1
```

Apply the Stateless Firewall Filter to a Logical Interface

Step-by-Step Procedure

To apply the filter **rl_dclass1** to a logical interface:

1. Configure the logical interface to which you will apply the stateless firewall filter.

```
[edit]
user@host# edit interfaces ge-0/0/1 unit 0 family inet
```

2. Configure the interface address for the logical interface.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
user@host# set address 10.1.2.3/30
```

3. Apply the stateless firewall filter to the logical interface.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
user@host# set filter input rl_dclass1
```

Confirm and Commit Your Candidate Configuration

Step-by-Step Procedure

To confirm and then commit your candidate configuration:

1. Confirm the configuration of the stateless firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show firewall
filter rl_dclass1 {
  policer police_class1 {
    if-exceeding {
      bandwidth-limit 25;
      burst-size-limit 1000;
    }
    then {
      discard;
    }
  }
  term term1 {
    from {
      destination-class class1;
    }
    then {
      policer police_class1;
    }
  }
}
```

2. Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      filter {
        input rl_dclass1;
      }
      address 10.1.2.3/30;
    }
  }
}
```

3. If you are done configuring the device, commit your candidate configuration.

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

Verification

To confirm that the configuration is working properly, enter the **show class-of-service ge-0/0/1** operational mode command.

Related Documentation

- [Understanding How to Use Standard Firewall Filters on page 3293](#)
- [Filtering Packets Received on an Interface Set Overview on page 4515](#)
- [Statement Hierarchy for Defining an Interface Set on page 4777](#)
- [Statement Hierarchy for Configuring a Filter to Match on an Interface Set on page 4777](#)
- [Example: Filtering Packets Received on an Interface Set on page 4638](#)

Examples of Standard Firewall Filters Configuration

- [Example: Applying Lists of Multiple Standard Firewall Filters on page 4621](#)
- [Example: Nesting References to Multiple Standard Firewall Filters on page 4626](#)
- [Example: Configuring Interface-Specific Firewall Filter Counters on page 4629](#)
- [Example: Filtering Packets Received on an Interface Group on page 4634](#)
- [Example: Filtering Packets Received on an Interface Set on page 4638](#)
- [Example: Configuring Filter-Based Forwarding on the Source Address on page 4644](#)
- [Example: Configuring Filter-Based Forwarding on Logical Systems on page 4652](#)
- [Example: Configuring Filter-Based Forwarding to a Specific Outgoing Interface or Destination IP Address on page 4662](#)

Example: Applying Lists of Multiple Standard Firewall Filters

This example shows how to apply lists of multiple stateless firewall filters.

- [Requirements on page 4621](#)
- [Overview on page 4621](#)
- [Configuration on page 4622](#)
- [Verification on page 4625](#)

Requirements

Before you begin, be sure that you have:

- Installed your router or switch, and supported PIC, DPC, or MPC and performed the initial router or switch configuration.
- Configured basic Ethernet in the topology.
- Configured a logical interface to run the IP version 4 (IPv4) protocol (**family inet**) and configured the logical interface with an interface address. This example uses logical interface **ge-1/3/0.0** configured with the IP address 1.1.1.2/30.



NOTE: For completeness, the configuration section of this example includes setting an IP address for logical interface **ge-1/3/0.0**.

- Verified that traffic is flowing in the topology and that ingress and egress IPv4 traffic is flowing through logical interface **ge-1/3/0.0**.
- Verified that you have access to the remote host that is connected to this router's or switch's logical interface **ge-1/3/0.0**.

Overview

In this example, you configure three IPv4 stateless firewall filters and apply each filter directly to the same logical interface by using a list.

Topology

This example applies the following firewall filters as a *list of input filters* at logical interface **ge-1/3/0.0**. Each filter contains a single term that evaluates IPv4 packets and accepts packets based on the value of the **destination port** field in the TCP header:

- Filter **filter_FTP** matches on the FTP port number (21).
- Filter **filter_SSH** matches on the SSH port number (22).
- Filter **filter_Telnet** matches on the Telnet port number (23).

If an inbound packet does not match any of the filters in the input list, the packet is discarded.



NOTE: The Junos OS uses filters in a list in the order in which the filter names appear in the list. In this simple example, the order is irrelevant because all of the filters specify the same action.

Any of the filters can be applied to other interfaces, either alone (using the **input** or **output** statement) or in combination with other filters (using the **input-list** or **output-list** statement). The objective is to configure multiple “minimalist” firewall filters that you can reuse in interface-specific filter lists.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

- [Configure Multiple IPv4 Stateless Firewall Filters on page 4622](#)
- [Apply the Filters to a Logical Interface as an Input List and an Output List on page 4623](#)
- [Confirm and Commit Your Candidate Configuration on page 4623](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set firewall family inet filter filter_FTP term 0 from protocol tcp
set firewall family inet filter filter_FTP term 0 from destination-port 21
set firewall family inet filter filter_FTP term 0 then count pkts_FTP
set firewall family inet filter filter_FTP term 0 then accept
set firewall family inet filter filter_SSH term 0 from protocol tcp
set firewall family inet filter filter_SSH term 0 from destination-port 23
set firewall family inet filter filter_SSH term 0 then count pkts_SSH
set firewall family inet filter filter_SSH term 0 then accept
set firewall family inet filter filter_Telnet term 0 from protocol tcp
set firewall family inet filter filter_Telnet term 0 from destination-port 22
set firewall family inet filter filter_Telnet term 0 then count pkts_Telnet
set firewall family inet filter filter_Telnet term 0 then accept
set firewall family inet filter filter_discard term 1 then count pkts_discarded
set firewall family inet filter filter_discard term 1 then discard
set interfaces ge-1/3/0 unit 0 family inet address 1.1.1.2/30
set interfaces ge-1/3/0 unit 0 family inet filter input-list filter_FTP
set interfaces ge-1/3/0 unit 0 family inet filter input-list filter_SSH
set interfaces ge-1/3/0 unit 0 family inet filter input-list filter_Telnet
set interfaces ge-1/3/0 unit 0 family inet filter input-list filter_discard
```

Configure Multiple IPv4 Stateless Firewall Filters

Step-by-Step Procedure

To configure the IPv4 stateless firewall filters:

1. Navigate the CLI to the hierarchy level at which you configure IPv4 firewall filters.

```
[edit]
user@host# edit firewall family inet
```
2. Configure the first firewall filter to count and accept packets for port 21.

```
[edit firewall family inet]
```

```

user@host# set filter filter_FTP term 0 from protocol tcp
user@host# set filter filter_FTP term 0 from destination-port 21
user@host# set filter filter_FTP term 0 then count pkts_FTP
user@host# set filter filter_FTP term 0 then accept

```

3. Configure the second firewall filter to count and accept packets for port 23.

```

[edit firewall family inet]
user@host# set filter filter_SSH term 0 from protocol tcp
user@host# set filter filter_SSH term 0 from destination-port 23
user@host# set filter filter_SSH term 0 then count pkt_SSH
user@host# set filter filter_SSH term 0 then accept

```

4. Configure the third firewall filter to count and accept packets from port 22.

```

[edit firewall family inet]
user@host# set filter filter_Telnet term 0 from protocol tcp
user@host# set filter filter_Telnet term 0 from destination-port 22
user@host# set filter filter_Telnet term 0 then count pkts_Telnet
user@host# set filter filter_Telnet term 0 then accept

```

5. Configure the last firewall filter to count the discarded packets.

```

[edit firewall family inet]
user@host# set filter filter_discard term 1 then count pkts_discarded
user@host# set filter filter_discard term 1 then discard

```

Apply the Filters to a Logical Interface as an Input List and an Output List

Step-by-Step Procedure To apply the six IPv4 stateless firewall filters as a list of input filters and a list of output filters:

1. Navigate the CLI to the hierarchy level at which you apply IPv4 firewall filters to logical interface **ge-1/3/0.0**.

```

[edit]
user@host# edit interfaces ge-1/3/0 unit 0 family inet

```

2. Configure the IPv4 protocol family for the logical interface.

```

[edit interfaces ge-1/3/0 unit 0 family inet]
user@host# set address 1.1.1.2/30

```

3. Apply the filters as a list of input filters.

```

[edit interfaces ge-1/3/0 unit 0 family inet]
user@host# set filter input-list [ filter_FTP filter_SSH filter_Telnet filter_discard ]

```

Confirm and Commit Your Candidate Configuration

Step-by-Step Procedure To confirm and then commit your candidate configuration:

1. Confirm the configuration of the stateless firewall filters by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

[edit]
user@host# show firewall

```

```
family inet {
  filter filter_FTP {
    term 0 {
      from {
        protocol tcp;
        destination-port 21;
      }
      then {
        count pkts_FTP;
        accept;
      }
    }
  }
  filter filter_SSH {
    term 0 {
      from {
        protocol tcp;
        destination-port 23;
      }
      then {
        count pkts_SSH;
        accept;
      }
    }
  }
  filter filter_Telnet {
    term 0 {
      from {
        protocol tcp;
        destination-port 22;
      }
      then {
        count pkts_Telnet;
        accept;
      }
    }
  }
  filter filter_discard {
    term 1 {
      then {
        count pkts_discarded;
        discard;
      }
    }
  }
}
```

2. Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-1/3/0 {
  unit 0 {
    family inet {
```

```

        filter {
            input-list [ filter_FTP filter_SSH filter_Telnet filter_discard ];
        }
        address 1.1.1.2/30;
    }
}

```

3. If you are done configuring the device, commit your candidate configuration.

```

[edit]
user@host# commit

```

Verification

Confirm that the configuration is working properly.

- [Verifying That Inbound Packets Are Accepted Only If Destined for the FTP, SSH or Telnet Port on page 4625](#)

Verifying That Inbound Packets Are Accepted Only If Destined for the FTP, SSH or Telnet Port

Purpose Verify that all three filters are active for the logical interface.

Action To verify that input packets are accepted according to the three filters:

1. From the remote host that is connected to this router's (or switch's) logical interface **ge-1/3/0.0**, send a packet with destination port number 21 in the header. The packet should be accepted.
2. From the remote host that is connected to this router's (or switch's) logical interface **ge-1/3/0.0**, send a packet with destination port number 23 in the header. The packet should be accepted.
3. From the remote host that is connected to this router's (or switch's) logical interface **ge-1/3/0.0**, send a packet with destination port number 22 in the header. The packet should be accepted.
4. From the remote host that is connected to this router's (or switch's) logical interface **ge-1/3/0.0**, send a packet with a destination port number *other than* 21, 22, or 23. The packet should be discarded.
5. To display counter information for the list of filters applied to the input at **ge-1/3/0.0-i** enter the **show firewall filter ge-1/3/0.0-i** operational mode command. The command output displays the number of bytes and packets that match filter terms associated with the following counters:
 - **pkts_FTP-ge-1/3/0.0-i**
 - **pkts_SSH-ge-1/3/0.0-i**
 - **pkts_Telnet-ge-1/3/0.0-i**
 - **pkts_discard-ge-1/3/0.0-i**

- Related Documentation**
- [Multiple Standard Firewall Filters Applied as a List Overview on page 4505](#)
 - [Guidelines for Applying Multiple Standard Firewall Filters as a List on page 4509](#)

Example: Nesting References to Multiple Standard Firewall Filters

This example shows how to configure nested references to multiple stateless firewall filters.

- [Requirements on page 4626](#)
- [Overview on page 4626](#)
- [Configuration on page 4626](#)
- [Verification on page 4629](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, you configure a stateless firewall filter for a match condition and action combination that can be shared among multiple firewall filters. You then configure two firewall filters that reference the first firewall filter. Later, if the common filtering criteria needs to be changed, you would modify only the one shared firewall filter configuration.

Topology

The **common_filter** firewall filter discards packets that have a UDP source or destination port field number of 69. Both of the two additional firewall filters, **filter1** and **filter2**, reference the **common_filter**.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

- [Configure the Nested Firewall Filters on page 4627](#)
- [Apply Both Nested Firewall Filters to Interfaces on page 4627](#)
- [Confirm and Commit Your Candidate Configuration on page 4628](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set firewall family inet filter common_filter term common_term from protocol udp
set firewall family inet filter common_filter term common_term from port tftp
set firewall family inet filter common_filter term common_term then discard
set firewall family inet filter filter1 term term1 filter common-filter
set firewall family inet filter filter1 term term2 from address 192.168.0.0/16
set firewall family inet filter filter1 term term2 then reject
set firewall family inet filter filter2 term term1 filter common-filter
set firewall family inet filter filter2 term term2 from protocol udp
set firewall family inet filter filter2 term term2 from port bootps
```



```

set firewall family inet filter filter2 term term2 then accept
set interfaces ge-0/0/0 unit 0 family inet address 10.1.0.1/24
set interfaces ge-0/0/0 unit 0 family inet filter input filter1
set interfaces ge-0/0/3 unit 0 family inet address 10.1.3.1/24
set interfaces ge-0/0/0 unit 0 family inet filter input filter2

```

Configure the Nested Firewall Filters

Step-by-Step Procedure

To configure two nested firewall filters that share a common filter:

1. Navigate the CLI to the hierarchy level at which you configure IPv4 firewall filters.

```

[edit]
user@host# edit firewall family inet

```

2. Configure the common filter that will be referenced by multiple other filters.

```

[edit firewall family inet]
user@host# set filter common_filter term common_term from protocol udp
user@host# set filter common_filter term common_term from port tftp
user@host# set filter common_filter term common_term then discard

```

3. Configure a filter that references the common filter.

```

[edit firewall family inet]
user@host# set filter filter1 term term1 filter common-filter
user@host# set filter filter1 term term2 from address 192.168.0.0/16
user@host# set filter filter1 term term2 then reject

```

4. Configure a second filter that references the common filter.

```

[edit firewall family inet]
user@host# set filter filter2 term term1 filter common-filter
user@host# set filter filter2 term term2 from protocol udp
user@host# set filter filter2 term term2 from port bootps
user@host# set filter filter2 term term2 then accept

```

Apply Both Nested Firewall Filters to Interfaces

Step-by-Step Procedure

To apply both nested firewall filters to logical interfaces:

1. Apply the first nested filter to a logical interface input.

```

[edit]
user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.1.0.1/24
user@host# set interfaces ge-0/0/0 unit 0 family inet filter input filter1

```

2. Apply the second nested filter to a logical interface input.

```

[edit]
user@host# set interfaces ge-0/0/3 unit 0 family inet address 10.1.3.1/24
user@host# set interfaces ge-0/0/0 unit 0 family inet filter input filter2

```

Confirm and Commit Your Candidate Configuration

Step-by-Step Procedure

To confirm and then commit your candidate configuration:

1. Confirm the configuration of the stateless firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
  filter common_filter {
    term common_term {
      from {
        protocol udp;
        port tftp;
      }
      then {
        discard;
      }
    }
  }
  filter filter1 {
    term term1 {
      filter common-filter;
    }
    term term2 {
      from {
        address 192.168/16;
      }
      then {
        reject;
      }
    }
  }
  filter filter2 {
    term term1 {
      filter common-filter;
    }
    term term2 {
      from {
        protocol udp;
        port bootps;
      }
      then {
        accept;
      }
    }
  }
}
```

2. Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
```

```

user@host# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      filter {
        input filter1;
      }
      address 10.1.0.1/24;
    }
  }
}
ge-0/0/3 {
  unit 0 {
    family inet {
      filter {
        input filter2;
      }
      address 10.1.3.1/24;
    }
  }
}

```

3. If you are done configuring the device, commit your candidate configuration.

```

[edit]
user@host# commit

```

Verification

To confirm that the configuration is working properly, enter the `show firewall filter filter1` and `show firewall filter filter2` operational mode commands.

Related Documentation

- [Multiple Standard Firewall Filters in a Nested Configuration Overview on page 4510](#)
- [Guidelines for Nesting References to Multiple Standard Firewall Filters on page 4511](#)

Example: Configuring Interface-Specific Firewall Filter Counters

This example shows how to configure and apply an interface-specific standard stateless firewall filter.

- [Requirements on page 4629](#)
- [Overview on page 4630](#)
- [Configuration on page 4630](#)
- [Verification on page 4632](#)

Requirements

Interface-specific stateless firewall filters are supported on T Series, M120, M320, MX Series routers, and EX Series switches only.

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, you create an interface-specific stateless firewall filter that counts and accepts packets with source or destination addresses in a specified prefix and the IP protocol type field set to a specific value.

Topology

You configure the interface-specific stateless firewall filter **filter_s_tcp** to count and accept packets with IP source or destination addresses in the **10.0.0.0/12** prefix and the IP protocol type field set to **tcp** (or the numeric value **6**).

The name of the firewall filter counter is **count_s_tcp**.

You apply the firewall filter to multiple logical interfaces:

- **at-1/1/1.0** input
- **ge-2/2/2.2** output

Applying the filter to these two interfaces results in two instances of the filter: **filter_s_tcp-at-1/1/1.0-i** and **filter_s_tcp-ge-2/2/2.2-o**, respectively.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configure the Interface-Specific Firewall Filter on page 4630](#)
- [Apply the Interface-Specific Firewall Filter to Multiple Interfaces on page 4631](#)
- [Confirm Your Candidate Configuration on page 4631](#)
- [Clear the Counters and Commit Your Candidate Configuration on page 4632](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set firewall family inet filter filter_s_tcp interface-specific
set firewall family inet filter filter_s_tcp term 1 from address 10.0.0.0/12
set firewall family inet filter filter_s_tcp term 1 from protocol tcp
set firewall family inet filter filter_s_tcp term 1 then count count_s_tcp
set firewall family inet filter filter_s_tcp term 1 then accept
set interfaces at-1/1/1 unit 0 family inet filter input filter_s_tcp
set interfaces ge-2/2/2 unit 2 family inet filter filter_s_tcp
```

Configure the Interface-Specific Firewall Filter

Step-by-Step Procedure

To configure the interface-specific firewall filter:

1. Create the IPv4 firewall filter **filter_s_tcp**.

[edit]
user@host# edit firewall family inet filter filter_s_tcp

2. Enable interface-specific instances of the filter.

```
[edit firewall family inet filter filter_s_tcp]
user@host# set interface-specific
```

3. Configure the match conditions for the term.

```
[edit firewall family inet filter filter_s_tcp]
user@host# set term 1 from address 10.0.0.0/12
user@host# set term 1 from protocol tcp
```

4. Configure the actions for the term.

```
[edit firewall family inet filter filter_s_tcp]
user@host# set term 1 then count count_s_tcp
user@host# set term 1 then accept
```

Apply the Interface-Specific Firewall Filter to Multiple Interfaces

Step-by-Step Procedure

To apply the filter **filter_s_tcp** to logical interfaces **at-1/1/1.0** and **ge-2/2/2.2**:

1. Apply the interface-specific filter to packets received on logical interface **at-1/1/1.0**.

```
[edit]
user@host# set interfaces at-1/1/1 unit 0 family inet filter input filter_s_tcp
```

2. Apply the interface-specific filter to packets transmitted from logical interface **ge-2/2/2.2**.

```
[edit]
user@host# set interfaces ge-2/2/2 unit 2 family inet filter filter_s_tcp
```

Confirm Your Candidate Configuration

Step-by-Step Procedure

To confirm your candidate configuration:

1. Confirm the configuration of the stateless firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
  filter filter_s_tcp {
    interface-specific;
    term 1 {
      from {
        address {
          10.0.0.0/12;
        }
        protocol tcp;
      }
      then {
        count count_s_tcp;
        accept;
      }
    }
  }
}
```

```
}  
}
```

2. Confirm the configuration of the interfaces by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]  
user@host# show interfaces  
at-1/1/1 {  
  unit 0  
    family inet {  
      filter {  
        input filter_s_tcp;  
      }  
    }  
  }  
}  
ge-2/2/2 {  
  unit 2  
    family inet {  
      filter {  
        output filter_s_tcp;  
      }  
    }  
  }  
}
```

Clear the Counters and Commit Your Candidate Configuration

Step-by-Step Procedure

To clear the counters and commit your candidate configuration:

1. From operational command mode, use the **clear firewall all** command to clear the statistics for all firewall filters.

To clear only the counters used in this example, include the interface-specific filter instance names:

```
[edit]  
user@host> clear firewall filter filter_s_tcp-at-1/1/1.0-i  
user@host> clear firewall filter filter_s_tcp-ge-2/2/2.2-o
```

2. Commit your candidate configuration.

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

Verification

Confirm that the configuration is working properly.

- [Verifying That the Filter Is Applied to Each of the Multiple Interfaces on page 4632](#)
- [Verifying That the Counters Are Collected Separately by Interface on page 4633](#)

Verifying That the Filter Is Applied to Each of the Multiple Interfaces

Purpose Verify that the filter is applied to each of the multiple interfaces.

Action Run the **show interfaces** command with the **detail** or **extensive** output level.

1. Verify that the filter is applied to the input for **at-1/1/1.0**:

```
user@host> show interfaces at-1/1/1 detail
Physical interface: at-1/1/1, Enabled, Physical link is Up
  Interface index: 300, SNMP ifIndex: 194, Generation: 183
...
  Logical interface at-1/1/1.0 (Index 64) (SNMP ifIndex 204) (Generation 5)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: ATM-SNAP
...
  Protocol inet, MTU: 4470, Generation: 13, Route table: 0
    Flags: Sendbcst-pkt-to-re
    Input Filters: filter_s_tcp-at-1/1/1.0-i,,,,,
```

2. Verify that the filter is applied to the output for **ge-2/2/2.2**:

```
user@host> show interfaces ge-2/2/2 detail
Physical interface: ge-2/2/2, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 502, Generation: 132
...
  Logical interface ge-2/2/2.2 (Index 70) (SNMP ifIndex 536) (Generation 135)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
...
  Protocol inet, MTU: 4470, Generation: 146, Route table: 0
    Flags: Sendbcst-pkt-to-re
    Output Filters: filter_s_tcp-ge-2/2/2.2-o,,,,,
```

Verifying That the Counters Are Collected Separately by Interface

Purpose Make sure that the **count_s_tcp** counters are collected separately for the two logical interfaces.

Action Run the **show firewall** command.

```
user@host> show firewall filter filter_s_tcp
Filter: filter_s_tcp
Counters:
Name                               Bytes      Packets
count_s_tcp-at-1/1/1.0-i           420         5
count_s_tcp-ge-2/2/2.2-o          8888        101
```

- Related Documentation**
- [Interface-Specific Firewall Filter Instances Overview](#)
 - [Statement Hierarchy for Configuring Interface-Specific Firewall Filters on page 4773](#)
 - [Statement Hierarchy for Applying Interface-Specific Firewall Filters on page 4774](#)

Example: Filtering Packets Received on an Interface Group

This example shows how to configure a standard stateless firewall filter to match packets tagged for a particular interface group.

- [Requirements on page 4634](#)
- [Overview on page 4634](#)
- [Configuration on page 4634](#)
- [Verification on page 4637](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, you configure two router or switch interfaces to belong to the interface group. You also configure a stateless firewall filter that matches packets that have been tagged as received on that interface group, contain a source or destination address within a particular prefix, and contain a TCP source or destination port of a specific type and port number. The filter counts, logs, and rejects packets that match this criteria. The filter counts, logs, and rejects all other packets. By applying this firewall filter to only one of the two interfaces in the interface group, you can apply the filtering mechanism to all packets input to the two interfaces.

Topology

You configure the interface group number 1 to consist of the management port **fxp0.0** and the loopback port **lo0.0**.

You configure the firewall filter **filter_if_group** to accept only packets from interface group 1, received from or destined for IP addresses in the 192.168.80.114/32 prefix, and received from or destined for TCP port number 79.

You apply the firewall filter **filter_if_group** to the router's (or switch's) loopback interface.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

- [Configure the Stateless Firewall Filter on page 4635](#)
- [Assign Interfaces to the Interface Group on page 4635](#)
- [Apply the Firewall Filter to the Loopback Interface on page 4636](#)
- [Confirm and Commit Your Candidate Configuration on page 4636](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set firewall family inet filter filter_if_group term term1 from interface-group 1
set firewall family inet filter filter_if_group term term1 from address 192.168.80.114/32
```



```

set firewall family inet filter filter_if_group term term1 from protocol tcp
set firewall family inet filter filter_if_group term term1 from port finger
set firewall family inet filter filter_if_group term term1 then count if_group_counter1
set firewall family inet filter filter_if_group term term1 then log
set firewall family inet filter filter_if_group term term1 then reject
set firewall family inet filter filter_if_group term term2 then count if_group_counter2
set firewall family inet filter filter_if_group term term2 then log
set firewall family inet filter filter_if_group term term2 then accept
set interfaces fxp0 unit 0 family inet filter group 1
set interfaces fxp0 unit 0 family inet address 192.168.5.38/24
set interfaces lo0 unit 0 family inet filter group 1
set interfaces lo0 unit 0 family inet address 10.0.0.1/32
set interfaces lo0 unit 0 family inet address 192.168.77.1/32
set interfaces lo0 unit 0 family inet filter input filter_if_group

```

Configure the Stateless Firewall Filter

Step-by-Step Procedure

To configure the stateless firewall filter `filter_if_group`:

1. Create the stateless firewall filter `filter_if_group`.


```

[edit]
user@host# edit firewall family inet filter filter_if_group

```
2. Configure term `term1` to match packets received on interface group 1, with the source or destination address field in the 192.168.80.114/32 prefix, and with the TCP source or destination port number 79.


```

[edit firewall family inet filter filter_if_group]
user@host# set term term1 from interface-group 1
user@host# set term term1 from address 192.168.80.114/32
user@host# set term term1 from protocol tcp
user@host# set term term1 from port finger

```
3. Configure term `term1` to count, log, and reject matching packets.


```

[edit firewall family inet filter filter_if_group]
user@host# set term term1 then count if_group_counter1
user@host# set term term1 then log
user@host# set term term1 then reject

```
4. Configure the term `term2` to count, log, and accept all other packets.


```

[edit firewall family inet filter filter_if_group]
user@host# set term term2 then count if_group_counter2
user@host# set term term2 then log
user@host# set term term2 then accept

```

Assign Interfaces to the Interface Group

Step-by-Step Procedure

To assign logical interfaces to the interface group number 1 referenced by the firewall filter match term `term1`:

1. Assign the management port to interface group number 1.


```

[edit]
user@host# set interfaces fxp0 unit 0 family inet filter group 1
user@host# set interfaces fxp0 unit 0 family inet address 192.168.5.38/24

```

2. Assign a second logical interface to interface group number 1.

```
[edit]
user@host# set interfaces lo0 unit 0 family inet filter group 1
user@host# set interfaces lo0 unit 0 family inet address 10.0.0.1/32
user@host# set interfaces lo0 unit 0 family inet address 192.168.77.1/32
```

Apply the Firewall Filter to the Loopback Interface

Step-by-Step Procedure

- To apply the firewall filter to the router's (or switch's) loopback interface:

```
[edit]
user@host# set interfaces lo0 unit 0 family inet filter input filter_if_group
```

Confirm and Commit Your Candidate Configuration

Step-by-Step Procedure

To confirm and then commit your candidate configuration:

1. Confirm the configuration of the stateless firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
  filter filter_if_group {
    term term1 {
      from {
        interface-group 1;
        address {
          192.168.80.114/32;
        }
        protocol tcp;
        port finger;
      }
      then {
        count if_group_counter1;
        log;
        reject;
      }
    }
    term term2 {
      then {
        count if_group_counter2;
        log;
        accept;
      }
    }
  }
}
```

2. Confirm the configuration of the interfaces by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
```

```

user@host# show interfaces
fxp0 {
  unit 0 {
    family inet {
      filter {
        group 1;
      }
      address 192.168.5.38/24;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      filter {
        input filter_if_group;
        group 1;
      }
      address 10.0.0.1/32;
      address 192.168.77.1/32;
    }
  }
}

```

3. If you are done configuring the device, commit your candidate configuration.

```

[edit]
user@host# commit

```

Verification

To display the firewall filter counters, enter the `show firewall filter filter_if_group` operational mode command:

```
user@host> show firewall filter filter_if_group
```

```
Filter: filter_if_group
```

```
Counters:
```

Name	Bytes	Packets
if_group_counter1	0	0
if_group_counter2	6452105	82667

To display the local log of packet headers for packets evaluated by the firewall filter, enter the `show firewall log` operational mode command.

Related Documentation

- [Filtering Packets Received on a Set of Interface Groups Overview on page 4514](#)
- [Statement Hierarchy for Assigning Interfaces to Interface Groups on page 4774](#)
- [Statement Hierarchy for Configuring a Filter to Match on a Set of Interface Groups on page 4775](#)

Example: Filtering Packets Received on an Interface Set

This example shows how to configure a standard stateless firewall filter to match packets tagged for a particular interface set.

- [Requirements on page 4638](#)
- [Overview on page 4638](#)
- [Configuration on page 4639](#)
- [Verification on page 4643](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, you apply a stateless firewall filter to the input of the router or switch loopback interface. The firewall filter includes a term that matches packets tagged for a particular interface set.

Topology

You create the firewall filter **L2_filter** to apply rate limits to the protocol-independent traffic received on the following interfaces:

- fe-0/0/0.0
- fe-1/0/0.0
- fe-1/1/0.0



NOTE: The interface type in this topic is just an example. The fe- interface type is not supported by EX Series switches.

First, for protocol-independent traffic received on **fe-0/0/0.0**, the firewall filter term **t1** applies policer **p1**.

For protocol-independent traffic received on any other Fast Ethernet interfaces, firewall filter term **t2** applies policer **p2**. To define an interface set that consists of all Fast Ethernet interfaces, you include the **interface-set interface-set-name interface-name** statement at the **[edit firewall]** hierarchy level. To define a packet-matching criteria based on the interface on which a packet arrives to a specified interface set, you configure a term that uses the **interface-set** firewall filter match condition.

Finally, for any other protocol-independent traffic, firewall filter term **t3** applies policer **p3**.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configuring the Interfaces for Which the Stateless Firewall Filter Terms Take Rate-Limiting Actions on page 4639](#)
- [Configuring the Stateless Firewall Filter That Rate-Limits Protocol-Independent Traffic Based on the Interfaces on Which Packets Arrive on page 4640](#)
- [Applying the Stateless Firewall Filter to the Routing Engine Input Interface on page 4643](#)

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces fe-0/0/0 unit 0 family inet address 10.1.1.1/30
set interfaces fe-1/0/0 unit 0 family inet address 10.2.2.1/30
set interfaces fe-1/1/0 unit 0 family inet address 10.4.4.1/30
set firewall policer p1 if-exceeding bandwidth-limit 5m
set firewall policer p1 if-exceeding burst-size-limit 10m
set firewall policer p1 then discard
set firewall policer p2 if-exceeding bandwidth-limit 40m
set firewall policer p2 if-exceeding burst-size-limit 100m
set firewall policer p2 then discard
set firewall policer p3 if-exceeding bandwidth-limit 600m
set firewall policer p3 if-exceeding burst-size-limit 1g
set firewall policer p3 then discard
set firewall interface-set ifset fe-*
set firewall family any filter L2_filter term t1 from interface fe-0/0/0.0
set firewall family any filter L2_filter term t1 then count c1
set firewall family any filter L2_filter term t1 then policer p1
set firewall family any filter L2_filter term t2 from interface-set ifset
set firewall family any filter L2_filter term t2 then count c2
set firewall family any filter L2_filter term t2 then policer p2
set firewall family any filter L2_filter term t3 then count c3
set firewall family any filter L2_filter term t3 then policer p3
set interfaces lo0 unit 0 family inet address 1.1.1.157/30
set interfaces lo0 unit 0 filter input L2_filter
```

Configuring the Interfaces for Which the Stateless Firewall Filter Terms Take Rate-Limiting Actions

Step-by-Step Procedure

To configure the interfaces for which the stateless firewall filter terms take rate-limiting actions:

1. Configure the logical interface whose input traffic will be matched by the first term of the firewall filter.

[edit]

user@host# **set interfaces fe-0/0/0 unit 0 family inet address 10.1.1.1/30**

2. Configure the logical interfaces whose input traffic will be matched by the second term of the firewall filter.

```
[edit ]
user@host# set interfaces fe-1/0/0 unit 0 family inet address 10.2.2.1/30
user@host# set interfaces fe-1/1/0 unit 0 family inet address 10.4.4.1/30
```

3. If you are done configuring the device, commit the configuration.

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

Results Confirm the configuration of the router (or switch) transit interfaces by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
fe-0/0/0 {
  unit 0 {
    family inet {
      address 10.1.1.1/30;
    }
  }
}
fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.2.2.1/30;
    }
  }
}
fe-1/1/0 {
  unit 0 {
    family inet {
      address 10.4.4.1/30;
    }
  }
}
```

Configuring the Stateless Firewall Filter That Rate-Limits Protocol-Independent Traffic Based on the Interfaces on Which Packets Arrive

Step-by-Step Procedure To configure the standard stateless firewall **L2_filter** that uses policers (**p1**, **p2**, and **p3**) to rate-limit protocol-independent traffic based on the interfaces on which the packets arrive:

1. Configure the firewall statements.

```
[edit]
user@host# edit firewall
```

2. Configure the policer **p1** to discard traffic that exceeds a traffic rate of **5m** bps or a burst size of **10m** bytes.

```
[edit firewall]
```

```

user@host# set policer p1 if-exceeding bandwidth-limit 5m
user@host# set policer p1 if-exceeding burst-size-limit 10m
user@host# set policer p1 then discard

```

3. Configure the policer **p2** to discard traffic that exceeds a traffic rate of **40m** bps or a burst size of **100m** bytes .

```

[edit firewall]
user@host# set policer p2 if-exceeding bandwidth-limit 40m
user@host# set policer p2 if-exceeding burst-size-limit 100m
user@host# set policer p2 then discard

```

4. Configure the policer **p3** to discard traffic that exceeds a traffic rate of **600m** bps or a burst size of **1g** bytes.

```

[edit firewall]
user@host# set policer p3 if-exceeding bandwidth-limit 600m
user@host# set policer p3 if-exceeding burst-size-limit 1g
user@host# set policer p3 then discard

```

5. Define the interface set **ifset** to be the group of all Fast Ethernet interfaces on the router.

```

[edit firewall]
user@host# set interface-set ifset fe-*

```

6. Create the stateless firewall filter **L2_filter**.

```

[edit firewall]
user@host# edit family any filter L2_filter

```

7. Configure filter term **t1** to match IPv4, IPv6, or MPLS packets received on interface **fe-0/0/0.0** and use policer **p1** to rate-limit that traffic.

```

[edit firewall family any filter L2_filter]
user@host# set term t1 from interface fe-0/0/0.0
user@host# set term t1 then count c1
user@host# set term t1 then policer p1

```

8. Configure filter term **t2** to match packets received on interface-set **ifset** and use policer **p2** to rate-limit that traffic.

```

[edit firewall family any filter L2_filter]
user@host# set term t2 from interface-set ifset
user@host# set term t2 then count c2
user@host# set term t2 then policer p2

```

9. Configure filter term **t3** to use policer **p3** to rate-limit all other traffic.

```

[edit firewall family any filter L2_filter]
user@host# set term t3 then count c3
user@host# set term t3 then policer p3

```

10. If you are done configuring the device, commit the configuration.

```

[edit]
user@host# commit

```

Results Confirm the configuration of the stateless firewall filter and the policers referenced as firewall filter actions by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
family any {
  filter L2_filter {
    term t1 {
      from {
        interface fe-0/0/0.0;
      }
      then {
        policer p1;
        count c1;
      }
    }
    term t2 {
      from {
        interface-set ifset;
      }
      then {
        policer p2;
        count c2;
      }
    }
    term t3 {
      then {
        policer p3;
        count c3;
      }
    }
  }
}
policer p1 {
  if-exceeding {
    bandwidth-limit 5m;
    burst-size-limit 10m;
  }
  then discard;
}
policer p2 {
  if-exceeding {
    bandwidth-limit 40m;
    burst-size-limit 100m;
  }
  then discard;
}
policer p3 {
  if-exceeding {
    bandwidth-limit 600m;
    burst-size-limit 1g;
  }
  then discard;
}
```



```

interface-set ifset {
  fe-*;
}

```

Applying the Stateless Firewall Filter to the Routing Engine Input Interface

Step-by-Step Procedure

To apply the stateless firewall filter to the Routing Engine input interface:

1. Apply the stateless firewall filter to the Routing Engine interface in the input direction.

```

[edit]
user@host# set interfaces lo0 unit 0 family inet address 1.1.1.157/30
user@host# set interfaces lo0 unit 0 filter input L2_filter

```

2. If you are done configuring the device, commit the configuration.

```

[edit]
user@host# commit

```

Results Confirm the application of the firewall filter to the Routing Engine input interface by entering the **show interfaces** command again. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```

user@host# show interfaces
fe-0/0/0 {
  ...
}
fe-1/0/0 {
  ...
}
fe-1/1/0 {
  ...
}
lo0 {
  unit 0 {
    filter {
      input L2_filter;
    }
    family inet {
      address 1.1.1.157/30;
    }
  }
}

```

Verification

To confirm that the configuration is working properly, use the **show firewall filter L2_filter** operational mode command to monitor traffic statistics about the firewall filter and three counters.

Related Documentation

- [Understanding How to Use Standard Firewall Filters on page 3293](#)
- [Filtering Packets Received on an Interface Set Overview on page 4515](#)
- [Statement Hierarchy for Defining an Interface Set on page 4777](#)

- [Statement Hierarchy for Configuring a Filter to Match on an Interface Set on page 4777](#)

Example: Configuring Filter-Based Forwarding on the Source Address

This example shows how to configure filter-based forwarding. The filter classifies packets to determine their forwarding path within the ingress routing device.

- [Requirements on page 4644](#)
- [Overview on page 4644](#)
- [Configuration on page 4646](#)
- [Verification on page 4651](#)

Requirements

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

Overview

Filter-based forwarding is supported for IP version 4 (IPv4) and IP version 6 (IPv6).

Use filter-based forwarding for service provider selection when customers have Internet connectivity provided by different ISPs yet share a common access layer. When a shared media (such as a cable modem) is used, a mechanism on the common access layer looks at Layer 2 or Layer 3 addresses and distinguishes between customers. You can use filter-based forwarding when the common access layer is implemented using a combination of Layer 2 switches and a single router.

With filter-based forwarding, all packets received on an interface are considered. Each packet passes through a filter that has match conditions. If the match conditions are met for a filter and you have created a routing instance, filter-based forwarding is applied to a packet. The packet is forwarded based on the next hop specified in the routing instance. For static routes, the next hop can be a specific LSP.



NOTE: Source-class usage filter matching and unicast reverse-path forwarding checks are not supported on an interface configured with filter-based forwarding (FBF).

To configure filter-based forwarding, perform the following tasks:

- Create a match filter on an ingress router or switch. To specify a match filter, include the **filter filter-name** statement at the **[edit firewall]** hierarchy level. A packet that passes through the filter is compared against a set of rules to classify it and to determine its membership in a set. Once classified, the packet is forwarded to a routing table specified in the accept action in the filter description language. The routing table then forwards the packet to the next hop that corresponds to the destination address entry in the table.
- Create routing instances that specify the routing table(s) to which a packet is forwarded, and the destination to which the packet is forwarded at the **[edit routing-instances]** hierarchy level. For example:

```
[edit]
routing-instances {
  routing-table-name1 {
    instance-type forwarding;
    routing-options {
      static {
        route 0.0.0.0/0 nexthop 10.0.0.1;
      }
    }
  }
  routing-table-name2 {
    instance-type forwarding;
    routing-options {
      static {
        route 0.0.0.0/0 nexthop 10.0.0.2;
      }
    }
  }
}
```

- Create a routing table group that adds interface routes to the forwarding routing instances used in filter-based forwarding (FBF), as well as to the default routing instance inet.0. This part of the configuration resolves the routes installed in the routing instances to directly connected next hops on that interface. Create the routing table group at the **[edit routing-options]** hierarchy level.




NOTE: Specify inet.0 as one of the routing instances that the interface routes are imported into. If the default instance inet.0 is not specified, interface routes are not imported into the default routing instance.

This example shows a packet filter that directs customer traffic to a next-hop router in the domains, SP 1 or SP 2, based on the packet's source address.

If the packet has a source address assigned to an SP 1 customer, destination-based forwarding occurs using the sp1-route-table.inet.0 routing table. If the packet has a source address assigned to an SP 2 customer, destination-based forwarding occurs using the sp2-route-table.inet.0 routing table. If a packet does not match either of these conditions, the filter accepts the packet, and destination-based forwarding occurs using the standard inet.0 routing table.

Figure 96 on page 4646 shows the topology used in this example.

On Device P1, an input filter classifies packets received from Device PE3 and Device PE4. The packets are routed based on the source addresses. Packets with source addresses in the 10.1.1.0/24 and 10.1.2.0/24 networks are routed to Device PE1. Packets with source addresses in the 10.2.1.0/24 and 10.2.2.0/24 networks are routed to Device PE2.

Figure 96: Filter-Based Forwarding

To establish connectivity, OSPF is configured on all of the interfaces. For demonstration purposes, loopback interface addresses are configured on the routing devices to represent networks in the clouds.

The [“CLI Quick Configuration” on page 4646](#) section shows the entire configuration for all of the devices in the topology. The [“Configuring the Routing Instances on the Device P1” on page 4648](#) section shows the step-by-step configuration of the ingress routing device, Device P1.

Configuration

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

```
Device P1
set firewall filter classify-customers term sp1-customers from source-address 10.1.1.0/24
set firewall filter classify-customers term sp1-customers from source-address 10.1.2.0/24
set firewall filter classify-customers term sp1-customers then log
set firewall filter classify-customers term sp1-customers then routing-instance
  sp1-route-table
set firewall filter classify-customers term sp2-customers from source-address 10.2.1.0/24
set firewall filter classify-customers term sp2-customers from source-address 10.2.2.0/24
set firewall filter classify-customers term sp2-customers then log
set firewall filter classify-customers term sp2-customers then routing-instance
  sp2-route-table
set firewall filter classify-customers term default then accept
set interfaces fe-1/2/0 unit 0 family inet filter input classify-customers
set interfaces fe-1/2/0 unit 0 family inet address 172.16.0.10/30
set interfaces fe-1/2/1 unit 0 family inet address 172.16.0.13/30
set interfaces fe-1/2/2 unit 0 family inet address 172.16.0.17/30
set protocols ospf rib-group bbf-group
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set routing-instances sp1-route-table instance-type forwarding
set routing-instances sp1-route-table routing-options static route 0.0.0.0/0 next-hop
  172.16.0.13
set routing-instances sp2-route-table instance-type forwarding
```

```

set routing-instances sp2-route-table routing-options static route 0.0.0.0/0 next-hop
  172.16.0.17
set routing-options rib-groups fbf-group import-rib inet.0
set routing-options rib-groups fbf-group import-rib sp1-route-table.inet.0
set routing-options rib-groups fbf-group import-rib sp2-route-table.inet.0

```

Device P2

```

set interfaces fe-1/2/0 unit 0 family inet address 172.16.0.2/30
set interfaces fe-1/2/1 unit 0 family inet address 172.16.0.6/30
set interfaces fe-1/2/2 unit 0 family inet address 172.16.0.9/30
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

Device PE1

```

set interfaces fe-1/2/0 unit 0 family inet address 172.16.0.14/30
set interfaces lo0 unit 0 family inet address 1.1.1.1/32
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

Device PE2

```

set interfaces fe-1/2/0 unit 0 family inet address 172.16.0.18/30
set interfaces lo0 unit 0 family inet address 2.2.2.2/32
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

Device PE3

```

set interfaces fe-1/2/0 unit 0 family inet address 172.16.0.1/30
set interfaces lo0 unit 0 family inet address 10.1.1.1/32
set interfaces lo0 unit 0 family inet address 10.1.2.1/32
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

Device PE4

```

set interfaces fe-1/2/0 unit 0 family inet address 172.16.0.5/30
set interfaces lo0 unit 0 family inet address 10.2.1.1/32
set interfaces lo0 unit 0 family inet address 10.2.2.1/32
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

Configuring the Firewall Filter

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure the firewall filter on the main router or switch:

1. Configure the source addresses for SP1 customers.

```

[edit firewall filter classify-customers term sp1-customers]
user@host# set from source-address 10.1.1.0/24
user@host# set from source-address 10.1.2.0/24

```
2. Configure the actions that are taken when packets are received with the specified source addresses.

To track the action of the firewall filter, a log action is configured. The sp1-route-table.inet.0 routing table on Device P1 routes the packets.

```

[edit firewall filter classify-customers term sp1-customers]
user@host# set then log

```

```
user@host# set then routing-instance sp1-route-table
```

3. Configure the source addresses for SP2 customers.

```
[edit firewall filter classify-customers term sp2-customers]
```

```
user@host# set from source-address 10.2.1.0/24
```

```
user@host# set from source-address 10.2.2.0/24
```

4. Configure the actions that are taken when packets are received with the specified source addresses.

To track the action of the firewall filter, a log action is configured. The sp2-route-table.inet.0 routing table on Device P1 routes the packet.

```
[edit firewall filter classify-customers term sp2-customers]
```

```
user@host# set then log
```

```
user@host# set then routing-instance sp2-route-table
```

5. Configure the action to take when packets are received from any other source address.

All of these packets are simply accepted and routed using the default IPv4 unicast routing table, inet.0.

```
[edit firewall filter classify-customers term default]
```

```
user@host# set then accept
```

Configuring the Routing Instances on the Device P1

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure the routing instances:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0]
```

```
user@host# set unit 0 family inet address 172.16.0.10/30
```

```
[edit interfaces fe-1/2/1]
```

```
user@host# set unit 0 family inet address 172.16.0.13/30
```

```
[edit interfaces fe-1/2/2]
```

```
user@host# set unit 0 family inet address 172.16.0.17/30
```

2. Assign the **classify-customers** firewall filter to router interface fe-1/2/0.0 as an input packet filter.

```
[edit interfaces fe-1/2/0]
```

```
user@host# set unit 0 family inet filter input classify-customers
```

3. Configure connectivity, using either a routing protocol or static routing.

As a best practice, disable routing on the management interface.

```
[edit protocols ospf area 0.0.0.0]
```

```
user@host# set interface all
```

```
user@host# set interface fxp0.0 disable
```

4. Create the routing instances.

These routing instances are referenced in the **classify-customers** firewall filter.

The forwarding instance type provides support for filter-based forwarding, where interfaces are not associated with instances. All interfaces belong to the default instance, in this case Device P1.

```
[edit routing-instances]
user@host# set sp1-route-table instance-type forwarding
```

```
user@host# set sp2-route-table instance-type forwarding
```

5. Resolve the routes installed in the routing instances to directly connected next hops.

```
[edit routing-instances sp1-route-table routing-options]
user@host# set static route 0.0.0.0/0 next-hop 172.16.0.13
```

```
user@host# set static route 0.0.0.0/0 next-hop 172.16.0.17
```

6. Group together the routing tables to form a routing table group.

The first routing table, inet.0, is the primary routing table, and the additional routing tables are the secondary routing tables.

The primary routing table determines the address family of the routing table group, in this case IPv4.

```
[edit routing-options]
user@host# set rib-groups fbf-group import-rib inet.0
user@host# set rib-groups fbf-group import-rib sp1-route-table.inet.0
user@host# set rib-groups fbf-group import-rib sp2-route-table.inet.0
```

7. Apply the routing table group to OSPF.

This causes the OSPF routes to be installed into all the routing tables in the group.

```
[edit protocols ospf]
user@host# set rib-group fbf-group
```

8. If you are done configuring the device, commit the configuration.

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

Results

Confirm your configuration by issuing the **show interfaces**, **show firewall**, **show protocols**, **show routing-instances**, and **show routing-options** commands.

```
user@host# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet {
      filter {
        input classify-customers;
      }
      address 172.16.0.10/30;
```

```
    }
  }
}
fe-1/2/1 {
  unit 0 {
    family inet {
      address 172.16.0.13/30;
    }
  }
}
fe-1/2/2 {
  unit 0 {
    family inet {
      address 172.16.0.17/30;
    }
  }
}
}

user@host# show firewall
filter classify-customers {
  term sp1-customers {
    from {
      source-address {
        10.1.1.0/24;
        10.1.2.0/24;
      }
    }
    then {
      log;
      routing-instance sp1-route-table;
    }
  }
  term sp2-customers {
    from {
      source-address {
        10.2.1.0/24;
        10.2.2.0/24;
      }
    }
    then {
      log;
      routing-instance sp2-route-table;
    }
  }
  term default {
    then accept;
  }
}

user@host# show protocols
ospf {
  rib-group fbf-group;
  area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}
```



```

    }
  }
user@host# show routing-instances
sp1-route-table {
  instance-type forwarding;
  routing-options {
    static {
      route 0.0.0.0/0 next-hop 172.16.0.13;
    }
  }
}
sp2-route-table {
  instance-type forwarding;
  routing-options {
    static {
      route 0.0.0.0/0 next-hop 172.16.0.17;
    }
  }
}

user@host# show routing-options
rib-groups {
  fbf-group {
    import-rib [ inet.0 sp1-route-table.inet.0 sp2-route-table.inet.0 ];
  }
}

```

Verification

Confirm that the configuration is working properly.

Pinging with Specified Source Addresses

Purpose Send some ICMP packets across the network to test the firewall filter.

Action 1. Run the **ping** command, pinging the lo0.0 interface on Device PE1.

The address configured on this interface is 1.1.1.1.

Specify the source address 10.1.2.1, which is the address configured on the lo0.0 interface on Device PE3.

```

user@PE3> ping 1.1.1.1 source 10.1.2.1
PING 1.1.1.1 (1.1.1.1): 56 data bytes
64 bytes from 1.1.1.1: icmp_seq=0 ttl=62 time=1.444 ms
64 bytes from 1.1.1.1: icmp_seq=1 ttl=62 time=2.094 ms
^C
--- 1.1.1.1 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.444/1.769/2.094/0.325 ms

```

2. Run the **ping** command, pinging the lo0.0 interface on Device PE2.

The address configured on this interface is 2.2.2.2.

Specify the source address 10.2.1.1, which is the address configured on the lo0.0 interface on Device PE4.

```

user@PE4> ping 2.2.2.2 source 10.2.1.1

```

```

PING 2.2.2.2 (2.2.2.2): 56 data bytes
64 bytes from 2.2.2.2: icmp_seq=0 ttl=62 time=1.473 ms
64 bytes from 2.2.2.2: icmp_seq=1 ttl=62 time=1.407 ms
^C
--- 2.2.2.2 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.407/1.440/1.473/0.033 ms

```

Meaning Sending these pings activates the firewall filter actions.

Verifying the Firewall Filter

Purpose Make sure the firewall filter actions take effect.

Action 1. Run the **show firewall log** command on Device P1.

```

user@P1> show firewall log
Log :
Time      Filter  Action Interface  Protocol  Src Addr
Dest Addr
13:52:20  pfe      A      fe-1/2/0.0    ICMP      10.2.1.1
2.2.2.2
13:52:19  pfe      A      fe-1/2/0.0    ICMP      10.2.1.1
2.2.2.2
13:51:53  pfe      A      fe-1/2/0.0    ICMP      10.1.2.1
1.1.1.1
13:51:52  pfe      A      fe-1/2/0.0    ICMP      10.1.2.1
1.1.1.1

```

- Related Documentation**
- *Example: Configuring Multitopology Routing Based on Applications*
 - *Configuring Filter-Based Forwarding*
 - *Copying and Redirecting Traffic with Port Mirroring and Filter-Based Forwarding*
 - *Using Filter-Based Forwarding to Export Monitored Traffic to Multiple Destinations*
 - [Filter-Based Forwarding Overview on page 4515](#)

Example: Configuring Filter-Based Forwarding on Logical Systems

This example shows how to configure filter-based forwarding within a logical system. The filter classifies packets to determine their forwarding path within the ingress routing device.

- [Requirements on page 4652](#)
- [Overview on page 4653](#)
- [Configuration on page 4655](#)
- [Verification on page 4661](#)

Requirements

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

Overview

Filter-based forwarding is supported for IP version 4 (IPv4) and IP version 6 (IPv6).

Use filter-based forwarding for service provider selection when customers have Internet connectivity provided by different ISPs yet share a common access layer. When a shared media (such as a cable modem) is used, a mechanism on the common access layer looks at Layer 2 or Layer 3 addresses and distinguishes between customers. You can use filter-based forwarding when the common access layer is implemented using a combination of Layer 2 switches and a single router.

With filter-based forwarding, all packets received on an interface are considered. Each packet passes through a filter that has match conditions. If the match conditions are met for a filter and you have created a routing instance, filter-based forwarding is applied to a packet. The packet is forwarded based on the next hop specified in the routing instance. For static routes, the next hop can be a specific LSP.



NOTE: Source-class usage filter matching and unicast reverse-path forwarding checks are not supported on an interface configured with filter-based forwarding (FBF).

To configure filter-based forwarding, perform the following tasks:

- Create a match filter on an ingress router or switch. To specify a match filter, include the **filter filter-name** statement at the **[edit firewall]** hierarchy level. A packet that passes through the filter is compared against a set of rules to classify it and to determine its membership in a set. Once classified, the packet is forwarded to a routing table specified in the accept action in the filter description language. The routing table then forwards the packet to the next hop that corresponds to the destination address entry in the table.
- Create routing instances that specify the routing table(s) to which a packet is forwarded, and the destination to which the packet is forwarded at the **[edit routing-instances]** or **[edit logical-systems logical-system-name routing-instances]** hierarchy level. For example:

```
[edit]
routing-instances {
  routing-table-name1 {
    instance-type forwarding;
    routing-options {
      static {
        route 0.0.0.0/0 nexthop 10.0.0.1;
      }
    }
  }
  routing-table-name2 {
    instance-type forwarding;
    routing-options {
      static {
        route 0.0.0.0/0 nexthop 10.0.0.2;
      }
    }
  }
}
```

```

    }
  }
}

```

- Create a routing table group that adds interface routes to the forwarding routing instances used in filter-based forwarding (FBF), as well as to the default routing instance **inet.0**. This part of the configuration resolves the routes installed in the routing instances to directly connected next hops on that interface. Create the routing table group at the **[edit routing-options]** or **[edit logical-systems *logical-system-name* routing-options]** hierarchy level.



NOTE: Specify **inet.0** as one of the routing instances that the interface routes are imported into. If the default instance **inet.0** is not specified, interface routes are not imported into the default routing instance.

This example shows a packet filter that directs customer traffic to a next-hop router in the domains, SP 1 or SP 2, based on the packet's source address.

If the packet has a source address assigned to an SP 1 customer, destination-based forwarding occurs using the **sp1-route-table.inet.0** routing table. If the packet has a source address assigned to an SP 2 customer, destination-based forwarding occurs using the **sp2-route-table.inet.0** routing table. If a packet does not match either of these conditions, the filter accepts the packet, and destination-based forwarding occurs using the standard **inet.0** routing table.

One way to make filter-based forwarding work within a logical system is to configure the firewall filter on the logical system that receives the packets. Another way is to configure the firewall filter on the main router or switch and then reference the logical system in the firewall filter. This example uses the second approach. The specific routing instances are configured within the logical system. Because each routing instance has its own routing table, you have to reference the routing instances in the firewall filter, as well. The syntax looks as follows:

```


[edit firewall filter filter-name term term-name]
user@host# set then logical-system logical-system-name routing-instance
routing-instance-name

```

Figure 59 on page 3300 shows the topology used in this example.

On Logical System P1, an input filter classifies packets received from Logical System PE3 and Logical System PE4. The packets are routed based on the source addresses. Packets with source addresses in the 10.1.1.0/24 and 10.1.2.0/24 networks are routed to Logical System PE1. Packets with source addresses in the 10.2.1.0/24 and 10.2.2.0/24 networks are routed to Logical System PE2.

Figure 97: Logical Systems with Filter-Based Forwarding



To establish connectivity, OSPF is configured on all of the interfaces. For demonstration purposes, loopback interface addresses are configured on the routing devices to represent networks in the clouds.

The [“CLI Quick Configuration” on page 3300](#) section shows the entire configuration for all of the devices in the topology. The [“Configuring the Routing Instances on the Logical System P1” on page 3303](#) and [“Configuring the Firewall Filter on the Main Router” on page 3302](#) sections shows the step-by-step configuration of the ingress routing device, Logical System P1.

Configuration

CLI Quick Configuration

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

```
set firewall filter classify-customers term sp1-customers from source-address 10.1.1.0/24
set firewall filter classify-customers term sp1-customers from source-address 10.1.2.0/24
set firewall filter classify-customers term sp1-customers then log
set firewall filter classify-customers term sp1-customers then logical-system P1
  routing-instance sp1-route-table
set firewall filter classify-customers term sp2-customers from source-address 10.2.1.0/24
set firewall filter classify-customers term sp2-customers from source-address 10.2.2.0/24
set firewall filter classify-customers term sp2-customers then log
set firewall filter classify-customers term sp2-customers then logical-system P1
  routing-instance sp2-route-table
set firewall filter classify-customers term default then accept
set logical-systems P1 interfaces lt-1/2/0 unit 10 encapsulation ethernet
set logical-systems P1 interfaces lt-1/2/0 unit 10 peer-unit 9
set logical-systems P1 interfaces lt-1/2/0 unit 10 family inet filter input classify-customers
set logical-systems P1 interfaces lt-1/2/0 unit 10 family inet address 172.16.0.10/30
set logical-systems P1 interfaces lt-1/2/0 unit 13 encapsulation ethernet
set logical-systems P1 interfaces lt-1/2/0 unit 13 peer-unit 14
set logical-systems P1 interfaces lt-1/2/0 unit 13 family inet address 172.16.0.13/30
set logical-systems P1 interfaces lt-1/2/0 unit 17 encapsulation ethernet
set logical-systems P1 interfaces lt-1/2/0 unit 17 peer-unit 18
set logical-systems P1 interfaces lt-1/2/0 unit 17 family inet address 172.16.0.17/30
set logical-systems P1 protocols ospf rib-group fbf-group
```

```
set logical-systems P1 protocols ospf area 0.0.0.0 interface all
set logical-systems P1 protocols ospf area 0.0.0.0 interface fxp0.0 disable
set logical-systems P1 routing-instances sp1-route-table instance-type forwarding
set logical-systems P1 routing-instances sp1-route-table routing-options static route
  0.0.0.0/0 next-hop 172.16.0.13
set logical-systems P1 routing-instances sp2-route-table instance-type forwarding
set logical-systems P1 routing-instances sp2-route-table routing-options static route
  0.0.0.0/0 next-hop 172.16.0.17
set logical-systems P1 routing-options rib-groups fbf-group import-rib inet.0
set logical-systems P1 routing-options rib-groups fbf-group import-rib
  sp1-route-table.inet.0
set logical-systems P1 routing-options rib-groups fbf-group import-rib
  sp2-route-table.inet.0
set logical-systems P2 interfaces lt-1/2/0 unit 2 encapsulation ethernet
set logical-systems P2 interfaces lt-1/2/0 unit 2 peer-unit 1
set logical-systems P2 interfaces lt-1/2/0 unit 2 family inet address 172.16.0.2/30
set logical-systems P2 interfaces lt-1/2/0 unit 6 encapsulation ethernet
set logical-systems P2 interfaces lt-1/2/0 unit 6 peer-unit 5
set logical-systems P2 interfaces lt-1/2/0 unit 6 family inet address 172.16.0.6/30
set logical-systems P2 interfaces lt-1/2/0 unit 9 encapsulation ethernet
set logical-systems P2 interfaces lt-1/2/0 unit 9 peer-unit 10
set logical-systems P2 interfaces lt-1/2/0 unit 9 family inet address 172.16.0.9/30
set logical-systems P2 protocols ospf area 0.0.0.0 interface all
set logical-systems P2 protocols ospf area 0.0.0.0 interface fxp0.0 disable
set logical-systems PE1 interfaces lt-1/2/0 unit 14 encapsulation ethernet
set logical-systems PE1 interfaces lt-1/2/0 unit 14 peer-unit 13
set logical-systems PE1 interfaces lt-1/2/0 unit 14 family inet address 172.16.0.14/30
set logical-systems PE1 interfaces lo0 unit 3 family inet address 1.1.1.1/32
set logical-systems PE1 protocols ospf area 0.0.0.0 interface all
set logical-systems PE1 protocols ospf area 0.0.0.0 interface fxp0.0 disable
set logical-systems PE2 interfaces lt-1/2/0 unit 18 encapsulation ethernet
set logical-systems PE2 interfaces lt-1/2/0 unit 18 peer-unit 17
set logical-systems PE2 interfaces lt-1/2/0 unit 18 family inet address 172.16.0.18/30
set logical-systems PE2 interfaces lo0 unit 4 family inet address 2.2.2.2/32
set logical-systems PE2 protocols ospf area 0.0.0.0 interface all
set logical-systems PE2 protocols ospf area 0.0.0.0 interface fxp0.0 disable
set logical-systems PE3 interfaces lt-1/2/0 unit 1 encapsulation ethernet
set logical-systems PE3 interfaces lt-1/2/0 unit 1 peer-unit 2
set logical-systems PE3 interfaces lt-1/2/0 unit 1 family inet address 172.16.0.1/30
set logical-systems PE3 interfaces lo0 unit 1 family inet address 10.1.1.1/32
set logical-systems PE3 interfaces lo0 unit 1 family inet address 10.1.2.1/32
set logical-systems PE3 protocols ospf area 0.0.0.0 interface all
set logical-systems PE3 protocols ospf area 0.0.0.0 interface fxp0.0 disable
set logical-systems PE4 interfaces lt-1/2/0 unit 5 encapsulation ethernet
set logical-systems PE4 interfaces lt-1/2/0 unit 5 peer-unit 6
set logical-systems PE4 interfaces lt-1/2/0 unit 5 family inet address 172.16.0.5/30
set logical-systems PE4 interfaces lo0 unit 2 family inet address 10.2.1.1/32
set logical-systems PE4 interfaces lo0 unit 2 family inet address 10.2.2.1/32
set logical-systems PE4 protocols ospf area 0.0.0.0 interface all
set logical-systems PE4 protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

Configuring the Firewall Filter on the Main Router

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure the firewall filter on the main router or switch:

1. Configure the source addresses for SP1 customers.

```
[edit firewall filter classify-customers term sp1-customers]
user@host# set from source-address 10.1.1.0/24
user@host# set from source-address 10.1.2.0/24
```

2. Configure the actions that are taken when packets are received with the specified source addresses.

To track the action of the firewall filter, a log action is configured. The sp1-route-table.inet.0 routing table on Logical System P1 routes the packets.

```
[edit firewall filter classify-customers term sp1-customers]
user@host# set then log
user@host# set then logical-system P1 routing-instance sp1-route-table
```

3. Configure the source addresses for SP2 customers.

```
[edit firewall filter classify-customers term sp2-customers]
user@host# set from source-address 10.2.1.0/24
user@host# set from source-address 10.2.2.0/24
```

4. Configure the actions that are taken when packets are received with the specified source addresses.

To track the action of the firewall filter, a log action is configured. The sp2-route-table.inet.0 routing table on Logical System P1 routes the packet.

```
[edit firewall filter classify-customers term sp2-customers]
user@host# set then log
user@host# set then logical-system P1 routing-instance sp2-route-table
```

5. Configure the action to take when packets are received from any other source address.

All of these packets are simply accepted and routed using the default IPv4 unicast routing table, inet.0.

```
[edit firewall filter classify-customers term default]
user@host# set then accept
```

Configuring the Routing Instances on the Logical System P1

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure the routing instances on a logical system:

1. Configure the interfaces on the logical system.

```
[edit logical-systems P1 interfaces lt-1/2/0]
user@host# set unit 10 encapsulation ethernet
user@host# set unit 10 peer-unit 9
user@host# set unit 10 family inet address 172.16.0.10/30
```

```
user@host# set unit 13 encapsulation ethernet
user@host# set unit 13 peer-unit 14
user@host# set unit 13 family inet address 172.16.0.13/30
```

```
user@host# set unit 17 encapsulation ethernet
user@host# set unit 17 peer-unit 18
user@host# set unit 17 family inet address 172.16.0.17/30
```

2. Assign the **classify-customers** firewall filter to router interface lt-1/2/0.10 as an input packet filter.

```
[edit logical-systems P1 interfaces lt-1/2/0]
user@host# set unit 10 family inet filter input classify-customers
```

3. Configure connectivity, using either a routing protocol or static routing.

As a best practice, disable routing on the management interface.

```
[edit logical-systems P1 protocols ospf area 0.0.0.0]
user@host# set interface all
user@host# set interface fxp0.0 disable
```

4. Create the routing instances.

These routing instances are referenced in the **classify-customers** firewall filter.

The forwarding instance type provides support for filter-based forwarding, where interfaces are not associated with instances. All interfaces belong to the default instance, in this case Logical System P1.

```
[edit logical-systems P1 routing-instances]
user@host# set sp1-route-table instance-type forwarding
```

```
user@host# set sp2-route-table instance-type forwarding
```

5. Resolve the routes installed in the routing instances to directly connected next hops.

```
[edit logical-systems P1 routing-instances]
user@host# set sp1-route-table routing-options static route 0.0.0.0/0 next-hop
172.16.0.13
```



```
user@host# set sp2-route-table routing-options static route 0.0.0.0/0 next-hop
172.16.0.17
```

6. Group together the routing tables to form a routing table group.

The first routing table, inet.0, is the primary routing table, and the additional routing tables are the secondary routing tables.

The primary routing table determines the address family of the routing table group, in this case IPv4.

```
[edit logical-systems P1 routing-options]
user@host# set rib-groups fbf-group import-rib inet.0
user@host# set rib-groups fbf-group import-rib sp1-route-table.inet.0
user@host# set rib-groups fbf-group import-rib sp2-route-table.inet.0
```

7. Apply the routing table group to OSPF.

This causes the OSPF routes to be installed into all the routing tables in the group.

```
[edit logical-systems P1 protocols ospf]
user@host# set rib-group fbf-group
```

8. If you are done configuring the device, commit the configuration.

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

Results

Confirm your configuration by issuing the **show firewall** and **show logical-systems P1** commands.

```
user@host# show firewall
filter classify-customers {
  term sp1-customers {
    from {
      source-address {
        10.1.1.0/24;
        10.1.2.0/24;
      }
    }
    then {
      log;
      logical-system P1 routing-instance sp1-route-table;
    }
  }
}
term sp2-customers {
  from {
    source-address {
      10.2.1.0/24;
      10.2.2.0/24;
    }
  }
  then {
    log;
    logical-system P1 routing-instance sp2-route-table;
  }
}
```

```
}
term default {
  then accept;
}
}

user@host# show logical-systems P1
interfaces {
  lt-1/2/0 {
    unit 10 {
      encapsulation ethernet;
      peer-unit 9;
      family inet {
        filter {
          input classify-customers;
        }
        address 172.16.0.10/30;
      }
    }
    unit 13 {
      encapsulation ethernet;
      peer-unit 14;
      family inet {
        address 172.16.0.13/30;
      }
    }
    unit 17 {
      encapsulation ethernet;
      peer-unit 18;
      family inet {
        address 172.16.0.17/30;
      }
    }
  }
}
protocols {
  ospf {
    rib-group fbf-group;
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
  }
}
routing-instances {
  sp1-route-table {
    instance-type forwarding;
    routing-options {
      static {
        route 0.0.0.0/0 next-hop 172.16.0.13;
      }
    }
  }
  sp2-route-table {
```

```

instance-type forwarding;
routing-options {
  static {
    route 0.0.0.0/0 next-hop 172.16.0.17;
  }
}
}
routing-options {
  rib-groups {
    fbf-group {
      import-rib [ inet.0 sp1-route-table.inet.0 sp2-route-table.inet.0 ];
    }
  }
}
}

```

Verification

Confirm that the configuration is working properly.

Pinging with Specified Source Addresses

Purpose Send some ICMP packets across the network to test the firewall filter.

Action 1. Log in to Logical System PE3.

```

user@host> set cli logical-system PE3
Logical system: PE3

```

2. Run the **ping** command, pinging the lo0.3 interface on Logical System PE1.

The address configured on this interface is 1.1.1.1.

Specify the source address 10.1.2.1, which is the address configured on the lo0.1 interface on Logical System PE3.

```

user@host:PE3> ping 1.1.1.1 source 10.1.2.1
PING 1.1.1.1 (1.1.1.1): 56 data bytes
64 bytes from 1.1.1.1: icmp_seq=0 ttl=62 time=1.444 ms
64 bytes from 1.1.1.1: icmp_seq=1 ttl=62 time=2.094 ms
^C
--- 1.1.1.1 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.444/1.769/2.094/0.325 ms

```

3. Log in to Logical System PE4.

```

user@host:PE3> set cli logical-system PE4
Logical system: PE4

```

4. Run the **ping** command, pinging the lo0.4 interface on Logical System PE2.

The address configured on this interface is 2.2.2.2.

Specify the source address 10.2.1.1, which is the address configured on the lo0.2 interface on Logical System PE4.

```

user@host:PE4> ping 2.2.2.2 source 10.2.1.1
PING 2.2.2.2 (2.2.2.2): 56 data bytes
64 bytes from 2.2.2.2: icmp_seq=0 ttl=62 time=1.473 ms
64 bytes from 2.2.2.2: icmp_seq=1 ttl=62 time=1.407 ms

```

```
^C
--- 2.2.2.2 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.407/1.440/1.473/0.033 ms
```

Meaning Sending these pings activates the firewall filter actions.

Verifying the Firewall Filter

Purpose Make sure the firewall filter actions take effect.

Action 1. Log in to Logical System P1.

```
user@host> set cli logical-system P1
Logical system: P1
```

2. Run the **show firewall log** command on Logical System P1.

```
user@host:P1> show firewall log
Log :
Time      Filter  Action Interface  Protocol  Src Addr
Dest Addr
13:52:20  pfe          A      1t-1/2/0.10  ICMP      10.2.1.1
2.2.2.2
13:52:19  pfe          A      1t-1/2/0.10  ICMP      10.2.1.1
2.2.2.2
13:51:53  pfe          A      1t-1/2/0.10  ICMP      10.1.2.1
1.1.1.1
13:51:52  pfe          A      1t-1/2/0.10  ICMP      10.1.2.1
1.1.1.1
```

- Related Documentation**
- [Example: Configuring Filter-Based Forwarding on the Source Address on page 4644](#)
 - [Example: Configuring Multitopology Routing Based on Applications](#)
 - [Configuring Filter-Based Forwarding](#)
 - [Copying and Redirecting Traffic with Port Mirroring and Filter-Based Forwarding](#)
 - [Using Filter-Based Forwarding to Export Monitored Traffic to Multiple Destinations](#)
 - [Filter-Based Forwarding Overview on page 4515](#)

Example: Configuring Filter-Based Forwarding to a Specific Outgoing Interface or Destination IP Address

- [Understanding Filter-Based Forwarding to a Specific Outgoing Interface or Destination IP Address on page 4663](#)
- [Example: Configuring Filter-Based Forwarding to a Specific Outgoing Interface on page 4664](#)

Understanding Filter-Based Forwarding to a Specific Outgoing Interface or Destination IP Address

Policy-based routing (also known as filter-based forwarding) refers to the use of firewall filters that are applied to an interface to match certain IP header characteristics and to route only those matching packets differently than the packets would normally be routed.

Starting in Junos OS Release 12.2, you can use **then next-interface**, **then next-ip**, or **then next-ip6** as an action in a firewall filter.

For example:

```
from {
  set of match conditions
}
then {
  IP-address (or)
  IPv6-address (or)
  Interface name
}
```

The set of match conditions can be as follows:

- Layer-3 properties (for example, the source or destination IP address or the TOS byte)
- Layer-4 properties (for example, the source or destination port)

The route for the given IPv4 or IPv6 address has to be present in the routing table for policy-based routing to take effect. Similarly, the route through the given interface has to be present in the forwarding table for **next-interface** action to take effect. This can be achieved by configuring an interior gateway protocol (IGP), such as OSPF or IS-IS, to advertise Layer 3 routes.

The firewall filter matches the conditions and forwards the packet to one of the following:

- An IPv4 address (using the **next-ip** firewall filter action)
- An IPv6 address (using the **next-ip6** firewall filter action)
- An interface (using the **next-interface** firewall filter action)

Suppose, for example, that you want to offer services to your customers, and the services reside on different servers. An example of a service might be hosted DNS or hosted FTP. As customer traffic arrives at the Juniper Networks routing device, you can use filter-based forwarding to send traffic to the servers by applying a match condition on a MAC address or an IP address or simply an incoming interface and send the packets to a certain outgoing interface that is associated with the appropriate server. Some of your destinations might be IPv4 or IPv6 addresses, in which case the **next-ip** or **next-ip6** action is useful.

Optionally, you can associate the outgoing interfaces or IP addresses with routing instances.

For example:

```
firewall {
```

```
filter filter1 {
  term t1 {
    from {
      source-address {
        10.1.1.3/32;
      }
    }
    then {
      next-interface {
        xe-0/1/0.1;
        routing-instance rins1;
      }
    }
  }
  term t2 {
    from {
      source-address {
        10.1.1.4/32;
      }
    }
    then {
      next-interface {
        xe-0/1/0.2;
        routing-instance rins2;
      }
    }
  }
}
routing-instances {
  rins1 {
    instance-type virtual-router;
    interface xe-0/1/0.1;
  }
  rins2 {
    instance-type virtual-router;
    interface xe-0/1/0.2;
  }
}
```

Example: Configuring Filter-Based Forwarding to a Specific Outgoing Interface

This example shows how to use **then next-interface** as an action in a firewall filter.

- [Requirements on page 4665](#)
- [Overview on page 4665](#)
- [Configuration on page 4665](#)
- [Verification on page 4668](#)

Requirements

This example has the following hardware and software requirements:

- MX Series 3D Universal Edge Router or an EX Series switch as the routing device with the firewall filter configured.
- Junos OS Release 12.3R2 running on the routing device with the firewall filter configured.
- The filter with the **next-interface** (or **next-ip**) action can only be applied to an interface that is hosted on a Trio MPC. If you apply the filter to an I-chip based DPC, the commit operation fails.
- The outgoing interface referred to in the **next-interface *interface-name*** action can be hosted on a Trio MPC or an I-chip based DPC.

Overview


In this example, Device R1 has two loopback interface addresses configured: 172.16.1.1 and 172.16.2.2.

On Device R2, a firewall filter has multiple terms configured. Each term matches one of the source addresses in incoming traffic, and routes the traffic to specified outgoing interfaces. The outgoing interfaces are configured as VLAN-tagged interfaces between Device R2 and Device R3.

IS-IS is used for connectivity among the devices.

Figure 98 on page 4665 shows the topology used in this example.

Figure 98: Filter-Based Forwarding to Specified Outgoing Interfaces



This example shows the configuration on Device R2.

Configuration

CLI Quick Configuration

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

```
Device R2
set interfaces ge-2/1/0 unit 0 family inet filter input filter1
set interfaces ge-2/1/0 unit 0 family inet address 10.0.0.10/30
set interfaces ge-2/1/0 unit 0 description to-R1
set interfaces ge-2/1/0 unit 0 family iso
set interfaces ge-2/1/1 vlan-tagging
set interfaces ge-2/1/1 description to-R3
set interfaces ge-2/1/1 unit 0 vlan-id 1001
```

```
set interfaces ge-2/1/1 unit 0 family inet address 10.0.0.13/30
set interfaces ge-2/1/1 unit 0 family iso
set interfaces ge-2/1/1 unit 1 vlan-id 1002
set interfaces ge-2/1/1 unit 1 family inet address 10.0.0.25/30
set interfaces ge-2/1/1 unit 1 family iso
set interfaces lo0 unit 0 family inet address 10.255.4.4/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0404.00
set firewall family inet filter filter1 term t1 from source-address 172.16.1.1/32
set firewall family inet filter filter1 term t1 then next-interface ge-2/1/1.0
set firewall family inet filter filter1 term t2 from source-address 172.16.2.2/32
set firewall family inet filter filter1 term t2 then next-interface ge-2/1/1.1
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure Device R2:

1. Configure the interfaces.

```
[edit interfaces]
```

```
user@R2# set ge-2/1/0 unit 0 family inet filter input filter1
user@R2# set ge-2/1/0 unit 0 family inet address 10.0.0.10/30
user@R2# set ge-2/1/0 unit 0 description to-R1
user@R2# set ge-2/1/0 unit 0 family iso
```

```
user@R2# set ge-2/1/1 vlan-tagging
user@R2# set ge-2/1/1 description to-R3
```

```
user@R2# set ge-2/1/1 unit 0 vlan-id 1001
user@R2# set ge-2/1/1 unit 0 family inet address 10.0.0.13/30
user@R2# set ge-2/1/1 unit 0 family iso
```

```
user@R2# set ge-2/1/1 unit 1 vlan-id 1002
user@R2# set ge-2/1/1 unit 1 family inet address 10.0.0.25/30
user@R2# set ge-2/1/1 unit 1 family iso
```

```
user@R2# set lo0 unit 0 family inet address 10.255.4.4/32
user@R2# set lo0 unit 0 family iso address 49.0001.0010.0000.0404.00
```

2. Configure the firewall filter.

```
[edit firewall family inet filter filter1]
```

```
user@R2# set term t1 from source-address 172.16.1.1/32
user@R2# set term t1 then next-interface ge-2/1/1.0
```

```
user@R2# set term t2 from source-address 172.16.2.2/32
user@R2# set term t2 then next-interface ge-2/1/1.1
```

3. Enable IS-IS on the interfaces.

```
[edit protocols is-is]
```



```

user@R2# set interface all level 1 disable
user@R2# set interface fxp0.0 disable
user@R2# set interface lo0.0

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show firewall**, and **show protocols** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

user@R2# show interfaces
ge-2/1/0 {
  unit 0 {
    description to-R1;
    family inet {
      filter {
        input filter1;
      }
      address 10.0.0.10/30;
    }
    family iso;
  }
}
ge-2/1/1 {
  description to-R3;
  vlan-tagging;
  unit 0 {
    vlan-id 1001;
    family inet {
      address 10.0.0.13/30;
    }
    family iso;
  }
  unit 1 {
    vlan-id 1002;
    family inet {
      address 10.0.0.25/30;
    }
    family iso;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.4.4/32;
    }
    family iso {
      address 49.0001.0010.0000.0404.00;
    }
  }
}
}

user@R2# show firewall
family inet {
  filter filter1 {
    term t1 {
      from {

```

```
        source-address {
            172.16.1.1/32;
        }
    }
    then {
        next-interface {
            ge-2/1/1.0;
        }
    }
}
term t2 {
    from {
        source-address {
            172.16.2.2/32;
        }
    }
    then {
        next-interface {
            ge-2/1/1.1;
        }
    }
}
}
}
}

user@R2# show protocols
isis {
    interface all {
        level 1 disable;
    }
    interface fxp0.0 {
        disable;
    }
    interface lo0.0;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Checking the Paths Used

Purpose Make sure that the expected paths are used when sending traffic from Device R1 to Device R4.

Action On Device R1, enter the **traceroute** command.

```
user@R1> traceroute 10.255.6.6 source 172.16.1.1
traceroute to 10.255.6.6 (10.255.6.6) from 172.16.1.1, 30 hops max, 40 byte packets
 1  10.0.0.10 (10.0.0.10)  0.976 ms  0.895 ms  0.815 ms
 2  10.0.0.14 (10.0.0.14)  0.868 ms  0.888 ms  0.813 ms
 3  10.255.6.6 (10.255.6.6)  1.715 ms  1.442 ms  1.382 ms

user@R1> traceroute 10.255.6.6 source 172.16.2.2
```

```
traceroute to 10.255.6.6 (10.255.6.6) from 172.16.2.2, 30 hops max, 40 byte packets
```

```
 1  10.0.0.10 (10.0.0.10)  0.973 ms  0.907 ms  0.782 ms
 2  10.0.0.26 (10.0.0.26)  0.844 ms  0.890 ms  0.852 ms
 3  10.255.6.6 (10.255.6.6)  1.384 ms  1.516 ms  1.462 ms
```

Meaning The output shows that the second hop changes, depending on the source address used in the **traceroute** command.

To verify this feature, a traceroute operation is performed on Device R1 to Device R4. When the source IP address is 172.16.1.1, packets are forwarded out the ge-2/1/1.0 interface on Device R2. When the source IP address is 172.16.2.2, packets are forwarded out the ge-2/1/1.1 interface on Device R2.

- Related Documentation**
- [Example: Configuring Filter-Based Forwarding on Logical Systems on page 3298](#)
 - [Example: Configuring Filter-Based Forwarding on the Source Address on page 4644](#)
 - [Standard Firewall Filter Nonterminating Actions on page 4744](#)

Examples of Standard Firewall Filters Configuration Options

- [Example: Configuring Statistics Collection for a Standard Firewall Filter on page 4669](#)
- [Example: Configuring Logging for a Stateless Firewall Filter Term on page 4674](#)

Example: Configuring Statistics Collection for a Standard Firewall Filter

This example shows how to configure and apply a stateless firewall filter that collects data according to parameters specified in an associated accounting profile.

- [Requirements on page 4669](#)
- [Overview on page 4669](#)
- [Configuration on page 4670](#)
- [Verification on page 4673](#)

Requirements

Firewall filter accounting profiles are supported for all traffic types except **family any**.

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, you create a firewall filter accounting profile and apply it to a stateless firewall filter. The accounting profile specifies how frequently to collect packet and byte count statistics and the name of the file to which the statistics are written. The profile also specifies that statistics are to be collected for three firewall filter counters.

Topology

The firewall filter accounting profile **filter_acctg_profile** specifies that statistics are collected every 60 minutes, and the statistics are written to the file

`/var/log/ff_accounting_file`. Statistics are collected for counters named `counter1`, `counter2`, and `counter3`.

The IPv4 stateless firewall filter named `my_firewall_filter` increments a counter for each of three filter terms. The filter is applied to logical interface `ge-0/0/1.0`.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configure an Accounting Profile on page 4670](#)
- [Configure a Firewall Filter That References the Accounting Profile on page 4671](#)
- [Apply the Firewall Filter to an Interface on page 4671](#)
- [Confirm Your Candidate Configuration on page 4672](#)
- [Clear the Counters and Commit Your Candidate Configuration on page 4673](#)

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the `[edit]` hierarchy level.

```
set accounting-options filter-profile filter_acctg_profile file ff_accounting_file
set accounting-options filter-profile filter_acctg_profile interval 60
set accounting-options filter-profile filter_acctg_profile counters counter1
set accounting-options filter-profile filter_acctg_profile counters counter2
set accounting-options filter-profile filter_acctg_profile counters counter3
set firewall family inet filter my_firewall_filter accounting-profile filter_acctg_profile
set firewall family inet filter my_firewall_filter term term1 from protocol ospf
set firewall family inet filter my_firewall_filter term term1 then count counter1
set firewall family inet filter my_firewall_filter term term1 then discard
set firewall family inet filter my_firewall_filter term term2 from source-address
  10.108.0.0/16
set firewall family inet filter my_firewall_filter term term2 then count counter2
set firewall family inet filter my_firewall_filter term term2 then discard
set firewall family inet filter my_firewall_filter term accept-all then count counter3
set firewall family inet filter my_firewall_filter term accept-all then accept
set interfaces ge-0/0/1 unit 0 family inet address 10.1.2.3/30
set interfaces ge-0/0/1 unit 0 family inet filter input my_firewall_filter
```

Configure an Accounting Profile

Step-by-Step Procedure

To configure an accounting profile:

1. Create the accounting profile `filter_acctg_profile`.

```
[edit]
user@host# edit accounting-options filter-profile filter_acctg_profile
```

2. Configure the accounting profile to filter and collect packet and byte count statistics every 60 minutes and write them to the `/var/log/ff_accounting_file` file.

```
[edit accounting-options filter-profile filter_acctg_profile]
```

```
user@host# set file ff_accounting_file
user@host# set interval 60
```

3. Configure the accounting profile to collect filter profile statistics (packet and byte counts) for three counters.

```
[edit accounting-options filter-profile filter_acctg_profile]
user@host# set counters counter1
user@host# set counters counter2
user@host# set counters counter3
```

Configure a Firewall Filter That References the Accounting Profile

Step-by-Step Procedure

To configure a firewall filter that references the accounting profile:

1. Create the stateless firewall filter **my_firewall_filter**.

```
[edit]
user@host# edit firewall family inet filter my_firewall_filter
```

2. Apply the filter-accounting profile **filter_acctg_profile** to the firewall filter.

```
[edit firewall family inet filter my_firewall_filter]
user@host# set accounting-profile filter_acctg_profile
```

3. Configure the first filter term and counter.

```
[edit firewall family inet filter my_firewall_filter]
user@host# set term term1 from protocol ospf
user@host# set term term1 then count counter1
user@host# set term term1 then discard
```

4. Configure the second filter term and counter.

```
[edit firewall family inet filter my_firewall_filter]
user@host# set term term2 from source-address 10.108.0.0/16
user@host# set term term2 then count counter2
user@host# set term term2 then discard
```

5. Configure the third filter term and counter.

```
[edit firewall family inet filter my_firewall_filter]
user@host# set term accept-all then count counter3
user@host# set term accept-all then accept
```

Apply the Firewall Filter to an Interface

Step-by-Step Procedure

To apply the stateless firewall filter to a logical interface:

1. Configure the logical interface to which you will apply the stateless firewall filter.

```
[edit]
user@host# edit interfaces ge-0/0/1 unit 0 family inet
```

2. Configure the interface address for the logical interface.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
user@host# set address 10.1.2.3/30
```

3. Apply the stateless firewall filter to the logical interface.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
user@host# set filter input my_firewall_filter
```

Confirm Your Candidate Configuration

Step-by-Step Procedure

To confirm your candidate configuration:

1. Confirm the configuration of the accounting profile by entering the **show accounting-options** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show accounting-options
filter-profile filter_acctg_profile {
  file ff_accounting_file;
  interval 60;
  counters {
    counter1;
    counter2;
    counter3;
  }
}
```

2. Confirm the configuration of the stateless firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
  filter my_firewall_filter {
    accounting-profile filter_acctg_profile;
    term term1 {
      from {
        protocol ospf;
      }
      then {
        count counter1;
        discard;
      }
    }
    term term2 {
      from {
        source-address {
          10.108.0.0/16;
        }
      }
      then {
        count counter2;
        discard;
      }
    }
    term accept-all {
```

```

        then {
            count counter3;
            accept;
        }
    }
}

```

3. Confirm the configuration of the interfaces by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

[edit]
user@host# show interfaces
ge-0/0/1 {
    unit 0 {
        family inet {
            filter {
                input my_firewall_filter;
            }
            address 10.1.2.3/30;
        }
    }
}

```

Clear the Counters and Commit Your Candidate Configuration

Step-by-Step Procedure

To clear the counters and commit your candidate configuration:

1. From operational command mode, use the **clear firewall all** command to clear the statistics for all firewall filters.

To clear only the counters incremented in this example, include the name of the firewall filter.

```

[edit]
user@host> clear firewall filter my_firewall_filter

```

2. Commit your candidate configuration.

```

[edit]
user@host# commit

```

Verification

To verify that the filter is applied to the logical interface, run the **show interfaces** command with the **detail** or **extensive** output level.

To verify that the three counters are collected separately, run the **show firewall filter my_firewall_filter** command.

```

user@host> show firewall filter my_firewall_filter

```

```

Filter: my_firewall_filter

```

```

Counters:

```

Name	Bytes	Packets
counter1	0	0

counter2	0	0
counter3	0	0

Related Documentation

- [Accounting for Standard Firewall Filters Overview on page 4517](#)
- [Statement Hierarchy for Configuring Firewall Filter Accounting Profiles on page 4783](#)
- [Statement Hierarchy for Applying Firewall Filter Accounting Profiles on page 4784](#)

Example: Configuring Logging for a Stateless Firewall Filter Term

This example shows how to configure a standard stateless firewall filter to log packet headers.

- [Requirements on page 4674](#)
- [Overview on page 4674](#)
- [Configuration on page 4674](#)
- [Verification on page 4677](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, you use a stateless firewall filter that logs and counts ICMP packets that have 192.168.207.222 as either their source or destination.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configure the Syslog Messages File for the Firewall Facility on page 4675](#)
- [Configure the Stateless Firewall Filter on page 4675](#)
- [Apply the Stateless Firewall Filter to a Logical Interface on page 4675](#)
- [Confirm and Commit Your Candidate Configuration on page 4676](#)

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set system syslog file messages_firewall_any firewall any
set system syslog file messages_firewall_any archive no-world-readable
set firewall family inet filter icmp_syslog term icmp_match from address
192.168.207.222/32
set firewall family inet filter icmp_syslog term icmp_match from protocol icmp
set firewall family inet filter icmp_syslog term icmp_match then count packets
set firewall family inet filter icmp_syslog term icmp_match then log
```



```

set firewall family inet filter icmp_syslog term icmp_match then accept
set firewall family inet filter icmp_syslog term default_term then accept
set interfaces ge-0/0/1 unit 0 family inet address 10.1.2.3/30
set interfaces ge-0/0/1 unit 0 family inet filter input icmp_syslog

```

Configure the Syslog Messages File for the Firewall Facility

Step-by-Step Procedure

To configure a syslog messages file for the **firewall** facility:

1. Configure a messages file for all syslog messages generated for the **firewall** facility.

```

user@host# set system syslog file messages_firewall_any firewall any

```

2. Restrict permission to the archived **firewall** facility syslog files to the root user and users who have the Junos OS maintenance permission.

```

user@host# set system syslog file messages_firewall_any archive no-world-readable

```

Configure the Stateless Firewall Filter

Step-by-Step Procedure

To configure the stateless firewall filter **icmp_syslog** that logs and counts ICMP packets that have **192.168.207.222** as either their source or destination:

1. Create the stateless firewall filter **icmp_syslog**.

```

[edit]
user@host# edit firewall family inet filter icmp_syslog

```

2. Configure matching on the ICMP protocol and an address.

```

[edit firewall family inet filter icmp_syslog]
user@host# set term icmp_match from address 192.168.207.222/32
user@host# set term icmp_match from protocol icmp

```

3. Count, log, and accept matching packets.

```

[edit firewall family inet filter icmp_syslog]
user@host# set term icmp_match then count packets
user@host# set term icmp_match then log
user@host# set term icmp_match then accept

```

4. Accept all other packets.

```

[edit firewall family inet filter icmp_syslog]
user@host# set term default_term then accept

```

Apply the Stateless Firewall Filter to a Logical Interface

Step-by-Step Procedure

To apply the stateless firewall filter to a logical interface:

1. Configure the logical interface to which you will apply the stateless firewall filter.

```

[edit]
user@host# edit interfaces ge-0/0/1 unit 0 family inet

```

2. Configure the interface address for the logical interface.

```

[edit interfaces ge-0/0/1 unit 0 family inet]
user@host# set address 10.1.2.3/30

```

3. Apply the stateless firewall filter to the logical interface.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
user@host# set filter input icmp_syslog
```

Confirm and Commit Your Candidate Configuration

Step-by-Step Procedure

To confirm and then commit your candidate configuration:

1. Confirm the configuration of the syslog message file for the **firewall** facility by entering the **show system** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show system
syslog {
  file messages_firewall_any {
    firewall any;
    archive no-world-readable;
  }
}
```

2. Confirm the configuration of the stateless firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
  filter icmp_syslog {
    term icmp_match {
      from {
        address {
          192.168.207.222/32;
        }
        protocol icmp;
      }
      then {
        count packets;
        log;
        accept;
      }
    }
    term default_term {
      then accept;
    }
  }
}
```

3. Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
```

```

unit 0 {
    family inet {
        filter {
            input icmp_syslog;
        }
        address 10.1.2.3/30;
    }
}

```

4. If you are done configuring the device, commit your candidate configuration.

```

[edit]
user@host# commit

```

Verification

To confirm that the configuration is working properly, enter the `show log filter` command:

```

user@host> show log messages_firewall_any
Mar 20 08:03:11 hostname feb FW: ge-0/1/0.0   A icmp 192.168.207.222
192.168.207.223      0      0 (1 packets)

```

This output file contains the following fields:

- **Date and Time**—Date and time at which the packet was received (not shown in the default).
- **Filter action**:
 - **A**—Accept (or next term)
 - **D**—Discard
 - **R**—Reject
- **Protocol**—Packet's protocol name or number.
- **Source address**—Source IP address in the packet.
- **Destination address**—Destination IP address in the packet.



NOTE: If the protocol is ICMP, the ICMP type and code are displayed. For all other protocols, the source and destination ports are displayed.

The last two fields (both zero) are the source and destination TCP/UDP ports, respectively, and are shown for TCP or UDP packets only. This log message indicates that only one packet for this match has been detected in about a 1-second interval. If packets arrive faster, the system log function compresses the information so that less output is generated, and displays an output similar to the following:

```

user@host> show log filter
Mar 20 08:08:45 hostname feb FW: ge-0/1/0.0   A icmp 192.168.207.222
192.168.207.223      0      0 (515 packets)

```

Related Documentation

- [System Logging Overview on page 4517](#)
- [Logging of Packet Headers Evaluated by a Firewall Filter Term on page 4520](#)
- System log messages with the **DFWD_** prefix, described in the *Junos OS System Log Messages Reference*
- System log messages with the **PFE_FW_*** prefix, described in the *Junos OS System Log Messages Reference*

Service Filters Configuration

- [Example: Configuring and Applying Service Filters on page 4678](#)

Example: Configuring and Applying Service Filters

This example shows how to configure and apply service filters.

- [Requirements on page 4678](#)
- [Overview on page 4678](#)
- [Configuration on page 4679](#)
- [Verification on page 4682](#)

Requirements

This example use the logical interface **xe-0/1/0.0** on any of the following hardware components:

- Adaptive Services (AS) PIC on an M Series or T Series router
- Multiservices (MS) PIC on an M Series or T Series router
- Multiservices (MS) DPC on an MX Series router
- EX Series switch

Before you begin, make sure that you have:

- Installed your supported router (or switch) and PICs or DPCs and performed the initial router (or switch) configuration.
- Configured basic Ethernet in the topology, and verified that traffic is flowing in the topology and that IPv4 traffic is flowing through logical interface **xe-0/1/0.0**.
- Configured the service set **vrf_svcs** with service input and output rules and default settings for services at a service interface.

For guidelines for configuring service sets, see “*Configuring Service Sets to be Applied to Services Interfaces*” in the *Junos Services Interfaces Configuration Release 11.2*.

Overview

In this example, you create three types of service filters for IPv4 traffic: one input service filter, one postservice input filter, and one output service filter.

Topology

You apply the input service filter and postservice input filter to input traffic at logical interface **xe-0/1/0.0**, and you apply the output service filter to the output traffic at the same logical interface.

- Filtering IPv4 traffic before it is accepted for input service processing—At logical interface **xe-0/1/0.0**, you use the service filter **in_filter_presvc** to filter IPv4 input traffic before the traffic can be accepted for processing by services associated with service set **vrf_svcs**. The **in_filter_presvc** service filter counts packets sent from ICMP port 179, directs these packets to the input services associated with the service set **vrf_svcs**, and discards all other packets.
- Filtering IPv4 traffic after it has completed input service processing—At logical interface **xe-0/1/0.0**, you use the service filter **in_filter_postsvc** to filter traffic that is returning to the services interface after the input service set **in_filter_presvc** is executed. The **in_filter_postsvc** service filter counts packets sent from ICMP port 179 and then discards them.
- Filtering IPv4 traffic before it is accepted for output service processing—At logical interface **xe-0/1/0.0**, you use the service-filter **out_filter_presvc** to filter IPv4 output traffic before the traffic can be accepted for processing by the services associated with service set **vrf_svcs**. The **out_filter_presvc** service filter counts packets destined for TCP port 179 and then directs the packets to the output services associated with the service set **vrf_svcs**.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configuring the Three Service Filters on page 4680](#)
- [Applying the Three Service Filters on page 4681](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set firewall family inet service-filter in_filter_presvc term t1 from protocol tcp
set firewall family inet service-filter in_filter_presvc term t1 from source-port bgp
set firewall family inet service-filter in_filter_presvc term t1 then count svc_in_pkts
set firewall family inet service-filter in_filter_postsvc term t2 from protocol tcp
set firewall family inet service-filter in_filter_postsvc term t2 from source-port bgp
set firewall family inet service-filter in_filter_postsvc term t2 then count svc_in_pkts_rtn
set firewall family inet service-filter in_filter_postsvc term t2 then skip
set firewall family inet service-filter out_filter_presvc term t3 from protocol icmp
set firewall family inet service-filter out_filter_presvc term t3 from destination-port bgp
set firewall family inet service-filter out_filter_presvc term t3 then count svc_out_pkts
set firewall family inet service-filter out_filter_presvc term t3 then service
```

```
set interfaces xe-0/1/0 unit 0 family inet service input service-set vrf_svcs service-filter
in_filter_presvc
set interfaces xe-0/1/0 unit 0 family inet service input post-service-filter in_filter_postsvc
set interfaces xe-0/1/0 unit 0 family inet service output service-set vrf_svcs service-filter
out_filter_presvc
```

Configuring the Three Service Filters

Step-by-Step Procedure

To configure the three service filters:

1. Configure the input service filter.

```
[edit]
user@host# edit firewall family inet service-filter in_filter_presvc
```

```
[edit firewall family inet service-filter in_filter_presvc]
user@host# set term t1 from protocol tcp
user@host# set term t1 from source-port bgp
user@host# set term t1 then count svc_in_pkts
user@host# set term t1 then service
```

2. Configure the postservice input filter.

```
[edit]
user@host# edit firewall family inet service-filter in_filter_postsvc
```

```
[edit firewall family inet service-filter in_filter_postsvc]
user@host# set term t2 from protocol tcp
user@host# set term t2 from source-port bgp
user@host# set term t2 then count svc_in_pkts_rtn
user@host# set term t2 then skip
```

3. Configure the output service filter.

```
[edit]
user@host# edit firewall family inet service-filter out_filter_presvc
```

```
[edit firewall family inet service-filter out_filter_presvc]
user@host# set term t3 from protocol icmp
user@host# set term t3 from destination-port bgp
user@host# set term t3 then count svc_out_pkts
user@host# set term t3 then service
```

Results Confirm the configuration of the input and output service filters and the postservice input filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
  service-filter in_filter_presvc {
    term t1 {
      from {
        protocol tcp;
        source-port bgp;
      }
    }
  }
}
```

```

        then {
            count svc_in_pkts;
            service;
        }
    }
}
service-filter in_filter_postsvc {
    term t2 {
        from {
            protocol tcp;
            source-port bgp;
        }
        then {
            count svc_in_pkts_rtn;
            skip;
        }
    }
}
service-filter out_filter_presvc {
    term t3 {
        from {
            protocol icmp;
            destination-port bgp;
        }
        then {
            count svc_out_pkts;
            service;
        }
    }
}
}

```

Applying the Three Service Filters

Step-by-Step Procedure

To apply the three service filters:

1. Access the IPv4 protocol on the input interface **xe-0/1/0.0**.

```

[edit]
user@host# edit interfaces xe-0/1/0 unit 0 family inet

```

2. Apply the input service filter and the postservice input filter.

```

[edit interfaces xe-0/1/0 unit 0 family inet]
user@host# set service input service-set vrf_svcs service-filter in_filter_presvc
user@host# set service input post-service-filter in_filter_postsvc
user@host# set service output service-set vrf_svcs service-filter out_filter_presvc

```

Results Confirm the configuration of the interfaces by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

[edit]
user@host# show interfaces
xe-0/1/0 {
    unit 0 {

```

```
family inet {
  service {
    input {
      service-set vrf_svcs service-filter in_filter_presvc;
      post-service-filter in_filter_postsvc;
    }
    output {
      service-set vrf_svcs service-filter out_filter_presvc;
    }
  }
}
```

When you are done configuring the device, commit your candidate configuration.

Verification

Confirm that the configuration is working properly.

- [Verifying That Inbound Traffic Is Filtered Before Input Service on page 4682](#)
- [Verifying That Inbound Traffic Is Filtered After Input Service Processing on page 4682](#)
- [Verifying That Outbound Traffic Is Filtered Before Output Service Processing on page 4682](#)

Verifying That Inbound Traffic Is Filtered Before Input Service

Purpose Verify that inbound packets sent from TCP port 179 are sent for processing by the *input* services associated with the service set **vrf_svcs**.

Action Display the count of packets sent for processing by the *input* services associated with the service set **vrf_svcs**.

[edit]

user@host> show firewall filter in_filter_presvc-vrf_svcs counter svc_in_pkts

Verifying That Inbound Traffic Is Filtered After Input Service Processing

Purpose Verify that inbound packets sent from TCP port 179 are returned from processing by the *input* services associated with the service set **vrf_svcs**.

Action Display the count of packets returned from processing by the *input* services associated with the service set **vrf_svcs**.

[edit]

user@host> show firewall filter in_filter_postsvc-vrf_svcs counter svc_in_pkts_rtn

Verifying That Outbound Traffic Is Filtered Before Output Service Processing

Purpose Verify that outbound packets sent to ICMP port 179 are sent for processing by the *output* services associated with the service set **vrf_svcs**.

Action Display the count of packets sent for processing by the *output* services associated with the service set **vrf_svcs**.

[edit]

```
user@host> show firewall filter out_filter_presvc-vrf_svcs counter svc_out_pkts
```

Related Documentation

- [Service Filter Overview on page 4522](#)
- [How Service Filters Evaluate Packets on page 4523](#)
- [Guidelines for Configuring Service Filters on page 4524](#)
- [Guidelines for Applying Service Filters on page 4526](#)

Simple Filters Configuration

- [Example: Configuring and Applying a Simple Filter on page 4683](#)

Example: Configuring and Applying a Simple Filter

This example shows how to configure a simple filter.

- [Requirements on page 4683](#)
- [Overview on page 4683](#)
- [Configuration on page 4684](#)
- [Verification on page 4686](#)

Requirements

This example uses one of the following hardware components:

- One Gigabit Ethernet intelligent queuing (IQ2) PIC installed on an M120, M320, or T Series router
- One Enhanced Queuing Dense Port Concentrator (EQ DPC) installed on an MX Series router or an EX Series switch

Before you begin, make sure that you have:

- Installed your supported router (or switch) and PIC or DPC and performed the initial router (or switch) configuration.
- Configured basic Ethernet in the topology, and verified that traffic is flowing in the topology and that ingress IPv4 traffic is flowing into logical interface **ge-0/0/1.0**.

Overview

This simple filter sets the loss priority to low for TCP traffic with source address **1.1.1.1**, sets the loss priority to high for HTTP (Web) traffic with source addresses in the **4.0.0.0/8** range, and sets the loss priority to low for all traffic with destination address **6.6.6.6**.

Topology

The simple filter is applied as an input filter (arriving packets are checking for destination address **6.6.6.6**, not queued output packets) on interface **ge-0/0/1.0**.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configuring the Simple Firewall Filter on page 4684](#)
- [Applying the Simple Filter to the Logical Interface Input on page 4685](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set firewall family inet simple-filter sf_classify_1 term 1 from source-address 1.1.1.1/32
set firewall family inet simple-filter sf_classify_1 term 1 from protocol tcp
set firewall family inet simple-filter sf_classify_1 term 1 then loss-priority low
set firewall family inet simple-filter sf_classify_1 term 2 from source-address 4.0.0.0/8
set firewall family inet simple-filter sf_classify_1 term 2 from protocol tcp
set firewall family inet simple-filter sf_classify_1 term 2 from source-port http
set firewall family inet simple-filter sf_classify_1 term 2 then loss-priority high
set firewall family inet simple-filter sf_classify_1 term 3 from destination-address
  6.6.6.6/32
set firewall family inet simple-filter sf_classify_1 term 3 then loss-priority low
set firewall family inet simple-filter sf_classify_1 term 3 then forwarding-class best-effort
set interfaces ge-0/0/1 unit 0 family inet simple-filter input sf_classify_1
set interfaces ge-0/0/1 unit 0 family inet address 10.1.2.3/30
```

Configuring the Simple Firewall Filter

Step-by-Step Procedure

To configure the simple filter:

1. Create the simple filter **sf_classify_1**.

```
[edit]
user@host# edit firewall family inet simple-filter sf_classify_1
```

2. Configure classification of TCP traffic based on the source IP address.

```
[edit firewall family inet simple-filter sf_classify_1]
user@host# set term 1 from source-address 1.1.1.1/32
user@host# set term 1 from protocol tcp
user@host# set term 1 then loss-priority low
```

3. Configure classification of HTTP traffic based on the source IP address.

```
[edit firewall family inet simple-filter sf_classify_1]
user@host# set term 2 from source-address 4.0.0.0/8
user@host# set term 2 from protocol tcp
user@host# set term 2 from source-port http
user@host# set term 2 then loss-priority high
```

4. Configure classification of other traffic based on the destination IP address.

```
[edit firewall family inet simple-filter sf_classify_1]
user@host# set term 3 from destination-address 6.6.6.6/32
user@host# set term 3 then loss-priority low
```

```
user@host# set term 3 then forwarding-class best-effort
```

Results Confirm the configuration of the simple filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
  simple-filter sf_classify_1 {
    term 1 {
      from {
        source-address {
          1.1.1.1/32;
        }
        protocol {
          tcp;
        }
      }
      then loss-priority low;
    }
    term 2 {
      from {
        source-address {
          4.0.0.0/8;
        }
        source-port {
          http;
        }
        protocol {
          tcp;
        }
      }
      then loss-priority high;
    }
    term 3 {
      from {
        destination-address {
          6.6.6.6/32;
        }
      }
      then {
        loss-priority low;
        forwarding-class best-effort;
      }
    }
  }
}
```

Applying the Simple Filter to the Logical Interface Input

Step-by-Step Procedure To apply the simple filter to the logical interface input:

1. Configure the logical interface to which you will apply the simple filter.

```
[edit]
```

```
user@host# edit interfaces ge-0/0/1 unit 0 family inet
```

2. Configure the interface address for the logical interface.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
```

```
user@host# set address 10.1.2.3/30
```

3. Apply the simple filter to the logical interface input.

```
[edit interfaces ge-0/0/1 unit 0 family inet]
```

```
user@host# set simple-filter input sf_classify_1
```

Results Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      simple-filter {
        input sf_classify_1;
      }
      address 10.1.2.3/30;
    }
  }
}
```

When you are done configuring the device, commit your candidate configuration.

Verification

Confirm that the configuration is working properly.

- [Displaying the Mapping of Forwarding Class Maps and Names to Queue Numbers on page 4686](#)
- [Displaying CoS Queue Counters for the Interface on page 4686](#)
- [Displaying CoS Queue Counter Details for the Physical Interface on page 4687](#)

Displaying the Mapping of Forwarding Class Maps and Names to Queue Numbers

Purpose Display the mapping of forwarding class names to queue numbers.

Action Enter the **show class-of-service forwarding-class** operational mode command.

```
[edit]
```

```
user@host> show class-of-service forwarding-class
```

For information about the command output, see “**show class-of-service forwarding-class**” in the *Junos OS Operational Mode Commands*.

Displaying CoS Queue Counters for the Interface

Purpose Verify that the class-of-service (CoS) queue counters for the interface reflect the simple filter applied to the logical interface.

Action Enter the **show interfaces** command for the physical interface on which the simple filter is applied, and specify **detail** or **extensive** output level.

[edit]

```
user@host> show interfaces ge-0/0/1 detail
```

In the **Physical interface** section, under **Ingress queues**, the **Queue counters** section displays ingress queue counters for each forwarding class.

For more detailed information about the command output, see “[show interfaces \(Gigabit Ethernet\)](#)” or “[show interfaces \(10-Gigabit Ethernet\)](#)” in the *Junos OS Operational Mode Commands*.

Displaying CoS Queue Counter Details for the Physical Interface

Purpose Verify that the CoS queue counter details for the physical interface reflect the simple filter applied to the logical interface.

Action Enter the **show interfaces queue** command for the physical interface on which the simple filter is applied, and specify the **ingress** option.

[edit]

```
user@host> show interfaces queue ge-0/0/1 ingress
```

For information about the command output, see “[show interfaces queue](#)” in the *Junos OS Operational Mode Commands*.

Related Documentation

- [Simple Filter Overview on page 4529](#)
- [How Simple Filters Evaluate Packets on page 4530](#)
- [Guidelines for Configuring Simple Filters on page 4531](#)
- [Guidelines for Applying Simple Filters on page 4534](#)

Firewall Filters Configuration in Logical Systems

- [Example: Configuring a Stateless Firewall Filter to Protect a Logical System Against ICMP Floods on page 4687](#)

Example: Configuring a Stateless Firewall Filter to Protect a Logical System Against ICMP Floods

This example shows how to configure a stateless firewall filter that protects against ICMP denial-of-service attacks on a logical system.

- [Requirements on page 4687](#)
- [Overview on page 4688](#)
- [Configuration on page 4688](#)
- [Verification on page 4690](#)

Requirements

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

Overview


This example shows a stateless firewall filter called **protect-RE** that polices ICMP packets. The **icmp-policer** limits the traffic rate of the ICMP packets to 1,000,000 bps and the burst size to 15,000 bytes. Packets that exceed the traffic rate are discarded.

The policer is incorporated into the action of a filter term called **icmp-term**.

In this example, a ping is sent from a directly connected physical router to the interface configured on the logical system. The logical system accepts the ICMP packets if they are received at a rate of up to 1 Mbps (bandwidth-limit). The logical system drops all ICMP packets when this rate is exceeded. The **burst-size-limit** statement accepts traffic bursts up to 15 Kbps. If bursts exceed this limit, all packets are dropped. When the flow rate subsides, ICMP packets are again accepted.

Figure 58 on page 3295 shows the topology used in this example.

Figure 99: Logical System with a Stateless Firewall



Configuration

CLI Quick Configuration

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

```
set logical-systems LS1 interfaces ge-0/0/2 unit 0 family inet policer input icmp-policer
set logical-systems LS1 interfaces ge-0/0/2 unit 0 family inet address 10.0.45.2/30
set logical-systems LS1 firewall family inet filter protect-RE term icmp-term from protocol icmp
set logical-systems LS1 firewall family inet filter protect-RE term icmp-term then policer icmp-policer
set logical-systems LS1 firewall policer icmp-policer if-exceeding bandwidth-limit 1m
set logical-systems LS1 firewall policer icmp-policer if-exceeding burst-size-limit 15k
set logical-systems LS1 firewall policer icmp-policer then discard
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure an ICMP firewall filter on a logical system:

1. Configure the interface on the logical system.

```
[edit]
user@host# set logical-systems LS1 interfaces ge-0/0/2 unit 0 family inet address
10.0.45.2/30
```

2. Explicitly enable ICMP packets to be received on the interface.

```
[edit]
user@host# set logical-systems LS1 firewall family inet filter protect-RE term
icmp-term from protocol icmp
user@host# set logical-systems LS1 firewall family inet filter protect-RE term
icmp-term then accept
```

3. Create the policer.

```
[edit]
user@host# set logical-systems LS1 firewall policer icmp-policer if-exceeding
bandwidth-limit 1m
user@host# set logical-systems LS1 firewall policer icmp-policer if-exceeding
burst-size-limit 15k
user@host# set logical-systems LS1 firewall policer icmp-policer then discard
```

4. Apply the policer to a filter term.

```
[edit]
user@host# set logical-systems LS1 firewall family inet filter protect-RE term
icmp-term then policer icmp-policer
```

5. Apply the policer to the logical system interface.

```
[edit]
user@host# set logical-systems LS1 interfaces ge-0/0/2 unit 0 family inet policer
input icmp-policer
```

6. If you are done configuring the device, commit the configuration.

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

Results

Confirm your configuration by issuing the **show logical-systems LS1** command.

```
user@host# show logical-systems LS1
interfaces {
  ge-0/0/2 {
    unit 0 {
      family inet {
        policer {
          input icmp-policer;
        }
      }
    }
  }
}
```

```
        address 10.0.45.2/30;
      }
    }
  }
}
firewall {
  family inet {
    filter protect-RE {
      term icmp-term {
        from {
          protocol icmp;
        }
        then {
          policer icmp-policer;
          accept;
        }
      }
    }
  }
}
policer icmp-policer {
  if-exceeding {
    bandwidth-limit 1m;
    burst-size-limit 15k;
  }
  then discard;
}
}
```

Verification

Confirm that the configuration is working properly.

Verifying That Ping Works Unless the Limits Are Exceeded

Purpose Make sure that the logical system interface is protected against ICMP-based DoS attacks.

Action Log in to a system that has connectivity to the logical system and run the **ping** command.

```
user@R2> ping 10.0.45.2
PING 10.0.45.2 (10.0.45.2): 56 data bytes
64 bytes from 10.0.45.2: icmp_seq=0 ttl=64 time=1.316 ms
64 bytes from 10.0.45.2: icmp_seq=1 ttl=64 time=1.277 ms
64 bytes from 10.0.45.2: icmp_seq=2 ttl=64 time=1.269 ms

user@R2> ping 10.0.45.2 size 20000
PING 10.0.45.2 (10.0.45.2): 20000 data bytes
^C
--- 10.0.45.2 ping statistics ---
4 packets transmitted, 0 packets received, 100% packet loss
```

Meaning When you send a normal ping, the packet is accepted. When you send a ping packet that exceeds the filter limit, the packet is discarded.

Related Documentation

- [Example: Creating an Interface on a Logical System on page 3275](#)

Configuration Statements

- [\[edit firewall\] Hierarchy Level on page 4691](#)

[edit firewall] Hierarchy Level

Several statements in the **[edit firewall]** hierarchy are valid at numerous locations within the hierarchy. To make the complete hierarchy easier to read, the repeated statements are listed in the following sections, which are referenced at the appropriate locations in [“Complete \[edit firewall\] Hierarchy” on page 325](#).

- [Common Firewall Actions on page 4691](#)
- [Common IP Firewall Actions on page 4691](#)
- [Common IPv4 Firewall Actions on page 4692](#)
- [Common IP Firewall Match Conditions on page 4692](#)
- [Common IPv4 Firewall Match Conditions on page 4693](#)
- [Common Layer 2 Firewall Match Conditions on page 4694](#)
- [Complete \[edit firewall\] Hierarchy on page 4695](#)

Common Firewall Actions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- **[edit firewall family (any | ccc | ethernet-switching | inet | inet6 | mpls | vpls) filter *filter-name* term *term-name* then]**
- **[edit firewall filter *filter-name* term *term-name* then]**

The common firewall actions are as follows:

```
count counter-name;
forwarding-class class-name;
loss-priority (high | low | medium-high | medium-low);
next term;
policer policer-name;
three-color-policer policer-name {
    (single-rate single-rate-policer-name | two-rate two-rate-policer-name);
}
```

Common IP Firewall Actions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- **[edit firewall family inet filter *filter-name* term *term-name* then]**
- **[edit firewall family inet6 filter *filter-name* term *term-name* then]**
- **[edit firewall filter *filter-name* term *term-name* then]**

The common IP firewall actions are as follows:

```
log;
logical-system logical-system-name <routing-instance routing-instance-name>
  <topology topology-name>;
port-mirror;
port-mirror-instance instance-name;
routing-instance routing-instance-name <topology topology-name>;
sample;
service-filter-hit;
syslog;
topology topology-name;
```

Common IPv4 Firewall Actions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- **[edit firewall family inet filter *filter-name* term *term-name* then]**
- **[edit firewall filter *filter-name* term *term-name* then]**

The common IP version 4 (IPv4) firewall actions are as follows:

```
(accept | discard <accounting collector-name> | reject <administratively-prohibited |
  bad-host-tos | bad-network-tos | fragmentation-needed | host-prohibited |
  host-unknown | host-unreachable | network-prohibited | network-unknown |
  network-unreachable | port-unreachable | precedence-cutoff | precedence-violation |
  protocol-unreachable | source-host-isolated | source-route-failed | tcp-reset>);
ipsec-sa sa-name;
load-balance sa-name;
next-hop-group group-name;
prefix-action action-name;
```

Common IP Firewall Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- **[edit firewall family inet dialer-filter *filter-name* term *term-name* from]** (with the exceptions noted at this level in [“Complete \[edit firewall\] Hierarchy” on page 325](#))
- **[edit firewall family inet filter *filter-name* term *term-name* from]**
- **[edit firewall family inet6 dialer-filter *filter-name* term *term-name* from]** (with the exceptions noted at this level in [“Complete \[edit firewall\] Hierarchy” on page 325](#))
- **[edit firewall family inet6 filter *filter-name* term *term-name* from]**
- **[edit firewall filter *filter-name* term *term-name* from]**

The common IP firewall match conditions are as follows:

```
address {
  ip-prefix/prefix-length> <except>;
}
```

```

destination-address {
    ip-prefix</prefix-length> <except>;
}
destination-class [ class-names ] | destination-class-except [ class-names ];
(destination-port [ port-names ] | destination-port-except [ port-names ]);
destination-prefix-list {
    list-name <except>;
}
(forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
icmp-code [ codes ] | icmp-code-except [ codes ];
icmp-type [ types ] | icmp-type-except [ types ];
interface interface-name;
interface-group [ group-names ] | interface-group-except [ group-names ];
interface-set set-name;
(loss-priority [ priorities ] | loss-priority-except [ priorities ]);
(packet-length [ values ] | packet-length-except [ values ]);
(port [ port-names ] | port-except [ port-names ]);
prefix-list {
    list-name <except>;
}
service-filter-hit;
source-address {
    ip-prefix</prefix-length> <except>;
}
(source-class [ class-names ] | source-class-except [ class-names ]);
(source-port [ port-names ] | source-port-except [ port-names ]);
source-prefix-list {
    list-name <except>;
}
tcp-established;
tcp-flags flag;
tcp-initial;

```

Common IPv4 Firewall Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- [\[edit firewall family inet dialer-filter filter-name term term-name from\]](#) (with the exceptions noted at this level in [“Complete \[edit firewall\] Hierarchy” on page 325](#))
- [\[edit firewall family inet filter filter-name term term-name from\]](#)
- [\[edit firewall filter filter-name term term-name from\]](#)

The common IPv4 firewall match conditions are as follows:

```

(ah-spi [ values ] | ah-spi-except [ values ]);
(dscp [ code-point-values ] | dscp-except [ code-point-values ]);
(esp-spi [ values ] | esp-spi-except [ values ]);
first-fragment;
fragment-flags flag;
(fragment-offset [ offsets ] | fragment-offset-except [ offsets ]);
(ip-options [ option-names ] | ip-options-except [ option-names ]);
is-fragment;
(precedence [ precedence-names ] | precedence-except [ precedence-names ]);

```

```
(protocol [ protocol-names ] | protocol-except [ protocol-names ] );  
(ttl [ tll-values ] | ttl-except [ tll-values ] );
```

Common Layer 2 Firewall Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit firewall\] Hierarchy](#)” on page 325 instead of the statements being repeated.

- [edit firewall family ethernet-switching filter *filter-name* term *term-name* from]
- [edit firewall family vpls filter *filter-name* term *term-name* from]

The common Layer 2 firewall match conditions are as follows:

```
destination-mac-address {  
    mac-address <except>;  
}  
(destination-port [ port-names ] | destination-port-except [ port-names ] );  
(dscp [ code-point-values ] | dscp-except [ code-point-values ] );  
(ether-type [ protocol-types ] | ether-type-except [ protocol-types ] );  
(forwarding-class [ class-names ] | forwarding-class-except [ class-names ] );  
(icmp-code [ codes ] | icmp-code-except [ codes ] );  
(icmp-type [ types ] | icmp-type-except [ types ] );  
(interface-group [ group-names ] | interface-group-except [ group-names ] );  
ip-address {  
    ip-prefix </prefix-length> <except>;  
}  
ip-destination-address {  
    ip-prefix </prefix-length> <except>;  
}  
(ip-precedence [ precedence-names ] | ip-precedence-except [ precedence-names ] );  
(ip-protocol [ protocol-names ] | ip-protocol-except [ protocol-names ] );  
ip-source-address ip-prefix </prefix-length>;  
(learn-vlan-lp-priority [ priorities ] | learn-vlan-lp-priority [ priorities ] );  
(learn-vlan-id [ vlan-ids ] | learn-vlan-id-except [ vlan-ids ] );  
(loss-priority [ priorities ] | loss-priority-except [ priorities ] );  
(port [ port-names ] | port-except [ port-names ] );  
source-mac-address {  
    mac-address <except>;  
}  
(source-port [ port-names ] | source-port-except [ port-names ] );  
tcp-flags flag;  
(traffic-type [ broadcast known-unicast multicast unknown-unicast ] |  
    traffic-type-except [ broadcast known-unicast multicast unknown-unicast ] );  
(user-vlan-lp-priority [ priorities ] | user-vlan-lp-priority [ priorities ] );  
(user-vlan-id [ vlan-ids ] | user-vlan-id-except [ vlan-ids ] );  
(vlan-ether-type [ protocol-types ] | vlan-ether-type-except [ protocol-types ] );
```

Complete [edit firewall] Hierarchy

```

firewall {
  family (any | ccc | ethernet-switching | inet | inet6 | mpls | vpls) {
    ... the family subhierarchies appear after the main [edit firewall] hierarchy ...
  }
  filter filter-name {
    accounting-profile [ profile-names ];
    enhanced-mode;
    interface-shared-with;
    interface-specific;
    physical-interface-policer;
    term term-name {
      filter filter-name;
      from {
        ... statements in Common IP Firewall Match Conditions on page 322 AND
        statements in Common IPv4 Firewall Match Conditions on page 323 ...
      }
      then {
        ... statements in Common Firewall Actions on page 320 AND
        statements in Common IP Firewall Actions on page 321 AND
        statements in Common IPv4 Firewall Actions on page 321 ...
      }
    }
  }
  hierarchical-policer policer-name {
    aggregate {
      if-exceeding {
        bandwidth-limit bps;
        burst-size-limit bytes;
      }
      then {
        discard;
        forwarding-class class-name;
        loss-priority (high | low | medium-high | medium-low);
      }
    }
    logical-interface-policer;
    physical-interface-policer;
    premium {
      if-exceeding {
        bandwidth-limit bps;
        burst-size-limit bytes;
      }
      then {
        discard;
      }
    }
  }
  shared-bandwidth-policer;
  interface-set interface-set-name {
    interface-name;
  }
  load-balance-group group-name {
    next-hop-group [ group-names ];
  }
}

```

```
}
policer policer-name {
  filter-specific;
  if-exceeding {
    (bandwidth-limit bps | bandwidth-percent percentage);
    burst-size-limit bytes;
  }
  logical-bandwidth-policer;
  logical-interface-policer;
  physical-interface-policer;
  then {
    discard;
    forwarding-class class-name;
    loss-priority (high | low | medium-high | medium-low);
  }
}
three-color-policer policer-name {
  action {
    loss-priority high then discard;
  }
  filter-specific;
  logical-interface-policer;
  physical-interface-policer;
  shared-bandwidth-policer;
  single-rate {
    (color-aware | color-blind);
    committed-burst-size bytes;
    committed-information-rate bps;
    excess-burst-size bytes;
  }
  two-rate {
    (color-aware | color-blind);
    committed-burst-size bytes;
    committed-information-rate bps;
    peak-burst-size bytes;
    peak-information-rate bps;
  }
}
}

firewall {
  family any {
    filter filter-name {
      interface-shared;
      term term-name {
        from {
          (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
          interface interface-name;
          interface-set set-name;
          (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
          (packet-length [ values ] | packet-length-except [ values ]);
        }
        then {
          ... statements in Common Firewall Actions on page 320 PLUS ...
          (accept | discard);
        }
      }
    }
  }
}
```

```

    }
  }
}

firewall {
  family ccc {
    filter filter-name {
      accounting-profile [ profile-names ];
      physical-interface-filter;
      interface-specific;
      term term-name {
        filter filter-name;
        from {
          (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
          (interface-group [ group-names ] | interface-group-except [ group-names ]);
          (learn-vlan-1p-priority [ priorities ] | learn-vlan-1p-priority [ priorities ]);
          (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
          (user-vlan-1p-priority [ priorities ] | user-vlan-1p-priority [ priorities ]);
        }
        then {
          ... statements in Common Firewall Actions on page 320 PLUS ...
          (accept | discard);
          port-mirror-instance instance-name;
        }
      }
    }
  }
}

firewall {
  family ethernet-switching {
    filter filter-name {
      interface-specific;
      term term-name {
        from {
          destination-address {
            ip-prefix</prefix-length>;
          }
          destination-mac-address {
            mac-address;
          }
          destination-port [ port-names ];
          destination-prefix-list {
            list-name;
          }
          dot1q-tag [ tag-values ];
          dot1q-user-priority [ priority-values ];
          dscp [ code-point-values ];
          ether-type [ protocol-names ];
          fragment-flags flag;
          icmp-code [ codes ];
          icmp-type [ types ];
          interface interface-name;
          is-fragment;

```

```

precedence [ precedence-names ];
protocol [ protocol-names ];
source-address {
    ip-prefix < / prefix-length >;
}
source-mac-address {
    mac-address;
}
source-port [ port-names ];
source-prefix-list {
    list-name;
}
tcp-established;
tcp-flags flag;
tcp-initial;
vlan [ vlan-names ];
}
then {
    (accept | discard);
    analyzer analyzer-name;
    count counter-name;
    forwarding-class class-name;
    interface interface-name;
    log;
    loss-priority (high | low);
    policer policer-name;
    syslog;
    vlan vlan-name;
}
}
}
}
}

firewall {
    family inet {
        dialer-filter filter-name {
            accounting-profile [ profile-names ];
            term term-name {
                from {
                    ... statements in Common IP Firewall Match Conditions on page 322 AND
                    statements in Common IPv4 Firewall Match Conditions on page 323 EXCEPT
                    FOR ...
                    (ah-spi [ values ] | ah-spi-except [ values ]); # NOT valid at this level
                    (destination-class [ class-names ] |
                     destination-class-except [ class-names ]); # NOT valid at this level
                    interface interface-name; # NOT valid at this level
                    (loss-priority [ priorities ] | loss-priority-except [ priorities ]); # NOT valid at
                     this level
                    service-filter-hit; # NOT valid at this level
                    (source-class [ class-names ] | source-class-except [ class-names ]); # NOT
                     valid at this level
                }
            }
            then {
                (ignore | note);
                log;
            }
        }
    }
}

```



```

        sample;
        syslog;
    }
}
filter filter-name {
    accounting-profile [ profile-names ];
    interface-specific;
    term term-name {
        filter filter-name;
        from {
            ... statements in Common IP Firewall Match Conditions on page 322 AND
               statements in Common IPv4 Firewall Match Conditions on page 323 ...
        }
        then {
            ... statements in Common Firewall Actions on page 320 AND
               statements in Common IP Firewall Actions on page 321 AND
               statements in Common IPv4 Firewall Actions on page 321 ...
        }
    }
}
prefix-action name {
    count;
    destination-prefix-length prefix-length;
    filter-specific;
    policer policer-name;
    source-prefix-length prefix-length;
    subnet-prefix-length prefix-length;
}
service-filter filter-name {
    term term-name {
        from {
            address {
                ip-prefix</prefix-length>;
            }
            (ah-spi [ values ] | ah-spi-except [ values ]);
            destination-address {
                ip-prefix</prefix-length>;
            }
            (destination-port [ port-names ] | destination-port-except [ port-names ]);
            destination-prefix-list {
                list-name;
            }
            (esp-spi [ values ] | esp-spi-except [ values ]);
            first-fragment;
            fragment-flags flag;
            (fragment-offset [ offsets ] | fragment-offset-except [ offsets ]);
            (interface-group [ group-names ] | interface-group-except [ group-names ]);
            (ip-options [ option-names ] | ip-options-except [ option-names ]);
            is-fragment;
            (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
            (port [ port-names ] | port-except [ port-names ]);
            prefix-list {
                list-name;
            }
            (protocol [ protocol-names ] | protocol-except [ protocol-names ]);
        }
    }
}

```

```

        source-address {
            ip-prefix </prefix-length>;
        }
        (source-port [ port-names ] | source-port-except [ port-names ]);
        source-prefix-list {
            list-name;
        }
        tcp-flags flag-name;
    }
    then {
        count counter-name;
        log;
        port-mirror;
        sample;
        (service | skip);
    }
}
}
simple-filter filter-name {
    term term-name {
        from {
            destination-address ip-prefix </prefix-length>;
            destination-port port-name;
            forwarding-class [ class-names ];
            protocol protocol-name;
            source-address ip-prefix </prefix-length>;
            source-port port-name;
        }
        then {
            forwarding-class class-name;
            loss-priority (high | low | medium-high | medium-low);
            policer policer-name;
        }
    }
}
}
}
}
firewall {
    family inet6 {
        dialer-filter filter-name {
            accounting-profile [ profile-names ];
            term term-name {
                from {
                    ... statements in Common IP Firewall Match Conditions on page 322 PLUS ...
                    (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
                    ... BUT NOT ...
                    (destination-class [ class-names ] |
                     destination-class-except [ class-names ]); # NOT valid at this level
                    (forwarding-class [ class-names ] |
                     forwarding-class-except [ class-names ]); # NOT valid at this level
                    interface interface-name; # NOT valid at this level
                    (interface-group [ group-names ] | interface-group-except [ group-names ]); #
                     NOT valid at this level
                    (loss-priority [ priorities ] | loss-priority-except [ priorities ]); # NOT valid at
                     this level
                }
            }
        }
    }
}

```

```

        service-filter-hit; # NOT valid at this level
        (source-class [ class-names ] | source-class-except [ class-names ]); # NOT
            valid at this level
        tcp-established; # NOT valid at this level
        tcp-flags flag; # NOT valid at this level
        tcp-initial; # NOT valid at this level
    }
    then {
        (ignore | note);
        log;
        sample;
        syslog;
    }
}
}
filter filter-name {
    accounting-profile [ profile-names ];
    interface-specific;
    term term-name {
        filter filter-name;
        from {
            ... statements in Common IP Firewall Match Conditions on page 322 PLUS ...
            (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
            (traffic-class [ code-point-values ] | traffic-class-except [ code-point-values ]);
        }
        then {
            ... statements in Common Firewall Actions on page 320 AND
            statements in Common IP Firewall Actions on page 321 PLUS ...
            (accept | discard | reject <address-unreachable | administratively-prohibited |
                beyond-scope | fragmentation-needed | no-route | port-unreachable |
                tcp-reset>);
        }
    }
}
}
service-filter filter-name {
    term term-name {
        from {
            address {
                ip-prefix</prefix-length>;
            }
            (ah-spi [ values ] | ah-spi-except [ values ]);
            destination-address {
                ip-prefix</prefix-length>;
            }
            (destination-port [ port-names ] | destination-port-except [ port-names ]);
            destination-prefix-list {
                list-name;
            }
            (esp-spi [ values ] | esp-spi-except [ values ]);
            (interface-group [ group-names ] | interface-group-except [ group-names ]);
            (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
            (port [ port-names ] | port-except [ port-names ]);
            prefix-list {
                list-name;
            }
            source-address {

```

```

        ip-prefix </prefix-length>;
    }
    (source-port [ port-names ] | source-port-except [ port-names ]);
    source-prefix-list {
        list-name;
    }
    tcp-flags flag-name;
}
then {
    count counter-name;
    log;
    port-mirror;
    sample;
    (service | skip);
}
}
}
}
}

firewall {
    family mpls {
        filter filter-name {
            accounting-profile [ profile-names ];
            interface-specific;
            physical-interface-filter;
            term term-name {
                from {
                    (exp [ exp-bits ] | exp-except [ exp-bits ]);
                }
                then {
                    (ignore | note);
                    log;
                    sample;
                    syslog;
                }
            }
        }
    }
    filter filter-name {
        accounting-profile [ profile-names ];
        interface-specific;
        physical-interface-filter;
        term term-name {
            filter filter-name;
            from {
                (exp [ exp-bits ] | exp-except [ exp-bits ]);
                (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
                interface interface-name;
                interface-set set-name;
                (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
            }
            then {
                ... statements in Common Firewall Actions on page 320 PLUS ...
                (accept | discard);
                sample;
            }
        }
    }
}

```

```

    }
  }
}

firewall {
  family vpls {
    filter filter-name {
      accounting-profile [ profile-names ];
      interface-specific;
      term term-name {
        filter filter-name;
        from {
          ... statements in Common Layer 2 Firewall Match Conditions on page 323 ...
        }
        then {
          ... statements in Common Firewall Actions on page 320 PLUS ...
          (accept | discard);
          port-mirror;
          port-mirror-instance instance-name;
        }
      }
    }
  }
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

Administration

- [Firewall Filters Standards on page 4703](#)
- [Firewall Filters Reference on page 4704](#)
- [Standard Firewall Filter Match Conditions and Actions on page 4706](#)
- [Standard Firewall Filter Match Conditions and Actions for ACX Series Routers on page 4751](#)
- [Service Filter Match Conditions and Actions on page 4760](#)
- [Reference Information for Firewall Filters in Logical Systems on page 4767](#)
- [Firewall Filters Statement Hierarchies on page 4773](#)
- [Summary of Firewall Filters Configuration Statements on page 4785](#)

Firewall Filters Standards

- [Supported Standards for Filtering on page 4703](#)

Supported Standards for Filtering

The Junos OS supports the following RFCs related to filtering:

- RFC 792, *Internet Control Message Protocol*
- RFC 2460, *Internet Protocol, Version 6 (IPv6)*
- RFC 2474, *Definition of the Differentiated Services (DS) Field*

- RFC 2475, *An Architecture for Differentiated Services*
- RFC 2597, *Assured Forwarding PHB Group*
- RFC 3246, *An Expedited Forwarding PHB (Per-Hop Behavior)*
- RFC 4291, *IP Version 6 Addressing Architecture*
- RFC 4443, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*

Related Documentation

- [Standard Stateless Firewall Filter Overview on page 4475](#)
- [Service Filter Overview on page 4522](#)
- [Simple Filter Overview on page 4529](#)
- [Stateless Firewall Filters in Logical Systems Overview on page 4535](#)

Firewall Filters Reference

- [Using the CLI Editor in Configuration Mode on page 4704](#)

Using the CLI Editor in Configuration Mode

This topic describes some of the basic commands that you must use to enter configuration mode in the command-line interface (CLI) editor, navigate through the configuration hierarchy, get help, and commit or revert the changes that you make during the configuration session.

Task	Command/Statement	Example
Edit Your Configuration		
Enter configuration mode. When you first log in to the device, the device is in operational mode. You must explicitly enter configuration mode. When you do, the CLI prompt changes from user@host> to user@host# and the hierarchy level appears in square brackets.	configure	user@host> configure [edit] user@host#
Create a statement hierarchy. You can use the edit command to simultaneously create a hierarchy and move to that new level in the hierarchy. You cannot use the edit command to change the value of identifiers.	edit hierarchy-level value	[edit] user@host# edit security zones security-zone myzone [edit security zones security-zone myzone] user@host#
Create a statement hierarchy and set identifier values. The set command is similar to edit except that your current level in the hierarchy does not change.	set hierarchy-level value	[edit] user@host# set security zones security-zone myzone [edit] user@host#

Task	Command/Statement	Example
Navigate the Hierarchy		
Navigate down to an existing hierarchy level.	<code>edit hierarchy-level</code>	<pre>[edit] user@host# edit security zones [edit security zones] user@host#</pre>
Navigate up one level in the hierarchy.	<code>up</code>	<pre>[edit security zones] user@host# up [edit security] user@host#</pre>
Navigate to the top of the hierarchy.	<code>top</code>	<pre>[edit security zones] user@host# top [edit] user@host#</pre>
Commit or Revert Changes		
Commit your configuration.	<code>commit</code>	<pre>[edit] user@host# commit commit complete</pre>
<p>Roll back changes from the current session.</p> <p>Use the rollback command to revert all changes from the current configuration session. When you run the rollback command before exiting your session or committing changes, the software loads the most recently committed configuration onto the device. You must enter the rollback statement at the edit level in the hierarchy.</p>	<code>rollback</code>	<pre>[edit] user@host# rollback load complete</pre>
Exit Configuration Mode		
Commit the configuration and exit configuration mode.	<code>commit and-quit</code>	<pre>[edit] user@host# commit and-quit user@host></pre>
<p>Exit configuration mode without committing your configuration.</p> <p>You must navigate to the top of the hierarchy using the up or top commands before you can exit configuration mode.</p>	<code>exit</code>	<pre>[edit] user@host# exit The configuration has been changed but not committed Exit with uncommitted changes? [yes,no] (yes)</pre>
Get Help		

Task	Command/Statement	Example
Display a list of valid options for the current hierarchy level.	?	<pre>[edit] user@host# edit security zones ?</pre> <p>Possible completions:</p> <pre><[Enter]> Execute this command > functional-zone Functional zone > security-zone Security zones Pipe through a command [edit]</pre>

Related Documentation

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

Standard Firewall Filter Match Conditions and Actions

- [Standard Firewall Filter Match Conditions for Protocol-Independent Traffic on page 4706](#)
- [Standard Firewall Filter Match Conditions for IPv4 Traffic on page 4707](#)
- [Standard Firewall Filter Match Conditions for IPv6 Traffic on page 4716](#)
- [Standard Firewall Filter Match Conditions for MPLS Traffic on page 4723](#)
- [Standard Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic on page 4725](#)
- [Standard Firewall Filter Match Conditions for VPLS Traffic on page 4727](#)
- [Standard Firewall Filter Match Conditions for Layer 2 CCC Traffic on page 4734](#)
- [Standard Firewall Filter Match Conditions for Layer 2 Bridging Traffic on page 4737](#)
- [Standard Firewall Filter Terminating Actions on page 4742](#)
- [Standard Firewall Filter Nonterminating Actions on page 4744](#)

Standard Firewall Filter Match Conditions for Protocol-Independent Traffic

You can configure a standard stateless firewall filter with match conditions for protocol-independent traffic (**family any**).



NOTE: Protocol-independent standard firewall filters—firewall filters configured at the [edit firewall family any] hierarchy level— are not supported on the router (or switch) loopback interface (lo0).

Table 340 on page 4707 describes the **match-conditions** you can configure at the [edit firewall family any filter *filter-name* term *term-name* from] hierarchy level.

Table 340: Standard Firewall Filter Match Conditions for Protocol-Independent Traffic

Match Condition	Description
forwarding-class <i>class</i>	<p>Match the forwarding class of the packet.</p> <p>Specify assured-forwarding, best-effort, expedited-forwarding, or network-control.</p> <p>For information about forwarding classes and router-internal output queues, see the <i>Junos OS Class of Service Configuration Guide</i>.</p> <p>NOTE: On T4000 Type 5 FPCs, a filter attached at the Layer 2 application point (that is, at the logical interface level) is unable to match with the forwarding class of a packet that is set by a Layer 3 classifier such as DSCP, DSCP V6, inet-precedence, and mpls-exp.</p>
forwarding-class-except <i>class</i>	Do not match on the forwarding class. For details, see the forwarding-class match condition.
interface <i>interface-name</i>	<p>Match the interface on which the packet was received.</p> <p>NOTE: If you configure this match condition with an interface that does not exist, the term does not match any packet.</p>
interface-set <i>interface-set-name</i>	<p>Match the interface on which the packet was received to the specified interface set.</p> <p>To define an interface set, include the interface-set statement at the [edit firewall] hierarchy level. For more information, see "Filtering Packets Received on an Interface Set Overview" on page 4515.</p>
packet-length <i>bytes</i>	Match the length of the received packet, in bytes. The length refers only to the IP packet, including the packet header, and does not include any Layer 2 encapsulation overhead.
packet-length-except <i>bytes</i>	Do not match on the received packet length, in bytes. For details, see the packet-length match type.

Related Documentation

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Standard Firewall Filter Terminating Actions on page 4742](#)
- [Standard Firewall Filter Nonterminating Actions on page 4744](#)

Standard Firewall Filter Match Conditions for IPv4 Traffic

You can configure a standard stateless firewall filter with match conditions for Internet Protocol version 4 (IPv4) traffic (**family inet**). [Table 341 on page 4707](#) describes the **match-conditions** you can configure at the **[edit firewall family inet filter filter-name term term-name from]** hierarchy level.

Table 341: Standard Firewall Filter Match Conditions for IPv4 Traffic

Match Condition	Description
address <i>address</i> [except]	Match the IPv4 source or destination address field unless the except option is included. If the option is included, do not match the IPv4 source or destination address field.

Table 341: Standard Firewall Filter Match Conditions for IPv4 Traffic (*continued*)

Match Condition	Description
ah-spi <i>spi-value</i>	(M Series routers, except M120 and M320) Match the IPsec authentication header (AH) security parameter index (SPI) value. NOTE: This match condition is not supported on PTX series packet transport switches.
ah-spi-except <i>spi-value</i>	(M Series routers, except M120 and M320) Do not match the IPsec AH SPI value. NOTE: This match condition is not supported on PTX series packet transport switches.
apply-groups	Specify which groups to inherit configuration data from. You can specify more than one group name. You must list them in order of inheritance priority. The configuration data in the first group takes priority over the data in subsequent groups.
apply-groups-except	Specify which groups not to inherit configuration data from. You can specify more than one group name.
destination-address <i>address</i> [except]	Match the IPv4 destination address field unless the except option is included. If the option is included, do not match the IPv4 destination address field.. You cannot specify both the address and destination-address match conditions in the same term.
destination-class <i>class-names</i>	Match one or more specified destination class names (sets of destination prefixes grouped together and given a class name). For more information, see “Firewall Filter Match Conditions Based on Address Classes” on page 4502 . NOTE: This match condition is not supported on PTX series packet transport switches.
destination-class-except <i>class-names</i>	Do not match one or more specified destination class names. For details, see the destination-class match condition. NOTE: This match condition is not supported on PTX series packet transport switches.
destination-port <i>number</i>	Match the UDP or TCP destination port field. You cannot specify both the port and destination-port match conditions in the same term. If you configure this match condition, we recommend that you also configure the protocol udp or protocol tcp match statement in the same term to specify which protocol is being used on the port. In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed): afs (1483), bgp (179), biff (512), bootpc (68), bootps (67), cmd (514), cvspserver (2401), dhcp (67), domain (53), eklogin (2105), ekshell (2106), exec (512), finger (79), ftp (21), ftp-data (20), http (80), https (443), ident (113), imap (143), kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544), ldap (389), ldp (646), login (513), mobileip-agent (434), mobilip-mn (435), msdp (639), netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123), pop3 (110), pptp (1723), printer (515), radacct (1813), radius (1812), rip (520), rkinit (2108), smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514), tacacs (49), tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525), who (513), or xdmcp (177).
destination-port-except <i>number</i>	Do not match the UDP or TCP destination port field. For details, see the destination-port match condition.

Table 341: Standard Firewall Filter Match Conditions for IPv4 Traffic (*continued*)

Match Condition	Description
destination-prefix-list <i>name</i> [except]	<p>Match destination prefixes in the specified list unless the except option is included. If the option is included, do not match the destination prefixes in the specified list.</p> <p>Specify the name of a prefix list defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.</p>
dscp <i>number</i>	<p>Match the Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant 6 bits of this byte form the DSCP. For more information, see the <i>Junos OS Class of Service Configuration Guide</i>.</p> <p>You can specify a numeric value from 0 through 63. To specify the value in hexadecimal form, include 0x as a prefix. To specify the value in binary form, include b as a prefix.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> • RFC 3246, <i>An Expedited Forwarding PHB (Per-Hop Behavior)</i>, defines one code point: ef (46). • RFC 2597, <i>Assured Forwarding PHB Group</i>, defines 4 classes, with 3 drop precedences in each class, for a total of 12 code points: <ul style="list-style-type: none"> • af11 (10), af12 (12), af13 (14) • af21 (18), af22 (20), af23 (22) • af31 (26), af32 (28), af33 (30) • af41 (34), af42 (36), af43 (38)
dscp-except <i>number</i>	Do not match on the DSCP number. For more information, see the dscp match condition.
esp-spi <i>spi-value</i>	<p>Match the IPsec encapsulating security payload (ESP) SPI value. Match on this specific SPI value. You can specify the ESP SPI value in hexadecimal, binary, or decimal form.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p>
esp-spi-except <i>spi-value</i>	<p>Match the IPsec ESP SPI value. Do not match on this specific SPI value.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p>
first-fragment	<p>Match if the packet is the first fragment of a fragmented packet. Do not match if the packet is a trailing fragment of a fragmented packet. The first fragment of a fragmented packet has a fragment offset value of 0.</p> <p>This match condition is an alias for the bit-field match condition fragment-offset 0 match condition.</p> <p>To match both first and trailing fragments, you can use two terms that specify different match conditions: first-fragment and is-fragment.</p>
forwarding-class <i>class</i>	<p>Match the forwarding class of the packet.</p> <p>Specify assured-forwarding, best-effort, expedited-forwarding, or network-control.</p> <p>For information about forwarding classes and router-internal output queues, see the <i>Junos OS Class of Service Configuration Guide</i>.</p>
forwarding-class-except <i>class</i>	Do not match the forwarding class of the packet. For details, see the forwarding-class match condition.

Table 341: Standard Firewall Filter Match Conditions for IPv4 Traffic (*continued*)

Match Condition	Description
fragment-flags <i>number</i>	<p>(Ingress only) Match the three-bit IP fragmentation flags field in the IP header.</p> <p>In place of the numeric field value, you can specify one of the following keywords (the field values are also listed): dont-fragment (0x4), more-fragments (0x2), or reserved (0x8).</p>
fragment-offset <i>value</i>	<p>Match the 13-bit fragment offset field in the IP header. The value is the offset, in 8-byte units, in the overall datagram message to the data fragment. Specify a numeric value, a range of values, or a set of values. An offset value of 0 indicates the first fragment of a fragmented packet.</p> <p>The first-fragment match condition is an alias for the fragment-offset 0 match condition.</p> <p>To match both first and trailing fragments, you can use two terms that specify different match conditions (first-fragment and is-fragment).</p>
fragment-offset-except <i>number</i>	Do not match the 13-bit fragment offset field.
icmp-code <i>number</i>	<p>Match the ICMP message code field.</p> <p>If you configure this match condition, we recommend that you also configure the protocol icmp match condition in the same term.</p> <p>If you configure this match condition, you must also configure the icmp-type <i>message-type</i> match condition in the same term. An ICMP message code provides more specific information than an ICMP message type, but the meaning of an ICMP message code is dependent on the associated ICMP message type.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> parameter-problem: ip-header-bad (0), required-option-missing (1) redirect: redirect-for-host (1), redirect-for-network (0), redirect-for-tos-and-host (3), redirect-for-tos-and-net (2) time-exceeded: ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0) unreachable: communication-prohibited-by-filtering (13), destination-host-prohibited (10), destination-host-unknown (7), destination-network-prohibited (9), destination-network-unknown (6), fragmentation-needed (4), host-precedence-violation (14), host-unreachable (1), host-unreachable-for-TOS (12), network-unreachable (0), network-unreachable-for-TOS (11), port-unreachable (3), precedence-cutoff-in-effect (15), protocol-unreachable (2), source-host-isolated (8), source-route-failed (5)
icmp-code-except <i>message-code</i>	Do not match the ICMP message code field. For details, see the icmp-code match condition.
icmp-type <i>number</i>	<p>Match the ICMP message type field.</p> <p>If you configure this match condition, we recommend that you also configure the protocol icmp match condition in the same term.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): echo-reply (0), echo-request (8), info-reply (16), info-request (15), mask-request (17), mask-reply (18), parameter-problem (12), redirect (5), router-advertisement (9), router-solicit (10), source-quench (4), time-exceeded (11), timestamp (13), timestamp-reply (14), or unreachable (3).</p>

Table 341: Standard Firewall Filter Match Conditions for IPv4 Traffic (*continued*)

Match Condition	Description
icmp-type-except <i>message-type</i>	Do not match the ICMP message type field. For details, see the icmp-type match condition.
interface <i>interface-name</i>	<p>Match the interface on which the packet was received.</p> <p>NOTE: If you configure this match condition with an interface that does not exist, the term does not match any packet.</p>
interface-group <i>group-number</i>	<p>Match the logical interface on which the packet was received to the specified interface group or set of interface groups. For <i>group-number</i>, specify a single value or a range of values from 0 through 255.</p> <p>To assign a logical interface to an interface group <i>group-number</i>, specify the <i>group-number</i> at the [interfaces interface-name unit number family family filter group] hierarchy level.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p> <p>For more information, see “Filtering Packets Received on a Set of Interface Groups Overview” on page 4514.</p>
interface-group-except <i>group-number</i>	<p>Do not match the logical interface on which the packet was received to the specified interface group or set of interface groups. For details, see the interface-group match condition.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p>
interface-set <i>interface-set-name</i>	<p>Match the interface on which the packet was received to the specified interface set.</p> <p>To define an interface set, include the interface-set statement at the [edit firewall] hierarchy level.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p> <p>For more information, see “Filtering Packets Received on an Interface Set Overview” on page 4515.</p>

Table 341: Standard Firewall Filter Match Conditions for IPv4 Traffic (*continued*)

Match Condition	Description
ip-options values	<p>Match the 8-bit IP option field, if present, to the specified value or list of values.</p> <p>In place of a numeric value, you can specify one of the following text synonyms (the option values are also listed): loose-source-route (131), record-route (7), router-alert (148), security (130), stream-id (136), strict-source-route (137), or timestamp (68).</p> <p>To match <i>any</i> value for the IP option, use the text synonym any. To match on <i>multiple</i> values, specify the list of values within square brackets ('[' and ']'). To match a <i>range</i> of values, use the value specification [<i>value1-value2</i>].</p> <p>For example, the match condition ip-options [0-147] matches on an IP options field that contains the loose-source-route, record-route, or security values, or any other value from 0 through 147. However, this match condition does not match on an IP options field that contains only the router-alert value (148).</p> <p>For most interfaces, a filter term that specifies an ip-option match on one or more <i>specific</i> IP option values (a value other than any) causes packets to be sent to the Routing Engine so that the kernel can parse the IP option field in the packet header.</p> <ul style="list-style-type: none"> For a firewall filter term that specifies an ip-option match on one or more specific IP option values, you cannot specify the count, log, or syslog nonterminating actions <i>unless</i> you also specify the discard terminating action in the same term. This behavior prevents double-counting of packets for a filter applied to a transit interface on the router or switch. Packets processed on the kernel might be dropped in case of a system bottleneck. To ensure that matched packets are instead sent to the Packet Forwarding Engine (where packet processing is implemented in hardware), use the ip-options any match condition. <p>The 10-Gigabit Ethernet Modular Port Concentrator (MPC), 100-Gigabit Ethernet MPC, 60-Gigabit Ethernet MPC, 60-Gigabit Queuing Ethernet MPC, and 60-Gigabit Ethernet Enhanced Queuing MPC on MX Series routers (and EX Series switches) are capable of parsing the IP option field of the IPv4 packet header. For interfaces configured on those MPCs, <i>all</i> packets that are matched using the ip-options match condition are sent to the Packet Forwarding Engine for processing.</p> <p>NOTE: On M and T series routers, firewall filters cannot count ip-options packets on a per option type and per interface basis. A limited work around is to use the show pfe statistics ip options command to see ip-options statistics on a per PFE basis. See <i>show pfe statistics ip</i> for sample output.</p>
ip-options-except values	<p>Do not match the IP option field to the specified value or list of values. For details about specifying the values, see the ip-options match condition.</p>
is-fragment	<p>Match if the packet is a trailing fragment of a fragmented packet. Do not match the first fragment of a fragmented packet.</p> <p>NOTE: To match both first and trailing fragments, you can use two terms that specify different match conditions (first-fragment and is-fragment).</p>

Table 341: Standard Firewall Filter Match Conditions for IPv4 Traffic (*continued*)

Match Condition	Description
loss-priority level	<p>Match the packet loss priority (PLP) level.</p> <p>Specify a single level or multiple levels: low, medium-low, medium-high, or high.</p> <p>Supported on M120 and M320 routers; M7i and M10i routers with the Enhanced CFEB (CFEB-E); and MX Series routers.</p> <p>For IP traffic on M320, MX Series and T Series routers with Enhanced II Flexible PIC Concentrators (FPCs) and EX Series switches, you must include the tri-color statement at the [edit class-of-service] hierarchy level to commit a PLP configuration with any of the four levels specified. If the tri-color statement is not enabled, you can only configure the high and low levels. This applies to all protocol families.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p> <p>For information about the tri-color statement and for information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>Junos OS Class of Service Configuration Guide</i>.</p>
loss-priority-except level	<p>Do not match the PLP level. For details, see the loss-priority match condition.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p>
packet-length bytes	Match the length of the received packet, in bytes. The length refers only to the IP packet, including the packet header, and does not include any Layer 2 encapsulation overhead.
packet-length-except bytes	Do not match the length of the received packet, in bytes. For details, see the packet-length match type.
port number	<p>Match the UDP or TCP source or destination port field.</p> <p>If you configure this match condition, you cannot configure the destination-port match condition or the source-port match condition in the same term.</p> <p>If you configure this match condition, we recommend that you also configure the protocol udp or protocol tcp match statement in the same term to specify which protocol is being used on the port.</p> <p>In place of the numeric value, you can specify one of the text synonyms listed under destination-port.</p>
port-except number	Do not match the UDP or TCP source or destination port field. For details, see the port match condition.
precedence ip-precedence-value	<p>Match the IP precedence field.</p> <p>In place of the numeric field value, you can specify one of the following text synonyms (the field values are also listed): critical-ecp (0xa0), flash (0x60), flash-override (0x80), immediate (0x40), internet-control (0xc0), net-control (0xe0), priority (0x20), or routine (0x00). You can specify precedence in hexadecimal, binary, or decimal form.</p>

Table 341: Standard Firewall Filter Match Conditions for IPv4 Traffic (*continued*)

Match Condition	Description
precedence-except ip-precedence-value	Do not match the IP precedence field. In place of the numeric field value, you can specify one of the following text synonyms (the field values are also listed): critical-ecp (0xa0), flash (0x60), flash-override (0x80), immediate (0x40), internet-control (0xc0), net-control (0xe0), priority (0x20), or routine (0x00). You can specify precedence in hexadecimal, binary, or decimal form.
prefix-list name	Match the prefixes of the source or destination address fields to the prefixes in the specified list unless the except option is included. If the option is included, do not match the prefixes of the source or destination address fields to the prefixes in the specified list. The prefix list is defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.
protocol number	Match the IP protocol type field. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): ah (51), dstopts (60), egp (8), esp (50), fragment (44), gre (47), hop-by-hop (0), icmp (1), icmp6 (58), icmpv6 (58), igmp (2), ipip (4), ipv6 (41), ospf (89), pim (103), rsvp (46), sctp (132), tcp (6), udp (17), or vrrp (112).
protocol-except number	Do not match the IP protocol type field. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): ah (51), dstopts (60), egp (8), esp (50), fragment (44), gre (47), hop-by-hop (0), icmp (1), icmp6 (58), icmpv6 (58), igmp (2), ipip (4), ipv6 (41), ospf (89), pim (103), rsvp (46), sctp (132), tcp (6), udp (17), or vrrp (112).
rat-type tech-type-value	Match the radio-access technology (RAT) type specified in the 8-bit Tech-Type field of Proxy Mobile IPv4 (PMIPv4) access technology type extension. The technology type specifies the access technology through which the mobile device is connected to the access network. Specify a single value, a range of values, or a set of values. You can specify a technology type as a numeric value from 0 through 255 or as a system keyword. <ul style="list-style-type: none"> The following numeric values are examples of well-known technology types: <ul style="list-style-type: none"> Numeric value 1 matches IEEE 802.3. Numeric value 2 matches IEEE 802.11a/b/g. Numeric value 3 matches IEEE 802.16e Numeric value 4 matches IEEE 802.16m. Text string eutran matches 4G. Text string geran matches 2G. Text string utran matches 3G.
rat-type-except tech-type-value	Do not match the RAT Type.
service-filter-hit	Match a packet received from a filter where a service-filter-hit action was applied. NOTE: This match condition is not supported on PTX series packet transport switches.
source-address address [except]	Match the IPv4 address of the source node sending the packet unless the except option is included. If the option is included, do not match the IPv4 address of the source node sending the packet. You cannot specify both the address and source-address match conditions in the same term.

Table 341: Standard Firewall Filter Match Conditions for IPv4 Traffic (*continued*)

Match Condition	Description
source-class <i>class-names</i>	<p>Match one or more specified source class names (sets of source prefixes grouped together and given a class name). For more information, see “Firewall Filter Match Conditions Based on Address Classes” on page 4502.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p>
source-class-except <i>class-names</i>	<p>Do not match one or more specified source class names. For details, see the source-class match condition.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p>
source-port <i>number</i>	<p>Match the UDP or TCP source port field.</p> <p>You cannot specify the port and source-port match conditions in the same term.</p> <p>If you configure this match condition for IPv4 traffic, we recommend that you also configure the protocol udp or protocol tcp match statement in the same term to specify which protocol is being used on the port.</p> <p>In place of the numeric value, you can specify one of the text synonyms listed with the destination-port <i>number</i> match condition.</p>
source-port-except <i>number</i>	Do not match the UDP or TCP source port field. For details, see the source-port match condition.
source-prefix-list <i>name</i> [except]	<p>Match source prefixes in the specified list unless the except option is included. If the option is included, do not match the source prefixes in the specified list.</p> <p>Specify the name of a prefix list defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.</p>
tcp-established	<p>Match TCP packets of an established TCP session (packets other than the first packet of a connection). This is an alias for tcp-flags "(ack rst)".</p> <p>This match condition does not implicitly check that the protocol is TCP. To check this, specify the protocol tcp match condition.</p>

Table 341: Standard Firewall Filter Match Conditions for IPv4 Traffic (*continued*)

Match Condition	Description
tcp-flags value	<p>Match one or more of the low-order 6 bits in the 8-bit TCP flags field in the TCP header.</p> <p>To specify individual bit fields, you can specify the following text synonyms or hexadecimal values:</p> <ul style="list-style-type: none"> • fin (0x01) • syn (0x02) • rst (0x04) • push (0x08) • ack (0x10) • urgent (0x20) <p>In a TCP session, the SYN flag is set only in the initial packet sent, while the ACK flag is set in all packets sent after the initial packet.</p> <p>You can string together multiple flags using the bit-field logical operators.</p> <p>For combined bit-field match conditions, see the tcp-established and tcp-initial match conditions.</p> <p>If you configure this match condition, we recommend that you also configure the protocol tcp match statement in the same term to specify that the TCP protocol is being used on the port.</p> <p>For IPv4 traffic only, this match condition does not implicitly check whether the datagram contains the first fragment of a fragmented packet. To check for this condition for IPv4 traffic only, use the first-fragment match condition.</p>
tcp-initial	<p>Match the initial packet of a TCP connection. This is an alias for tcp-flags "(lack & syn)".</p> <p>This condition does not implicitly check that the protocol is TCP. If you configure this match condition, we recommend that you also configure the protocol tcp match condition in the same term.</p>
ttl number	<p>Match the IPv4 time-to-live number. Specify a TTL value or a range of TTL values. For number, you can specify one or more values from 0 through 255. This match condition is supported only on M120, M320, MX Series, and T Series routers.</p>
ttl-except number	<p>Do not match on the IPv4 TTL number. For details, see the ttl match condition.</p>

Related Documentation

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Standard Firewall Filter Terminating Actions on page 4742](#)
- [Standard Firewall Filter Nonterminating Actions on page 4744](#)

Standard Firewall Filter Match Conditions for IPv6 Traffic

You can configure a standard stateless firewall filter with match conditions for Internet Protocol version 6 (IPv6) traffic (**family inet6**). [Table 342 on page 4717](#) describes the **match-conditions** you can configure at the **[edit firewall family inet6 filter filter-name term term-name from]** hierarchy level.

Table 342: Standard Firewall Filter Match Conditions for IPv6 Traffic

Match Condition	Description
address <i>address</i> [except]	Match the IPv6 source or destination address field unless the except option is included. If the option is included, do not match the IPv6 source or destination address field.
apply-groups	Specify which groups to inherit configuration data from. You can specify more than one group name. You must list them in order of inheritance priority. The configuration data in the first group takes priority over the data in subsequent groups.
apply-groups-except	Specify which groups not to inherit configuration data from. You can specify more than one group name.
destination-address <i>address</i> [except]	Match the IPv6 destination address field unless the except option is included. If the option is included, do not match the IPv6 destination address field. You cannot specify both the address and destination-address match conditions in the same term.
destination-class <i>class-names</i>	Match one or more specified destination class names (sets of destination prefixes grouped together and given a class name). NOTE: This match condition is not supported on PTX series packet transport switches. For more information, see “Firewall Filter Match Conditions Based on Address Classes” on page 4502 .
destination-class-except <i>class-names</i>	Do not match one or more specified destination class names. For details, see the destination-class match condition. NOTE: This match condition is not supported on PTX series packet transport switches.
destination-port <i>number</i>	Match the UDP or TCP destination port field. You cannot specify both the port and destination-port match conditions in the same term. If you configure this match condition, we recommend that you also configure the next-header udp or next-header tcp match condition in the same term to specify which protocol is being used on the port. In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed): afs (1483), bgp (179), biff (512), bootpc (68), bootps (67), cmd (514), cvspserver (2401), dhcp (67), domain (53), eklogin (2105), ekshell (2106), exec (512), finger (79), ftp (21), ftp-data (20), http (80), https (443), ident (113), imap (143), kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544), ldap (389), ldp (646), login (513), mobileip-agent (434), mobileip-mn (435), msdp (639), netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123), pop3 (110), pptp (1723), printer (515), radacct (1813), radius (1812), rip (520), rkinit (2108), smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514), tacacs (49), tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525), who (513), or xdmcp (177).
destination-port-except <i>number</i>	Do not match the UDP or TCP destination port field. For details, see the destination-port match condition.
destination-prefix-list <i>prefix-list-name</i> [except]	Match the IPv6 destination prefix to the specified list unless the except option is included. If the option is included, do not match the IPv6 destination prefix to the specified list. The prefix list is defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.

Table 342: Standard Firewall Filter Match Conditions for IPv6 Traffic (*continued*)

Match Condition	Description
forwarding-class <i>class</i>	<p>Match the forwarding class of the packet.</p> <p>Specify assured-forwarding, best-effort, expedited-forwarding, or network-control.</p> <p>For information about forwarding classes and router-internal output queues, see the <i>Junos OS Class of Service Configuration Guide</i>.</p>
forwarding-class-except <i>class</i>	Do not match the forwarding class of the packet. For details, see the forwarding-class match condition.
hop-limit <i>hop-limit</i>	Match the hop limit to the specified hop limit or set of hop limits. For hop-limit , specify a single value or a range of values from 0 through 255..
hop-limit-except <i>message-code</i>	Do not match the hop limit to the specified hop limit or set of hop limits. For details, see the hop-limit match condition.
icmp-code <i>message-code</i>	<p>Match the ICMP message code field.</p> <p>If you configure this match condition, we recommend that you also configure the next-header icmp or next-header icmp6 match condition in the same term.</p> <p>If you configure this match condition, you must also configure the icmp-type message-type match condition in the same term. An ICMP message code provides more specific information than an ICMP message type, but the meaning of an ICMP message code is dependent on the associated ICMP message type.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> parameter-problem: ip6-header-bad (0), unrecognized-next-header (1), unrecognized-option (2) time-exceeded: ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0) destination-unreachable: administratively-prohibited (1), address-unreachable (3), no-route-to-destination (0), port-unreachable (4)
icmp-code-except <i>message-code</i>	Do not match the ICMP message code field. For details, see the icmp-code match condition.
icmp-type <i>message-type</i>	<p>Match the ICMP message type field.</p> <p>If you configure this match condition, we recommend that you also configure the next-header icmp or next-header icmp6 match condition in the same term.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): destination-unreachable (1), echo-reply (129), echo-request (128), membership-query (130), membership-report (131), membership-termination (132), neighbor-advertisement (136), neighbor-solicit (135), node-information-reply (140), node-information-request (139), packet-too-big (2), parameter-problem (4), redirect (137), router-advertisement (134), router-renumbering (138), router-solicit (133), or time-exceeded (3).</p>
icmp-type-except <i>message-type</i>	Do not match the ICMP message type field. For details, see the icmp-type match condition.

Table 342: Standard Firewall Filter Match Conditions for IPv6 Traffic (*continued*)

Match Condition	Description
interface <i>interface-name</i>	<p>Match the interface on which the packet was received.</p> <p>NOTE: If you configure this match condition with an interface that does not exist, the term does not match any packet.</p>
interface-group group-number	<p>Match the logical interface on which the packet was received to the specified interface group or set of interface groups. For group-number, specify a single value or a range of values from 0 through 255.</p> <p>To assign a logical interface to an interface group group-number, specify the group-number at the [interfaces interface-name unit number family family filter group] hierarchy level.</p> <p>For more information, see “Filtering Packets Received on a Set of Interface Groups Overview” on page 4514.</p>
interface-group-except group-number	<p>Do not match the logical interface on which the packet was received to the specified interface group or set of interface groups. For details, see the interface-group match condition.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p>
interface-set interface-set-name	<p>Match the interface on which the packet was received to the specified interface set.</p> <p>To define an interface set, include the interface-set statement at the [edit firewall] hierarchy level.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p> <p>For more information, see “Filtering Packets Received on an Interface Set Overview” on page 4515.</p>
loss-priority level	<p>Match the packet loss priority (PLP) level.</p> <p>Specify a single level or multiple levels: low, medium-low, medium-high, or high.</p> <p>Supported on M120 and M320 routers; M7i and M10i routers with the Enhanced CFEB (CFEB-E); and MX Series routers and EX Series switches.</p> <p>For IP traffic on M320, MX Series, T Series routers and EX Series switches with Enhanced II Flexible PIC Concentrators (FPCs), you must include the tri-color statement at the [edit class-of-service] hierarchy level to commit a PLP configuration with any of the four levels specified. If the tri-color statement is not enabled, you can only configure the high and low levels. This applies to all protocol families.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p> <p>For information about the tri-color statement and for information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>Junos OS Class of Service Configuration Guide</i>.</p>
loss-priority-except level	<p>Do not match the PLP level. For details, see the loss-priority match condition.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p>

Table 342: Standard Firewall Filter Match Conditions for IPv6 Traffic (*continued*)

Match Condition	Description
next-header <i>header-type</i>	<p>Match the 8-bit Next Header field that identifies the type of header between the IPv6 header and payload.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): ah (51), dstopts (60), egp (8), esp (50), fragment (44), gre (47), hop-by-hop (0), icmp (1), icmp6 (58), icmpv6 (58), igmp (2), ipip (4), ipv6 (41), no-next-header (59), ospf (89), pim (103), routing (43), rsvp (46), sctp (132), tcp (6), udp (17), or vrrp (112).</p> <p>NOTE: next-header icmp6 and next-header icmpv6 match conditions perform the same function. next-header icmp6 is the preferred option. next-header icmpv6 is hidden in the Junos OS CLI.</p>
next-header-except <i>header-type</i>	Do not match the 8-bit Next Header field that identifies the type of header between the IPv6 header and payload. For details, see the next-header match type.
packet-length <i>bytes</i>	Match the length of the received packet, in bytes. The length refers only to the IP packet, including the packet header, and does not include any Layer 2 encapsulation overhead.
packet-length-except <i>bytes</i>	Do not match the length of the received packet, in bytes. For details, see the packet-length match type.
payload-protocol <i>protocol-type</i>	<p>Match the payload protocol type.</p> <p>In place of the protocol-type numeric value, you can specify one of the following text synonyms (the field values are also listed): specify one or a set of of the following: ah (51), dstopts (60), egp (8), esp (50), fragment (44), gre (47), hop-by-hop (0), icmp (1), icmp6 (58), igmp (2), ipip (4), ipv6 (41), no-next-header, ospf (89), pim (103), routing, rsvp (46), sctp (132), tcp (6), udp (17), or vrrp (112).</p>
payload-protocol-except <i>protocol-type</i>	Do not match the payload protocol type. For details, see the payload-protocol match type.
port <i>number</i>	<p>Match the UDP or TCP source or destination port field.</p> <p>If you configure this match condition, you cannot configure the destination-port match condition or the source-port match condition in the same term.</p> <p>If you configure this match condition, we recommend that you also configure the next-header udp or next-header tcp match condition in the same term to specify which protocol is being used on the port.</p> <p>In place of the numeric value, you can specify one of the text synonyms listed under destination-port.</p>
port-except <i>number</i>	Do not match the UDP or TCP source or destination port field. For details, see the port match condition.
prefix-list <i>prefix-list-name</i> [except]	<p>Match the prefixes of the source or destination address fields to the prefixes in the specified list unless the except option is included. If the option is included, do not match the prefixes of the source or destination address fields to the prefixes in the specified list.</p> <p>The prefix list is defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.</p>

Table 342: Standard Firewall Filter Match Conditions for IPv6 Traffic (*continued*)

Match Condition	Description
service-filter-hit	<p>Match a packet received from a filter where a service-filter-hit action was applied.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p>
source-address <i>address</i> [except]	<p>Match the IPv6 address of the source node sending the packet unless the except option is included. If the option is included, do not match the IPv6 address of the source node sending the packet.</p> <p>You cannot specify both the address and source-address match conditions in the same term.</p>
source-class <i>class-names</i>	<p>Match one or more specified source class names (sets of source prefixes grouped together and given a class name).</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p> <p>For more information, see “Firewall Filter Match Conditions Based on Address Classes” on page 4502.</p>
source-class-except <i>class-names</i>	<p>Do not match one or more specified source class names. For details, see the source-class match condition.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p>
source-port <i>number</i>	<p>Match the UDP or TCP source port field.</p> <p>You cannot specify the port and source-port match conditions in the same term.</p> <p>If you configure this match condition, we recommend that you also configure the next-header udp or next-header tcp match condition in the same term to specify which protocol is being used on the port.</p> <p>In place of the numeric value, you can specify one of the text synonyms listed with the destination-port number match condition.</p>
source-port-except <i>number</i>	<p>Do not match the UDP or TCP source port field. For details, see the source-port match condition.</p>
source-prefix-list <i>name</i> [except]	<p>Match the IPv6 address prefix of the packet source field unless the except option is included. If the option is included, do not match the IPv6 address prefix of the packet source field.</p> <p>Specify a prefix list name defined at the [edit policy-options prefix-list prefix-list-name] hierarchy level.</p>
tcp-established	<p>Match TCP packets other than the first packet of a connection. This is a text synonym for tcp-flags "(ack rst)" (0x14).</p> <p>NOTE: This condition does not implicitly check that the protocol is TCP. To check this, specify the protocol tcp match condition.</p> <p>If you configure this match condition, we recommend that you also configure the next-header tcp match condition in the same term.</p>

Table 342: Standard Firewall Filter Match Conditions for IPv6 Traffic (*continued*)

Match Condition	Description
tcp-flags flags	<p>Match one or more of the low-order 6 bits in the 8-bit TCP flags field in the TCP header.</p> <p>To specify individual bit fields, you can specify the following text synonyms or hexadecimal values:</p> <ul style="list-style-type: none"> • fin (0x01) • syn (0x02) • rst (0x04) • push (0x08) • ack (0x10) • urgent (0x20) <p>In a TCP session, the SYN flag is set only in the initial packet sent, while the ACK flag is set in all packets sent after the initial packet.</p> <p>You can string together multiple flags using the bit-field logical operators.</p> <p>For combined bit-field match conditions, see the tcp-established and tcp-initial match conditions.</p> <p>If you configure this match condition, we recommend that you also configure the next-header tcp match condition in the same term to specify that the TCP protocol is being used on the port.</p>
tcp-initial	<p>Match the initial packet of a TCP connection. This is a text synonym for tcp-flags "(!ack & syn)".</p> <p>This condition does not implicitly check that the protocol is TCP. If you configure this match condition, we recommend that you also configure the next-header tcp match condition in the same term.</p>
traffic-class number	<p>Match the 8-bit field that specifies the class-of-service (CoS) priority of the packet.</p> <p>This field was previously used as the type-of-service (ToS) field in IPv4.</p> <p>You can specify a numeric value from 0 through 63. To specify the value in hexadecimal form, include 0x as a prefix. To specify the value in binary form, include b as a prefix.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> • RFC 3246, <i>An Expedited Forwarding PHB (Per-Hop Behavior)</i>, defines one code point: ef (46). • RFC 2597, <i>Assured Forwarding PHB Group</i>, defines 4 classes, with 3 drop precedences in each class, for a total of 12 code points: <ul style="list-style-type: none"> • af11 (10), af12 (12), af13 (14) • af21 (18), af22 (20), af23 (22) • af31 (26), af32 (28), af33 (30) • af41 (34), af42 (36), af43 (38)
traffic-class-except number	<p>Do not match the 8-bit field that specifies the CoS priority of the packet. For details, see the traffic-class match description.</p>



NOTE: If you specify an IPv6 address in a match condition (the *address*, *destination-address*, or *source-address* match conditions), use the syntax for text representations described in RFC 2373, *IP Version 6 Addressing Architecture*. For more information about IPv6 addresses, see “IPv6 Overview” and “IPv6 Standards” in the *Junos OS Routing Protocols Configuration Guide*.

Related Documentation

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Standard Firewall Filter Terminating Actions on page 4742](#)
- [Standard Firewall Filter Nonterminating Actions on page 4744](#)

Standard Firewall Filter Match Conditions for MPLS Traffic

You can configure a standard stateless firewall filter with match conditions for MPLS traffic (*family mpls*).



NOTE: The input-list *filter-names* and output-list *filter-names* statements for firewall filters for the *mpls* protocol family are supported on all interfaces with the exception of management interfaces and internal Ethernet interfaces (fxp or em0), loopback interfaces (lo0), and USB modem interfaces (umd).

Table 270 on page 3566 describes the *match-conditions* you can configure at the [edit firewall family mpls filter *filter-name* term *term-name* from] hierarchy level.

Table 343: Standard Firewall Filter Match Conditions for MPLS Traffic

Match Condition	Description
apply-groups	Specify which groups to inherit configuration data from. You can specify more than one group name. You must list them in order of inheritance priority. The configuration data in the first group takes priority over the data in subsequent groups.
apply-groups-except	Specify which groups not to inherit configuration data from. You can specify more than one group name.
exp <i>number</i>	Experimental (EXP) bit number or range of bit numbers in the MPLS header. For <i>number</i> , you can specify one or more values from 0 through 7 in decimal, binary, or hexadecimal format. NOTE: This match condition is not supported on PTX series packet transport switches.
exp-except <i>number</i>	Do not match on the EXP bit number or range of bit numbers in the MPLS header. For <i>number</i> , you can specify one or more values from 0 through 7. NOTE: This match condition is not supported on PTX series packet transport switches.
forwarding-class <i>class</i>	Forwarding class. Specify assured-forwarding , best-effort , expedited-forwarding , or network-control .
forwarding-class-except <i>class</i>	Do not match on the forwarding class. Specify assured-forwarding , best-effort , expedited-forwarding , or network-control .

Table 343: Standard Firewall Filter Match Conditions for MPLS Traffic (*continued*)

Match Condition	Description
interface <i>interface-name</i>	<p>Interface on which the packet was received. You can configure a match condition that matches packets based on the interface on which they were received.</p> <p>NOTE: If you configure this match condition with an interface that does not exist, the term does not match any packet.</p>
interface-set <i>interface-set-name</i>	<p>Match the interface on which the packet was received to the specified interface set.</p> <p>To define an interface set, include the interface-set statement at the [edit firewall] hierarchy level.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p> <p>For more information, see “Filtering Packets Received on an Interface Set Overview” on page 4515.</p>
ip-version <i>number</i>	<p>(Interfaces on Enhanced Scaling flexible PIC concentrators [FPCs] on supported T Series routers only) Inner IP version. To match MPLS-tagged IPv4 packets, match on the text synonym ipv4.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p>
loss-priority <i>level</i>	<p>Match the packet loss priority (PLP) level.</p> <p>Specify a single level or multiple levels: low, medium-low, medium-high, or high.</p> <p>Supported on M120 and M320 routers; M7i and M10i routers with the Enhanced CFEB (CFEB-E); and MX Series routers and EX Series switches.</p> <p>For IP traffic on M320, MX Series, T Series routers with Enhanced II Flexible PIC Concentrators (FPCs) and EX Series switches, you must include the tri-color statement at the [edit class-of-service] hierarchy level to commit a PLP configuration with any of the four levels specified. If the tri-color statement is not enabled, you can only configure the high and low levels. This applies to all protocol families.</p> <p>For information about the tri-color statement and for information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>Junos OS Class of Service Configuration Guide</i>.</p>
loss-priority-except <i>level</i>	<p>Do not match the PLP level. For details, see the loss-priority match condition.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p>

**Related
Documentation**

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Standard Firewall Filter Terminating Actions on page 4742](#)
- [Standard Firewall Filter Nonterminating Actions on page 4744](#)

Standard Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic

This topic covers the following information:

- [Matching on IPv4 or IPv6 Packet Header Address or Port Fields in MPLS Flows on page 4725](#)
- [IP Address Match Conditions for MPLS Traffic on page 4726](#)
- [IP Port Match Conditions for MPLS Traffic on page 4726](#)

Matching on IPv4 or IPv6 Packet Header Address or Port Fields in MPLS Flows

To support network-based service in a core network, you can configure a standard firewall filter that matches Internet Protocol version 4 (IPv4) or version 6 (IPv6) packet header fields in MPLS traffic (**family mpls**). The firewall filter can match IPv4 or IPv6 packets as an inner payload of an MPLS packet that has a single MPLS label or up to five MPLS labels stacked together. You can configure match conditions based on IPv4 addresses and IPv4 port numbers or IPv6 addresses and IPv6 port numbers in the header.

Standard firewall filters based on MPLS-tagged IPv4 headers are supported for interfaces on Enhanced Scaling flexible PIC concentrators (FPCs) on T320, T640, T1600, TX Matrix, and TX Matrix Plus routers only. However, the standard firewall filters based on MPLS-tagged IPv6 headers are supported for interfaces on the Type 5 FPC on T4000 Core Routers only. The feature is not supported for the router loopback interface (**lo0**), the router management interface (**fxp0** or **em0**), or USB modem interfaces (**umd**).

To configure a stateless firewall filter term that matches an address or port fields in the Layer 4 header of packets in an MPLS flow, you use the **ip-version ipv4** match condition to specify that the term is to match packets based on inner IP fields:

- To match an MPLS-tagged IPv4 packet on the source or destination address field in the IPv4 header, specify the match condition at the **[edit firewall family mpls filter filter-name term term-name from ip-version ipv4]** hierarchy level.
- To match an MPLS-tagged IPv4 packet on the source or destination port field in the Layer 4 header, specify the match condition at the **[edit firewall family mpls filter filter-name term term-name from ip-version ipv4 protocol (udp | tcp)]** hierarchy level.

To configure a stateless firewall filter term that matches an address or port fields in the IPv6 header of packets in an MPLS flow, you use the **ip-version ipv6** match condition to specify that the term is to match packets based on inner IP fields:

- To match an MPLS-tagged IPv6 packet on the source or destination address field in the IPv6 header, specify the match condition at the **[edit firewall family mpls filter filter-name term term-name from ip-version ipv6]** hierarchy level.
- To match an MPLS-tagged IPv6 packet on the source or destination port field in the Layer 4 header, specify the match condition at the **[edit firewall family mpls filter filter-name term term-name from ip-version ipv6 protocol (udp | tcp)]** hierarchy level.

IP Address Match Conditions for MPLS Traffic

Table 344 on page 4726 describes the IP address-specific match conditions you can configure at the `[edit firewall family mpls filter filter-name term term-name from ip-version ip-version]` hierarchy level.

Table 344: IP Address-Specific Firewall Filter Match Conditions for MPLS Traffic

Match Condition	Description
<code>destination-address <i>address</i></code>	Match the address of the destination node to receive the packet.
<code>destination-address <i>address</i> except</code>	Do not match the address of the destination node to receive the packet.
<code>protocol <i>number</i></code>	Match the IP protocol type field. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): <code>ah</code> (51), <code>dstopts</code> (60), <code>egp</code> (8), <code>esp</code> (50), <code>fragment</code> (44), <code>gre</code> (47), <code>hop-by-hop</code> (0), <code>icmp</code> (1), <code>icmp6</code> (58), <code>icmpv6</code> (58), <code>igmp</code> (2), <code>ipip</code> (4), <code>ipv6</code> (41), <code>ospf</code> (89), <code>pim</code> (103), <code>rsvp</code> (46), <code>sctp</code> (132), <code>tcp</code> (6), <code>udp</code> (17), or <code>vrrp</code> (112).
<code>source-address <i>address</i></code>	Match the address of the source node sending the packet.
<code>source-address <i>address</i> except</code>	Do not match the address of the source node sending the packet.

IP Port Match Conditions for MPLS Traffic

Table 345 on page 4727 describes the IP port-specific *match-conditions* you can configure at the `[edit firewall family mpls filter filter-name term term-name from ip-version ip-version protocol (udp | tcp)]` hierarchy level.

Table 345: IP Port-Specific Firewall Filter Match Conditions for MPLS Traffic

Match Condition	Description
destination-port <i>number</i>	<p>Match on the UDP or TCP destination port field.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed): afs (1483), bgp (179), biff (512), bootpc (68), bootps (67), cmd (514), cvspserver (2401), dhcp (67), domain (53), eklogin (2105), ekshell (2106), exec (512), finger (79), ftp (21), ftp-data (20), http (80), https (443), ident (113), imap (143), kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544), ldap (389), ldp (646), login (513), mobileip-agent (434), mobilip-mn (435), msdp (639), netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nnntp (119), ntalk (518), ntp (123), pop3 (110), pptp (1723), printer (515), radacct (1813), radius (1812), rip (520), rkinit (2108), smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514), tacacs (49), tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525), who (513), or xmcp (177).</p>
destination-port-except <i>number</i>	<p>Do not match on the UDP or TCP destination port field.</p> <p>In place of the numeric value, you can specify one of the text synonyms listed with the destination-port match condition.</p>
source-port <i>number</i>	<p>Match on the TCP or UDP source port field.</p> <p>In place of the numeric field, you can specify one of the text synonyms listed under destination-port.</p>
source-port-except <i>number</i>	Do not match on the TCP or UDP source port field.

Related Documentation

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Standard Firewall Filter Terminating Actions on page 4742](#)
- [Standard Firewall Filter Nonterminating Actions on page 4744](#)

Standard Firewall Filter Match Conditions for VPLS Traffic

In the **from** statement in the VPLS filter term, you specify conditions that the packet must match for the action in the **then** statement to be taken. All conditions in the **from** statement must match for the action to be taken. The order in which you specify match conditions is not important, because a packet must match all the conditions in a term for a match to occur.

If you specify no match conditions in a term, that term matches all packets.

An individual condition in a **from** statement can contain a list of values. For example, you can specify numeric ranges. You can also specify multiple source addresses or destination addresses. When a condition defines a list of values, a match occurs if one of the values in the list matches the packet.

Individual conditions in a **from** statement can be negated. When you negate a condition, you are defining an explicit mismatch. For example, the negated match condition for **forwarding-class** is **forwarding-class-except**. If a packet matches a negated condition, it is immediately considered not to match the **from** statement, and the next term in the filter is evaluated, if there is one. If there are no more terms, the packet is discarded.

You can configure a standard firewall filter with match conditions for Virtual Private LAN Service (VPLS) traffic (**family vpls**). [Table 346 on page 4728](#) describes the **match-conditions** you can configure at the **[edit firewall family vpls filter *filter-name* term *term-name* from]** hierarchy level.



NOTE: Not all match conditions for VPLS traffic are supported on all routing platforms or switching platforms. A number of match conditions for VPLS traffic are supported only on MX Series 3D Universal Edge Routers.

In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Table 346: Standard Firewall Filter Match Conditions for VPLS Traffic

Match Condition	Description
destination-mac-address address	Match the destination media access control (MAC) address of a VPLS packet.
destination-port number	<p>(MX Series routers and EX Series switches only) Match the UDP or TCP destination port field.</p> <p>You cannot specify both the port and destination-port match conditions in the same term.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed): afs (1483), bgp (179), biff (512), bootpc (68), bootps (67), cmd (514), cvspserver (2401), dhcp (67), domain (53), eklogin (2105), ekshell (2106), exec (512), finger (79), ftp (21), ftp-data (20), http (80), https (443), ident (113), imap (143), kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544), ldap (389), ldp (646), login (513), mobileip-agent (434), mobileip-mn (435), msdp (639), netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123), pop3 (110), pptp (1723), printer (515), radacct (1813), radius (1812), rip (520), rkit (2108), smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514), tacacs (49), tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525), who (513), or xdmcp (177).</p>
destination-port-except number	<p>(MX Series routers and EX Series switches only) Do not match on the TCP or UDP destination port field. You cannot specify both the port and destination-port match conditions in the same term.</p>
destination-prefix-list name	<p>(MX Series routers and EX Series switches only) Match destination prefixes in the specified list. Specify the name of a prefix list defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.</p> <p>NOTE: VPLS prefix lists support only IPv4 addresses. IPv6 addresses included in a VPLS prefix list will be discarded.</p>
destination-prefix-list name except	<p>(MX Series routers and EX Series switches only) Do not match destination prefixes in the specified list. For more information, see the destination-prefix-list match condition.</p>

Table 346: Standard Firewall Filter Match Conditions for VPLS Traffic (*continued*)

Match Condition	Description
dscp number	<p>(MX Series routers and EX Series switches only) Match the Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant 6 bits of this byte form the DSCP. For more information, see the <i>Junos OS Class of Service Configuration Guide</i>.</p> <p>You can specify a numeric value from 0 through 63. To specify the value in hexadecimal form, include 0x as a prefix. To specify the value in binary form, include b as a prefix.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> • RFC 3246, <i>An Expedited Forwarding PHB (Per-Hop Behavior)</i>, defines one code point: ef (46). • RFC 2597, <i>Assured Forwarding PHB Group</i>, defines 4 classes, with 3 drop precedences in each class, for a total of 12 code points: <p>af11 (10), af12 (12), af13 (14),</p> <p>af21 (18), af22 (20), af23 (22),</p> <p>af31 (26), af32 (28), af33 (30),</p> <p>af41 (34), af42 (36), af43 (38)</p>
dscp-except number	<p>(MX Series routers and EX Series switches only) Do not match on the DSCP. For details, see the dscp match condition.</p>
ether-type values	<p>Match the 2-octet IEEE 802.3 Length/EtherType field to the specified value or list of values.</p> <p>You can specify decimal or hexadecimal values from 0 through 65535 (0xFFFF). A value from 0 through 1500 (0x05DC) specifies the length of an Ethernet Version 1 frame. A value from 1536 (0x0600) through 65535 specifies the EtherType (nature of the MAC client protocol) of an Ethernet Version 2 frame.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the hexadecimal values are also listed): aarp (0x80F3), appletalk (0x809B), arp (0x0806), ipv4 (0x0800), ipv6 (0x86DD), mpls-multicast (0x8848), mpls-unicast (0x8847), oam (0x8902), ppp (0x880B), pppoe-discovery (0x8863), pppoe-session (0x8864), or sna (0x80D5).</p>
ether-type-except values	<p>Do not match the 2-octet Length/EtherType field to the specified value or list of values.</p> <p>For details about specifying the values, see the ether-type match condition.</p>
forwarding-class class	<p>Match the forwarding class. Specify assured-forwarding, best-effort, expedited-forwarding, or network-control.</p>
forwarding-class-except class	<p>Do not match the forwarding class. For details, see the forwarding-class match condition.</p>

Table 346: Standard Firewall Filter Match Conditions for VPLS Traffic (*continued*)

Match Condition	Description
icmp-code message-code	<p>Match the ICMP message code field.</p> <p>If you configure this match condition, we recommend that you also configure the next-header icmp or next-header icmp6 match condition in the same term.</p> <p>If you configure this match condition, you must also configure the icmp-type message-type match condition in the same term. An ICMP message code provides more specific information than an ICMP message type, but the meaning of an ICMP message code is dependent on the associated ICMP message type.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> parameter-problem: ip6-header-bad (0), unrecognized-next-header (1), unrecognized-option (2) time-exceeded: ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0) destination-unreachable: address-unreachable (3), administratively-prohibited (1), no-route-to-destination (0), port-unreachable (4)
icmp-code-except message-code	Do not match the ICMP message code field. For details, see the icmp-code match condition.
icmp-code number	<p>(MX Series routers and EX Series switches only) Match the ICMP message code field.</p> <p>If you configure this match condition, we recommend that you also configure the ip-protocol icmp or ip-protocol icmp6 match condition in the same term.</p> <p>If you configure this match condition, you must also configure the icmp-type message-type match condition in the same term. An ICMP message code provides more specific information than an ICMP message type, but the meaning of an ICMP message code is dependent on the associated ICMP message type.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> parameter-problem: ip6-header-bad (0), unrecognized-next-header (1), unrecognized-option (2) time-exceeded: ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0) destination-unreachable: address-unreachable (3), administratively-prohibited (1), no-route-to-destination (0), port-unreachable (4)
icmp-code-except number	(MX Series routers and EX Series switches only) Do not match on the ICMP code field. For details, see the icmp-code match condition.
icmp-type number	<p>(MX Series routers and EX Series switches only) Match the ICMP message type field.</p> <p>If you configure this match condition, we recommend that you also configure the ip-protocol icmp, ip-protocol icmp6, or ip-protocol icmpv6 match condition in the same term.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): destination-unreachable (1), echo-reply (129), echo-request (128), membership-query (130), membership-report (131), membership-termination (132), neighbor-advertisement (136), neighbor-solicit (135), node-information-reply (140), node-information-request (139), packet-too-big (2), parameter-problem (4), redirect (137), router-advertisement (134), router-renumbering (138), router-solicit (133), or time-exceeded (3).</p>

Table 346: Standard Firewall Filter Match Conditions for VPLS Traffic (*continued*)

Match Condition	Description
icmp-type-except <i>number</i>	(MX Series routers and EX Series switches only) Do not match the ICMP message type field. For details, see the icmp-type match condition.
interface <i>interface-name</i>	Interface on which the packet was received. You can configure a match condition that matches packets based on the interface on which they were received. NOTE: If you configure this match condition with an interface that does not exist, the term does not match any packet.
interface-group <i>group-number</i>	Match the logical interface on which the packet was received to the specified interface group or set of interface groups. For <i>group-number</i> , specify a single value or a range of values from 0 through 255. To assign a logical interface to an interface group <i>group-number</i> , specify the <i>group-number</i> at the [interfaces <i>interface-name</i> unit <i>number</i> family <i>family</i> filter group] hierarchy level. For more information, see “Filtering Packets Received on a Set of Interface Groups Overview” on page 4514. NOTE: This match condition is not supported on T4000 Type 5 FPCs.
interface-group-except <i>group-name</i>	Do not match the logical interface on which the packet was received to the specified interface group or set of interface groups. For details, see the interface-group match condition. NOTE: This match condition is not supported on T4000 Type 5 FPCs.
interface-set <i>interface-set-name</i>	Match the interface on which the packet was received to the specified interface set. To define an interface set, include the interface-set statement at the [edit firewall] hierarchy level. For more information, see “Filtering Packets Received on an Interface Set Overview” on page 4515.
ip-address <i>address</i>	(MX Series routers and EX Series switches only) 32-bit address that supports the standard syntax for IPv4 addresses.
ip-destination-address <i>address</i>	(MX Series routers and EX Series switches only) 32-bit address that is the final destination node address for the packet.
ip-precedence <i>ip-precedence-field</i>	(MX Series routers and EX Series switches only) IP precedence field. In place of the numeric field value, you can specify one of the following text synonyms (the field values are also listed): critical-ecp (0xa0), flash (0x60), flash-override (0x80), immediate (0x40), internet-control (0xc0), net-control (0xe0), priority (0x20), or routine (0x00).
ip-precedence-except <i>ip-precedence-field</i>	(MX Series routers and EX Series switches only) Do not match on the IP precedence field.
ip-protocol <i>number</i>	(MX Series routers and EX Series switches only) IP protocol field.
ip-protocol-except <i>number</i>	(MX Series routers and EX Series switches only) Do not match on the IP protocol field.
ip-source-address <i>address</i>	(MX Series routers and EX Series switches only) IP address of the source node sending the packet.

Table 346: Standard Firewall Filter Match Conditions for VPLS Traffic (*continued*)

Match Condition	Description
learn-vlan-1p-priority <i>number</i>	<p>(MX Series routers and EX Series switches only) Match on the IEEE 802.1p learned VLAN priority bits in the provider VLAN tag (the only tag in a single-tag frame with 802.1Q VLAN tags or the outer tag in a dual-tag frame with 802.1Q VLAN tags). Specify a single value or multiple values from 0 through 7.</p> <p>Compare with the user-vlan-1p-priority match condition.</p>
learn-vlan-1p-priority-except <i>number</i>	(MX Series routers and EX Series switches only) Do not match on the IEEE 802.1p learned VLAN priority bits. For details, see the learn-vlan-1p-priority match condition.
learn-vlan-dei	(MX Series routers and EX Series switches only) Match the user VLAN ID drop eligibility indicator (DEI) bit.
learn-vlan-dei-except	(MX Series routers and EX Series switches only) Do not match the user VLAN ID DEI bit.
learn-vlan-id <i>number</i>	(MX Series routers and EX Series switches only) VLAN identifier used for MAC learning.
learn-vlan-id-except <i>number</i>	(MX Series routers and EX Series switches only) Do not match on the VLAN identifier used for MAC learning.
loss-priority <i>level</i>	<p>Packet loss priority (PLP) level. Specify a single level or multiple levels: low, medium-low, medium-high, or high.</p> <p>Supported on M120 and M320 routers; M7i and M10i routers with the Enhanced CFEB (CFEB-E); and MX Series routers.</p> <p>For IP traffic on M320, MX Series, and T Series routers with Enhanced II Flexible PIC Concentrators (FPCs) and EX Series switches, you must include the tri-color statement at the [edit class-of-service] hierarchy level to commit a PLP configuration with any of the four levels specified. If the tri-color statement is not enabled, you can only configure the high and low levels. This applies to all protocol families.</p> <p>For information about the tri-color statement and about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>Junos OS Class of Service Configuration Guide</i>.</p>
loss-priority-except <i>level</i>	<p>Do not match on the packet loss priority level. Specify a single level or multiple levels: low, medium-low, medium-high, or high.</p> <p>For information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>Junos OS Class of Service Configuration Guide</i>.</p>
port <i>number</i>	(MX Series routers and EX Series switches only) TCP or UDP source or destination port. You cannot specify both the port match condition and either the destination-port or source-port match condition in the same term.
port-except <i>number</i>	(MX Series routers and EX Series switches only) Do not match on the TCP or UDP source or destination port. You cannot specify both the port match condition and either the destination-port or source-port match condition in the same term.

Table 346: Standard Firewall Filter Match Conditions for VPLS Traffic (*continued*)

Match Condition	Description
prefix-list <i>name</i>	<p>(MX Series routers and EX Series switches only) Match the destination or source prefixes in the specified list. Specify the name of a prefix list defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.</p> <p>NOTE: VPLS prefix lists support only IPv4 addresses. IPv6 addresses included in a VPLS prefix list will be discarded.</p>
prefix-list <i>name</i> except	<p>(MX Series routers and EX Series switches only) Do not match the destination or source prefixes in the specified list. For more information, see the destination-prefix-list match condition.</p>
source-mac-address <i>address</i>	Source MAC address of a VPLS packet.
source-port <i>number</i>	<p>(MX Series routers and EX Series switches only) TCP or UDP source port field. You cannot specify the port and source-port match conditions in the same term.</p>
source-port-except <i>number</i>	<p>(MX Series routers and EX Series switches only) Do not match on the TCP or UDP source port field. You cannot specify the port and source-port match conditions in the same term.</p>
source-prefix-list <i>name</i>	<p>(MX Series routers and EX Series switches only) Match the source prefixes in the specified prefix list. Specify a prefix list name defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.</p> <p>NOTE: VPLS prefix lists support only IPv4 addresses. IPv6 addresses included in a VPLS prefix list will be discarded.</p>
source-prefix-list <i>name</i> except	<p>(MX Series routers and EX Series switches only) Do not match the source prefixes in the specified prefix list. For more information, see the source-prefix-list match condition.</p>
tcp-flags <i>flags</i>	<p>Match one or more of the low-order 6 bits in the 8-bit TCP flags field in the TCP header.</p> <p>To specify individual bit fields, you can specify the following text synonyms or hexadecimal values:</p> <ul style="list-style-type: none"> • fin (0x01) • syn (0x02) • rst (0x04) • push (0x08) • ack (0x10) • urgent (0x20) <p>In a TCP session, the SYN flag is set only in the initial packet sent, while the ACK flag is set in all packets sent after the initial packet.</p> <p>You can string together multiple flags using the bit-field logical operators.</p> <p>If you configure this match condition for IPv6 traffic, we recommend that you also configure the next-header tcp match condition in the same term to specify that the TCP protocol is being used on the port.</p>
traffic-type <i>type-name</i>	<p>(MX Series routers and EX Series switches only) Traffic type. Specify broadcast, multicast, unknown-unicast, or known-unicast.</p>

Table 346: Standard Firewall Filter Match Conditions for VPLS Traffic (*continued*)

Match Condition	Description
traffic-type-except <i>type-name</i>	(MX Series routers and EX Series switches only) Do not match on the traffic type. Specify broadcast , multicast , unknown-unicast , or known-unicast .
user-vlan-1p-priority <i>number</i>	(MX Series routers and EX Series switches only) Match on the IEEE 802.1p user priority bits in the customer VLAN tag (the inner tag in a dual-tag frame with 802.1Q VLAN tags). Specify a single value or multiple values from 0 through 7. Compare with the learn-vlan-1p-priority match condition.
user-vlan-1p-priority-except <i>number</i>	(MX Series routers and EX Series switches only) Do not match on the IEEE 802.1p user priority bits. For details, see the user-vlan-1p-priority match condition.
user-vlan-id <i>number</i>	(MX Series routers and EX Series switches only) Match the first VLAN identifier that is part of the payload.
user-vlan-id-except <i>number</i>	(MX Series routers and EX Series switches only) Do not match on the first VLAN identifier that is part of the payload.
vlan-ether-type <i>value</i>	VLAN Ethernet type field of a VPLS packet.
vlan-ether-type-except <i>value</i>	Do not match on the VLAN Ethernet type field of a VPLS packet.

**Related
Documentation**

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Standard Firewall Filter Terminating Actions on page 4742](#)
- [Standard Firewall Filter Nonterminating Actions on page 4744](#)

Standard Firewall Filter Match Conditions for Layer 2 CCC Traffic

You can configure a standard stateless firewall filter with match conditions for Layer 2 circuit cross-connect (CCC) traffic (**family ccc**).

The following restrictions apply to firewall filters for Layer 2 CCC traffic:

- The **input-list** *filter-names* and **output-list** *filter-names* statements for firewall filters for the **ccc** protocol family are supported on all interfaces with the exception of management interfaces and internal Ethernet interfaces (**fxp** or **em0**), loopback interfaces (**lo0**), and USB modem interfaces (**umd**).
- Only on MX Series routers and EX Series switches, you cannot apply a Layer 2 CCC stateless firewall filter (a firewall filter configured at the **[edit firewall filter family ccc]** hierarchy level) as an output filter. On MX Series routers and EX Series switches, firewall filters configured for the **family ccc** statement can be applied only as input filters.

[Table 347 on page 4735](#) describes the *match-conditions* you can configure at the **[edit firewall family ccc filter filter-name term term-name from]** hierarchy level.

Table 347: Standard Firewall Filter Match Conditions for Layer 2 CCC Traffic

Match Condition	Description
apply-groups	Specify which groups to inherit configuration data from. You can specify more than one group name. You must list them in order of inheritance priority. The configuration data in the first group takes priority over the data in subsequent groups.
apply-groups-except	Specify which groups not to inherit configuration data from. You can specify more than one group name.
destination-mac-address address	<p>(MX Series routers and EX Series switches only) Match the destination media access control (MAC) address of a virtual private LAN service (VPLS) packet.</p> <p>To have packets correctly evaluated by this match condition when applied to egress traffic flowing over a CCC circuit from a logical interface on an I-chip DPC in a Layer 2 virtual private network (VPN) routing instance, you must make a configuration change to the Layer 2 VPN routing instance. You must explicitly disable the use of a control word for traffic flowing out over a Layer 2 circuit. The use of a control word is enabled by default for Layer 2 VPN routing instances to support the emulated virtual circuit (VC) encapsulation for Layer 2 circuits.</p> <p>To explicitly disable the use of a control word for Layer 2 VPNs, include the no-control-word statement at either of the following hierarchy levels:</p> <ul style="list-style-type: none"> • [edit routing-instances <i>routing-instance-name</i> protocols l2vpn] • [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn] <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p> <p>For more information, see “Disabling the Control Word for Layer 2 VPNs” in the <i>Junos OS VPNs Configuration Guide</i>.</p>
forwarding-class class	Forwarding class. Specify assured-forwarding , best-effort , expedited-forwarding , or network-control .
forwarding-class-except class	Do not match on the forwarding class. Specify assured-forwarding , best-effort , expedited-forwarding , or network-control .
interface-group group-number	<p>Match the logical interface on which the packet was received to the specified interface group or set of interface groups. For <i>group-number</i>, specify a single value or a range of values from 0 through 255.</p> <p>To assign a logical interface to an interface group <i>group-number</i>, specify the <i>group-number</i> at the [interfaces <i>interface-name</i> unit <i>number</i> family <i>family</i> filter group] hierarchy level.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p> <p>For more information, see “Filtering Packets Received on a Set of Interface Groups Overview” on page 4514.</p>
interface-group-except number	<p>Do not match the logical interface on which the packet was received to the specified interface group or set of interface groups. For details, see the interface-group match condition.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p>

Table 347: Standard Firewall Filter Match Conditions for Layer 2 CCC Traffic (*continued*)

Match Condition	Description
learn-vlan-1p-priority <i>number</i>	<p>(MX Series routers and EX Series switches only) Match on the IEEE 802.1p learned VLAN priority bits in the provider VLAN tag (the only tag in a single-tag frame with 802.1Q VLAN tags or the outer tag in a dual-tag frame with 802.1Q VLAN tags). Specify a single value or multiple values from 0 through 7.</p> <p>Compare with the user-vlan-1p-priority match condition.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p>
learn-vlan-1p-priority-except <i>number</i>	<p>(MX Series routers and EX Series switches only) Do not match on the IEEE 802.1p learned VLAN priority bits. For details, see the learn-vlan-1p-priority match condition.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p>
loss-priority <i>level</i>	<p>Packet loss priority (PLP) level. Specify a single level or multiple levels: low, medium-low, medium-high, or high.</p> <p>Supported on M120 and M320 routers; M7i and M10i routers with the Enhanced CFEB (CFEB-E); and MX Series routers and EX Series switches.</p> <p>For IP traffic on M320, MX Series, T Series routers with Enhanced II Flexible PIC Concentrators (FPCs) and EX Series switches, you must include the tri-color statement at the [edit class-of-service] hierarchy level to commit a PLP configuration with any of the four levels specified. If the tri-color statement is not enabled, you can only configure the high and low levels. This applies to all protocol families.</p> <p>For information about the tri-color statement and for information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>Junos OS Class of Service Configuration Guide</i>.</p>
loss-priority-except <i>level</i>	<p>Do not match on the packet loss priority level. Specify a single level or multiple levels: low, medium-low, medium-high, or high.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p> <p>For information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>Junos OS Class of Service Configuration Guide</i>.</p>
user-vlan-1p-priority <i>number</i>	<p>(MX Series routers and EX Series switches only) Match on the IEEE 802.1p user priority bits in the customer VLAN tag (the inner tag in a dual-tag frame with 802.1Q VLAN tags). Specify a single value or multiple values from 0 through 7.</p> <p>Compare with the learn-vlan-1p-priority match condition.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p>
user-vlan-1p-priority-except <i>number</i>	<p>(MX Series routers and EX Series switches only) Do not match on the IEEE 802.1p user priority bits. For details, see the user-vlan-1p-priority match condition.</p> <p>NOTE: This match condition is not supported on PTX series packet transport switches.</p>

Related Documentation

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Standard Firewall Filter Terminating Actions on page 4742](#)

- [Standard Firewall Filter Nonterminating Actions on page 4744](#)

Standard Firewall Filter Match Conditions for Layer 2 Bridging Traffic

Only on MX Series routers and EX Series switches, you can configure a standard stateless firewall filter with match conditions for Layer 2 bridging traffic (**family bridge**).

[Table 348 on page 4737](#) describes the **match-conditions** you can configure at the **[edit firewall family bridge filter filter-name term term-name from]** hierarchy level.

Table 348: Standard Firewall Filter Match Conditions for Layer 2 Bridging (MX Series 3D Universal Edge Routers and EX Series switches Only)

Match Condition	Description
destination-mac-address address	Destination media access control (MAC) address of a Layer 2 packet in a bridging environment.
destination-port number	TCP or UDP destination port field. You cannot specify both the port and destination-port match conditions in the same term.
destination-port-except	Do not match the TCP/UDP destination port.
destination-prefix-list	Match the IP destination prefixes in a named list.
dscp number	<p>Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant 6 bits of this byte form the DSCP. For more information, see the <i>Junos OS Class of Service Configuration Guide</i>.</p> <p>You can specify a numeric value from 0 through 63. To specify the value in hexadecimal form, include 0x as a prefix. To specify the value in binary form, include b as a prefix.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> • RFC 3246, <i>An Expedited Forwarding PHB (Per-Hop Behavior)</i>, defines one code point: ef (46). • RFC 2597, <i>Assured Forwarding PHB Group</i>, defines 4 classes, with 3 drop precedences in each class, for a total of 12 code points: <p>af11 (10), af12 (12), af13 (14),</p> <p>af21 (18), af22 (20), af23 (22),</p> <p>af31 (26), af32 (28), af33 (30),</p> <p>af41 (34), af42 (36), af43 (38)</p>
dscp-except number	Do not match on the DSCP number. For more information, see the dscp-except match condition.

Table 348: Standard Firewall Filter Match Conditions for Layer 2 Bridging (MX Series 3D Universal Edge Routers and EX Series switches Only) (continued)

Match Condition	Description
ether-type value	<p>Match the 2-octet IEEE 802.3 Length/EtherType field to the specified value or list of values.</p> <p>You can specify decimal or hexadecimal values from 0 through 65535 (0xFFFF). A value from 0 through 1500 (0x05DC) specifies the length of an Ethernet Version 1 frame. A value from 1536 (0x0600) through 65535 specifies the EtherType (nature of the MAC client protocol) of an Ethernet Version 2 frame.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the hexadecimal values are also listed): aarp (0x80F3), appletalk (0x809B), arp (0x0806), ipv4 (0x0800), ipv6 (0x86DD), mpls-multicast (0x8848), mpls-unicast (0x8847), oam (0x8902), ppp (0x880B), pppoe-discovery (0x8863), pppoe-session (0x8864), sna (0x80D5).</p>
ether-type-except value	<p>Do not match the 2-octet IEEE 802.3 Length/EtherType field to the specified value or list of values.</p> <p>For details about specifying the values, see the ether-type match condition.</p>
forwarding class class	Forwarding class. Specify assured-forwarding , best-effort , expedited-forwarding , or network-control .
forwarding-class-except class	Ethernet type field of a Layer 2 packet environment. Specify assured-forwarding , best-effort , expedited-forwarding , or network-control .
icmp-code message-code	<p>Match the ICMP message code field.</p> <p>If you configure this match condition, we recommend that you also configure the ip-protocol icmp, ip-protocol icmp6, or ip-protocol icmpv6 match condition in the same term.</p> <p>If you configure this match condition, you must also configure the icmp-type message-type match condition in the same term. An ICMP message code provides more specific information than an ICMP message type, but the meaning of an ICMP message code is dependent on the associated ICMP message type.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> parameter-problem: ip6-header-bad (0), unrecognized-next-header (1), unrecognized-option (2) time-exceeded: ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0) destination-unreachable: address-unreachable (3), administratively-prohibited (1), no-route-to-destination (0), port-unreachable (4)
icmp-code-except message-code	Do not match the ICMP message code field. For details, see the icmp-code match condition.

Table 348: Standard Firewall Filter Match Conditions for Layer 2 Bridging (MX Series 3D Universal Edge Routers and EX Series switches Only) (*continued*)

Match Condition	Description
icmp-type <i>message-type</i>	<p>Match the ICMP message type field.</p> <p>If you configure this match condition, we recommend that you also configure the ip-protocol icmp, ip-protocol icmp6, or ip-protocol icmpv6 match condition in the same term.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): destination-unreachable (1), echo-reply (129), echo-request (128), membership-query (130), membership-report (131), membership-termination (132), neighbor-advertisement (136), neighbor-solicit (135), node-information-reply (140), node-information-request (139), packet-too-big (2), parameter-problem (4), redirect (137), router-advertisement (134), router-renumbering (138), router-solicit (133), or time-exceeded (3).</p>
icmp-type-except <i>message-type</i>	Do not match the ICMP message type field. For details, see the icmp-type match condition.
interface <i>interface-name</i>	<p>Interface on which the packet was received. You can configure a match condition that matches packets based on the interface on which they were received.</p> <p>NOTE: If you configure this match condition with an interface that does not exist, the term does not match any packet.</p>
interface-group <i>group-number</i>	<p>Match the logical interface on which the packet was received to the specified interface group or set of interface groups. For <i>group-number</i>, specify a single value or a range of values from 0 through 255.</p> <p>To assign a logical interface to an interface group <i>group-number</i>, specify the <i>group-number</i> at the [interfaces <i>interface-name</i> unit <i>number</i> family <i>family</i> filter <i>group</i>] hierarchy level.</p> <p>For more information, see “Filtering Packets Received on a Set of Interface Groups Overview” on page 4514.</p>
interface-group-except <i>number</i>	Do not match the logical interface on which the packet was received to the specified interface group or set of interface groups. For details, see the interface-group match condition.
interface-set <i>interface-set-name</i>	<p>Match the interface on which the packet was received to the specified interface set.</p> <p>To define an interface set, include the interface-set statement at the [edit firewall] hierarchy level. For more information, see “Filtering Packets Received on an Interface Set Overview” on page 4515.</p>
ip-address <i>address</i>	32-bit address that supports the standard syntax for IPv4 addresses.
ip-destination-address <i>address</i>	32-bit address that is the final destination node address for the packet.
ip-precedence <i>ip-precedence-field</i>	<p>IP precedence field. In place of the numeric field value, you can specify one of the following text synonyms (the field values are also listed): critical-ecp (0xa0), flash (0x60), flash-override (0x80), immediate (0x40), internet-control (0xc0), net-control (0xe0), priority (0x20), or routine (0x00).</p>
ip-precedence-except <i>ip-precedence-field</i>	Do not match on the IP precedence field.

Table 348: Standard Firewall Filter Match Conditions for Layer 2 Bridging (MX Series 3D Universal Edge Routers and EX Series switches Only) (*continued*)

Match Condition	Description
ip-protocol <i>number</i>	IP protocol field.
ip-protocol-except	Do not match the IP protocol type.
ip-source-address <i>address</i>	IP address of the source node sending the packet.
isid <i>number</i>	(Supported with Provider Backbone Bridging [PBB]) Match internet service identifier.
isid-dei <i>number</i>	(Supported with PBB) Match the Internet service identifier drop eligibility indicator (DEI) bit.
isid-dei-except <i>number</i>	(Supported with PBB) Do not match the Internet service identifier DEI bit.
isid-priority-code-point <i>number</i>	(Supported with PBB) Match the Internet service identifier priority code point.
isid-priority-code-point-except <i>number</i>	(Supported with PBB) Do not match the Internet service identifier priority code point.
learn-vlan-1p-priority <i>value</i>	(MX Series routers and EX Series switches only) Match on the IEEE 802.1p learned VLAN priority bits in the provider VLAN tag (the only tag in a single-tag frame with 802.1Q VLAN tags or the outer tag in a dual-tag frame with 802.1Q VLAN tags). Specify a single value or multiple values from 0 through 7. Compare with the user-vlan-1p-priority match condition.
learn-vlan-1p-priority-except <i>value</i>	(MX Series routers and EX Series switches only) Do not match on the IEEE 802.1p learned VLAN priority bits. For details, see the learn-vlan-1p-priority match condition.
learn-vlan-dei <i>number</i>	(Supported with bridging) Match user virtual LAN (VLAN) identifier DEI bit.
learn-vlan-dei-except <i>number</i>	(Supported with bridging) Do not match user VLAN identifier DEI bit.
learn-vlan-id <i>number</i>	VLAN identifier used for MAC learning.
learn-vlan-id-except <i>number</i>	Do not match on the VLAN identifier used for MAC learning.

Table 348: Standard Firewall Filter Match Conditions for Layer 2 Bridging (MX Series 3D Universal Edge Routers and EX Series switches Only) (*continued*)

Match Condition	Description
loss-priority <i>level</i>	<p>Packet loss priority (PLP) level. Specify a single level or multiple levels: low, medium-low, medium-high, or high.</p> <p>Supported on M120 and M320 routers; M7i and M10i routers with the Enhanced CFEB (CFEB-E); and MX Series routers and EX Series switches.</p> <p>For IP traffic on M320, MX Series, T Series routers with Enhanced II Flexible PIC Concentrators (FPCs), and EX Series switches, you must include the tri-color statement at the [edit class-of-service] hierarchy level to commit a PLP configuration with any of the four levels specified. If the tri-color statement is not enabled, you can only configure the high and low levels. This applies to all protocol families.</p> <p>For information about the tri-color statement and for information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>Junos OS Class of Service Configuration Guide</i>.</p>
loss-priority-except <i>level</i>	<p>Do not match on the packet loss priority level. Specify a single level or multiple levels: low, medium-low, medium-high, or high.</p> <p>For information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>Junos OS Class of Service Configuration Guide</i>.</p>
port <i>number</i>	TCP or UDP source or destination port. You cannot specify both the port match condition and either the destination-port or source-port match conditions in the same term.
source-mac-address <i>address</i>	Source MAC address of a Layer 2 packet.
source-port <i>number</i>	TCP or UDP source port field. You cannot specify the port and source-port match conditions in the same term.
source-port-except	Do not match the TCP/UDP source port.
tcp-flags <i>flags</i>	<p>Match one or more of the low-order 6 bits in the 8-bit TCP flags field in the TCP header.</p> <p>To specify individual bit fields, you can specify the following text synonyms or hexadecimal values:</p> <ul style="list-style-type: none"> • fin (0x01) • syn (0x02) • rst (0x04) • push (0x08) • ack (0x10) • urgent (0x20) <p>In a TCP session, the SYN flag is set only in the initial packet sent, while the ACK flag is set in all packets sent after the initial packet.</p> <p>You can string together multiple flags using the bit-field logical operators.</p> <p>Configuring the tcp-flags match condition requires that you configure the next-header-tcp match condition.</p>
traffic-type <i>type</i>	Traffic type. Specify broadcast , multicast , unknown-unicast , or known-unicast .

Table 348: Standard Firewall Filter Match Conditions for Layer 2 Bridging (MX Series 3D Universal Edge Routers and EX Series switches Only) (*continued*)

Match Condition	Description
traffic-type-except <i>type</i>	Do not match on the traffic type.
user-vlan-1p-priority <i>value</i>	(MX Series routers and EX Series switches only) Match on the IEEE 802.1p user priority bits in the customer VLAN tag (the inner tag in a dual-tag frame with 802.1Q VLAN tags). Specify a single value or multiple values from 0 through 7. Compare with the learn-vlan-1p-priority match condition.
user-vlan-1p-priority-except <i>value</i>	(MX Series routers and EX Series switches only) Do not match on the IEEE 802.1p user priority bits. For details, see the user-vlan-1p-priority match condition.
user-vlan-id <i>number</i>	(MX Series routers and EX Series switches only) Match the first VLAN identifier that is part of the payload.
user-vlan-id-except <i>number</i>	(MX Series routers and EX Series switches only) Do not match on the first VLAN identifier that is part of the payload.
vlan-ether-type <i>value</i>	VLAN Ethernet type field of a Layer 2 bridging packet.
vlan-ether-type-except <i>value</i>	Do not match on the VLAN Ethernet type field of a Layer 2 bridging packet.

Related Documentation

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Standard Firewall Filter Terminating Actions on page 4742](#)
- [Standard Firewall Filter Nonterminating Actions on page 4744](#)

Standard Firewall Filter Terminating Actions

Standard stateless firewall filters support different sets of terminating actions for each protocol family.



NOTE: You cannot configure the next term action with a *terminating* action in the same filter term. However, you can configure the next term action with another *nonterminating* action in the same filter term.

[Table 349 on page 4743](#) describes the terminating actions you can specify in a standard firewall filter term.

Table 349: Terminating Actions for Standard Firewall Filters

Terminating Action	Description	Protocols
accept	Accept the packet.	<ul style="list-style-type: none"> • family any • family inet • family inet6 • family mpls • family vpls • family ccc • family bridge <p>family ethernet-switching (for EX Series switches only)</p>
discard	Discard a packet silently, without sending an Internet Control Message Protocol (ICMP) message. Discarded packets are available for logging and sampling.	<ul style="list-style-type: none"> • family any • family inet • family inet6 • family mpls • family vpls • family ccc • family bridge <p>family ethernet-switching (for EX Series switches only)</p>
logical-system <i>logical-system-name</i>	Direct the packet to the specified logical system. NOTE: This action is not supported on PTX series packet transport switches.	<ul style="list-style-type: none"> • family inet • family inet6
reject <i>message-type</i>	Reject the packet and return an ICMPv4 or ICMPv6 message: <ul style="list-style-type: none"> • If no <i>message-type</i> is specified, a destination unreachable message is returned by default. • If tcp-reset is specified as the <i>message-type</i>, tcp-reset is returned only if the packet is a TCP packet. Otherwise, the administratively-prohibited message, which has a value of 13, is returned. • If any other <i>message-type</i> is specified, that message is returned. <p>NOTE: Rejected packets can be sampled or logged if you configure the sample or syslog action.</p> <p>The <i>message-type</i> can be one of the following values: address-unreachable, administratively-prohibited, bad-host-tos, bad-network-tos, beyond-scope, fragmentation-needed, host-prohibited, host-unknown, host-unreachable, network-prohibited, network-unknown, network-unreachable, no-route, port-unreachable, precedence-cutoff, precedence-violation, protocol-unreachable, source-host-isolated, source-route-failed, or tcp-reset.</p>	<ul style="list-style-type: none"> • family inet • family inet6

Table 349: Terminating Actions for Standard Firewall Filters (*continued*)

Terminating Action	Description	Protocols
routing-instance <i>routing-instance-name</i>	Direct the packet to the specified routing instance. NOTE: This action is not supported on PTX series packet transport switches.	<ul style="list-style-type: none"> family inet family inet6
topology <i>topology-name</i>	Direct the packet to the specified topology. NOTE: This action is not supported on PTX series packet transport switches. Each routing instance (master or virtual-router) supports one default topology to which all forwarding classes are forwarded. For Multitopology Routing, you can configure a firewall filter on the ingress interface to match a specific forwarding class, such as expedited forwarding, with a specific topology. The traffic that matches the specified forwarding class is then added to the routing table for that topology.	<ul style="list-style-type: none"> family inet family inet6

Related Documentation

- Guidelines for Configuring Standard Firewall Filters on page 4478
- Standard Firewall Filter Nonterminating Actions on page 4744

Standard Firewall Filter Nonterminating Actions

Standard stateless firewall filters support different sets of nonterminating actions for each protocol family.



NOTE: You cannot configure the next term action with a *terminating* action in the same filter term. However, you can configure the next term action with another *nonterminating* action in the same filter term.

Table 350 on page 4745 describes the nonterminating actions you can configure for a standard firewall filter term.

Table 350: Nonterminating Actions for Standard Firewall Filters

Nonterminating Action	Description	Protocol Families
<code>count</code> <i>counter-name</i>	Count the packet in the named counter.	<ul style="list-style-type: none">• <code>family any</code>• <code>family inet</code>• <code>family inet6</code>• <code>family mpls</code>• <code>family vpls</code>• <code>family ccc</code>• <code>family bridge</code>• <code>family ethernet-switching</code> (for EX Series switches only)

Table 350: Nonterminating Actions for Standard Firewall Filters (*continued*)

Nonterminating Action	Description	Protocol Families
dscp value	<p>Set the IPv4 Differentiated Services code point (DSCP) bit. You can specify a numerical value from 0 through 63. To specify the value in hexadecimal form, include 0x as a prefix. To specify the value in binary form, include b as a prefix.</p> <p>The default DSCP value is best effort, that is, be or 0.</p> <p>You can also specify on the following text synonyms:</p> <ul style="list-style-type: none"> • af11—Assured forwarding class 1, low drop precedence • af12—Assured forwarding class 1, medium drop precedence • af13—Assured forwarding class 1, high drop precedence • af21—Assured forwarding class 2, low drop precedence • af22—Assured forwarding class 2, medium drop precedence • af23—Assured forwarding class 2, high drop precedence • af31—Assured forwarding class 3, low drop precedence • af32—Assured forwarding class 3, medium drop precedence • af33—Assured forwarding class 3, high drop precedence • af41—Assured forwarding class 4, low drop precedence • af42—Assured forwarding class 4, medium drop precedence • af43—Assured forwarding class 4, high drop precedence • be—Best effort • cs0—Class selector 0 • cs1—Class selector 1 • cs2—Class selector 2 • cs3—Class selector 3 • cs4—Class selector 4 • cs5—Class selector 5 • cs6—Class selector 6 • cs7—Class selector 7 • ef—Expedited forwarding <p>NOTE: This action is not supported on PTX Series packet transport switches.</p> <p>NOTE: The actions dscp 0 or dscp be are supported only on T320, T640, T1600, TX Matrix, TX Matrix Plus, and M320 routers and on the 10-Gigabit Ethernet Modular Port Concentrators (MPC), 60-Gigabit Ethernet MPC, 60-Gigabit Ethernet Queuing MPC, and 60-Gigabit Ethernet Enhanced Queuing MPC on MX Series routers (and EX Series switches). However, these actions are not supported on Enhanced III Flexible PIC Concentrators (FPCs) on M320 routers.</p> <p>NOTE: On T4000 routers, the dscp 0 action is not supported during the interoperation between a T1600 Enhanced Scaling Type 4 FPC and a T4000 Type 5 FPC.</p>	family inet

Table 350: Nonterminating Actions for Standard Firewall Filters (*continued*)

Nonterminating Action	Description	Protocol Families
forwarding-class <i>class-name</i>	Classify the packet to the named forwarding class: <ul style="list-style-type: none"> <i>forwarding-class-name</i> assured-forwarding best-effort expedited-forwarding network-control 	<ul style="list-style-type: none"> family any family inet family inet6 family mpls family vpls family ccc family bridge family ethernet-switching (for EX Series switches only)
ipsec-sa <i>ipsec-sa</i>	Use the specified IPsec security association. NOTE: This action is not supported on MX Series routers and EX Series switches, Type 5 FPCs on T4000 routers, and PTX Series packet transport switches.	family inet
load-balance <i>group-name</i>	Use the specified load-balancing group. NOTE: This action is not supported on MX Series routers, EX Series switches, or PTX Series packet transport switches.	family inet
log	Log the packet header information in a buffer within the Packet Forwarding Engine. You can access this information by issuing the show firewall log command at the command-line interface (CLI).	<ul style="list-style-type: none"> family inet family inet6
logical-system <i>logical-system-name</i>	Direct packets to a specific logical system.	<ul style="list-style-type: none"> family inet family inet6
loss-priority (high medium-high medium-low low)	<p>Set the packet loss priority (PLP) level.</p> <p>You cannot also configure the three-color-policer nonterminating action for the same firewall filter term. These two nonterminating actions are mutually exclusive.</p> <p>Supported on M120 and M320 routers; M7i and M10i routers with the Enhanced CFEB (CFEB-E); and MX Series routers.</p> <p>For IP traffic on M320, MX Series, T Series routers with Enhanced II Flexible PIC Concentrators (FPCs), and EX Series switches, you must include the tri-color statement at the [edit class-of-service] hierarchy level to commit a PLP configuration with any of the four levels specified. If the tri-color statement is not enabled, you can only configure the high and low levels. This applies to all protocol families.</p> <p>For information about the tri-color statement and for information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>Junos OS Class of Service Configuration Guide</i>.</p>	<ul style="list-style-type: none"> family any family inet family inet6 family mpls family vpls family ccc family bridge family ethernet-switching (for EX Series switches only)
next-hop-group <i>group-name</i>	Use the specified next-hop group.	family inet

Table 350: Nonterminating Actions for Standard Firewall Filters (*continued*)

Nonterminating Action	Description	Protocol Families
next-interface <i>interface-name</i>	(MX Series routers and EX Series switches) Direct packets to the specified outgoing interface.	<ul style="list-style-type: none"> • family inet • family inet6
next-ip <i>ip-address</i>	(MX Series routers and EX Series switches) Direct packets to the specified destination IPv4 address.	family inet
next-ip6 <i>ipv6-address</i>	(MX Series routers and EX Series switches) Direct packets to the specified destination IPv6 address.	family inet6
packet-mode	Updates a bit field in the packet key buffer, which specifies traffic that will bypass flow-based forwarding. Packets with the packet-mode action modifier follow the packet-based forwarding path and bypass flow-based forwarding completely. For more information about selective stateless packet-based services, see the <i>Junos OS Security Configuration Guide</i> .	family any
policer <i>policer-name</i>	<p>Name of policer to use to rate-limit traffic.</p> <p>NOTE: For IPv6, applies to SRX100, SRX210, SRX220, SRX240, and SRX650 devices only.</p>	<ul style="list-style-type: none"> • family any • family inet • family inet6 • family mpls • family vpls • family ccc • family bridge • family ethernet-switching (for EX Series switches only)
port-mirror	Port-mirror the packet based on the specified family. Supported on M120 routers, M320 routers configured with Enhanced III FPCs, MX Series routers and EX Series switches only.	<ul style="list-style-type: none"> • family inet • family inet6 • family vpls • family ccc • family bridge • family ethernet-switching (for EX Series switches only)
prefix-action <i>action-name</i>	<p>Count or police packets based on the specified action name.</p> <p>NOTE: This action is not supported on PTX Series packet transport switches.</p>	family inet
routing-instance <i>routing-instance-name</i>	Direct packets to the specified routing instance.	<ul style="list-style-type: none"> • family inet • family inet6

Table 350: Nonterminating Actions for Standard Firewall Filters (*continued*)

Nonterminating Action	Description	Protocol Families
sample	<p>Sample the packet.</p> <p>NOTE: The Junos OS does not sample packets originating from the router or switch. If you configure a filter and apply it to the output side of an interface, then only the transit packets going through that interface are sampled. Packets that are sent from the Routing Engine to the Packet Forwarding Engine are not sampled.</p>	<ul style="list-style-type: none"> • family inet • family inet6 • family mpls
service-accounting	<p>Count the packet for service accounting. The count is applied to a specific named counter (<code>_junos-dyn-service-counter</code>) that RADIUS can obtain.</p> <p>NOTE: This action is not supported on T4000 Type 5 FPCs and PTX Series packet transport switches.</p>	<ul style="list-style-type: none"> • family inet • family inet6
service-filter-hit	<p>(Only if the service-filter-hit flag is marked by a previous filter in the current type of chained filters) Direct the packet to the next type of filters.</p> <p>Indicate to subsequent filters in the chain that the packet was already processed. This action, coupled with the service-filter-hit match condition in receiving filters, helps to streamline filter processing.</p> <p>NOTE: This action is not supported on T4000 Type 5 FPCs and PTX Series packet transport switches.</p>	<ul style="list-style-type: none"> • family inet • family inet6
syslog	Log the packet to the system log file.	<ul style="list-style-type: none"> • family inet • family inet6
three-color-policer (single-rate two-rate) policer-name	<p>Police the packet using the specified single-rate or two-rate three-color-policer.</p> <p>You cannot also configure the loss-priority action for the same firewall filter term. These two actions are mutually exclusive.</p>	<ul style="list-style-type: none"> • family inet • family inet6 • family mpls • family vpls • family ccc • family bridge • family ethernet-switching (for EX Series switches only)

Table 350: Nonterminating Actions for Standard Firewall Filters (*continued*)

Nonterminating Action	Description	Protocol Families
traffic-class value	<p>Specify the traffic-class code point. You can specify a numerical value from 0 through 63. To specify the value in hexadecimal form, include 0x as a prefix. To specify the value in binary form, include b as a prefix.</p> <p>The default traffic-class value is best effort, that is, be or 0.</p> <p>In place of the numeric value, you can specify one of the following text synonyms:</p> <ul style="list-style-type: none"> • af11—Assured forwarding class 1, low drop precedence • af12—Assured forwarding class 1, medium drop precedence • af13—Assured forwarding class 1, high drop precedence • af21—Assured forwarding class 2, low drop precedence • af22—Assured forwarding class 2, medium drop precedence • af23—Assured forwarding class 2, high drop precedence • af31—Assured forwarding class 3, low drop precedence • af32—Assured forwarding class 3, medium drop precedence • af33—Assured forwarding class 3, high drop precedence • af41—Assured forwarding class 4, low drop precedence • af42—Assured forwarding class 4, medium drop precedence • af43—Assured forwarding class 4, high drop precedence • be—Best effort • cs0—Class selector 0 • cs1—Class selector 1 • cs2—Class selector 2 • cs3—Class selector 3 • cs4—Class selector 4 • cs5—Class selector 5 • cs6—Class selector 6 • cs7—Class selector 7 • ef—Expedited forwarding <p>NOTE: The actions traffic-class 0 or traffic-class be are supported only on T Series and M320 routers and on the 10-Gigabit Ethernet Modular Port Concentrator (MPC), 60-Gigabit Ethernet MPC, 60-Gigabit Ethernet Queuing MPC, and 60-Gigabit Ethernet Enhanced Queuing MPC on MX Series routers (and EX Series switches). However, these actions are not supported on Enhanced III Flexible PIC Concentrators (FPCs) on M320 routers.</p>	family inet6

Related Documentation

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Standard Firewall Filter Terminating Actions on page 4742](#)

Standard Firewall Filter Match Conditions and Actions for ACX Series Routers

- [Standard Firewall Filter Match Conditions and Actions on ACX Series Routers Overview on page 4751](#)
- [Standard Firewall Filter Match Conditions for IPv4 Traffic on ACX Series Routers on page 4752](#)
- [Standard Firewall Filter Match Conditions for MPLS Traffic on ACX Series Routers on page 4756](#)
- [Standard Firewall Filter Terminating Actions on ACX Series Routers on page 4756](#)
- [Standard Firewall Filter Nonterminating Actions on ACX Series Routers on page 4757](#)

Standard Firewall Filter Match Conditions and Actions on ACX Series Routers Overview

On ACX Series Universal Access Routers, you can configure firewall filters to filter packets and to perform an action on packets that match the filter. The match conditions specified to filter the packets are specific to the type of traffic being filtered.



NOTE: On ACX Series routers, the filter for the exiting traffic (egress filter) can be applied only for interface-specific instances of the firewall filter.

[Table 351 on page 4751](#) describes the types of traffic for which you can configure standard stateless firewall filters.

Table 351: Standard Firewall Filter Match Conditions by Protocol Family for ACX Series Routers

Traffic Type	Hierarchy Level at Which Match Conditions Are Specified
Protocol-independent	<p>[edit firewall family any filter <i>filter-name</i> term <i>term-name</i>]</p> <p>No match conditions are supported for this traffic type on ACX Series routers.</p>
IPv4	<p>[edit firewall family inet filter <i>filter-name</i> term <i>term-name</i>]</p> <p>For the complete list of match conditions, see “Standard Firewall Filter Match Conditions for IPv4 Traffic on ACX Series Routers” on page 4752.</p>
MPLS	<p>[edit firewall family mpls filter <i>filter-name</i> term <i>term-name</i>]</p> <p>For the complete list of match conditions, see “Standard Firewall Filter Match Conditions for MPLS Traffic on ACX Series Routers” on page 4756.</p>
Layer 2 CCC	<p>[edit firewall family ccc filter <i>filter-name</i> term <i>term-name</i>]</p> <p>No match conditions are supported for this traffic type on ACX Series routers.</p>

Under the **then** statement for a standard stateless firewall filter term, you can specify the actions to be taken on a packet that matches the term.

[Table 352 on page 4752](#) summarizes the types of actions you can specify in a standard stateless firewall filter term.

Table 352: Standard Firewall Filter Action Categories for ACX Series Routers

Type of Action	Description	Comment
Terminating	<p>Halts all evaluation of a firewall filter for a specific packet. The router performs the specified action, and no additional terms are used to examine the packet.</p> <p>You can specify only one <i>terminating action</i> in a standard firewall filter. You can, however, specify one terminating action with one or more <i>nonterminating actions</i> in a single term. For example, within a term, you can specify accept with count and syslog.</p>	See “ Standard Firewall Filter Terminating Actions on ACX Series Routers ” on page 4756.
Nonterminating	Performs other functions on a packet (such as incrementing a counter, logging information about the packet header, sampling the packet data, or sending information to a remote host using the system log functionality), but any additional terms are used to examine the packet.	See “ Standard Firewall Filter Nonterminating Actions on ACX Series Routers ” on page 4757.

Related Documentation

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Interface-Specific Firewall Filter Instances Overview](#)

Standard Firewall Filter Match Conditions for IPv4 Traffic on ACX Series Routers

On ACX Series routers, you can configure a standard stateless firewall filter with match conditions for IP version 4 (IPv4) traffic (**family inet**). [Table 353 on page 4752](#) describes the match conditions you can configure at the **[edit firewall family inet filter *filter-name* term *term-name* from]** hierarchy level.

Table 353: Standard Firewall Filter Match Conditions for IPv4 Traffic on ACX Series Routers

Match Condition	Description
destination-address <i>address</i>	<p>Match the IPv4 destination address field.</p> <p>NOTE: On ACX Series routers, you can specify only one destination address. A list of IPv4 destination addresses is not supported.</p>

Table 353: Standard Firewall Filter Match Conditions for IPv4 Traffic on ACX Series Routers (*continued*)

Match Condition	Description
destination-port <i>number</i>	<p>Match the UDP or TCP destination port field.</p> <p>If you configure this match condition, we recommend that you also configure the protocol udp or protocol tcp match statement in the same term to specify which protocol is being used on the port.</p> <p>NOTE: On ACX Series routers, you can specify only one destination port number. A list of port numbers is not supported.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed): afs (1483), bgp (179), biff (512), bootpc (68), bootps (67), cmd (514), cvspserver (2401), dhcp (67), domain (53), eklogin (2105), ekshell (2106), exec (512), finger (79), ftp (21), ftp-data (20), http (80), https (443), ident (113), imap (143), kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544), ldap (389), ldp (646), login (513), mobileip-agent (434), mobileip-mn (435), msdp (639), netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123), pop3 (110), pptp (1723), printer (515), radacct (1813), radius (1812), rip (520), rkinit (2108), smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514), tacacs (49), tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525), who (513), or xmcp (177).</p>
dscp <i>number</i>	<p>Match the Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant 6 bits of this byte form the DSCP. For more information, see the <i>Junos OS Class of Service Configuration Guide</i>.</p> <p>You can specify a numeric value from 0 through 63. To specify the value in hexadecimal form, include 0x as a prefix. To specify the value in binary form, include b as a prefix.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> • RFC 3246, <i>An Expedited Forwarding PHB (Per-Hop Behavior)</i>, defines one code point: ef (46). • RFC 2597, <i>Assured Forwarding PHB Group</i>, defines 4 classes, with 3 drop precedences in each class, for a total of 12 code points: <ul style="list-style-type: none"> • af11 (10), af12 (12), af13 (14) • af21 (18), af22 (20), af23 (22) • af31 (26), af32 (28), af33 (30) • af41 (34), af42 (36), af43 (38)
fragment-flags <i>number</i>	<p>(Ingress only) Match the three-bit IP fragmentation flags field in the IP header.</p> <p>In place of the numeric field value, you can specify one of the following keywords (the field values are also listed): dont-fragment (0x4), more-fragments (0x2), or reserved (0x8).</p>

Table 353: Standard Firewall Filter Match Conditions for IPv4 Traffic on ACX Series Routers (*continued*)

Match Condition	Description
icmp-code <i>number</i>	<p>Match the ICMP message code field.</p> <p>If you configure this match condition, we recommend that you also configure the protocol icmp match condition in the same term.</p> <p>If you configure this match condition, you must also configure the icmp-type <i>message-type</i> match condition in the same term. An ICMP message code provides more specific information than an ICMP message type, but the meaning of an ICMP message code is dependent on the associated ICMP message type.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> parameter-problem: ip-header-bad (0), required-option-missing (1) redirect: redirect-for-host (1), redirect-for-network (0), redirect-for-tos-and-host (3), redirect-for-tos-and-net (2) time-exceeded: ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0) unreachable: communication-prohibited-by-filtering (13), destination-host-prohibited (10), destination-host-unknown (7), destination-network-prohibited (9), destination-network-unknown (6), fragmentation-needed (4), host-precedence-violation (14), host-unreachable (1), host-unreachable-for-TOS (12), network-unreachable (0), network-unreachable-for-TOS (11), port-unreachable (3), precedence-cutoff-in-effect (15), protocol-unreachable (2), source-host-isolated (8), source-route-failed (5)
icmp-type <i>number</i>	<p>Match the ICMP message type field.</p> <p>If you configure this match condition, we recommend that you also configure the protocol icmp match condition in the same term.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): echo-reply (0), echo-request (8), info-reply (16), info-request (15), mask-request (17), mask-reply (18), parameter-problem (12), redirect (5), router-advertisement (9), router-solicit (10), source-quench (4), time-exceeded (11), timestamp (13), timestamp-reply (14), or unreachable (3).</p>
ip-options <i>values</i>	<p>Match the 8-bit IP option field, if present, to the specified value.</p> <p>ACX Series routers support only the ip-options_any match condition, which ensures that the packets are sent to the Packet Forwarding Engine for processing.</p> <p>NOTE: On ACX Series routers, you can specify only one IP option value. Configuring multiple values is not supported.</p>
precedence ip-precedence-field	<p>Match the IP precedence field.</p> <p>In place of the numeric field value, you can specify one of the following text synonyms (the field values are also listed): critical-ecp (0xa0), flash (0x60), flash-override (0x80), immediate (0x40), internet-control (0xc0), net-control (0xe0), priority (0x20), or routine (0x00). You can specify precedence in hexadecimal, binary, or decimal form.</p>

Table 353: Standard Firewall Filter Match Conditions for IPv4 Traffic on ACX Series Routers (*continued*)

Match Condition	Description
protocol <i>number</i>	Match the IP protocol type field. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): ah (51), dstopts (60), egp (8), esp (50), fragment (44), gre (47), hop-by-hop (0), icmp (1), icmp6 (58), icmpv6 (58), igmp (2), ipip (4), ipv6 (41), ospf (89), pim (103), rsvp (46), sctp (132), tcp (6), udp (17), or vrrp (112).
source-address <i>address</i>	Match the IPv4 address of the source node sending the packet.
source-port <i>number</i>	Match the UDP or TCP source port field. If you configure this match condition for IPv4 traffic, we recommend that you also configure the protocol udp or protocol tcp match statement in the same term to specify which protocol is being used on the port. In place of the numeric value, you can specify one of the text synonyms listed with the destination-port <i>number</i> match condition.
tcp-flags <i>value</i>	Match one or more of the low-order 6 bits in the 8-bit TCP flags field in the TCP header. To specify individual bit fields, you can specify the following text synonyms or hexadecimal values: <ul style="list-style-type: none"> • fin (0x01) • syn (0x02) • rst (0x04) • push (0x08) • ack (0x10) • urgent (0x20) <p>In a TCP session, the SYN flag is set only in the initial packet sent, while the ACK flag is set in all packets sent after the initial packet.</p> <p>You can string together multiple flags using the bit-field logical operators.</p> <p>For combined bit-field match conditions, see the tcp-initial match conditions.</p> <p>If you configure this match condition, we recommend that you also configure the protocol tcp match statement in the same term to specify that the TCP protocol is being used on the port.</p>
tcp-initial	Match the initial packet of a TCP connection. This is an alias for tcp-flags "(ack & syn)" . This condition does not implicitly check that the protocol is TCP. If you configure this match condition, we recommend that you also configure the protocol tcp match condition in the same term.
ttl <i>number</i>	Match the IPv4 time-to-live number. Specify a TTL value or a range of TTL values. For <i>number</i> , you can specify one or more values from 2 through 255.

Related Documentation

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Standard Firewall Filter Match Conditions and Actions on ACX Series Routers Overview on page 4751](#)

- [Standard Firewall Filter Terminating Actions on ACX Series Routers on page 4756](#)
- [Standard Firewall Filter Nonterminating Actions on ACX Series Routers on page 4757](#)

Standard Firewall Filter Match Conditions for MPLS Traffic on ACX Series Routers

On ACX Series routers, you can configure a standard stateless firewall filter with match conditions for MPLS traffic (**family mpls**).



NOTE: The input-list *filter-names* and output-list *filter-names* statements for firewall filters for the mpls protocol family are supported on all interfaces with the exception of management interfaces and internal Ethernet interfaces (fxp or em0), loopback interfaces (lo0), and USB modem interfaces (umd).

[Table 354 on page 4756](#) describes the match conditions you can configure at the `[edit firewall family mpls filter filter-name term term-name from]` hierarchy level.

Table 354: Standard Firewall Filter Match Conditions for MPLS Traffic on ACX Series Routers

Match Condition	Description
<i>exp number</i>	Experimental (EXP) bit number or range of bit numbers in the MPLS header. For <i>number</i> , you can specify one or more values from 0 through 7 in decimal, binary, or hexadecimal format.

Related Documentation

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Standard Firewall Filter Match Conditions and Actions on ACX Series Routers Overview on page 4751](#)
- [Standard Firewall Filter Terminating Actions on ACX Series Routers on page 4756](#)
- [Standard Firewall Filter Nonterminating Actions on ACX Series Routers on page 4757](#)

Standard Firewall Filter Terminating Actions on ACX Series Routers

Standard stateless firewall filters support different sets of terminating actions for each protocol family.



NOTE: ACX Series routers do not support the `next term` action.

[Table 355 on page 4757](#) describes the terminating actions you can specify in a standard firewall filter term.

Table 355: Terminating Actions for Standard Firewall Filters on ACX Series Routers

Terminating Action	Description	Protocols
accept	Accept the packet.	<ul style="list-style-type: none"> family any family inet family mpls family ccc
discard	Discard a packet silently, without sending an Internet Control Message Protocol (ICMP) message. Discarded packets are available for logging and sampling.	<ul style="list-style-type: none"> family any family inet family mpls family ccc
reject <i>message-type</i>	<p>Reject the packet and return an ICMPv4 or ICMPv6 message:</p> <ul style="list-style-type: none"> If no message type is specified, a destination-unreachable message is returned by default. If tcp-reset is specified as the message type, tcp-reset is returned only if the packet is a TCP packet. Otherwise, the administratively-prohibited message, which has a value of 13, is returned. If any other message type is specified, that message is returned. <p>NOTE:</p> <ul style="list-style-type: none"> Rejected packets can be sampled or logged if you configure the sample or syslog action. This action is supported on ingress only. <p>The message-type option can have one of the following values: address-unreachable, administratively-prohibited, bad-host-tos, bad-network-tos, beyond-scope, fragmentation-needed, host-prohibited, host-unknown, host-unreachable, network-prohibited, network-unknown, network-unreachable, no-route, port-unreachable, precedence-cutoff, precedence-violation, protocol-unreachable, source-host-isolated, source-route-failed, or tcp-reset.</p>	family inet
routing-instance <i>routing-instance-name</i>	Direct the packet to the specified routing instance.	<ul style="list-style-type: none"> family inet

Related Documentation

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Standard Firewall Filter Match Conditions and Actions on ACX Series Routers Overview on page 4751](#)
- [Standard Firewall Filter Nonterminating Actions on ACX Series Routers on page 4757](#)

Standard Firewall Filter Nonterminating Actions on ACX Series Routers

Standard stateless firewall filters support different sets of nonterminating actions for each protocol family.



NOTE: ACX Series routers do not support the next term action.

Table 356 on page 4758 describes the nonterminating actions you can configure for a standard firewall filter term.

Table 356: Nonterminating Actions for Standard Firewall Filters on ACX Series Routers

Nonterminating Action	Description	Protocol Families
count <i>counter-name</i>	Count the packet in the named counter.	<ul style="list-style-type: none"> family any family inet family mpls family ccc
forwarding-class <i>class-name</i>	Classify the packet based on the specified forwarding class: <ul style="list-style-type: none"> assured-forwarding best-effort expedited-forwarding network-control <p>NOTE: This action is supported on ingress only.</p>	<ul style="list-style-type: none"> family inet family any family mpls family ccc
log	Log the packet header information in a buffer within the Packet Forwarding Engine. You can access this information by issuing the show firewall log command at the command-line interface (CLI). <p>NOTE: This action is supported on ingress only.</p>	family inet

Table 356: Nonterminating Actions for Standard Firewall Filters on ACX Series Routers (*continued*)

Nonterminating Action	Description	Protocol Families
loss-priority (high medium-high low)	<p>Set the packet loss priority (PLP) level.</p> <p>You cannot also configure the three-color-policer nonterminating action for the same firewall filter term. These two nonterminating actions are mutually exclusive.</p> <p>You must include the tri-color statement at the [edit class-of-service] hierarchy level to commit a PLP configuration with any of the four levels specified. If the tri-color statement is not enabled, you can configure only the high and low levels. This applies to all protocol families.</p> <p>For information about the tri-color statement and for information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>Junos OS Class of Service Configuration Guide</i>.</p> <p>NOTE: This action is supported on ingress only.</p>	<ul style="list-style-type: none"> • family any • family inet • family mpls • family ccc
policer <i>policer-name</i>	Name of policer to use to rate-limit traffic.	<ul style="list-style-type: none"> • family any • family inet • family mpls • family ccc
port-mirror	<p>Port-mirror the packet based on the specified family.</p> <p>NOTE: This action is supported on ingress only.</p>	family inet
syslog	<p>Log the packet to the system log file.</p> <p>NOTE: This action is supported on ingress only.</p>	family inet

Table 356: Nonterminating Actions for Standard Firewall Filters on ACX Series Routers (*continued*)

Nonterminating Action	Description	Protocol Families
three-color-policer (single-rate two-rate) <i>policer-name</i>	Police the packet using the specified single-rate or two-rate three-color policer. You cannot also configure the loss-priority action for the same firewall filter term. These two actions are mutually exclusive.	<ul style="list-style-type: none"> • family any • family inet • family mpls • family ccc

Related Documentation

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Standard Firewall Filter Match Conditions and Actions on ACX Series Routers Overview on page 4751](#)
- [Standard Firewall Filter Terminating Actions on ACX Series Routers on page 4756](#)

Service Filter Match Conditions and Actions

- [Service Filter Match Conditions for IPv4 or IPv6 Traffic on page 4760](#)
- [Service Filter Terminating Actions on page 4766](#)
- [Service Filter Nonterminating Actions on page 4766](#)

Service Filter Match Conditions for IPv4 or IPv6 Traffic

Service filters support only a subset of the stateless firewall filter match conditions for IPv4 and IPv6 traffic. [Table 357 on page 4760](#) describes the service filter match conditions.

Table 357: Service Filter Match Conditions for IPv4 or IPv6 Traffic

Match Condition	Description	Protocol Families	
address <i>address</i>	Match the IP source or destination address field.	family inet	family inet6
address <i>address</i> except	Do not match the IP source or destination address field.	family inet	family inet6
ah-spi <i>spi-value</i>	(M Series routers, except M120 and M320) Match on the IPsec authentication header (AH) security parameter index (SPI) value.	family inet	—
ah-spi-except <i>spi-value</i>	(M Series routers, except M120 and M320) Do not match on the IPsec AH SPI value.	family inet	—
destination-address <i>address</i>	Match the IP destination address field. You cannot specify both the address and destination-address match conditions in the same term.	family inet	family inet6

Table 357: Service Filter Match Conditions for IPv4 or IPv6 Traffic (continued)

Match Condition	Description	Protocol Families	
destination-address address address	Do not match the IP destination address field. You cannot specify both the address and destination-address match conditions in the same term.	family inet	family inet6
destination-port number	Match the UDP or TCP destination port field. You cannot specify both the port and destination-port match conditions in the same term. If you configure this match condition for IPv4 traffic, we recommend that you also configure the protocol udp or protocol tcp match statement in the same term to specify which protocol is being used on the port. If you configure this match condition for IPv6 traffic, we recommend that you also configure the next-header udp or next-header tcp match condition in the same term to specify which protocol is being used on the port. In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed): afs (1483), bgp (179), biff (512), bootpc (68), bootps (67), cmd (514), cvspserver (2401), dhcp (67), domain (53), eklogin (2105), ekshell (2106), exec (512), finger (79), ftp (21), ftp-data (20), http (80), https (443), ident (113), imap (143), kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544), ldap (389), ldp (646), login (513), mobileip-agent (434), mobileip-mn (435), msdp (639), netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123), pop3 (110), pptp (1723), printer (515), radacct (1813), radius (1812), rip (520), rkinet (2108), smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514), tacacs (49), tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525), who (513), or xmcp (177).	family inet	family inet6
destination-port-except number	Do not match the UDP or TCP destination port field. For details, see the destination-port match description.	family inet	family inet6
destination-prefix-list name	Match the list of destination prefixes. The prefix list is defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.	family inet	family inet6
esp-spi value	Match the IPsec encapsulating security payload (ESP) SPI value. Specify a single value or a range of values. You can specify a <i>value</i> in hexadecimal, binary, or decimal form. To specify the value in hexadecimal form, include 0x as a prefix. To specify the value in binary form, include b as a prefix.	family inet	family inet6
esp-spi-except value	Do not match the IPsec ESP SPI value or range of values. For details, see the esp-spi match condition.	family inet	family inet6

Table 357: Service Filter Match Conditions for IPv4 or IPv6 Traffic (*continued*)

Match Condition	Description	Protocol Families	
first-fragment	<p>Match if the packet is the first fragment of a fragmented packet. Do not match if the packet is a trailing fragment of a fragmented packet. The first fragment of a fragmented packet has a fragment offset value of 0.</p> <p>This match condition is an alias for the bit-field match condition fragment-offset 0 match condition.</p> <p>To match both first and trailing fragments, you can use two terms that specify different match conditions: first-fragment and is-fragment.</p>	family inet	—
fragment-flags <i>number</i>	<p>(Ingress only) Match the three-bit IP fragmentation flags field in the IP header.</p> <p>In place of the numeric field value, you can specify one of the following keywords (the field values are also listed): dont-fragment (0x4), more-fragments (0x2), or reserved (0x8).</p>	family inet	—
fragment-offset <i>number</i>	<p>Match the 13-bit fragment offset field in the IP header. The value is the offset, in 8-byte units, in the overall datagram message to the data fragment. Specify a numeric value, a range of values, or a set of values. An offset value of 0 indicates the first fragment of a fragmented packet.</p> <p>The first-fragment match condition is an alias for the fragment-offset 0 match condition.</p> <p>To match both first and trailing fragments, you can use two terms that specify different match conditions (first-fragment and is-fragment).</p>	family inet	—
fragment-offset-except <i>number</i>	Do not match the 13-bit fragment offset field.	family inet	—
interface-group <i>group-number</i>	<p>Match the interface group (set of one or more logical interfaces) on which the packet was received. For <i>group-number</i>, specify a value from 0 through 255.</p> <p>For information about configuring interface groups, see “Filtering Packets Received on a Set of Interface Groups Overview” on page 4514.</p>	family inet	family inet6
interface-group-except <i>group-number</i>	Do not match the interface group on which the packet was received. for details, see the interface-group match condition.	family inet	family inet6

Table 357: Service Filter Match Conditions for IPv4 or IPv6 Traffic (*continued*)

Match Condition	Description	Protocol Families
ip-options values	<p>Match the 8-bit IP option field, if present, to the specified value or list of values.</p> <p>In place of a numeric value, you can specify one of the following text synonyms (the option values are also listed): loose-source-route (131), record-route (7), router-alert (148), security (130), stream-id (136), strict-source-route (137), or timestamp (68).</p> <p>To match <i>any</i> value for the IP option, use the text synonym any. To match on <i>multiple</i> values, specify the list of values within square brackets ('[' and ']'). To match a <i>range</i> of values, use the value specification [<i>value1-value2</i>].</p> <p>For example, the match condition ip-options [0-147] matches on an IP options field that contains the loose-source-route, record-route, or security values, or any other value from 0 through 147. However, this match condition does not match on an IP options field that contains only the router-alert value (148).</p> <p>For most interfaces, a filter term that specifies an ip-option match on one or more <i>specific</i> IP option values (a value other than any) causes packets to be sent to the Routing Engine so that the kernel can parse the IP option field in the packet header.</p> <ul style="list-style-type: none"> For a firewall filter term that specifies an ip-option match on one or more specific IP option values, you cannot specify the count, log, or syslog nonterminating actions <i>unless</i> you also specify the discard terminating action in the same term. This behavior prevents double-counting of packets for a filter applied to a transit interface on the router (or switch). Packets processed on the kernel might be dropped in case of a system bottleneck. To ensure that matched packets are instead sent to the Packet Forwarding Engine (where packet processing is implemented in hardware), use the ip-options any match condition. <p>The 10-Gigabit Ethernet Modular Port Concentrator (MPC), 60-Gigabit Ethernet MPC, 60-Gigabit Queuing Ethernet MPC, 60-Gigabit Ethernet Enhanced Queuing MPC on MX Series routers and EX Series switches are capable of parsing the IP option field of the IPv4 packet header. This capability is supported on EX Series switches also. For interfaces configured on those MPCs, <i>all</i> packets that are matched using the ip-options match condition are sent to the Packet Forwarding Engine for processing.</p>	family inet —
ip-options-except values	Do not match the IP option field to the specified value or list of values. For details about specifying the values , see the ip-options match condition.	family inet —
is-fragment	<p>Match if the packet is a trailing fragment of a fragmented packet. Do not match the first fragment of a fragmented packet.</p> <p>This match condition is an alias for the bit-field match condition fragment-offset 0 except bits.</p> <p>NOTE: To match both first and trailing fragments, you can use two terms that specify different match conditions (first-fragment and is-fragment).</p>	family inet —

Table 357: Service Filter Match Conditions for IPv4 or IPv6 Traffic (*continued*)

Match Condition	Description	Protocol Families	
port number	<p>Match the UDP or TCP source or destination port field.</p> <p>If you configure this match condition, you cannot configure the destination-port match condition or the source-port match condition in the same term.</p> <p>If you configure this match condition for IPv4 traffic, we recommend that you also configure the protocol udp or protocol tcp match statement in the same term to specify which protocol is being used on the port.</p> <p>If you configure this match condition for IPv6 traffic, we recommend that you also configure the next-header udp or next-header tcp match condition in the same term to specify which protocol is being used on the port.</p> <p>In place of the numeric value, you can specify one of the text synonyms listed under destination-port.</p>	family inet	family inet6
port-except number	Do not match the UDP or TCP source or destination port field. For details, see the port match condition.	family inet	family inet6
prefix-list prefix-list-name	Match the prefixes of the source or destination address fields to the prefixes in the specified list. The prefix list is defined at the [edit policy-options prefix-list prefix-list-name] hierarchy level.	family inet	family inet6
protocol number	<p>Match the IP protocol type field.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): ah (51), dstopts (60), egp (8), esp (50), fragment (44), gre (47), hop-by-hop (0), icmp (1), icmp6 (58), icmpv6 (58), igmp (2), ipip (4), ipv6 (41), ospf (89), pim (103), rsvp (46), sctp (132), tcp (6), udp (17), or vrrp (112).</p>	family inet	—
protocol-except number	Do not match the IP protocol type field. For details, see the protocol match condition.	family inet	—
source-address address	<p>Match the IP source address.</p> <p>You cannot specify both the address and source-address match conditions in the same term.</p>	family inet	family inet6
source-address address except	<p>Do not match the IP source address.</p> <p>You cannot specify both the address and source-address match conditions in the same term.</p>	family inet	family inet6

Table 357: Service Filter Match Conditions for IPv4 or IPv6 Traffic (*continued*)

Match Condition	Description	Protocol Families	
source-port <i>number</i>	<p>Match the UDP or TCP source port field.</p> <p>You cannot specify the port and source-port match conditions in the same term.</p> <p>If you configure this match condition for IPv4 traffic, we recommend that you also configure the protocol udp or protocol tcp match statement in the same term to specify which protocol is being used on the port.</p> <p>If you configure this match condition for IPv6 traffic, we recommend that you also configure the next-header udp or next-header tcp match condition in the same term to specify which protocol is being used on the port.</p> <p>In place of the numeric value, you can specify one of the text synonyms listed with the destination-port <i>number</i> match condition.</p>	family inet	family inet6
source-port-except <i>number</i>	Do not match the UDP or TCP source port field. For details, see the source-port match condition.	family inet	family inet6
source-prefix-list <i>name</i>	Match source prefixes in the specified list. Specify the name of a prefix list defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.	family inet	family inet6
tcp-flags <i>value</i>	<p>Match one or more of the low-order 6 bits in the 8-bit TCP flags field in the TCP header.</p> <p>To specify individual bit fields, you can specify the following text synonyms or hexadecimal values:</p> <ul style="list-style-type: none"> • fin (0x01) • syn (0x02) • rst (0x04) • push (0x08) • ack (0x10) • urgent (0x20) <p>In a TCP session, the SYN flag is set only in the initial packet sent, while the ACK flag is set in all packets sent after the initial packet.</p> <p>You can string together multiple flags using the bit-field logical operators.</p> <p>For combined bit-field match conditions, see the tcp-established and tcp-initial match conditions.</p> <p>If you configure this match condition for IPv4 traffic, we recommend that you also configure the protocol tcp match statement in the same term to specify that the TCP protocol is being used on the port.</p> <p>If you configure this match condition for IPv6 traffic, we recommend that you also configure the next-header tcp match condition in the same term to specify that the TCP protocol is being used on the port.</p>	family inet	family inet6



NOTE: If you specify an IPv6 address in a match condition (the address, destination-address, or source-address match conditions), use the syntax for text representations described in RFC 2373, *IP Version 6 Addressing Architecture*. For more information about IPv6 addresses, see “IPv6 Overview” and “IPv6 Standards” in the *Junos OS Routing Protocols Configuration Guide*.

Related Documentation

- [Service Filter Overview on page 4522](#)
- [Guidelines for Configuring Service Filters on page 4524](#)
- [Example: Configuring and Applying Service Filters on page 4678](#)
- [Service Filter Terminating Actions on page 4766](#)
- [Service Filter Nonterminating Actions on page 4766](#)

Service Filter Terminating Actions

Service filters support different sets of terminating actions than standard stateless firewall filters or simple filters.



NOTE: Service filters do not support the next term action.

[Table 358 on page 4766](#) describes the terminating actions you can configure in a service filter term.

Table 358: Terminating Actions for Service Filters

Terminating Action	Description	Protocol Families
service	Direct the packet to service processing.	<ul style="list-style-type: none"> • inet • inet6
skip	Let the packet bypass service processing.	<ul style="list-style-type: none"> • inet • inet6

Related Documentation

- [Service Filter Overview on page 4522](#)
- [Guidelines for Configuring Service Filters on page 4524](#)
- [Example: Configuring and Applying Service Filters on page 4678](#)
- [Service Filter Match Conditions for IPv4 or IPv6 Traffic on page 4760](#)
- [Service Filter Nonterminating Actions on page 4766](#)

Service Filter Nonterminating Actions

Service filters support different sets of terminating actions for each protocol family.



NOTE: Service filters do not support the `next term` action.

Table 359 on page 4767 describes the nonterminating actions you can configure in a service filter term.

Table 359: Nonterminating Actions for Service Filters

Nonterminating Action	Description	Protocol Families
<code>count counter-name</code>	Count the packet in the named counter.	<ul style="list-style-type: none"> • <code>inet</code> • <code>inet6</code>
<code>log</code>	Log the packet header information in a buffer within the Packet Forwarding Engine. You can access this information by issuing the <code>show firewall log</code> command at the command-line interface (CLI).	<ul style="list-style-type: none"> • <code>inet</code> • <code>inet6</code>
<code>port-mirror</code>	Port-mirror the packet based on the specified family. Supported on M120 routers, M320 routers configured with Enhanced III FPCs, MX Series routers, and EX Series switches only.	<ul style="list-style-type: none"> • <code>inet</code> • <code>inet6</code>
<code>sample</code>	Sample the packet.	<ul style="list-style-type: none"> • <code>inet</code> • <code>inet6</code>

Related Documentation

- [Service Filter Overview on page 4522](#)
- [Guidelines for Configuring Service Filters on page 4524](#)
- [Example: Configuring and Applying Service Filters on page 4678](#)
- [Service Filter Match Conditions for IPv4 or IPv6 Traffic on page 4760](#)
- [Service Filter Terminating Actions on page 4766](#)

Reference Information for Firewall Filters in Logical Systems

- [Unsupported Firewall Filter Statements for Logical Systems on page 4767](#)
- [Unsupported Actions for Firewall Filters in Logical Systems on page 4769](#)

Unsupported Firewall Filter Statements for Logical Systems

Table 360 on page 4768 shows statements that are supported at the `[edit firewall]` hierarchy level but not at the `[edit logical-systems logical-system-name firewall]` hierarchy level.

Table 360: Unsupported Firewall Statements for Logical Systems

Statement	Example	Description
accounting-profile	<pre>[edit] logical-systems { ls1 { firewall { family inet { filter myfilter { accounting-profile fw-profile; ... } } } } }</pre>	In this example, the accounting-profile statement is not allowed because the accounting profile fw-profile is configured under the [edit accounting-options] hierarchy.
hierarchical-policer	<pre>[edit] logical-systems { lr1 { firewall { hierarchical-policer { ... } } } }</pre>	In this example, the hierarchical policer statement requires a class-of-service configuration, which is not supported under logical systems.
load-balance-group	<pre>[edit] logical-systems { ls1 { firewall { load-balance-group lb-group { next-hop-group nh-group; } } } }</pre>	<p>This configuration is not allowed because the next-hop-group nh-group statement must be configured at the [edit forwarding-options next-hop-group] hierarchy level—outside the [edit logical-systems logical-system-name firewall] hierarchy.</p> <p>Currently, the forwarding-options dhcp-relay statement is the only forwarding option supported for logical systems.</p>

Table 360: Unsupported Firewall Statements for Logical Systems (*continued*)

Statement	Example	Description
virtual-channel	<pre>[edit] logical-systems { ls1 { firewall { family inet { filter foo { term one { from { source-address 10.1.0.0/16; } then { virtual-channel sammy; } } } } } } }</pre>	<p>This configuration is not allowed because the virtual channel sammy refers to an object defined at the [edit class-of-service] hierarchy level, and class of service is not supported for logical systems.</p> <p>NOTE:</p> <p>The virtual-channel statement is supported for J Series devices only, provided the firewall filter is configured outside of a logical-system.</p>

Related Documentation

- [Stateless Firewall Filters in Logical Systems Overview on page 4535](#)
- [Guidelines for Configuring and Applying Firewall Filters in Logical Systems on page 4536](#)
- [Unsupported Actions for Firewall Filters in Logical Systems on page 4769](#)
- “[Introduction to Logical Systems on page 3259](#)” in the *Logical Systems Configuration Guide*
- “[Logical Systems Operations and Restrictions on page 3263](#)” in the *Logical Systems Configuration Guide*

Unsupported Actions for Firewall Filters in Logical Systems

[Table 361 on page 4769](#) describes the firewall filter actions that are supported at the **[edit firewall]** hierarchy level, but not supported at the **[edit logical-systems logical-system-name firewall]** hierarchy level.

Table 361: Unsupported Actions for Firewall Filters in Logical Systems

Firewall Filter Action	Example	Description
------------------------	---------	-------------

Terminating Actions Not Supported in a Logical System

Table 361: Unsupported Actions for Firewall Filters in Logical Systems (*continued*)

Firewall Filter Action	Example	Description
logical-system	<pre>[edit] logical-systems { ls1 { firewall { family inet { filter foo { term one { from { source-address 10.1.0.0/16; } then { logical-system fred; } } } } } } }</pre>	Because the logical-system action refers to fred —a logical system defined outside the local logical system—, this action is not supported.

Nonterminating Actions Not Supported in a Logical System

ipsec-sa	<pre>[edit] logical-systems { ls1 { firewall { family inet { filter foo { term one { from { source-address 10.1.0.0/16; } then { ipsec-sa barney; } } } } } } }</pre>	Because the ipsec-sa action modifier references barney —a security association defined outside the local logical system—this action is not supported.
-----------------	---	---

Table 361: Unsupported Actions for Firewall Filters in Logical Systems (*continued*)

Firewall Filter Action	Example	Description
next-hop-group	<pre> [edit] logical-systems { ls1 { firewall { family inet { filter foo { term one { from { source-address 10.1.0.0/16; } then { next-hop-group fred; } } } } } } } </pre>	Because the next-hop-group action refers to fred —an object defined at the [edit forwarding-options next-hop-group] hierarchy level—this action is not supported.
port-mirror	<pre> [edit] logical-systems { ls1 { firewall { family inet { filter foo { term one { from { source-address 10.1.0.0/16; } then { port-mirror; } } } } } } } </pre>	Because the port-mirror action relies on a configuration defined at the [edit forwarding-options port-mirroring] hierarchy level, this action is not supported.

Table 361: Unsupported Actions for Firewall Filters in Logical Systems (*continued*)

Firewall Filter Action	Example	Description
sample	<pre> [edit] logical-systems { ls1 { firewall { family inet { filter foo { term one { from { source-address 10.1.0.0/16; } then { sample; } } } } } } } </pre>	<p>In this example, the sample action depends on the sampling configuration defined under the [edit forwarding-options] hierarchy. Therefore, the sample action is not supported.</p>
syslog	<pre> [edit] logical-systems { ls1 { firewall { family inet { filter icmp-syslog { term icmp-match { from { address { 192.168.207.222/32; } protocol icmp; } then { count packets; syslog; accept; } } term default { then accept; } } } } } } </pre>	<p>In this example, there must be at least one system log (system syslog file filename) with the firewall facility enabled for the icmp-syslog filter's logs to be stored.</p> <p>Because this firewall configuration relies on a configuration outside the logical system, the syslog action modifier is not supported.</p>

- Related Documentation**
- [Stateless Firewall Filters in Logical Systems Overview on page 4535](#)
 - [Guidelines for Configuring and Applying Firewall Filters in Logical Systems on page 4536](#)
 - [Unsupported Firewall Filter Statements for Logical Systems on page 4767](#)
 - “Introduction to Logical Systems on page 3259” in the *Logical Systems Configuration Guide*

- "Logical Systems Operations and Restrictions on page 3263" in the *Logical Systems Configuration Guide*

Firewall Filters Statement Hierarchies

- Statement Hierarchy for Configuring Interface-Specific Firewall Filters on page 4773
- Statement Hierarchy for Applying Interface-Specific Firewall Filters on page 4774
- Statement Hierarchy for Assigning Interfaces to Interface Groups on page 4774
- Statement Hierarchy for Configuring a Filter to Match on a Set of Interface Groups on page 4775
- Statement Hierarchy for Applying Filters to an Interface Group on page 4776
- Statement Hierarchy for Defining an Interface Set on page 4777
- Statement Hierarchy for Configuring a Filter to Match on an Interface Set on page 4777
- Statement Hierarchy for Configuring FBF for IPv4 or IPv6 Traffic on page 4778
- Statement Hierarchy for Configuring FBF for IPv4 Traffic on ACX Series Routers on page 4778
- Statement Hierarchy for Configuring FBF for MPLS-Tagged IPv4 Traffic on page 4779
- Statement Hierarchy for Configuring Routing Instances for FBF on page 4781
- Statement Hierarchy for Applying FBF Filters to Interfaces on page 4782
- Statement Hierarchy for Configuring Firewall Filter Accounting Profiles on page 4783
- Statement Hierarchy for Applying Firewall Filter Accounting Profiles on page 4784

Statement Hierarchy for Configuring Interface-Specific Firewall Filters

To enable interface-specific instances for stateless firewall filters, include the **interface-specific** statement in the **filter** *filter-name* or **service-filter** *service-filter-name* stanza. Any counters specified as actions in an interface-specific filter are maintained separately per filter instance. Any policers specified as actions in an interface-specific filter are applied per filter instance.

```

firewall {
  family family-name {
    (filter filter-name | service-filter service-filter-name) {
      ...
      interface-specific;
      ...
      term term-name {
        from {
          match-conditions;
        }
        then {
          count counter-name;
          policer policer-name;
        }
      }
      ...
    }
  }
}

```

```
]
```

You can include the firewall configuration at one of the following hierarchy levels:

- `[edit]`
- `[edit logical-systems logical-system-name]`

**Related
Documentation**

- *Interface-Specific Firewall Filter Instances Overview*
- [Statement Hierarchy for Applying Interface-Specific Firewall Filters on page 4774](#)
- *Example: Configuring Interface-Specific Firewall Filter Counters*

Statement Hierarchy for Applying Interface-Specific Firewall Filters

To apply an interface-specific stateless firewall filter to a logical interface, include the **input *filter-name*** or **output *filter-name*** statement in the **filter** or **service-filter** stanza of the interfaces configuration:

```
interfaces {  
  interface-name {  
    unit unit-number {  
      family family-name {  
        filter {  
          input filter-name-1;  
          output filter-name-2;  
        }  
        service-filter {  
          input service-filter-name-1;  
          output service-filter-name-2;  
        }  
      }  
    }  
  }  
}
```

You can include the interface configuration at one of the following hierarchy levels:

- `[edit]`
- `[edit logical-systems logical-system-name]`

**Related
Documentation**

- *Interface-Specific Firewall Filter Instances Overview*
- [Statement Hierarchy for Configuring Interface-Specific Firewall Filters on page 4773](#)
- *Example: Configuring Interface-Specific Firewall Filter Counters*

Statement Hierarchy for Assigning Interfaces to Interface Groups

To assign a logical interface to an interface group, specify the group number by including the **group *interface-group-number*** statement in the **filter** stanza:

```
interfaces {  
  interface-name {  
    unit unit-number {
```

```

family ( inet | inet6 | vpls | ccc | bridge ) {
  filter {
    group interface-group-number;
  }
}

```



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

You can configure the firewall filter at one of the following hierarchy levels:

- [edit]
- [edit logical-systems *logical-system-name*]

Related Documentation

- [Filtering Packets Received on a Set of Interface Groups Overview on page 4514](#)
- [Statement Hierarchy for Configuring a Filter to Match on a Set of Interface Groups on page 4775](#)
- [Example: Filtering Packets Received on an Interface Group on page 4634](#)

Statement Hierarchy for Configuring a Filter to Match on a Set of Interface Groups

You can configure a standard stateless firewall filter or a service filter term that matches packets tagged for a specified interface group or set of interface groups.

To configure a standard stateless firewall filter that matches packets tagged for a specified interface group or set of interface groups, configure a filter term that uses the **interface-group** *interface-group-number* match condition:

```

firewall {
  family (inet | inet6 | vpls | ccc | bridge) {
    filter filter-name {
      term term-name {
        from {
          interface-group interface-group-number;
        }
        then {
          filter-actions;
        }
      }
    }
  }
}

```

To configure a service filter that matches packets tagged for a specified interface group or set of interface groups, configure a filter term that uses the **interface-group** *interface-group-name* match condition:

```

firewall {
  family (inet | inet6) {
    service-filter filter-name {

```

```
term term-name {  
  from {  
    interface-group interface-group-number;  
  }  
  then {  
    service-filter-actions;  
  }  
}  
}  
}
```

You can configure the firewall filter at one of the following hierarchy levels:

- [edit]
- [edit logical-systems *logical-system-name*]

**Related
Documentation**

- [Filtering Packets Received on a Set of Interface Groups Overview on page 4514](#)
- [Statement Hierarchy for Assigning Interfaces to Interface Groups on page 4774](#)
- [Example: Filtering Packets Received on an Interface Group on page 4634](#)

Statement Hierarchy for Applying Filters to an Interface Group

To apply a standard stateless firewall filter to an interface group, include the input *filter-name* or output *filter-name* in the filter stanza:

```
interfaces {  
  interface-name {  
    unit unit-number {  
      family family-name {  
        ...  
        filter {  
          input filter-name;  
          output filter-name;  
        }  
      }  
    }  
  }  
}
```

You can include the interface configuration at one of the following hierarchy levels:

- [edit]
- [edit logical-systems *logical-system-name*]

**Related
Documentation**

- [Interface-Specific Firewall Filter Instances Overview](#)
- [Statement Hierarchy for Assigning Interfaces to Interface Groups on page 4774](#)
- [Statement Hierarchy for Configuring a Filter to Match on a Set of Interface Groups on page 4775](#)
- [Example: Filtering Packets Received on an Interface Group on page 4634](#)

Statement Hierarchy for Defining an Interface Set

To configure a named group of interfaces that can be referenced in a stateless firewall filter match condition, use the **interface-set** statement to define the interface-set name and two or more interfaces:

```
firewall {
  interface-set interface-set-name {
    interface-name;
  }
}
```

You can include the statements at one of the following hierarchy levels:

- [\[edit firewall\]](#)
- [\[edit logical-systems logical-system-name firewall\]](#)

To specify that the interface set contains all interfaces of a particular type, you can use the '*' (asterisk) wildcard character. For example, use **fe-*** to specify all Fast Ethernet interfaces.

Related Documentation

- [Filtering Packets Received on an Interface Set Overview on page 4515](#)
- [Statement Hierarchy for Configuring a Filter to Match on an Interface Set on page 4777](#)
- [Example: Configuring a Rate-Limiting Filter Based on Destination Class on page 4617](#)
- [Example: Filtering Packets Received on an Interface Set on page 4638](#)

Statement Hierarchy for Configuring a Filter to Match on an Interface Set

To configure a standard stateless firewall filter that matches packets tagged for a specified interface group or set of interface groups, configure a filter term that uses the **interface-group interface-group-name** match condition:

```
firewall {
  family (any | inet | inet6 | mpls | vpls | bridge) {
    filter filter-name {
      term term-name {
        from {
          interface-set interface-set-name;
        }
        then {
          filter-actions;
        }
      }
    }
  }
}
```

Related Documentation

- [Filtering Packets Received on an Interface Set Overview on page 4515](#)
- [Statement Hierarchy for Defining an Interface Set on page 4777](#)
- [Example: Configuring a Rate-Limiting Filter Based on Destination Class on page 4617](#)
- [Example: Filtering Packets Received on an Interface Set on page 4638](#)

Statement Hierarchy for Configuring FBF for IPv4 or IPv6 Traffic

You can configure stateless firewall filters for filter-based forwarding by configuring filter terms that specify the **forwarding-class** *class-name* nonterminating action or the **routing-instance** *routing-instance-name* terminating action:

```
firewall {  
  family (inet | inet6) {  
    filter filter-name {  
      term term-name {  
        from {  
          ipv4-or-ipv6-match-conditions;  
        }  
        then {  
          forwarding-class class-name; #optional  
          other-optional-nonterminating-actions;  
          routing-instance routing-instance-name <topology topology-name>;  
        }  
      }  
    }  
  }  
}
```

You can include the firewall configuration at one of the following hierarchy levels:

- [edit]
- [edit logical-systems *logical-system-name*]

Related Documentation

- [Filter-Based Forwarding Overview on page 4515](#)
- [Standard Firewall Filter Match Conditions for IPv4 Traffic on page 4707](#)
- [Standard Firewall Filter Match Conditions for IPv6 Traffic on page 4716](#)
- [Statement Hierarchy for Configuring Routing Instances for FBF on page 4781](#)
- [Example: Configuring Filter-Based Forwarding on the Source Address on page 4644](#)

Statement Hierarchy for Configuring FBF for IPv4 Traffic on ACX Series Routers

On ACX Series routers, you can configure stateless firewall filters for filter-based forwarding (FBF) by configuring filter terms that specify the optional nonterminating actions or the **routing-instance** *routing-instance-name* terminating action:

```
firewall {  
  family inet {  
    filter filter-name {  
      term term-name {  
        from {  
          ipv4-match-conditions;  
        }  
        then {  
          optional-nonterminating-actions;  
          routing-instance routing-instance-name ;  
        }  
      }  
    }  
  }  
}
```



```

    }
  }
}

```

You can include the firewall configuration at the **[edit]** hierarchy level:

The **[edit logical-systems *logical-system-name*]** hierarchy level is not supported on the ACX Series routers.

Related Documentation

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Standard Firewall Filter Match Conditions and Actions on ACX Series Routers Overview on page 4751](#)
- [Standard Firewall Filter Terminating Actions on ACX Series Routers on page 4756](#)
- [Standard Firewall Filter Nonterminating Actions on ACX Series Routers on page 4757](#)
- [Standard Firewall Filter Match Conditions for IPv4 Traffic on ACX Series Routers on page 4752](#)
- [Standard Firewall Filter Match Conditions for MPLS Traffic on ACX Series Routers on page 4756](#)

Statement Hierarchy for Configuring FBF for MPLS-Tagged IPv4 Traffic

- [Matching on IPv4 Address and TCP/UDP Port Fields on page 4779](#)
- [Configuration Example on page 4780](#)

Matching on IPv4 Address and TCP/UDP Port Fields

To configure a firewall filter term that matches on IP source and destination address fields, and TCP and UDP ports in the IPv4 header of packets in an MPLS flow, you can specify supported match conditions as shown here:

```

[edit]
interfaces {
  interface-name {
    unit logical-unit-number {
      family family {
        address ip-address;
      }
      family mpls {
        filter {
          input filter-name;
        }
      }
    }
  }
}
firewall {
  family mpls {
    filter filter-name {
      term term-name {
        from {
          ip-version ipv4 {
            destination-address {

```

```
        ip-address;
      }
      source-address {
        ip-address;
      }
      protocol tcp {
        destination-port tcp-port;
        destination-port-except tcp-port;
        source-port tcp-port;
        source-port-except tcp-port;
      }
      protocol udp {
        destination-port udp-port;
        destination-port-except udp-port;
        source-port udp-port;
        source-port-except udp-port;
      }
    }
  }
  then {
    ...
  }
}
}
```

Configuration Example

```
interfaces {
  ge-6/0/0 {
    unit 0 {
      family inet {
        address 20.20.20.1/30;
      }
      family mpls {
        filter {
          input mpls-ipv4-filter1;
        }
      }
    }
  }
}
firewall {
  family mpls {
    filter mpls-ipv4-filter1 {
      term term1 {
        from {
          ip-version ipv4 {
            source-address {
              10.0.0.1/32;
            }
          }
          protocol tcp {
            destination-port ftp;
          }
        }
      }
    }
  }
}
```

```

    }
    then {
        count counter1;
        discard;
    }
}
}
}
}
}

```

Related Documentation

- [Filter-Based Forwarding Overview on page 4515](#)
- [Standard Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic on page 4725](#)
- [Statement Hierarchy for Configuring Routing Instances for FBF on page 4781](#)
- [Example: Configuring Filter-Based Forwarding on the Source Address on page 4644](#)

Statement Hierarchy for Configuring Routing Instances for FBF

A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The set of interfaces belongs to the routing tables, and the routing protocol parameters control the information in the routing tables.

To configure a routing instance for filter-based forwarding:

1. The **instance-type** must be **forwarding**. The **forwarding** routing instance type supports filter-based forwarding, where interfaces are not associated with instances. For this instance type, there is no one-to-one mapping between an interface and a routing instance. All interfaces belong to the default instance **inet.0**.
2. The name of the routing instance name must be the one referenced in the firewall filter action.



NOTE: In Junos OS Release 9.0 and later, you can no longer specify a routing-instance name of **default** or include special characters within the name of a routing instance.

You must also create a routing table group that adds interface routes to the following routing instances:

- Routing instance named in the action
- Default routing table **inet.0**

You create a routing table group to resolve the routes installed in the routing instance to directly connected next hops on that interface. For more information on routing table groups and interface routes, see the *Junos OS Routing Protocols Configuration Guide*.

```

routing-instances {
    routing-table-name {
        instance-type forwarding;
    }
}

```

```
routing-options {
  static {
    route destination-prefix nexthop address;
  }
}
```

You can include the **forwarding** routing instance at one of the following hierarchy levels:

- **[edit]**
- **[edit logical-systems *logical-system-name*]**

Related Documentation

- [Filter-Based Forwarding Overview on page 4515](#)
- [Statement Hierarchy for Configuring FBF for IPv4 or IPv6 Traffic on page 4778](#)
- [Statement Hierarchy for Configuring FBF for IPv4 Traffic on ACX Series Routers on page 4778](#)
- [Statement Hierarchy for Configuring FBF for MPLS-Tagged IPv4 Traffic on page 4779](#)
- [Example: Configuring Filter-Based Forwarding on the Source Address on page 4644](#)

Statement Hierarchy for Applying FBF Filters to Interfaces

To apply filter-based forwarding to a logical interface, include the **input** or **output** statement in the **filter** stanza.



NOTE: An interface configured with filter-based forwarding does not support source-class usage (SCU) filter matching and unicast reverse-path forwarding (RPF) check filters.

```
interfaces {
  interface-name {
    unit unit-number {
      family (inet | inet6 | mpls) {
        filter {
          input filter-name;
          output filter-name;
        }
        address address;
      }
    }
  }
}
```

You can include the interfaces configuration at one of the following hierarchy levels:

- **[edit]**
- **[edit logical-systems *logical-system-name*]**

- Related Documentation**
- [Filter-Based Forwarding Overview on page 4515](#)
 - [Statement Hierarchy for Configuring FBF for IPv4 or IPv6 Traffic on page 4778](#)
 - [Statement Hierarchy for Configuring FBF for MPLS-Tagged IPv4 Traffic on page 4779](#)
 - [Statement Hierarchy for Configuring Routing Instances for FBF on page 4781](#)
 - [Example: Configuring Filter-Based Forwarding on the Source Address on page 4644](#)

Statement Hierarchy for Configuring Firewall Filter Accounting Profiles

To configure an accounting profile that you can apply to a firewall filter, include the **filter-profile** *filter-profile-name* statement in the **accounting-options** stanza.

```
accounting-options {
  filter-profile filter-profile-name {
    file log-filename {
      archive-sites {
        site-urls;
      }
      files number;
      size bytes;
      start-time time;
      transfer-interval minutes;
    }
    interval minutes;
    counters {
      counter-name-1;
      counter-name-2;
    }
  }
}
```

You can include the accounting options configuration at one of the following hierarchy levels:

- **[edit]**
- **[edit logical-systems *logical-system-name*]**

To specify the name of the accounting data log file in the **/var/log** directory to be used in conjunction with the accounting profile, include the **file *log-filename*** statement.

To specify how often statistics are collected for the accounting profile, include the **interval *minutes*** statement.

To specify the names of the firewall filter counters for which filter profile statistics are collected, include the **counters** statement.

- Related Documentation**
- [Accounting for Standard Firewall Filters Overview on page 4517](#)
 - [Statement Hierarchy for Applying Firewall Filter Accounting Profiles on page 4784](#)
 - [Example: Configuring Statistics Collection for a Standard Firewall Filter on page 4669](#)

Statement Hierarchy for Applying Firewall Filter Accounting Profiles

You can apply an accounting profile to a standard stateless firewall filter for any supported protocol family except **family any**.

To apply a filter-accounting profile to a stateless firewall filter, include the **accounting-profile *accounting-profile-name*** statement at the firewall **filter** stanza:

```
firewall {  
  family family-name {  
    filter filter-name {  
      accounting-profile accounting-profile-name;  
      interface-specific;  
      physical-interface-policer;  
      term {  
        filter filter-profile-name;  
      }  
      term term-name {  
        from {  
          match-conditions;  
        }  
        then {  
          actions;  
        }  
      }  
    }  
  }  
}
```

You can include the firewall configuration at one of the following hierarchy levels:

- **[edit]**
- **[edit logical-systems *logical-system-name*]**

Related Documentation

- [Accounting for Standard Firewall Filters Overview on page 4517](#)
- [Statement Hierarchy for Configuring Firewall Filter Accounting Profiles on page 4783](#)
- [Example: Configuring Statistics Collection for a Standard Firewall Filter on page 4669](#)

Summary of Firewall Filters Configuration Statements

accounting-profile

Syntax	<code>accounting-profile <i>name</i>;</code>
Hierarchy Level	<code>[edit firewall family <i>family-name</i> filter <i>filter-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Enable collection of accounting data for the specified filter.
Options	<i>name</i> —Name assigned to the accounting profile.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Accounting for Standard Firewall Filters Overview on page 4517


destination-address

Syntax	<code>destination-address <i>address</i>;</code>
Hierarchy Level	<code>[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> from ip-version <i>ip-version</i>]</code>
Release Information	Statement introduced in Junos OS Release 10.1R1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the IPv4 or IPv6 address of the destination node that receives the packet.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Standard Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic on page 4725

destination-port

Syntax	<code>destination-port <destination-port>;</code>
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> from ip-version <i>ip-version</i> protocol (tcp udp)]
Release Information	Statement introduced in Junos OS Release 10.1R1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the destination port of the Layer 4 header.
Options	<i>destination-port</i> —The destination port of the Layer 4 header. Range: 0 through 65,535
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Standard Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic on page 4725

enhanced-mode

Syntax	enhanced-mode;
Hierarchy Level	[edit firewall filter <i>filter-name</i>], [edit firewall family <i>family-name</i> filter <i>filter-name</i>], [edit logical-systems <i>logical-system-name</i> firewall filter <i>filter-name</i>], [edit logical-systems <i>logical-system-name</i> firewall family <i>family-name</i> filter <i>filter-name</i>]
Release Information	Statement introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>Limit static service filters or API-client filters to term-based filter format only for inet or inet6 families when enhanced network services mode is configured at the [edit chassis network-services] hierarchy level. When used with one of the chassis enhanced network services modes, firewall filters are generated in term-based format for use with MPC modules.</p> <p>If enhanced network services are not configured for the chassis, the enhanced-mode statement is ignored and any enhanced mode firewall filters are generated in both term-based and compiled format (the default).</p>
	<div>  <p>NOTE: You cannot attach enhanced mode filters to local loopback, management, or MS-DPC interfaces. These interfaces are processed by the Routing Engine kernel and DPC modules and can accept only compiled firewall filter format.</p> </div>
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Network Services Mode Overview</i> in the <i>Junos OS System Basics Configuration Guide</i> • <i>Firewall Filters and Enhanced Network Services Mode Overview</i> in the <i>Junos OS Subscriber Management, Release 12.3</i> • <i>Configuring a Filter for Use with Enhanced Network Services Mode</i> in the <i>Junos OS Subscriber Management, Release 12.3</i>

family (Firewall)

Syntax family *family-name* {
 filter *filter-name* {
 accounting-profile *name*;
 enhanced-mode;
 interface-specific;
 physical-interface-filter;
 }
 prefix-action *name* {
 count;
 destination-prefix-length *prefix-length*;
 policer *policer-name*;
 source-prefix-length *prefix-length*;
 subnet-prefix-length *prefix-length*;
 }
 simple-filter *filter-name* {
 term *term-name* {
 from {
 match-conditions;
 }
 then {
 action;
 action-modifiers;
 }
 }
 }
 }
 }

Hierarchy Level [edit [firewall](#)],
 [edit logical-systems *logical-system-name* [firewall](#)]

Release Information Statement introduced before Junos OS Release 7.4.
 Logical systems support introduced in Junos OS Release 9.3.
 simple-filter statement introduced in Junos OS Release 7.6.
 any family type introduced in Junos OS Release 8.0.
 bridge family type introduced in Junos OS Release 8.4 (MX Series routers only).
 Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure a firewall filter for IP version 4 (IPv4) or IP version 6 (IPv6) traffic. Only on MX Series routers and EX Series switches, configure a firewall filter for Layer 2 traffic in a bridging environment.

Options *family-name*—Version or type of addressing protocol:

- **any**—Protocol-independent match conditions.
- **bridge**—(MX Series routers only) Layer 2 packets that are part of bridging domain.
- **ethernet-switching**—(EX Series switches) Filter Layer 2 (Ethernet) packets and Layer 3 (IP) packets.
- **ccc**—Layer 2 switching cross-connects.

- **inet**—IPv4 addressing protocol.
- **inet6**—IPv6 addressing protocol.
- **mpls**—MPLS.
- **vpls**—Virtual private LAN service (VPLS).

The remaining statements are explained separately.



NOTE: The packet lengths that a policer considers depends on the address family of the firewall filter. See [“Understanding the Frame Length for Policing Packets” on page 4811](#).

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Guidelines for Configuring Standard Firewall Filters on page 4478 • Guidelines for Configuring Service Filters on page 4524 • Guidelines for Configuring Simple Filters on page 4531

filter (Applying to a Logical Interface)

Syntax	<pre>filter { group <i>filter-group-number</i>; input <i>filter-name</i>; input-list [<i>filter-names</i>]; output <i>filter-name</i>; output-list [<i>filter-names</i>]; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Apply a stateless firewall filter to a logical interface at a specific protocol level.
Options	<p>group <i>filter-group-number</i>—Number of the group to which the interface belongs. Range: 1 through 255</p> <p>input <i>filter-name</i>—Name of one filter to evaluate when packets are received on the interface.</p> <p>input-list [<i>filter-names</i>]—Names of filters to evaluate when packets are received on the interface. Up to 16 filters can be included in a filter input list.</p> <p>output <i>filter-name</i>—Name of one filter to evaluate when packets are transmitted on the interface.</p> <p>output-list [<i>filter-names</i>]—Names of filters to evaluate when packets are transmitted on the interface. Up to 16 filters can be included in a filter output list.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Guidelines for Configuring Standard Firewall Filters on page 4478• Guidelines for Applying Standard Firewall Filters on page 4483

filter (Configuring)

Syntax	<pre> filter <i>filter-name</i> { accounting-profile <i>name</i>; enhanced-mode; interface-shared; interface-specific; physical-interface-filter; term <i>term-name</i> { filter <i>filter-name</i>; from { match-conditions; } then { actions; } } } </pre>
Hierarchy Level	<p>[edit dynamic-profiles <i>profile-name</i> firewall family <i>family-name</i>], [edit firewall family <i>family-name</i>], [edit logical-systems <i>logical-system-name</i> firewall family <i>family-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Logical systems support introduced in Junos OS Release 9.3. physical-interface-filter statement introduced in Junos OS Release 9.6. Support at the [edit dynamic-profiles ... family <i>family-name</i>] hierarchy level introduced in Junos OS Release 11.4. Support for the interface-shared> statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Configure firewall filters.
Options	<p><i>filter-name</i>—Name that identifies the filter. This must be a non-reserved string of not more than 64 characters. To include spaces in the name, enclose it in quotation marks (" "). In Junos OS Release 9.0 and later, you can no longer use special characters within the name of a firewall filter. Firewall filter names are restricted from having the form _.* (beginning and ending with underscores) or _.* (beginning with an underscore).</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Guidelines for Configuring Standard Firewall Filters on page 4478 • Guidelines for Applying Standard Firewall Filters on page 4483 • Configuring Multifield Classifiers on page 1390 • Using Multifield Classifiers to Set PLP on page 1461

- [simple-filter on page 4798](#)

firewall

Syntax	firewall { ... }
Hierarchy Level	[edit], [edit logical-systems <i>logical-system-name</i>] [edit dynamic-profiles <i>profile-name</i>],
Release Information	Statement introduced before Junos OS Release 7.4. Logical systems support introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure firewall filters. The statements are explained separately.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Guidelines for Configuring Standard Firewall Filters on page 4478• Guidelines for Configuring Service Filters on page 4524• Guidelines for Configuring Simple Filters on page 4531• Configuring Multifield Classifiers on page 1390• Using Multifield Classifiers to Set PLP on page 1461

interface-set

Syntax	<code>interface-set <i>interface-set-name</i> { <i>interface-name</i>; }</code>
Hierarchy Level	[edit firewall], [edit logical-systems <i>logical-system-name</i> firewall]
Release Information	Statement introduced before Junos OS Release 7.4. Logical systems support introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure an interface set.
Options	<i>interface-name</i> —Names of each interface to include in the interface set. You must specify more than one name.
Required Privilege Level	<code>firewall</code> —To view this statement in the configuration. <code>firewall-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Filtering Packets Received on an Interface Set Overview on page 4515

interface-specific (Firewall Filters)

Syntax	<code>interface-specific;</code>
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i>], [edit logical-systems <i>logical-system-name</i> firewall family <i>family-name</i> filter <i>filter-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Logical systems support introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure interface-specific names for firewall counters.
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Interface-Specific Firewall Filter Instances Overview

ip-version

Syntax	<code>ip-version <i>ip-version</i>;</code>
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in Junos OS Release 10.1R1. Option ipv6 introduced in Junos OS Release 12.2R1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the IP version for the firewall filter.
Options	<i>ip-version</i> —Version of the IP addressing. <ul style="list-style-type: none">• ipv4—IP version 4• ipv6—IP version 6
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Standard Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic on page 4725

prefix-list

Syntax	<pre>prefix-list <i>name</i> { <i>ip-addresses</i>; apply-path <i>path</i>; }</pre>
Hierarchy Level	[edit dynamic policy-options], [edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for configuration in the dynamic database introduced in Junos OS Release 9.5.</p> <p>Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Support for the vpls protocol family introduced in Junos OS Release 10.2.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Define a list of IPv4 or IPv6 address prefixes for use in a routing policy statement or firewall filter statement.</p> <p>You can configure up to 85,325 prefixes in each prefix list. To configure more than 85,325 prefixes, configure multiple prefix lists and apply them to multiple firewall filter terms.</p>
Options	<p><i>name</i>—Name that identifies the list of IPv4 or IPv6 address prefixes.</p> <p><i>ip-addresses</i>—List of IPv4 or IPv6 address prefixes, one IP address per line in the configuration.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Prefix Lists for Use in Routing Policy Match Conditions • Firewall Filter Match Conditions Based on Address Fields on page 4493 • Example: Configuring a Filter to Limit TCP Access to a Port Based On a Prefix List on page 4555

protocol

Syntax	<code>protocol (tcp udp);</code>
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> from ip-version <i>ip-version</i>]
Release Information	Statement introduced in Junos OS Release 10.1R1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the protocol field of IPv4 or the next-header field of the IPv6 address.
Options	tcp —Transmission Control Protocol. udp —User Datagram Protocol.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Standard Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic on page 4725

service-filter (Firewall)

Syntax	<pre> service-filter <i>filter-name</i> { term <i>term-name</i> { from { <i>match-conditions</i>; } then { <i>actions</i>; } } } </pre>
Hierarchy Level	[edit firewall family (inet inet6), [edit logical-systems <i>logical-system-name</i> firewall family (inet inet6)]
Release Information	Statement introduced before Junos OS Release 7.4. Logical systems support introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure service filters.
Options	<p><i>filter-name</i>—Name that identifies the service filter. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	firewall —To view this statement in the configuration. firewall-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Guidelines for Configuring Service Filters on page 4524 • Guidelines for Applying Service Filters on page 4526

simple-filter

Syntax	<pre>simple-filter <i>filter-name</i> { term <i>term-name</i> { from { <i>match-conditions</i>; } then { <i>actions</i>; } } }</pre>
Hierarchy Level	[edit firewall family inet], [edit logical-systems <i>logical-system-name</i> firewall family inet]
Release Information	Statement introduced in Junos OS Release 7.6. Logical systems support introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure simple filters.
Options	<p><i>filter-name</i>—Name that identifies the simple filter. This must be a non-reserved string of not more than 64 characters. No special characters are restricted. However, to include spaces in the name, enclose it in quotation marks (" ").</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Simple Filter Overview on page 4529• How Simple Filters Evaluate Packets on page 4530• Guidelines for Configuring Simple Filters on page 4531• Guidelines for Applying Simple Filters on page 4534

source-address

Syntax	source-address <i>address</i> ;
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> from ip-version <i>ip-version</i>]
Release Information	Statement introduced in Junos OS Release 10.1R1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the IPv4 or IPv6 address of the source node that sends the packet.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Standard Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic on page 4725

source-checking

Syntax	source-checking;
Hierarchy Level	[edit forwarding-options family inet6]
Description	<p>(MX Series 3D Universal Edge Routers Only) Discard IPv6 packets when the source address type is unspecified, loopback, multicast or link-local</p> <p>RFC 4291, <i>IP Version 6 Addressing Architecture</i>, refers to four address types that require special treatment when they are used as source addresses. The four address types are:</p> <ul style="list-style-type: none"> • Unspecified • Loopback • Multicast • Link-Local Unicast <p>The loopback and multicast addresses must never be used as a source address in IPv6 packets. The unspecified and link-local addresses can be used as source addresses but routers must never forward packets that have these addresses as source addresses. Typically, packets that contain unspecified or link-local addresses as source addresses are delivered to the local host. If the destination is not the local host, then the packet must not be forwarded. Configuring this statement filters or discards IPv6 packets of these four address types.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Applying Filters to Forwarding Tables

source-port

Syntax	<code>source-port <source-port>;</code>
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> from ip-version <i>ip-version</i> protocol (tcp udp)]
Release Information	Statement introduced in Junos OS Release 10.1R1. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure the source port of the Layer 4 header.
Options	source-port —The source port of the Layer 4 header. Range: 0 through 65,535
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Standard Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic on page 4725

term (Firewall Filter)

Syntax	<pre> term <i>term-name</i> { from { <i>match-conditions</i>; ip-version ipv4 { <i>match-conditions-mpls-ipv4-address</i>; protocol (tcp udp) { <i>match conditions-mpls-ipv4-port</i>; } } } then { <i>actions</i>; } } </pre>
Hierarchy Level	<pre> [edit firewall family <i>family-name</i> filter <i>filter-name</i>], [edit firewall family <i>family-name</i> service-filter <i>filter-name</i>], [edit firewall family <i>family-name</i> simple-filter <i>filter-name</i>], [edit logical-systems <i>logical-system-name</i> firewall family <i>family-name</i> filter <i>filter-name</i>], [edit logical-systems <i>logical-system-name</i> firewall family <i>family-name</i> service-filter <i>filter-name</i>], [edit logical-systems <i>logical-system-name</i> firewall family <i>family-name</i> simple-filter <i>filter-name</i>] </pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>filter option introduced in Junos OS Release 7.6.</p> <p>Logical systems support introduced in Junos OS Release 9.3.</p> <p>ip-version ipv4 support introduced in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Define a firewall filter term.
Options	<p>actions—(Optional) Actions to perform on the packet if conditions match. You can specify one <i>terminating action</i> supported for the specified filter type. If you do not specify a terminating action, the packets that match the conditions in the from statement are accepted by default. As an option, you can specify one or more <i>nonterminating actions</i> supported for the specified filter type.</p> <p>filter-name—(Optional) For family <i>family-name</i> filter <i>filter-name</i> only, reference another standard stateless firewall filter from within this term.</p> <p>from—(Optional) Match packet fields to values. If not included, all packets are considered to match and the actions and action modifiers in the then statement are taken.</p> <p>match-conditions—One or more conditions to use to make a match on a packet.</p> <p>match-conditions-mpls-ipv4-address—(MPLS-tagged IPv4 traffic only) One or more IP address match conditions to match on the IPv4 packet header. Supports network-based service in a core network with IPv4 packets as an inner payload of an MPLS packet with labels stacked up to five deep.</p>

match-conditions-mpls-ipv4-port—(MPLS-tagged IPv4 traffic only) One or more UDP or TCP port match conditions to use to match a packet in an MPLS flow. Supports network-based service in a core network with IPv4 packets as an inner payload of an MPLS packet with labels stacked up to five deep.

term-name—Name that identifies the term. The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks (" ").

then—(Optional) Actions to take on matching packets. If not included and a packet matches all the conditions in the **from** statement, the packet is accepted.

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

Related Documentation

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Guidelines for Configuring Service Filters on page 4524](#)
- [Guidelines for Configuring Simple Filters on page 4531](#)
- [Guidelines for Configuring and Applying Firewall Filters in Logical Systems on page 4536](#)

Traffic Policers

- [Overview on page 4802](#)
- [Configuration on page 4838](#)
- [Administration on page 5013](#)

Overview

- [Introduction to Traffic Policing on page 4802](#)
- [Introduction to Configuring Policers on page 4811](#)
- [Policer Rate Limits and Actions on page 4821](#)
- [Policer Implementation on page 4828](#)

Introduction to Traffic Policing

- [Traffic Policing Overview on page 4802](#)
- [Traffic Policer Types on page 4807](#)
- [Order of Policer and Firewall Filter Operations on page 4810](#)
- [Understanding the Frame Length for Policing Packets on page 4811](#)

Traffic Policing Overview

This topic covers the following information:

- [Congestion Management for IP Traffic Flows on page 4803](#)
- [Traffic Limits on page 4803](#)
- [Traffic Color Marking on page 4804](#)

- [Forwarding Classes and PLP Levels on page 4805](#)
- [Policer Application to Traffic on page 4806](#)

Congestion Management for IP Traffic Flows

Traffic policing, also known *rate limiting*, is an essential component of network access security that is designed to thwart denial-of-service (DoS) attacks. Traffic policing enables you to control the maximum rate of IP traffic sent or received on an interface and also to partition network traffic into multiple priority levels, also known as *classes of service*. A policer defines a set of traffic rate limits and sets consequences for traffic that does not conform to the configured limits. Packets in a traffic flow that does not conform to traffic limits are either discarded or marked with a different forwarding class or packet loss priority (PLP) level.

With the exception of policers configured to rate-limit aggregate traffic (all protocol families and logical interfaces configured on a physical interface), you can apply a policer to all IP packets in a Layer 2 or Layer 3 traffic flow at a logical interface.

With the exception of policers configured to rate-limit based on physical interface media rate, you can apply a policer to specific IP packets in a Layer 3 traffic flow at a logical interface by using a stateless firewall filter.

You can apply a policer to inbound or outbound interface traffic. Policers applied to inbound traffic help to conserve resources by dropping traffic that does not need to be routed through a network. Dropping inbound traffic also helps to thwart denial-of-service (DoS) attacks. Policers applied to outbound traffic control the bandwidth used.



NOTE: Traffic policers are instantiated on a per-PIC basis. Traffic policing does not work when the traffic for one local policy decision function (L-PDF) subscriber is distributed over multiple Multiservices PICs in an AMS group.

Traffic Limits

Junos[®] OS policers use the *token-bucket* algorithm to enforce a limit on average transmit or receive rate of IP traffic at an interface while allowing bursts of traffic up to a maximum value based on the overall traffic load. The token-bucket algorithm offers more flexibility than the *leaky-bucket* algorithm in that you can allow a specified amount of bursting before starting to discard packets or apply a penalty to packet output-queuing priority or packet drop priority.


In the token-bucket model, the bucket represents the policing function. Tokens are added to the bucket at a fixed rate, but only up to the specified depth of the bucket. Each token represents a “credit” for some number of bits, and tokens in the bucket are “cashed in” for the ability to transmit or receive traffic at the interface. When sufficient tokens are present in the bucket, a traffic flow continues unrestricted. Otherwise, packets might be dropped or else re-marked with a lower forwarding class, a higher packet loss priority (PLP) level, or both.

- The rate at which tokens are added to the bucket represents the highest average transmit or receive rate in bits per second allowed for a given service level. You specify

this highest average traffic rate as the *bandwidth limit* of the policer. If the traffic arrival rate is so high that at some point insufficient tokens are present in the bucket, then the traffic flow is no longer conforming to the traffic limit.

- The depth of the bucket in bytes controls the amount of back-to-back bursting allowed. You specify this factor as the *burst-size limit* of the policer. This second limit affects the average transmit or receive rate by limiting the number of bytes permitted in a transmission burst for a given interval of time. Bursts exceeding the current burst-size limit are dropped until there are sufficient tokens available to permit the burst to proceed.

Figure 100: Network Traffic and Burst Rates



As shown in the figure above, a UPC bar code is a good facsimile of what traffic looks like on the line; an interface is either transmitting (bursting at full rate) or it is not. The black lines represent periods of data transmission and the white space represents periods of silence when the token bucket can replenish.

Depending on the type of policer used, packets in a policed traffic flow that surpasses the defined limits might be implicitly set to a higher PLP level, assigned to a configured forwarding class or set to a configured PLP level (or both), or simply discarded. If packets encounter downstream congestion, packets with a **low** PLP level are less likely to be discarded than those with a **medium-low**, **medium-high**, or **high** PLP level.

Traffic Color Marking

Based on the particular set of traffic limits configured, a policer identifies a traffic flow as belonging to one of either two or three categories that are similar to the colors of a traffic light used to control automobile traffic.

A *two-color-marking* policer categorizes traffic as either conforming to the traffic limits (green) or violating the traffic limits (red):

- **Green**—Two-color-marking policers implicitly set the packets in a green flow to the low PLP level, and you cannot configure any policer actions for conforming traffic.
- **Red**—Two-color-marking policers do not perform any implicit actions on packets in a red flow. Instead, those packets are handled according to the actions specified in the policer configuration. You can configure a two-color-marking policer to simply discard packets if the traffic flow is red. Alternatively, you can configure a two-color-marking policer to handle the packets in a red flow by setting the PLP level to either **low** or **high**, assigning the packets to any forwarding class already configured, or both.

On MX Series, M120, and M320 routers and M7i and M10i routers with the Enhanced CFEB (CFEB-E) and EX Series switches only, you can specify two additional PLP levels for packets in a red flow: **medium-low** or **medium-high**.

Three-color-marking policers categorize traffic as conforming to the traffic limits (green), violating the traffic limits (red), or exceeding the traffic limits but within an allowed range (yellow):

- **Green**—Like two-color-marking policers, three-color-marking policers implicitly set the packets in a green flow to the low PLP level, and you cannot configure any policer actions for conforming traffic.
- **Yellow**—Unlike two-color-marking policers, three-color-marking policers categorize a second type of nonconforming traffic: yellow.

Single-rate three-color policing categorizes as yellow traffic that exceeds the traffic limits while conforming to a second defined burst-size limit. Two-rate three-color policing categorizes as yellow traffic that exceeds the traffic limits while conforming to both a second defined burst-size limit and a second defined bandwidth limit.

Three-color-marking policers implicitly set the packets in a yellow flow to the medium-high PLP level so that the packets incur a less severe penalty than those in a red flow. You cannot configure any policer actions for yellow traffic.

- **Red**—Unlike two-color-marking policers, three-color-marking policers implicitly set the packets in a red flow to the high PLP level, which is the highest PLP value. You can also configure a three-color-marking policer to discard the packets in a red flow instead of forwarding them with a high PLP setting.

Two-color-marking policers allows bursts of traffic for short periods, whereas three-color-marking policers allow more sustained bursts of traffic.

Forwarding Classes and PLP Levels

A packet's forwarding class assignment and PLP level are used by the Junos OS class of service (CoS) features. The Junos CoS features include a set of mechanisms that you can use to provide differentiated services when best-effort traffic delivery is insufficient. For router (and switch) interfaces that carry IPv4, IPv6, and MPLS traffic, you can configure CoS features to take in a single flow of traffic entering at the edge of your network and provide different levels of service across the network—internal forwarding and scheduling (queuing) for output—based on the forwarding class assignments and PLP levels of the individual packets.



NOTE: Forwarding-class or loss-priority assignments performed by a policer or a stateless firewall filter override any such assignments performed on the ingress by the CoS default IP precedence classification at all logical interfaces or by any configured behavior aggregate (BA) classifier that is explicitly mapped to a logical interface.

Based on CoS configurations, packets of a given forwarding class are transmitted through a specific output queue, and each output queue is associated with a transmission service level defined in a *scheduler*.

Based on other CoS configurations, when packets in an output queue encounter congestion, packets with higher loss-priority values are more likely to be dropped by the random early detection (RED) algorithm. Packet loss priority values affect the scheduling of a packet without affecting the packet's relative ordering within the traffic flow.

Policer Application to Traffic

After you have defined and named a policer, it is stored as a template. You can later use the same policer name to provide the same policer configuration each time you want to use it. This eliminates the need to define the same policer values more than once.

You can apply a policer to a traffic flow in either of two ways:

- You can configure a standard stateless firewall filter that specifies the **policer *policer-name*** nonterminating action or the **three-color-policer (single-rate | two-rate) *policer-name*** nonterminating action. When you apply the standard filter to the input or output at a logical interface, the policer is applied to all packets of the filter-specific protocol family that match the conditions specified in the filter configuration.

With this method of applying a policer, you can define specific classes of traffic on an interface and apply traffic rate-limiting to each class.

- You can apply a policer directly to an interface so that traffic rate-limiting applies to all traffic on that interface, regardless of protocol family or any match conditions.

You can configure policers at the queue, logical interface, or Layer 2 (MAC) level. Only a single policer is applied to a packet at the egress queue, and the search for policers occurs in this order:

- Queue level
- Logical interface level
- Layer 2 (MAC) level

Related Documentation

- [Stateless Firewall Filter Overview on page 4463.](#)
- [Traffic Policer Types](#)
- [Order of Policer and Firewall Filter Operations on page 4810](#)
- [Packet Flow Through the CoS Process Overview on page 1273](#)

Traffic Policer Types

This topic covers the following information:

- [Single-Rate Two-Color Policers on page 4807](#)
- [Three-Color Policers on page 4808](#)
- [Two-Color and Three-Color Policer Options on page 4808](#)

Single-Rate Two-Color Policers

You can use a single-rate two-color policer, or “policer” when used without qualification, to rate-limit a traffic flow to an average bits-per-second arrival rate (specified by the single specified bandwidth limit) while allowing bursts of traffic for short periods (controlled by the single specified burst-size limit). This type of policer categorizes a traffic flow as either green (conforming) or red (nonconforming). Packets in a green flow are implicitly set to a **low** loss priority and then transmitted. Packets in a red flow are handled according to actions specified in the policer configuration. Packets in a red flow can be marked—set to a specified forwarding class, set to a specified loss priority, or both—or they can be discarded.

A single-rate two-color policer is most useful for metering traffic at the port (physical interface) level.

Basic Single-Rate Two-Color Policer

You can apply a basic single-rate two-color policer to Layer 3 traffic in either of two ways: as an interface policer or as a firewall filter policer. You can apply the policer as an *interface policer*, meaning that you apply the policer directly to a logical interface at the protocol family level. If you want to apply the policer to selected packets only, you can apply the policer as a *firewall filter policer*, meaning that you reference the policer in a stateless firewall filter term and then apply the filter to a logical interface at the protocol family level.

Bandwidth Policer

A bandwidth policer is simply a single-rate two-color policer that is defined using a bandwidth limit specified as a percentage value rather than as an absolute number of bits per second. When you apply the policer (as an interface policer or as a firewall filter policer) to a logical interface at the protocol family level, the effective bandwidth limit is calculated based on either the physical interface media rate or the logical interface configured shaping rate.

Logical Bandwidth Policer

A logical bandwidth policer is a bandwidth policer for which the effective bandwidth limit is calculated based on the logical interface configured shaping rate. You can apply the policer as a firewall filter policer only, and the firewall filter must be configured as an interface-specific filter. When you apply an interface-specific filter to multiple logical interfaces on supported routing platforms, any **count** or **policer** actions act on the traffic stream entering or exiting each individual interface, regardless of the sum of traffic on the multiple interfaces.

Three-Color Policers

The Junos OS supports two types of three-color policers: single-rate and two-rate. The main difference between a single-rate and a two-rate policer is that the single-rate policer allows bursts of traffic for short periods, while the two-rate policer allows more sustained bursts of traffic. Single-rate policing is implemented using a single token-bucket model, so that periods of relatively low traffic must occur between traffic bursts to allow the token bucket to refill. Two-rate policing is implemented using a dual token-bucket model, which allows bursts of traffic for longer periods.

Single-Rate Three-Color Policers

The single-rate three-color type of policer is defined in RFC 2697, *A Single Rate Three Color Marker*. You use this type of policer to rate-limit a traffic flow to a single rate and three traffic categories (green, yellow, and red). A single-rate three-color policer defines a *committed* bandwidth limit and burst-size limit plus an excess burst-size limit. Traffic that conforms to the committed traffic limits is categorized as green (conforming). Traffic that conforms to the bandwidth limit while allowing bursts of traffic as controlled by the excess burst-size limit is categorized as yellow. All other traffic is categorized as red.

A single-rate three-color policer is most useful when a service is structured according to packet length, not peak arrival rate.

Two-Rate Three-Color Policers

The two-rate three-color type of policer is defined in RFC 2698, *A Two Rate Three Color Marker*. You use this type of policer to rate-limit a traffic flow to two rates and three traffic categories (green, yellow, and red). A two-rate three-color policer defines a *committed* bandwidth limit and burst-size limit plus a *peak* bandwidth limit and burst-size limit. Traffic that conforms to the committed traffic limits is categorized as green (conforming). Traffic that exceeds the committed traffic limits but remains below the peak traffic limits is categorized as yellow. Traffic that exceeds the peak traffic limits is categorized as red.

A two-rate three-color policer is most useful when a service is structured according to arrival rates and not necessarily packet length.

Two-Color and Three-Color Policer Options

Both two-color and three-color policers can be configured with the following options:

- [Logical Interface \(Aggregate\) Policers on page 4809](#)
- [Physical Interface Policers on page 4809](#)
- [Policers Applied to Layer 2 Traffic on page 4809](#)
- [Multifield Classification on page 4809](#)

Logical Interface (Aggregate) Policers

A logical interface policer can be a two-color policer, not a three-color policer. When you apply a logical interface policer to multiple protocol families on the same logical interface, multiple instances of the policer are created, meaning that traffic for each protocol family is policed separately. You apply a logical interface policer directly to a logical interface configuration (and not by referencing the policer in a stateless firewall filter and then applying the filter to the logical interface).

- You can apply the policer at the interface logical unit level to rate-limit all traffic types, regardless of the protocol family.

When applied in this manner, the logical interface policer will be used by all traffic types (inet, inet6, etc.) and across all layers (layer 2, layer 3) no matter where the policer is attached on the logical interface.

- You can also apply the policer at the logical interface protocol family level, to rate-limit traffic for a specific protocol family.

You can apply a logical interface policer to unicast traffic only. For information about configuring a stateless firewall filter for flooded traffic, see “*Applying Filters to Forwarding Tables*” in the “Traffic Sampling, Forwarding, and Monitoring” section of the *Routing Policy Configuration Guide*.

Physical Interface Policers

A physical interface policer can be a two-color or three-color policer. When you apply physical interface policer, to different protocol families on the same logical interface, the protocol families share the same policer instance. This means that rate limiting is performed aggregately for the protocol families for which the policer is applied. This feature enables you to use a single policer instance to perform aggregate policing for different protocol families on the same physical interface. If you want a policer instance to be associated with a protocol family, the corresponding physical interface filter needs to be applied to that protocol family. The policer is not automatically applied to all protocol families configured on the physical interface.

In contrast, with logical interface policers there are multiple separate policer instances.

Policers Applied to Layer 2 Traffic

In addition to hierarchical policing, you can also apply single-rate two-color policers and three-color policers (both single-rate and two-rate) to Layer 2 input or output traffic. You must configure the two-color or three-color policer as a logical interface policer and reference the policer in the interface configuration at the logical unit level, and not at the protocol level. You cannot apply a two-color or three-color policer to Layer 2 traffic as a stateless firewall filter action.

Multifield Classification

Like behavior aggregate (BA) classification, which is sometimes referred to as class-of-service (CoS) value traffic classification, multifield classification is a method of classifying incoming traffic by associating each packet with a forwarding class, a packet loss priority level, or both. The CoS scheduling configuration assigns packets to output queues based on forwarding class. The CoS random early detection (RED) process uses

the drop probability configuration, output queue fullness percentage, and packet loss priority to drop packets as needed to control congestion at the output stage.

BA classification and multifield classification use different fields of a packet to perform traffic classification. BA classification is based on a *CoS value* in the IP packet header. Multifield classification can be based on *multiple fields* in the IP packet header, including CoS values. Multifield classification is used instead of BA classification when you need to classify packets based on information in the packet other than the CoS values only. Multifield classification is configured using a stateless firewall filter term that matches on any packet header fields and associates matched packets with a forwarding class, a loss priority, or both. The forwarding class or loss priority can be set by a firewall filter action or by a policer referenced as a firewall filter action.

Related Documentation

- [Traffic Policing Overview on page 1440](#)
- [Order of Policer and Firewall Filter Operations on page 4810](#)
- [Two-Color Policer Configuration Overview on page 4813](#)
- [Three-Color Policer Configuration Overview on page 4817](#)
- [Two-Color Policing at Layer 2 Overview on page 4947](#)
- [Three-Color Policing at Layer 2 Overview on page 4949](#)


Order of Policer and Firewall Filter Operations

You can apply both a traffic policer and a stateless firewall filter (with or without policing actions) to a single logical interface at the same time. In this case, the order of precedence of operations is such that policers applied directly to the logical interface are evaluated before input filters but after output filters.

- If an input firewall filter is configured on the same logical interface as a policer, the policer is executed first.
- If an output firewall filter is configured on the same logical interface as a policer, the firewall filter is executed first.

[Figure 101 on page 4810](#) illustrates the order of policer and firewall filter processing at the same interface.

Figure 101: Incoming and Outgoing Policers and Firewall Filters



Related Documentation

- [Two-Color Policer Configuration Overview on page 4813](#)
- [Three-Color Policer Configuration Overview on page 4817](#)
- [Hierarchical Policer Configuration Overview](#)

Understanding the Frame Length for Policing Packets

[Table 362 on page 4811](#) describes the packet lengths that are considered when you use a traffic policer.

Table 362: Packet Lengths Considered for Traffic Policers

Protocol	Policing Packet Lengths
Any	L3 frame including header
IPv4	L3 frame including header
IPv6	L3 frame including header
MPLS	L3 frame including header
VPLS	L2 frame including header + FCS
ethernet-switching	L2 frame including header + FCS
CCC	L2 frame including header + FCS

Related Documentation

- [Policer Overhead to Account for Rate Shaping in the Traffic Manager on page 4904](#)

Introduction to Configuring Policers

- [Statement Hierarchy for Configuring Policers on page 4811](#)
- [Two-Color Policer Configuration Overview on page 4813](#)
- [Three-Color Policer Configuration Overview on page 4817](#)
- [Guidelines for Applying Traffic Policers on page 4821](#)

Statement Hierarchy for Configuring Policers

```

firewall {
  family (any | ccc | ethernet-switching | inet | inet6 | mpls | vpls) {
    filter filter-name {
      ... protocol-family-specific-firewall-filter-configuration ...
      prefix-action name {
        count;
        destination-prefix-length prefix-length;
        policer policer-name;
        source-prefix-length prefix-length;
        subnet-prefix-length prefix-length;
      }
    }
  }
  filter filter-name {
    accounting-profile [ profile-names ];
    interface-specific;
    physical-interface-filter;
    term term-name {

```

```
filter filter-name;
from {
    ... ipv4-firewall-filter-match-conditions ...
}
then {
    ... ipv4-firewall-filter-terminating-actions ...
    ... ipv4-firewall-filter-nonterminating-actions ...
    next term;
}
}
}
hierarchical-policer policer-name {
    aggregate {
        if-exceeding {
            bandwidth-limit-limit bps;
            burst-size-limit bytes;
        }
        then {
            discard;
            forwarding-class class-name;
            loss-priority (high | low | medium-high | medium-low);
        }
    }
    premium {
        if-exceeding {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
        then {
            discard;
        }
    }
}
}
interface-set interface-set-name {
    interface-name;
}
}
load-balance-group group-name {
    next-hop-group [ group-names ];
}
}
policer policer-name {
    filter-specific;
    if-exceeding {
        (bandwidth-limit bps | bandwidth-percent percentage);
        burst-size-limit bytes;
    }
    logical-bandwidth-policer;
    logical-interface-policer;
    physical-interface-policer;
    then {
        discard;
        forwarding-class class-name;
        loss-priority (high | low | medium-high | medium-low);
    }
}
}
three-color-policer policer-name {
    action {
```

```
        loss-priority high then discard;
    }
    logical-interface-policer;
    physical-interface-policer;
    single-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        excess-burst-size bytes;
    }
    two-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        peak-burst-size bytes;
        peak-information-rate bps;
    }
}
}
```

Related Documentation

- [Two-Color Policer Configuration Overview on page 4813](#)
- [Three-Color Policer Configuration Overview on page 4817](#)
- [Hierarchical Policer Configuration Overview](#)
- [Guidelines for Applying Traffic Policers on page 4821](#)

Two-Color Policer Configuration Overview

Table 363 on page 4813 describes the hierarchy levels at which you can configure and apply single-rate two-color policers to Layer 3 traffic. For information about applying single-rate two-color policers to Layer 2 traffic, see “Two-Color Policing at Layer 2 Overview” on page 4947.

Table 363: Two-Color Policer Configuration and Application Overview

Policy Configuration	Layer 3 Application	Key Points
Single-Rate Two-Color Policer <i>Defines traffic rate limiting that you can apply to Layer 3 protocol-specific traffic at a logical interface. Can be applied as an interface policer or as a firewall filter policer.</i>		
Basic policer configuration: [edit firewall] policer <i>policer-name</i> { if-exceeding { bandwidth-limit <i>bps</i> ; burst-size-limit <i>bytes</i> ; } then { discard; forwarding-class <i>class-name</i> ; loss-priority <i>supported-value</i> ; } }	Method A—Apply as an interface policer at the protocol family level: [edit interfaces] <i>interface-name</i> { unit <i>unit-number</i> { family <i>family-name</i> { policer { input <i>policer-name</i> ; output <i>policer-name</i> ; } } } }	Policy configuration: <ul style="list-style-type: none">• Use bandwidth-limit <i>bps</i> to specify an absolute value. Firewall filter configuration (*) <ul style="list-style-type: none">• If applying to multiple interfaces, include the interface-specific statement to create unique policers and counters for each interface.

Table 363: Two-Color Policer Configuration and Application Overview (*continued*)

Policer Configuration	Layer 3 Application	Key Points
<pre> }</pre>	<pre> } Method B—Apply as a firewall filter policer at the protocol family level: [edit firewall] family <i>family-name</i> { filter <i>filter-name</i> { <i>interface-specific</i>; # (*) from { ... <i>match-conditions</i> ... } then { policer <i>policer-name</i>; } } } [edit interfaces] <i>interface-name</i> { unit <i>unit-number</i> { family <i>family-name</i> { filter { input <i>filter-name</i>; output <i>filter-name</i>; } ... <i>protocol-configuration</i> ... } } } }</pre>	<p>Interface policer verification:</p> <ul style="list-style-type: none"> • Use the show interfaces (detail extensive) operational mode command. • Use the show policer operational mode command. <p>Firewall filter policer verification:</p> <ul style="list-style-type: none"> • Use the show interfaces (detail extensive) operational mode command. • Use the show firewall filter <i>filter-name</i> operational mode command.

Table 363: Two-Color Policer Configuration and Application Overview (*continued*)

Policy Configuration	Layer 3 Application	Key Points
Bandwidth Policer <i>Defines traffic rate limiting that you can apply to Layer 3 protocol-specific traffic at a logical interface, but the bandwidth limit is specified as a percentage value. Bandwidth can be based on physical interface line rate (the default) or the logical interface shaping rate. Can be applied as an interface policer or as a firewall filter policer where the filter is either interface-specific or a physical interface filter.</i>		
Bandwidth policer configuration: <pre>[edit firewall] policer <i>policer-name</i> { logical-bandwidth-policer; if-exceeding { bandwidth-percent (1..100); burst-size-limit <i>bytes</i>; } then { discard; forwarding-class <i>class-name</i>; loss-priority <i>supported-value</i>; } }</pre>	Method A—Apply as an interface policer at the protocol family level: <pre>[edit interfaces] interface-name { unit <i>unit-number</i> { family <i>family-name</i> { policer { input <i>policer-name</i>; output <i>policer-name</i>; } } } }</pre> Method B—Apply as a firewall filter policer at the protocol family level: <pre>[edit firewall] family <i>family-name</i> { filter <i>filter-name</i> { interface-specific; from { ... <i>match-conditions</i> ... } then { policer <i>policer-name</i>; } } }</pre> <pre>[edit interfaces] interface-name { unit <i>unit-number</i> { family <i>family-name</i> { filter { input <i>filter-name</i>; output <i>filter-name</i>; } ... <i>protocol-configuration</i> ... } } }</pre>	Policer configuration: <ul style="list-style-type: none"> Use the bandwidth-percent <i>percentage</i> statement instead of the bandwidth-limit <i>bps</i> statement. By default, bandwidth policing rate-limits traffic based on a percentage of the physical interface media rate. To rate-limit traffic based on a percentage of the logical interface configured shaping rate, also include the logical-bandwidth-policer statement. Firewall filter configuration: <ul style="list-style-type: none"> Percentage bandwidth policers can only be referenced by filters configured with the interface-specific statement. Interface policer verification: <ul style="list-style-type: none"> Use the show interfaces (<i>detail</i> <i>extensive</i>) operational mode command. Use the show policer operational mode command. Firewall filter policer verification: <ul style="list-style-type: none"> Use the show interfaces (<i>detail</i> <i>extensive</i>) operational mode command. Use the show firewall filter <i>filter-name</i> operational mode command.

Table 363: Two-Color Policer Configuration and Application Overview (*continued*)

Policy Configuration	Layer 3 Application	Key Points
Logical Interface (Aggregate) Policer <i>Defines traffic rate limiting that you can apply to multiple protocol families on the same logical interface without creating multiple instances of the policer. Can be applied directly to a logical interface configuration only.</i>		
Logical interface policer configuration: <pre>[edit firewall] policer <i>policer-name</i> { <i>logical-interface-policer</i>; if-exceeding { bandwidth-limit <i>bps</i>; burst-size-limit <i>bytes</i>; } then { discard; forwarding-class <i>class-name</i>; loss-priority <i>supported-value</i>; } }</pre>	Apply as an interface policer only: <pre>[edit interfaces] interface-name { unit <i>unit-number</i> { policer { # All protocols input <i>policer-name</i>; output <i>policer-name</i>; } } family <i>family-name</i> { policer { # One protocol input <i>policer-name</i>; output <i>policer-name</i>; } } }</pre>	Policer configuration: <ul style="list-style-type: none"> • Include the logical-interface-policer statement. Two options for interface policer application: <ul style="list-style-type: none"> • To rate-limit all traffic types, regardless of the protocol family, apply the logical interface policer at the logical unit level. • To rate-limit traffic of a specific protocol family, apply the logical interface policer at the protocol family level. Interface policer verification: <ul style="list-style-type: none"> • Use the show interfaces (detail extensive) operational mode command. • Use the show policer operational mode command.

Table 363: Two-Color Policer Configuration and Application Overview (*continued*)

Policy Configuration	Layer 3 Application	Key Points
Physical Interface Policer Defines traffic rate limiting that applies to all logical interfaces and protocol families configured on a physical interface, even if the interfaces belong to different routing instances. Can be applied as a firewall filter policer referenced from a physical interface filter only.		
Physical interface policer configuration: [edit firewall] policer <i>policer-name</i> { physical-interface-policer; if-exceeding { bandwidth-limit <i>bps</i> ; burst-size-limit <i>bytes</i> ; } then { discard; forwarding-class <i>class-name</i> ; loss-priority <i>supported-value</i> ; } }	Apply as a firewall filter policer referenced from a physical interface filter that you apply at the protocol family level: [edit firewall] family <i>family-name</i> { filter <i>filter-name</i> { physical-interface-filter; from { ... <i>match-conditions</i> ... } then { policer <i>policer-name</i> ; } } } [edit interfaces] interface-name { unit <i>number</i> { family <i>family-name</i> { filter { input <i>filter-name</i> ; output <i>filter-name</i> ; } ... <i>protocol-configuration</i> ... } } }	Policer configuration: <ul style="list-style-type: none">• Include the physical-interface-policer statement. Firewall filter configuration: <ul style="list-style-type: none">• Include the physical-interface-filter statement. Application: <ul style="list-style-type: none">• Apply the filter to the input or output of a logical interface at the protocol family level. Firewall filter policer verification: <ul style="list-style-type: none">• Use the show interfaces (detail extensive) operational mode command.• Use the show firewall filter filter-name operational mode command.

- Related Documentation
- [Basic Single-Rate Two-Color Policers on page 4838](#)
 - [Bandwidth Policers on page 4854](#)
 - [Filter-Specific Counters and Policers on page 4863](#)
 - [Prefix-Specific Counting and Policing Actions on page 4874](#)
 - [Multifield Classification on page 4889](#)
 - [Policer Overhead to Account for Rate Shaping in the Traffic Manager on page 4904](#)
 - [Two-Color and Three-Color Physical Interface Policers on page 4940](#)

Three-Color Policer Configuration Overview

Table 364 on page 4818 describes the hierarchy levels at which you can configure and apply single-rate tricolor-marking (single-rate TCM) policers and two-rate tricolor-marking (two-rate TCM) policers to Layer 3 traffic. For information about applying three-color policers to Layer 2 traffic, see “Three-Color Policing at Layer 2 Overview” on page 4949.

Table 364: Three-Color Policer Configuration and Application Overview

Policer Configuration	Layer 3 Application	Key Points
Single-Rate Three-Color Policer Defines traffic rate limiting that you can apply to Layer 3 protocol-specific traffic at a logical interface. Can be applied as a firewall filter policer only. Provides moderate allowances for short periods of traffic that exceed the committed burst size.		
Basic single-rate TCM policer configuration: <pre>[edit firewall] three-color-policer <i>policer-name</i> { single-rate { (color-aware color-blind); committed-information-rate <i>bps</i>; committed-burst-size <i>bytes</i>; excess-burst-size <i>bytes</i>; } action { loss-priority high then discard; } }</pre>	Reference the policer from a firewall filter, and apply the filter to a protocol family on a logical interface: <pre>[edit firewall] family <i>family-name</i> { filter <i>filter-name</i> { term <i>term-name</i> { from { ... <i>match-conditions</i> ... } then { three-color-policer { single-rate <i>policer-name</i>; } } } } }</pre> Apply the filter to a logical interface at the protocol family level: <pre>[edit interfaces] interface-name { unit <i>unit-number</i> { family <i>family-name</i> { filter { input <i>filter-name</i>; output <i>filter-name</i>; } } } }</pre>	Policer configuration: <ul style="list-style-type: none"> Include the single-rate (color-aware color-blind) statement. Firewall filter configuration: <ul style="list-style-type: none"> Include the three-color-policer single-rate <i>policer-name</i> action. Applying the firewall filter to the logical interface: <ul style="list-style-type: none"> Include the filter (input output) <i>filter-name</i> statement.

Table 364: Three-Color Policer Configuration and Application Overview (*continued*)

Policy Configuration	Layer 3 Application	Key Points
Single-Rate Three-Color Physical Interface Policer Defines traffic rate limiting that applies to all logical interfaces and protocol families configured on a physical interface, even if the interfaces belong to different routing instances. Can be applied as a firewall filter policer only.		
Physical interface single-rate TCM policer: <pre>[edit firewall] three-color-policer <i>policer-name</i> { physical-interface-policer; single-rate { (color-aware color-blind); committed-information-rate <i>bps</i>; committed-burst-size <i>bytes</i>; excess-burst-size <i>bytes</i>; } action { loss-priority high then discard; } }</pre>	Reference the policer from a physical interface filter only, and apply the filter to a protocol family on a logical interface: <pre>[edit firewall] family <i>family-name</i> { filter <i>filter-name</i> { physical-interface-filter term <i>term-name</i> { from { ... <i>match-conditions</i> ... } then { three-color-policer { single-rate <i>policer-name</i>; } } } } }</pre> <pre>[edit interfaces] interface-name { unit <i>number</i> { family <i>family-name</i> { filter { input <i>filter-name</i>; output <i>filter-name</i>; } } } }</pre>	Policer configuration: <ul style="list-style-type: none"> • Include the physical-interface-policer statement. Firewall filter configuration: <ul style="list-style-type: none"> • Include the physical-interface-filter statement. Application: <ul style="list-style-type: none"> • Include the filter (input output) <i>filter-name</i> statement. Verification <ul style="list-style-type: none"> • To verify, use the show firewall filter <i>filter-name</i> operational mode command.

Table 364: Three-Color Policer Configuration and Application Overview (*continued*)

Policy Configuration	Layer 3 Application	Key Points
Basic Two-Rate Three-Color Policer Defines traffic rate limiting that you can apply to Layer 3 protocol-specific traffic at a logical interface. Can be applied as a firewall filter policer only. Provides moderate allowances for sustained periods of traffic that exceed the committed bandwidth limit or burst size.		
Basic two-rate TCM policer configuration: <pre>[edit firewall] three-color-policer <i>policer-name</i> { two-rate { (color-aware color-blind); committed-information-rate <i>bps</i>; committed-burst-size <i>bytes</i>; peak-information-rate <i>bps</i>; peak-burst-size <i>bytes</i>; } action { loss-priority high then discard; } }</pre>	Reference the policer from a firewall filter, and apply the filter to a protocol family on a logical interface: <pre>[edit firewall] family <i>family-name</i> { filter <i>filter-name</i> { term <i>term-name</i> { from { ... <i>match-conditions</i> ... } then { three-color-policer { two-rate <i>policer-name</i>; } } } } } [edit interfaces] interface <i>interface-name</i> { unit <i>unit-number</i> { family <i>family-name</i> { filter { input <i>filter-name</i>; output <i>filter-name</i>; } } } }</pre>	Policer configuration: <ul style="list-style-type: none"> Include the two-rate (color-aware color-blind) statement. Firewall filter configuration: <ul style="list-style-type: none"> Include the three-color-policer two-rate <i>policer-name</i> action. Applying the firewall filter to the logical interface: <ul style="list-style-type: none"> Include the filter (input output) <i>filter-name</i> statement.

Related Documentation

- [Three-Color Policer Configuration Guidelines on page 4912](#)
- [Basic Single-Rate Three-Color Policers on page 4915](#)
- [Basic Two-Rate Three-Color Policers on page 4921](#)
- [Two-Color and Three-Color Logical Interface Policers on page 4927](#)
- [Two-Color and Three-Color Physical Interface Policers on page 4940](#)

Guidelines for Applying Traffic Policers

The following general guidelines pertain to applying traffic policers:

- Only one type of policer can be applied to the input or output of the same physical or logical interface. For example, you are not allowed to apply a policer and a hierarchical policer in the same direction at the same logical interface.
- Chaining of policers—that is, applying policers to both a port and the logical interfaces of that port—is not allowed.
- A maximum of 64 policers is supported per physical or logical interface, provided no behavior aggregate (BA) classification—traffic classification based on CoS values in the packet headers—is applied to the logical interface.
- The policer should be independent of BA classification. Without BA classification, all traffic on an interface is treated either as expedited forwarding (EF) or non-EF, based on the configuration. With BA classification, a physical or logical interface can support up to 64 policers. The interface might be a physical interface or logical interface.
- With BA classification, the miscellaneous traffic (the traffic *not* matching any of the BA classification DSCP/EXP bits) is policed as non-EF traffic. No separate policers are installed for this traffic.
- Policers can be applied to unicast packets only. For information about configuring a filter for flooded traffic, see “*Applying Filters to Forwarding Tables*” in the “Traffic Sampling, Forwarding, and Monitoring” section of the *Routing Policy Configuration Guide*.

Related Documentation

- [Statement Hierarchy for Configuring Policers on page 4811](#)
- [Two-Color Policer Configuration Overview on page 4813](#)
- [Three-Color Policer Configuration Overview on page 4817](#)
- [Hierarchical Policer Configuration Overview](#)

Policer Rate Limits and Actions

- [Policer Bandwidth and Burst-Size Limits on page 4821](#)
- [Policer Color-Marking and Actions on page 4822](#)
- [Single Token Bucket Algorithm on page 4824](#)
- [Dual Token Bucket Algorithms on page 4826](#)

Policer Bandwidth and Burst-Size Limits

[Table 365 on page 4822](#) lists each of the Junos OS policer types supported. For each policer type, the table summarizes the bandwidth limits and burst-size limits used to rate-limit traffic.

Table 365: Policer Bandwidth Limits and Burst-Size Limits

Policer Type	Bandwidth Limits	Burst-Size Limits
Single-Rate Two-Color Policer		
Defines a single rate limit: a bandwidth limit and an allowed burst size for conforming traffic.	bandwidth-limit <i>bps</i>; M, MX, T Series routers and EX Series switches: 8000..500000000000	burst-size-limit <i>bytes</i>; M, MX, T Series routers, and EX Series switches: 1500..100000000000
For a single-rate two-color policer only, you can specify the bandwidth limit as a percentage value from 1 through 100 instead of as an absolute number of bits per second. The effective bandwidth limit is calculated as a percentage of either the physical interface media rate or the logical interface configured shaping rate.	bandwidth-percent 1..100 percent	
Single-Rate Three-Color Policer		
Defines a single rate limit: a bandwidth limit and an allowed burst size for conforming traffic.	committed-information-rate <i>bps</i>; M, MX, T Series routers, and EX Series switches: 1500..100000000000	committed-burst-size <i>bytes</i>; M, MX, T Series routers, and EX Series switches: 1500..100000000000
Also defines a second, larger burst size. This second burst size is used to differentiate between two categories of nonconforming traffic (yellow or red).		excess-burst-size <i>bytes</i>; M, MX, T Series routers, and EX Series switches: 1500..100000000000
Two-Rate Three-Color Policer		
Defines a committed rate limit: a bandwidth limit and an allowed burst size for conforming traffic.	committed-information-rate <i>bps</i>; M, MX, T Series routers, and EX Series switches: 1500..100000000000	committed-burst-size <i>bytes</i>; M, MX, T Series routers, and EX Series switches: 1500..100000000000
Also defines a peak rate limit: a second, larger burst size and a second, higher bandwidth limit. These additional rate-limit components are used to differentiate between two categories of nonconforming traffic (yellow or red).	peak-information-rate <i>bps</i>; M, MX, T Series routers, and EX Series switches: 1500..100000000000	peak-burst-size <i>bytes</i>; M, MX, T Series routers, and EX Series switches: 1500..100000000000

- Related Documentation**
- [Policer Color-Marking and Actions on page 4822](#)
 - [Determining Proper Burst Size for Traffic Policers on page 4834](#)

Policer Color-Marking and Actions

[Table 366 on page 4823](#) lists each of the Junos OS policer types supported. For each policer type, the table summarizes the color-marking criteria used to categorize a traffic flow and, for each color, the actions taken on packets in that type of traffic flow.

Table 366: Implicit and Configurable Policer Actions Based on Color Marking

Policer Rate Limits and Color Marking	Implicit Action	Configurable Actions
Single-Rate Two-Color Policer <ul style="list-style-type: none"> Bandwidth limit Burst size 		
Green Conforms to rate and burst size limits	Set PLP to low	–
Red Exceeds rate and burst size limits	–	<ul style="list-style-type: none"> Discard the packet. Assign to a forwarding class. Set PLP to low or high. On some platforms, you can also set the PLP to medium-low or medium-high.
Single-Rate Three-Color Policer <ul style="list-style-type: none"> Committed information rate (CIR) Committed burst size (CBS) Excess burst size (EBS) 		
Green Conforms to the CIR and CBS	Set PLP to low	–
Yellow Exceeds the CIR and CBS but conforms to the EBS	Set PLP to medium-high	–
Red Exceeds the EBS	Set PLP to high	<ul style="list-style-type: none"> Discard the packet.
Two-Rate Three-Color Policer <ul style="list-style-type: none"> Committed information rate (CIR) Committed burst size (CBS) Peak information rate (PIR) Peak burst size (PBS) 		
Green Conforms to the CIR and CBS	Set PLP to low	–
Yellow Exceeds the CIR and CBS, but conforms to the PIR	Set PLP to medium-high	–
Red Exceeds the PIR and PBS	Set PLP to high	<ul style="list-style-type: none"> Discard the packet.
Hierarchical Policer		

Table 366: Implicit and Configurable Policer Actions Based on Color Marking (*continued*)

Policer Rate Limits and Color Marking	Implicit Action	Configurable Actions
Aggregate policer <ul style="list-style-type: none"> Bandwidth limit Burst size 		
Green Conforms to rate limits	Set PLP to low	–
Red Exceeds rate limits	–	<ul style="list-style-type: none"> Discard the packet. Assign to a forwarding class. Set PLP to low or high. On some platforms, you can also set the PLP to medium-low or medium-high.
Premium policer <ul style="list-style-type: none"> Bandwidth limit Burst size 		
Green Conforms to rate limits	Set PLP to low	–
Red Exceeds rate limits	–	<ul style="list-style-type: none"> Discard the packet.

- Related Documentation**
- [Policer Bandwidth and Burst-Size Limits](#)
 - [Determining Proper Burst Size for Traffic Policers on page 4834](#)

Single Token Bucket Algorithm

This topic covers the following information:

- [Token Bucket Concepts on page 4825](#)
- [Single Token Bucket Algorithm on page 4825](#)
- [Conformance Measurement for Two-Color Marking on page 4826](#)

Token Bucket Concepts

When you apply traffic policing to the input or output traffic at an interface, the rate limits and actions specified in the policer configuration are used to enforce a limit on the average throughput rate at the interface while also allowing bursts of traffic up to a maximum number of bytes based on the overall traffic load. Junos OS policers measure traffic-flow conformance to a policing rate limit by using a *token bucket algorithm*:

- The *bucket* represents a rate-limiting function of the policer on the interface input or output traffic.
- Each *token* in the bucket represents a “credit” for some number of *bits*, and tokens in the bucket are “cashed in” for the ability to receive or transmit traffic that conforms to a rate limit configured for the policer.
- The *token arrival rate* is the fixed *bits-per-second* rate at which tokens are added to the token bucket, but only up to the specified depth of the bucket.
- The *token bucket depth* defines the capacity of the bucket in *bytes*.

An algorithm based on a single token bucket allows burst of traffic for short periods, whereas an algorithm based dual token buckets allows more sustained bursts of traffic.

Single Token Bucket Algorithm

A single-rate two-color policer limits traffic throughput at an interface based on how the traffic conforms to rate-limit values specified in the policer configuration. Similarly, a hierarchical policer limits traffic throughput at an interface based on how aggregate and premium traffic subflows conform to aggregate and premium rate-limit values specified in the policer configuration. For both two-color policer types, packets in a conforming traffic flow are categorized as *green*, and packets in a non-conforming traffic flow are categorized as *red*.

The single token bucket algorithm measures traffic-flow conformance to a two-color policer rate limit as follows:

- The token arrival rate represents the single *bandwidth limit* configured for the policer. You can specify the bandwidth limit as an absolute number of bits per second by including the **bandwidth-limit *bps*** statement. Alternatively, for single-rate two-color policers only, you can use the **bandwidth-percent *percentage*** statement to specify the bandwidth limit as a percentage of either the physical interface port speed or the configured logical interface shaping rate.
- The token bucket depth represents the single *burst size* configured for the policer. You specify the burst size by including the **burst-size-limit *bytes*** statement.
- If the bucket is filled to capacity, arriving tokens “overflow” the bucket and are lost.

When the bucket contains insufficient tokens for receiving or transmitting the traffic at the interface, packets might be dropped or else re-marked with a lower forwarding class, a higher packet loss priority (PLP) level, or both.

Conformance Measurement for Two-Color Marking

In two-color-marking policing, a traffic flow whose average arrival or departure rate does not exceed the token arrival rate (bandwidth limit) is considered *conforming traffic*. Packets in a conforming traffic flow (categorized as green traffic) are implicitly marked with a packet loss priority (PLP) level of **low** and then passed through the interface.

For a traffic flow whose average arrival or departure rate exceeds the token arrival rate, conformance to a two-color policer rate limit depends on the tokens in the bucket. If sufficient tokens remain in the bucket, the flow is considered conforming traffic. If the bucket does not contain sufficient tokens, the flow is considered *non-conforming traffic*. Packets in a non-conforming traffic flow (categorized as red traffic) are handled according to policing actions. Depending on the configuration of the two-color policer, packets might be implicitly discarded; or the packets might be re-marked with a specified forwarding class, a specified PLP, or both, and then passed through the interface.



NOTE: The number of tokens remaining in the bucket at any given time is a function of the token bucket depth and the overall traffic load.

The token bucket is initially filled to capacity, and so the policer allows an initial traffic burst (back-to-back traffic at average rates that exceed the token arrival rate) up to the size of the token bucket depth.

During periods of relatively low traffic (traffic that arrives at or departs from the interface at average rates below the token arrival rate), unused tokens accumulate in the bucket, but only up to the configured token bucket depth.

Related Documentation

- [Two-Color Policer Configuration Overview on page 4813](#)
- [Hierarchical Policer Configuration Overview](#)
- [Policer Color-Marking and Actions on page 4822](#)
- [bandwidth-limit \(Hierarchical Policer\)](#)
- [bandwidth-limit \(Policer\) on page 4971](#)
- [bandwidth-percent on page 4973](#)
- [burst-size-limit \(Hierarchical Policer\)](#)
- [burst-size-limit \(Policer\) on page 4975](#)

Dual Token Bucket Algorithms

This topic covers the following information:

- [Token Bucket Concepts on page 4827](#)
- [Guaranteed Bandwidth for Three-Color Marking on page 4827](#)
- [Nonconformance Measurement for Single-Rate Three-Color Marking on page 4827](#)
- [Nonconformance Measurement for Two-Rate Three-Color Marking on page 4828](#)

Token Bucket Concepts

When you apply traffic policing to the input or output traffic at an interface, the rate limits and actions specified in the policer configuration are used to enforce a limit on the average throughput rate at the interface while also allowing bursts of traffic up to a maximum number of bytes based on the overall traffic load. Junos OS policers measure traffic-flow conformance to a policing rate limit by using a *token bucket algorithm*:

- The *bucket* represents a rate-limiting function of the policer on the interface input or output traffic.
- Each *token* in the bucket represents a “credit” for some number of *bits*, and tokens in the bucket are “cashed in” for the ability to receive or transmit traffic that conforms to a rate limit configured for the policer.
- The *token arrival rate* is the fixed *bits-per-second* rate at which tokens are added to the token bucket, but only up to the specified depth of the bucket.
- The *token bucket depth* defines the capacity of the bucket in *bytes*.

An algorithm based on a single token bucket allows burst of traffic for short periods, whereas an algorithm based dual token buckets allows more sustained bursts of traffic.

Guaranteed Bandwidth for Three-Color Marking

A committed information rate (CIR) defines the guaranteed bandwidth for traffic arriving at or departing from the interface under normal line conditions. A flow of traffic at an average rate that conforms to the CIR is categorized green, and packets in a green flow are implicitly marked with **low** packet loss priority (PLP) and then passed through the interface. During periods of relatively low traffic (traffic that arrives at or departs from the interface at average rates below the CIR), any unused bandwidth capacity accumulates in the first token bucket, but only up to a configured number of bytes. If any unused bandwidth capacity overflows the first bucket, the excess accumulates in a second token bucket.

The committed burst size (CBS) defines the maximum number of bytes for which unused amounts of the guaranteed bandwidth can be accumulated in the first token bucket. A burst of traffic at an average rate that exceeds the CIR is also categorized as green provided that sufficient unused bandwidth capacity is available in the first token bucket.

Nonconformance Measurement for Single-Rate Three-Color Marking

Single-rate three-color policer configurations specify a second burst size—the excess burst size (EBS)—that defines the maximum number of bytes for which the second token bucket can accumulate unused bandwidth that overflows from the first bucket.

A traffic flow is categorized yellow if its average rate exceeds the CIR and the available bandwidth capacity accumulated in the first bucket if sufficient unused bandwidth capacity is available in the second token bucket. Packets in a yellow flow are implicitly marked with **medium-high** PLP and then passed through the interface.

A traffic flow is categorized red its average rate exceeds the CIR and the available bandwidth capacity accumulated in the second bucket. Packets in a red flow are implicitly marked with **high** PLP and then either passed through the interface or optionally discarded.

Nonconformance Measurement for Two-Rate Three-Color Marking

Two-rate three-color policer configurations include a second rate limit—the peak-information-rate (PIR)—that you set to the expected average data rate for traffic arriving at or departing from the interface under peak conditions.

Two-rate three-color policer configurations also include a second burst size—the peak burst size (PBS)—that defines the maximum number of bytes for which the second token bucket can accumulate unused peak bandwidth capacity. During periods of relatively little peak traffic (traffic that arrives at or departs from the interface at average rates that exceed the PIR), any unused peak bandwidth capacity accumulates in the second token bucket, but only up to the maximum number of bytes specified by the PBS.

A traffic flow is categorized yellow if it exceeds the CIR and the available committed bandwidth capacity accumulated in the first token bucket but conforms to the PIR. Packets in a yellow flow are implicitly marked with **medium-high** PLP and then passed through the interface.

A traffic flow is categorized red if it exceeds the PIR and the available peak bandwidth capacity accumulated in the second token bucket. Packets in a red flow are implicitly marked with **high** PLP and then either passed through the interface or optionally discarded.

Related Documentation

- [Three-Color Policer Configuration Overview on page 4817](#)
- [Policer Color-Marking and Actions on page 4822](#)
- [committed-burst-size on page 4980](#)
- [committed-information-rate on page 4982](#)
- [excess-burst-size on page 4984](#)
- [peak-burst-size on page 4999](#)
- [peak-information-rate on page 5001](#)

Policer Implementation

- [Policer Implementation Overview on page 4828](#)
- [Understanding the Benefits of Policers and Token Bucket Algorithms on page 4833](#)
- [Determining Proper Burst Size for Traffic Policers on page 4834](#)

Policer Implementation Overview

Traffic policing enables you to control the maximum rate of traffic sent or received on an interface and also to partition network traffic into multiple priority levels, also known as classes of service. A policer defines a set of traffic rate limits and sets consequences for traffic that does not conform to the configured limits. Packets in a traffic flow that do not conform to traffic limits are either discarded or marked with a different forwarding class or packet loss priority (PLP) level.

You can apply a policer to inbound or outbound traffic. Policers applied to inbound traffic help to conserve resources by dropping traffic that does not need to be routed through a network. Policers applied to outbound traffic control the bandwidth used.

The Juniper Networks® Junos® operating system (Junos OS) supports three types of policers:

- *Single-rate two-color policer* — The most common policer. Single-rate means that there is only a single bandwidth and burst rate referenced in the policer. The two colors associated with this policer are red (nonconforming) and green (conforming).
- *Single-rate three-color policer* — Similar to the single-rate two-color policer with the addition of the color yellow. This type also introduces the *committed information rate* (CIR) and a *committed burst rate* (CBR).
- *Two-rate three-color policer* — Builds off of the single-rate three-color policer by adding a second rate tier. *Two-rate* means there is an upper bandwidth limit and associated burst size as well as a *peak information rate* (PIR) and a *peak burst rate* (PBS).



NOTE: The remainder of this topic covers the single-rate two-color policer. For more information about the other types of policers, see the *Junos OS Traffic Policers Configuration Guide*.

Junos OS policers use a *token bucket algorithm* to enforce a limit on an average transmit or receive rate of traffic at an interface while allowing bursts of traffic up to a maximum value based on the configured bandwidth limit and configured burst size. The token bucket algorithm offers more flexibility than a *leaky bucket algorithm* in that you can allow a specified traffic burst before starting to discard packets or apply a penalty such as packet output-queuing priority or packet-drop priority.

There are two types of token bucket algorithms that can be used, depending on the type of policer that is applied to network traffic. Single-rate two-color policers use the *single token bucket algorithm* to measure traffic flow conformance to a two-color policer rate limit. Single-rate three-color policers and two-rate three-color policers both use the *dual token bucket algorithm* to measure traffic flow conformance to a three-color policer rate. The main difference between these two token bucket algorithms is that the single token bucket algorithm allows bursts of traffic for short periods, whereas the dual token bucket algorithm allows more sustained bursts of traffic.



NOTE: The remainder of this topic discusses the single token bucket algorithm. For more information about the dual token bucket algorithm, see the *Junos OS Traffic Policers Configuration Guide*.

Following are the main components of the token bucket algorithm:

- The *bucket* represents a rate-limiting function of the policer on the interface input or output traffic.
- Each *token* in the bucket represents a “credit” for some number of *bits*, and tokens in the bucket are “cashed in” for the ability to receive or transmit traffic that conforms to a rate limit configured for the policer.

- The *token arrival rate* is a periodic allocation of tokens into the token bucket that is calculated from the configured bandwidth limit.
- The *token bucket depth* defines the capacity of the bucket in *bytes*. Tokens that are allocated after the bucket reaches capacity are not able to be stored and used.

In the token bucket model, the bucket represents the policing function. Tokens are added to the bucket at a fixed rate, but once the specified depth of the bucket is reached, tokens allocated after cannot be stored and used. Each token represents a “credit” for some number of bits, and the tokens in the bucket are “cashed in” for the ability to transmit or receive traffic at the interface. When sufficient tokens are present in the bucket, a traffic flow continues unrestricted.

- The rate at which tokens are added to the bucket represents the highest average transmit or receive rate in bits per second allowed for a given service level. You specify this highest average traffic rate as the *bandwidth limit* of the policer. If the traffic arrival rate is so high that at some point insufficient tokens are present in the bucket, then the traffic flow is no longer conforming to the traffic limit. During periods of relatively low traffic (traffic that arrives at or departs from the interface at average rates below the token arrival rate), unused tokens accumulate in the bucket.
- The depth of the bucket in bytes controls the amount of back-to-back bursting allowed. You specify this factor as the *burst-size limit* of the policer. This second limit affects the average transmit or receive rate by limiting the number of bytes permitted in a transmission burst for a given interval of time. Bursts exceeding the current burst-size limit are dropped until there are sufficient tokens available to permit the burst to proceed.

To configure a policer, you need to set two parameters:

- Bandwidth limit configured in bps (using the **bandwidth-limit** statement)
- Burst size configured in bytes (using the **burst-size-limit** statement)



NOTE: For single-rate two-color policers only, you can also specify the bandwidth limit as a percentage of either the physical interface port speed or the configured logical interface shaping rate by using the **bandwidth-percent *percentage*** statement. You cannot configure a policer to use bandwidth percentage for aggregate, tunnel, or software interfaces.

Use the following command to set the policer conditions:

```
user@router# set firewall policer <policer name> if-exceeding ?
Possible completions:
  <[Enter]>          Execute this command
+ apply-groups       Groups from which to inherit configuration data
+ apply-groups-except Don't inherit configuration data from these groups
  bandwidth-limit    Bandwidth limit (8000..1000000000000 bits per second)
  bandwidth-percent  Bandwidth limit in percentage (1..100 percent)
  burst-size-limit    Burst size limit (1500..1000000000000 bytes)
  |                  Pipe through a command
```

The bandwidth limit parameter is used to determine the average rate limit applied to the traffic, while the burst-size parameter is used to allow for short periods of traffic bursting (back-to-back traffic at average rates that exceed the configured bandwidth limit). Once you apply a set of policer configuration settings (bandwidth limit and burst size), the configured values are adjusted to hardware programmable values. The conversion adjustment introduced is normally less than 1 percent of the configured bandwidth limit. This adjustment is needed because the software allows you to configure the bandwidth limit and burst size to any value within the specified ranges, but those values must be adjusted to the nearest value that can be programmed in the hardware.

The policer bandwidth limit configuration in the hardware is represented by two values: the *credit update frequency* and the *credit size*. The credit update frequency is used by the hardware to determine how frequently tokens (bits of unused bandwidth) are added to the token bucket. The credit size is based on the number of tokens that can fit in the token bucket. The MX Series, M120, M320 routers, and EX Series switches contain a set of credit update frequencies instead of having a single credit update frequency to minimize the adjustment difference from the configured bandwidth limit and to support a wide range of policer bandwidth rates (from 40 Kbps to 40 Gbps). One of the frequencies is used to program the policer (bandwidth limit and burst size) in the hardware.

The burst size is based on the overall traffic load and allows bursts of traffic to exceed the configured bandwidth limit. A policer with a large burst size effectively disables the configured bandwidth limit function, so the burst size must be relative to the configured bandwidth limit. You need to consider the traffic patterns in your network before determining the burst size. For more information about determining burst size, see [“Determining Proper Burst Size for Traffic Policers” on page 4834](#).


The configured burst size is adjusted in the hardware to a value that is based on the configured bandwidth limit. The burst size extends the configured bandwidth limit for bursty traffic that exceeds the configured bandwidth limit.

When a policer is applied to the traffic at an interface, the initial capacity for traffic bursting is equal to the number of bytes specified in the **burst-size-limit** statement.

[Figure 102 on page 4832](#) represents how a policer is implemented using the token bucket algorithm. The token allocator allocates tokens to the policer based on the configured bandwidth limit, which is the token size multiplied by the token arrival rate.

token size x token arrival rate = policer rate (configured bandwidth limit)

Figure 102: Token Bucket Algorithm




When a packet arrives at an interface configured with a policer, tokens that represent the number of bits that correspond to the length of the packet are used (or “cached in”) from the token bucket. If the token arrival rate is higher than the rate of traffic so that there are tokens not being used, the token bucket is filled to capacity, and arriving tokens “overflow” the bucket and are lost. The token bucket depth represents the single user-configured burst size for the policer.

If there are tokens in the token bucket and the incoming traffic rate is higher than the token rate (the configured policer rate, bandwidth limit), the traffic can use the tokens until the bucket is empty. The token consumption rate can be as high as the incoming traffic rate, which creates the burst of traffic shown in [Figure 103 on page 4832](#).

By using the token bucket algorithm, the average bandwidth rate being allowed is close to the configured bandwidth limit while simultaneously supporting bursty traffic, as shown in [Figure 103 on page 4832](#).

Figure 103: Traffic Behavior Using Policer and Burst Size





NOTE: The measured length of a packet changes according to the family type that the policer applies to. If the policer is applied under the family inet hierarchy, the policer considers only the IPv4 packet length. If the policer is applied under the family vpls hierarchy, the entire Ethernet frame (including the Ethernet MAC header) is included in the packet length.

The major factor that affects the policer shaping result is not the conversion adjustment, but the traffic pattern since most network traffic is not consistent and is not sent at a constant rate. Due to the fluctuation of the incoming traffic rate, some of the allocated tokens are not used. As a result, the shaped traffic rate is lower than you might expect, and the TCP connection behavior discussed in “[Understanding the Benefits of Policers and Token Bucket Algorithms](#)” on page 4833 is a typical example of this. To alleviate this effect of the lower shaped traffic rate, a proper burst size configuration is required.

Related Documentation

- [Understanding the Benefits of Policers and Token Bucket Algorithms on page 4833](#)
- [Determining Proper Burst Size for Traffic Policers on page 4834](#)

Understanding the Benefits of Policers and Token Bucket Algorithms

This topic describes some scenarios that demonstrate how difficult it is to control traffic that comes into your network without the help of policers and the token bucket algorithm. These scenarios assume that traffic is coming from a TCP-based connection. Depending on the number of TCP connections, policers can have different affects on rate limits.


This topic presents the following scenarios:

- [Scenario 1: Single TCP Connection on page 4833](#)
- [Scenario 2: Multiple TCP Connections on page 4834](#)

Scenario 1: Single TCP Connection

[Figure 104 on page 4833](#) shows the traffic loading on an interface with a policer configured. When the traffic rate reaches the configured bandwidth limit (which results in a packet drop), a TCP slow-start mechanism reduces the traffic rate down to half of what it was. When the traffic rate rises again, the same cycle repeats.

Figure 104: Policer Behavior With a Single TCP Connection




The problem presented in this scenario is that some bandwidth is available, but it is not being used by the traffic. The unused bandwidth shown in [Figure 104 on page 4833](#) is the result of an overall data throughput that is lower than the configured bandwidth value. This example is an extreme case because there is only a single TCP connection.

Scenario 2: Multiple TCP Connections

With multiple TCP connections or some background non-TCP-based traffic, there is less unused bandwidth, as depicted in [Figure 105 on page 4834](#). However, the same issue of unused bandwidth still exists if all the TCP connections experience a drop when the aggregated traffic rate exceeds the configured bandwidth limit.

Figure 105: Policer Behavior With Background Traffic (Multiple TCP Connections)



To reduce the problem of unused bandwidth in your network, you can configure a burst size.

Related Documentation

- [Policer Implementation Overview on page 4828](#)
- [Determining Proper Burst Size for Traffic Policers on page 4834](#)

Determining Proper Burst Size for Traffic Policers

A policer burst-size limit controls the number of bytes of traffic that can pass unrestricted through a policed interface when a burst of traffic pushes the average transmit or receive rate above the configured bandwidth limit. The actual number of bytes of bursty traffic allowed to pass through a policed interface can vary from zero to the configured burst-size limit, depending on the overall traffic load.

By configuring a proper burst size, the effect of a lower shaped rate is alleviated. Use the **burst-size-limit** statement to configure the burst size.



NOTE: If you set the burst-size limit too low, too many packets will be subjected to rate limiting. If you set the burst-size limit too high, too few packets will be rate-limited.

Consider these two main factors when determining the burst size to use:

- The allowed duration of a blast of traffic on the line.

- The burst size is large enough to handle the maximum transmission unit (MTU) size of the packets.

The following general guidelines apply to choosing a policer burst-size limit:


- A burst-size limit should not be set lower than 10 times the MTU of the traffic on the interface to be policed.
- The amount of time to allow a burst of traffic at the full line rate of a policed interface should not be lower than 5 ms.
- The minimum and maximum values you can specify for a policer burst-size limit depends on the policer type (two-color or three-color).




BEST PRACTICE: The preferred method for choosing a burst-size limit is based on the line rate of the interface on which you apply the policer and the amount of time you want to allow a burst of traffic at the full line rate.

Bursty traffic requires a relatively large burst size so that extra tokens can be allocated into the token bucket for upcoming traffic to use. [Figure 106 on page 4835](#) shows an extreme case of bursty traffic where the opportunity to allocate tokens is missed, and the bandwidth goes unused because a large burst size is not configured.

Figure 106: Bursty Traffic Without Configured Burst Size (Excessive Unused Bandwidth)



[Figure 107 on page 4836](#) depicts how bandwidth usage changes when a large burst size is configured to handle bursty traffic. The large burst size minimizes the amount of unused bandwidth because tokens are being allocated in between the bursts of traffic that can be used during traffic peaks. The burst size determines the depth of the token bucket.

Figure 107: Bursty Traffic With Configured Burst Size (Less Unused Bandwidth)

Configuring a large burst size for the unused tokens creates another issue. If the burst size is set to a very large value, the burst of traffic can be transmitted from the interface at line rate until all the accumulated tokens in the token bucket are used up. This means that configuring a large burst size can allow too many packets to avoid rate limiting, which can lead to a traffic rate that exceeds the bandwidth limit for an extended period of time.

If the average rate is considered within 1 second, the rate is still below the configured bandwidth limit. However, the downstream device might not be able to handle bursty traffic, so some packets might be dropped. As a result, the way to determine the best burst size configuration is to perform experimental configurations, since one burst size is not suitable for every traffic pattern.


To program the burst size on the MX Series, M120, M320 routers, and EX Series switches, the user-configured burst size value is restricted to 1 ms through 600 ms of the policer rate (configured bandwidth limit).

Before performing burst size test configurations, select the initial burst size value using one of the following:

- The recommended formula for calculating burst size for bandwidth described as bits per second is: **burst size = bandwidth x allowable time for burst traffic / 8** (5 ms of the policer rate)
- For policers where the interface bandwidth is unknown, use the MTU method of calculating burst size: **burst size = interface MTU x 10**

[Figure 108 on page 4837](#) depicts a comparison between the two methods.

Figure 108: Comparing Burst Size Configuration Methods: 5 ms V.S. 10 MTU



In [Figure 108 on page 4837](#), with the 5 ms method, the burst size decreases with the policer rate (the configured bandwidth limit). As a result, the burst size can become so small that it might not be able to hold 2 MTU packets. This can severely affect the policer performance, which makes the 10 MTU method appear to be a better choice.

For example, with a 100 Kbps bandwidth limit configured on a Gigabit Ethernet interface and a burst size configured to 100 ms, the burst size becomes **100 Kbps x 100 ms = 1250 bytes**. This burst size is smaller than one standard MTU payload size on an Ethernet-type interface. A 10 MTU burst size provides a burst size of **1500 bytes x 10 = 15000 bytes**. However, since the maximum burst size is 600 ms of the bandwidth limit, the maximum configured burst size is **100 Kbps x 600 ms = 7500 bytes**. On the Gigabit Ethernet interface, the burst duration is **7500 bytes / 1 Gbps = 60 μs** at Gigabit Ethernet line rate. If the burst size is too large for the downstream device, the burst size can be further reduced until the result is acceptable.

However, if the bandwidth limit is very high, the 10 MTU method might not be able to create a token bucket large enough to accommodate the unused credits. As a result, the average bandwidth limit is lower than what you configured. In this case, the 5 ms method is the better choice for configuring the burst size.

If a 200 Mbps bandwidth limit is configured with a 5 ms burst size, the calculation becomes **200 Mbps x 5 ms = 125 Kbytes**, which is approximately 83 1500-byte packets. If the 200 Mbps bandwidth limit is configured on a Gigabit Ethernet interface, the burst duration is **125000 bytes / 1 Gbps = 1 ms** at the Gigabit Ethernet line rate.

If a large burst size is configured at 600 ms with the bandwidth limit configured at 200 Mbps, the calculation becomes **200 Mbps x 600 ms = 15 Mbytes**. This creates a burst duration of 120 ms at the Gigabit Ethernet line rate. The average bandwidth rate in 1 second becomes **200 Mbps + 15 Mbytes = 320 Mbps**, which is much higher than the configured bandwidth limit at 200 Mbps. This example shows that a larger burst size can affect the measured bandwidth rate.

Related Documentation

- [Policer Implementation Overview on page 4828](#)
- [Understanding the Benefits of Policers and Token Bucket Algorithms on page 4833](#)

Configuration

- [Configuring Single-Rate Two-Color Policers on page 4838](#)
- [Configuring Three-Color Policers on page 4912](#)
- [Configuring Logical and Physical Interface Policers on page 4927](#)
- [Configuring Layer 2 Policers on page 4947](#)
- [Configuration Statements on page 4956](#)

Configuring Single-Rate Two-Color Policers

- [Basic Single-Rate Two-Color Policers on page 4838](#)
- [Bandwidth Policers on page 4854](#)
- [Filter-Specific Counters and Policers on page 4863](#)
- [Prefix-Specific Counting and Policing Actions on page 4874](#)
- [Multifield Classification on page 4889](#)
- [Policer Overhead to Account for Rate Shaping in the Traffic Manager on page 4904](#)

Basic Single-Rate Two-Color Policers

- [Single-Rate Two-Color Policer Overview on page 4838](#)
- [Example: Configuring a Single-Rate Two-Color Policer on page 4839](#)
- [Example: Configuring Interface and Firewall Filter Policers at the Same Interface on page 4845](#)

Single-Rate Two-Color Policer Overview

Single-rate two color policing enforces a configured rate of traffic flow for a particular service level by applying implicit or configured actions to traffic that does not conform to the limits. When you apply a single-rate two-color policer to the input or output traffic at an interface, the policer meters the traffic flow to the rate limit defined by the following components:

- **Bandwidth limit**—The average number of bits per second permitted for packets received or transmitted at the interface. You can specify the bandwidth limit as an absolute number of bits per second or as a percentage value from 1 through 100. If a percentage value is specified, the effective bandwidth limit is calculated as a percentage of either the physical interface media rate or the logical interface configured shaping rate.
- **Burst-size limit**—The maximum size permitted for bursts of data.

For a traffic flow that conforms to the configured limits (categorized as green traffic), packets are implicitly marked with a packet loss priority (PLP) level of **low** and are allowed to pass through the interface unrestricted.

For a traffic flow that exceeds the configured limits (categorized as red traffic), packets are handled according to the traffic-policing actions configured for the policer. The action might be to discard the packet, or the action might be to re-mark the packet with a specified forwarding class, a specified PLP, or both, and then transmit the packet.

To rate-limit Layer 3 traffic, you can apply a two-color policer in the following ways:

- Directly to a logical interface, at a specific protocol level.
- As the action of a standard stateless firewall filter that is applied to a logical interface, at a specific protocol level.

To rate-limit Layer 2 traffic, you can apply a two-color policer as a *logical interface policer* only. You cannot apply a two-color policer to Layer 2 traffic through a firewall filter.

Example: Configuring a Single-Rate Two-Color Policer

This example shows how to configure a single-rate two-color policer that you apply to a logical interface as a firewall filter action.

- [Requirements on page 4839](#)
- [Overview on page 4839](#)
- [Configuration on page 4839](#)
- [Verification on page 4844](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

To set the maximum burst size of the policer, you should know the line rate of the interface (preferred) or else the maximum transmission unit (MTU) of the traffic on the target interface.

Overview

In this example, you configure a single-rate two-color policer and then use an IPv4 firewall filter to apply the policer to all traffic except for BGP messages. You apply the firewall filter to the input and output of the same logical interface.

Topology

Assume that you know the traffic flow needs to be rate-limited to an average bandwidth of 9 Mbps. Also assume that you do not know the line rate of the target interface, but you do know that the MTU of traffic on the interface is 4700 bytes. As described in [“Determining Proper Burst Size for Traffic Policers” on page 4834](#), you can estimate a starting value of 10 times the MTU.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configuring the Logical Interface and Monitoring the Traffic Rate Without Policing on page 4840](#)
- [Configuring the Basic Single-Rate Two-Color Policer on page 4841](#)

- [Referencing the Policer from a Term in a Stateless Firewall Filter on page 4842](#)
- [Applying the Firewall Filter to the Input and Output of the Logical Interface on page 4843](#)

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces so-1/3/0 unit 0 family inet address 10.39.1.1/16
set firewall policer p-9m-47k-discard if-exceeding bandwidth-limit 9m
set firewall policer p-9m-47k-discard if-exceeding burst-size-limit 47k
set firewall policer p-9m-47k-discard then discard
set firewall family inet filter rate-limit-in term t1 from protocol tcp
set firewall family inet filter rate-limit-in term t1 from port bgp
set firewall family inet filter rate-limit-in term t1 then accept
set firewall family inet filter rate-limit-in term t2 then policer p-9m-47k-discard
set firewall family inet filter rate-limit-in term t2 then accept
set interfaces so-1/3/0 unit 0 family inet filter input rate-limit-in
set interfaces so-1/3/0 unit 0 family inet filter output rate-limit-in
```

Configuring the Logical Interface and Monitoring the Traffic Rate Without Policing**Step-by-Step Procedure**

To configure the logical interface and monitor the traffic rate:

1. Configure IPv4 on the logical interface.

```
[edit]
user@host# edit interfaces so-1/3/0 unit 0 family inet address 10.39.1.1/16
```

2. Commit the configuration.

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

3. Monitor the traffic flow at the interface (press 'q' to terminate monitoring).

```
[edit]
user@host# run monitor interface so-1/3/0.0
```

The traffic statistics report the input and output rates before applying interface rate limiting.

Results

Confirm the configuration of the logical interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
so-3/3/0 {
  unit 0 {
    family inet {
      address 10.39.1.1/16
    }
  }
}
```

*Configuring the Basic Single-Rate Two-Color Policer***Step-by-Step
Procedure**

To configure the basic single-rate two-color policer:

1. Enable configuration of a single-rate two-color policer.

[edit]

user@host# edit **firewall policer** p-9m-47k-discard

**NOTE:**

This example illustrates a basic two-color policer, which does not require any of the following statements:

- **filter-specific**—You can include this statement so that, when this policer is applied as a firewall filter policer, a single instance of the policer is used by all terms within the same firewall filter.
- **logical-bandwidth-policer**—Include this statement only if the policer bandwidth limit is specified as a percentage value (using the **bandwidth-percent percent** statement) and you want the bandwidth limit to be based on the configured shaping rate for the target logical interface.
- **logical-interface-policer**—Include this statement only if you apply the policer directly to a physical or logical interface to rate-limit traffic for any configured protocol family based on the physical interface media rate.
- **physical-interface-policer**—Include this statement if you apply the policer to packets filtered by a standard stateless firewall filter term and you want the policer to meter the aggregate traffic (all protocol families and logical interfaces configured on the underlying physical interface).

2. Configure the policer to rate-limit to an average bandwidth of 9 Mbps and a burst size of 47 KB.

[edit firewall policer p-9m-47k-discard]

user@host# set **if-exceeding bandwidth-limit** 9m

user@host# set **if-exceeding burst-size-limit** 47k

For information about configuring the burst size, see [“Determining Proper Burst Size for Traffic Policers”](#) on page 4834.

When the traffic flow conforms to these limits, the flow is categorized as green. For a single-rate two-color policer, packets in a green flow are implicitly set to a **low** packet loss priority (PLP) level.

3. Configure the policer to discard packets in a red traffic flow.

[edit firewall policer p-9m-47k-discard]

user@host# set **then discard**

Instead of discarding the traffic that exceeds the traffic limits, you can set the PLP level, the forwarding class assignment, or both.

Results Confirm the configuration of the policer by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
policer p-9m-47k-discard {
  if-exceeding {
    bandwidth-limit 9m;
    burst-size-limit 47k;
  }
  then discard;
}
```

Referencing the Policer from a Term in a Stateless Firewall Filter

Step-by-Step Procedure To reference the policer from a term in a stateless firewall filter:

1. Enable configuration of the stateless firewall filter.

```
[edit]
user@host# edit firewall family inet filter rate-limit-in
```

2. Configure the first term to accept all BGP traffic.

BGP messages are transported over TCP port 179.

```
[edit firewall family inet filter rate-limit-in]
user@host# set term t1 from protocol tcp
user@host# set term t1 from port bgp
user@host# set term t1 then accept
```

The BGP traffic is not rate-limited to avoid having the BGP session time out because the packet is discarded by the policer.

3. Configure the second term to rate-limit all other traffic using the policer.

```
[edit firewall family inet filter rate-limit-in]
user@host# set term t2 then policer p-9m-47k-discard
user@host# set term t2 then accept
```

Results Confirm the configuration of the firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
  filter rate-limit-in {
    term t1 {
      from {
        protocol tcp;
        port bgp;
      }
    }
  }
}
```



```

        then accept;
    }
    term t2 {
        then {
            policer p-9m-47k-discard;
            accept;
        }
    }
}
policer p-9m-47k-discard {
    if-exceeding {
        bandwidth-limit 9m;
        burst-size-limit 47k;
    }
    then discard;
}

```

Applying the Firewall Filter to the Input and Output of the Logical Interface

Step-by-Step Procedure

To apply the stateless firewall filter to the input and output of the logical interface:

1. Enable configuration of IPv4 on the logical interface.

```

[edit]
user@host# edit interfaces so-1/3/0 unit 0 family inet

```

2. Apply the stateless firewall filter.

```

[edit interfaces so-1/3/0 unit 0 family inet]
user@host# set filter input rate-limit-in
user@host# set filter output rate-limit-in

```

Results

Confirm the configuration of the logical interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```

[edit]
user@host# show interfaces
so-3/3/0 {
    unit 0 {
        family inet {
            filter {
                input rate-limit-in;
                output rate-limit-in;
            }
            address 10.39.1.1/16;
        }
    }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Monitoring and Adjusting Policing of the Logical Interface on page 4844](#)
- [Displaying the Number of Packets Processed by the Policer at the Logical Interface on page 4844](#)

Monitoring and Adjusting Policing of the Logical Interface

Purpose Ensure that the traffic limits configured in the policer are throttling the input and output rate as you intended.

Action To ensure that the configured traffic limits are throttling the traffic rate as intended:

1. Monitor the traffic flow at the interface (press 'q' to terminate monitoring).

[edit]

```
user@host# run monitor interface so-1/3/0.0
```

The traffic statistics report the input and output rate after applying interface rate limiting.

2. Adjust the burst-size limit if necessary.

If the policer appears to be throttling the input and output rate more than you intended, increase the burst size (possibly doubling it) and then check again.

Displaying the Number of Packets Processed by the Policer at the Logical Interface

Purpose Verify the traffic flow through the logical interface and that the policer is evaluated when packets are received on the logical interface.

Action Use the **show firewall** operational mode command for the filter you applied to the logical interface:

[edit]

```
user@host# run show firewall filter rate-limit-in
```

```
Filter: rate-limit-in
```

```
Policers:
```

Name	Packets
p-9m-47k-discard-t2	32863

The command output displays the name of policer (**p-9m-47k-discard**), the name of the filter term (**t2**) under which the policer action is specified, and the number of packets that matched the filter term. This is only the number of out-of-specification (out-of-spec) packet counts, not all packets policed by the policer.

Example: Configuring Interface and Firewall Filter Policers at the Same Interface

This example shows how to configure three single-rate two-color policers and apply the policers to the IPv4 input traffic at the same single-tag virtual LAN (VLAN) logical interface.

- [Requirements on page 4845](#)
- [Overview on page 4845](#)
- [Configuration on page 4846](#)
- [Verification on page 4852](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, you configure three single-rate two-color policers and apply the policers to the IPv4 input traffic at the same single-tag VLAN logical interface. Two policers are applied to the interface through a firewall filter, and one policer is applied directly to the interface.

You configure one policer, named **p-all-1m-5k-discard**, to rate-limit traffic to 1 Mbps with a burst size of 5000 bytes. You apply this policer directly to IPv4 input traffic at the logical interface. When you apply a policer directly to protocol-specific traffic at a logical interface, the policer is said to be applied as an *interface policer*.

You configure the other two policers to allow burst sizes of 500 KB, and you apply these policers to IPv4 input traffic at the logical interface by using an IPv4 standard stateless firewall filter. When you apply a policer to protocol-specific traffic at a logical interface through a firewall filter action, the policer is said to be applied as a *firewall-filter policer*.

- You configure the policer named **p-icmp-500k-500k-discard** to rate-limit traffic to 500 Kbps with a burst size of 500 K bytes by discarding packets that do not conform to these limits. You configure one of the firewall filter terms to apply this policer to Internet Control Message Protocol (ICMP) packets.
- You configure the policer named **p-ftp-10p-500k-discard** to rate-limit traffic to a 10 percent bandwidth with a burst size of 500 KB by discarding packets that do not conform to these limits. You configure another firewall-filter term to apply this policer to File Transfer Protocol (FTP) packets.

A policer that you configure with a bandwidth limit expressed as a percentage value (rather than as an absolute bandwidth value) is called a *bandwidth policer*. Only single-rate two-color policers can be configured with a percentage bandwidth specification. By default, a bandwidth policer rate-limits traffic to the specified percentage of the line rate of the physical interface underlying the target logical interface.

Topology

You configure the target logical interface as a single-tag VLAN logical interface on a Fast Ethernet interface operating at 100 Mbps. This means that the policer you configure with

the 10-percent bandwidth-limit (the policer that you apply to FTP packets) rate-limits the FTP traffic on this interface to 10 Mbps.



NOTE: In this example, you do not configure the bandwidth policer as a *logical-bandwidth policer*. Therefore, the percentage is based on the physical media rate rather than on the configured shaping rate of the logical interface.

The firewall filter that you configure to reference two of the policers must be configured as an *interface-specific filter*. Because the policer that is used to rate-limit FTP packets specifies the bandwidth limit as a percentage value, the firewall filter that references this policer must be configured as an interface-specific filter. Thus, if this firewall filter were to be applied to multiple interfaces instead of just the Fast Ethernet interface in this example, unique policers and counters would be created for each interface to which the filter is applied.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configuring the Single-Tag VLAN Logical Interface on page 4847](#)
- [Configuring the Three Policers on page 4848](#)
- [Configuring the IPv4 Firewall Filter on page 4849](#)
- [Applying the Interface Policer and Firewall Filter Policers to the Logical Interface on page 4851](#)

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces fe-0/1/1 vlan-tagging
set interfaces fe-0/1/1 unit 0 vlan-id 100
set interfaces fe-0/1/1 unit 0 family inet address 10.20.15.1/24
set interfaces fe-0/1/1 unit 1 vlan-id 101
set interfaces fe-0/1/1 unit 1 family inet address 10.20.240.1/24
set firewall policer p-all-1m-5k-discard if-exceeding bandwidth-limit 1m
set firewall policer p-all-1m-5k-discard if-exceeding burst-size-limit 5k
set firewall policer p-all-1m-5k-discard then discard
set firewall policer p-ftp-10p-500k-discard if-exceeding bandwidth-percent 10
set firewall policer p-ftp-10p-500k-discard if-exceeding burst-size-limit 500k
set firewall policer p-ftp-10p-500k-discard then discard
set firewall policer p-icmp-500k-500k-discard if-exceeding bandwidth-limit 500k
set firewall policer p-icmp-500k-500k-discard if-exceeding burst-size-limit 500k
set firewall policer p-icmp-500k-500k-discard then discard
set firewall family inet filter filter-ipv4-with-limits interface-specific
set firewall family inet filter filter-ipv4-with-limits term t-ftp from protocol tcp
set firewall family inet filter filter-ipv4-with-limits term t-ftp from port ftp
set firewall family inet filter filter-ipv4-with-limits term t-ftp from port ftp-data
```

```

set firewall family inet filter filter-ipv4-with-limits term t-ftp then policer
  p-ftp-10p-500k-discard
set firewall family inet filter filter-ipv4-with-limits term t-icmp from protocol icmp
set firewall family inet filter filter-ipv4-with-limits term t-icmp then policer
  p-icmp-500k-500k-discard
set firewall family inet filter filter-ipv4-with-limits term catch-all then accept
set interfaces fe-0/1/1 unit 1 family inet filter input filter-ipv4-with-limits
set interfaces fe-0/1/1 unit 1 family inet policer input p-all-1m-5k-discard

```

Configuring the Single-Tag VLAN Logical Interface

Step-by-Step Procedure To configure the single-tag VLAN logical interface:

1. Enable configuration of the Fast Ethernet interface.


```

[edit]
user@host# edit interfaces fe-0/1/1

```
2. Enable single-tag VLAN framing.


```

[edit interfaces fe-0/1/1]
user@host# set vlan-tagging

```
3. Bind VLAN IDs to the logical interfaces.


```

[edit interfaces fe-0/1/1]
user@host# set unit 0 vlan-id 100
user@host# set unit 1 vlan-id 101

```
4. Configure IPv4 on the single-tag VLAN logical interfaces.


```

[edit interfaces fe-0/1/1]
user@host# set unit 0 family inet address 10.20.15.1/24
user@host# set unit 1 family inet address 10.20.240.1/24

```

Results Confirm the configuration of the VLAN by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```

[edit]
user@host# show interfaces
fe-0/1/1 {
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    family inet {
      address 10.20.15.1/24;
    }
  }
  unit 1 {
    vlan-id 101;
    family inet {
      address 10.20.240.1/24;
    }
  }
}

```

Configuring the Three Policers

Step-by-Step Procedure

To configure the three policers:

1. Enable configuration of a two-color policer that discards packets that do not conform to a bandwidth of 1 Mbps and a burst size of 5000 bytes.



NOTE: You apply this policer directly to all IPv4 input traffic at the single-tag VLAN logical interface, so the packets will not be filtered before being subjected to rate limiting.

[edit]

user@host# edit **firewall policer p-all-1m-5k-discard**

2. Configure the first policer.

[edit firewall policer p-all-1m-5k-discard]

user@host# set **if-exceeding bandwidth-limit 1m**

user@host# set **if-exceeding burst-size-limit 5k**

user@host# set then discard

3. Enable configuration of a two-color policer that discards packets that do not conform to a bandwidth specified as "10 percent" and a burst size of 500,000 bytes.

You apply this policer only to the FTP traffic at the single-tag VLAN logical interface.

You apply this policer as the action of an IPv4 firewall filter term that matches FTP packets from TCP.

[edit firewall policer p-all-1m-5k-discard]

user@host# up

[edit]

user@host# edit **firewall policer p-ftp-10p-500k-discard**

4. Configure policing limits and actions.

[edit firewall policer p-ftp-10p-500k-discard]

user@host# set **if-exceeding bandwidth-percent 10**

user@host# set **if-exceeding burst-size-limit 500k**

user@host# set then discard

Because the bandwidth limit is specified as a percentage, the firewall filter that references this policer must be configured as an interface-specific filter.



NOTE: If you wanted this policer to rate-limit to 10 percent of the logical interface configured shaping rate (rather than to 10 percent of the physical interface media rate), you would need to include the **logical-bandwidth-policer** statement at the [edit firewall policer p-all-1m-5k-discard] hierarchy level. This type of policer is called a *logical-bandwidth policer*.

5. Enable configuration of the IPv4 firewall filter policer for ICMP packets.

```
[edit firewall policer p-ftp-10p-500k-discard]
user@host# up

[edit]
user@host# edit firewall policer p-icmp-500k-500k-discard
```

6. Configure policing limits and actions.

```
[edit firewall policer p-icmp-500k-500k-discard]
user@host# set if-exceeding bandwidth-limit 500k
user@host# set if-exceeding burst-size-limit 500k
user@host# set then discard
```

Results Confirm the configuration of the policers by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
policer p-all-1m-5k-discard {
  if-exceeding {
    bandwidth-limit 1m;
    burst-size-limit 5k;
  }
  then discard;
}
policer p-ftp-10p-500k-discard {
  if-exceeding {
    bandwidth-percent 10;
    burst-size-limit 500k;
  }
  then discard;
}
policer p-icmp-500k-500k-discard {
  if-exceeding {
    bandwidth-limit 500k;
    burst-size-limit 500k;
  }
  then discard;
}
```

Configuring the IPv4 Firewall Filter

Step-by-Step Procedure To configure the IPv4 firewall filter:

1. Enable configuration of the IPv4 firewall filter.

```
[edit]
user@host# edit firewall family inet filter filter-ipv4-with-limits
```

2. Configure the firewall filter as interface-specific.

```
[edit firewall family inet filter filter-ipv4-with-limits]
user@host# set interface-specific
```

The firewall filter must be interface-specific because one of the policers referenced is configured with a bandwidth limit expressed as a percentage value.

3. Enable configuration of a filter term to rate-limit FTP packets.

```
[edit firewall family inet filter filter-ipv4-with-limits]
user@host# edit term t-ftp
```

```
[edit firewall family inet filter filter-ipv4-with-limits term t-ftp]
user@host# set from protocol tcp
user@host# set from port [ ftp ftp-data ]
```

FTP messages are sent over TCP port 20 (**ftp**) and received over TCP port 21 (**ftp-data**).

4. Configure the filter term to match FTP packets.

```
[edit firewall family inet filter filter-ipv4-with-limits term t-ftp]
user@host# set then policer p-ftp-10p-500k-discard
```

5. Enable configuration of a filter term to rate-limit ICMP packets.

```
[edit firewall family inet filter filter-ipv4-with-limits term t-ftp]
user@host# up
```

```
[edit firewall family inet filter filter-ipv4-with-limits]
user@host# edit term t-icmp
```

6. Configure the filter term for ICMP packets

```
[edit firewall family inet filter filter-ipv4-with-limits term t-icmp]
user@host# set from protocol icmp
user@host# set then policer p-icmp-500k-500k-discard
```

7. Configure a filter term to accept all other packets without policing.

```
[edit firewall family inet filter filter-ipv4-with-limits term t-icmp]
user@host# up
```

```
[edit firewall family inet filter filter-ipv4-with-limits]
user@host# set term catch-all then accept
```

Results Confirm the configuration of the firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
  filter filter-ipv4-with-limits {
    interface-specific;
    term t-ftp {
      from {
        protocol tcp;
        port [ ftp ftp-data ];
      }
      then policer p-ftp-10p-500k-discard;
    }
  }
}
```



```

term t-icmp {
    from {
        protocol icmp;
    }
    then policer p-icmp-500k-500k-discard;
}
term catch-all {
    then accept;
}
}
}
policer p-all-1m-5k-discard {
    if-exceeding {
        bandwidth-limit 1m;
        burst-size-limit 5k;
    }
    then discard;
}
policer p-ftp-10p-500k-discard {
    if-exceeding {
        bandwidth-percent 10;
        burst-size-limit 500k;
    }
    then discard;
}
policer p-icmp-500k-500k-discard {
    if-exceeding {
        bandwidth-limit 500k;
        burst-size-limit 500k;
    }
    then discard;
}
}

```

Applying the Interface Policer and Firewall Filter Policers to the Logical Interface

Step-by-Step Procedure

To apply the three policers to the VLAN:

1. Enable configuration of IPv4 on the logical interface.

```
[edit]
user@host# edit interfaces fe-0/1/1 unit 1 family inet
```
2. Apply the firewall filter policers to the interface.

```
[edit interfaces fe-0/1/1 unit 1 family inet]
user@host# set filter input filter-ipv4-with-limits
```
3. Apply the interface policer to the interface.

```
[edit interfaces fe-0/1/1 unit 1 family inet]
user@host# set policer input p-all-1m-5k-discard
```

Input packets at **fe-0/1/1.0** are evaluated against the interface policer before they are evaluated against the firewall filter policers. For more information, see [“Order of Policer and Firewall Filter Operations” on page 4810](#).

Results Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
fe-0/1/1 {
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    family inet {
      address 10.20.15.1/24;
    }
  }
  unit 1 {
    vlan-id 101;
    family inet {
      filter {
        input filter-ipv4-with-limits;
      }
      policer {
        input p-all-1m-5k-discard;
      }
      address 10.20.240.1/24;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Displaying Policers Applied Directly to the Logical Interface on page 4852](#)
- [Displaying Statistics for the Policer Applied Directly to the Logical Interface on page 4853](#)
- [Displaying the Policers and Firewall Filters Applied to an Interface on page 4853](#)
- [Displaying Statistics for the Firewall Filter Policers on page 4854](#)

Displaying Policers Applied Directly to the Logical Interface

Purpose Verify that the interface policer is evaluated when packets are received on the logical interface.

Action Use the **show interfaces policers** operational mode command for logical interface **fe-0/1/1.1**. The command output section for the **Proto** column and **Input Policer** column shows that the policer **p-all-1m-5k-discard** is evaluated when packets are received on the logical interface.

```
user@host> show interfaces policers fe-0/1/1.1
Interface      Admin Link Proto Input Policer      Output Policer
fe-0/1/1.1     up    up
```

```
inet p-all-1m-5k-discard-fe-0/1/1.1-inet-i
```

In this example, the interface policer is applied to logical interface traffic in the input direction only.

Displaying Statistics for the Policer Applied Directly to the Logical Interface

Purpose Verify the number of packets evaluated by the interface policer.

Action Use the `show policer` operational mode command and optionally specify the name of the policer. The command output displays the number of packets evaluated by each configured policer (or the specified policer), in each direction.

```
user@host> show policer p-all-1m-5k-discard-fe-0/1/1.1-inet-i
Policers:
Name                                     Bytes      Packets
p-all-1m-5k-discard-fe-0/1/1.1-inet-i    200         5
```

Displaying the Policers and Firewall Filters Applied to an Interface

Purpose Verify that the firewall filter `filter-ipv4-with-limits` is applied to the IPv4 input traffic at logical interface `fe-0/1/1.1`.

Action Use the `show interfaces statistics` operational mode command for logical interface `fe-0/1/1.1`, and include the `detail` option. Under the **Protocol inet** section of the command output section, the **Input Filters** and **Policer** lines display the names of filter and policer applied to the logical interface in the input direction.

```
user@host> show interfaces statistics fe-0/1/1.1 detail
Logical interface fe-0/1/1.1 (Index 83) (SNMP ifIndex 545) (Generation 153)
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ] Encapsulation: ENET2
Traffic statistics:
  Input bytes : 0
  Output bytes : 46
  Input packets: 0
  Output packets: 1
Local statistics:
  Input bytes : 0
  Output bytes : 46
  Input packets: 0
  Output packets: 1
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
Protocol inet, MTU: 1500, Generation: 176, Route table: 0
Flags: Sendbcst-pkt-to-re
Input Filters: filter-ipv4-with-limits-fe-0/1/1.1-i
Policer: Input: p-all-1m-5k-discard-fe-0/1/1.1-inet-i
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.20.130/24, Local: 10.20.130.1, Broadcast: 10.20.130.255,
```

Generation: 169

In this example, the two firewall filter policers are applied to logical interface traffic in the input direction only.

Displaying Statistics for the Firewall Filter Policers

Purpose Verify the number of packets evaluated by the firewall filter policers.

Action Use the **show firewall** operational mode command for the filter you applied to the logical interface.

[edit]

```
user@host> show firewall filter filter-ipv4-with-limits-fe-0/1/1.1-i
```

Filter: filter-ipv4-with-limits-fe-0/1/1.1-i

Policers:

Name	Bytes	Packets
p-ftp-10p-500k-discard-t-ftp-fe-0/1/1.1-i	0	0
p-icmp-500k-500k-discard-t-icmp-fe-0/1/1.1-i	0	0

The command output displays the names of the policers (**p-ftp-10p-500k-discard** and **p-icmp-500k-500k-discard**), combined with the names of the filter terms (**t-ftp** and **t-icmp**, respectively) under which the policer action is specified. The policer-specific output lines display the number of packets that matched the filter term. This is only the number of out-of-specification (out-of-spec) packet counts, not all packets policed by the policer.

Related Documentation

- [Order of Policer and Firewall Filter Operations on page 4810](#)
- [Statement Hierarchy for Configuring Policers on page 4811](#)
- [Two-Color Policer Configuration Overview on page 4813](#)
- [Guidelines for Applying Traffic Policers on page 4821](#)
- [Determining Proper Burst Size for Traffic Policers on page 4834](#)

Bandwidth Policers

- [Bandwidth Policer Overview on page 4854](#)
- [Example: Configuring a Logical Bandwidth Policer on page 4856](#)

Bandwidth Policer Overview

For a single-rate two-color policer only, you can specify the bandwidth limit as a percentage value from 1 through 100 instead of as an absolute number of bits per second. This type of two-color policer, called a *bandwidth policer*, rate-limits traffic to a bandwidth limit that is calculated as a percentage of either the physical interface media rate or the logical interface configured shaping rate.

Guidelines for Configuring a Bandwidth Policer

The following guidelines apply to configuring a bandwidth policer:

- To specify a percentage bandwidth limit, you include the **bandwidth-percent *percentage*** statement in place of the **bandwidth-limit *bps*** statement.
- By default, a bandwidth policer calculates the percentage bandwidth limit based on the physical interface port speed. To configure a bandwidth policer to calculate the percentage bandwidth limit based on the configured logical interface shaping rate instead, include the **logical-bandwidth-policer** statement at the **[edit firewall policer *policer-name*]** hierarchy level. This type of bandwidth policer is called a *logical bandwidth policer*.

You can configure a logical interface shaping rate by including the **shaping-rate *bps*** statement at the **[edit class-of-service interfaces interface *interface-name* unit *logical-unit-number*]** hierarchy level. A logical interface shaping rate causes the specified amount of bandwidth to be allocated to the logical interface.



NOTE: If you configure a logical-bandwidth policer and then apply the policer to a logical interface that is not configured with a shaping rate, then the policer rate-limits traffic on that logical interface to calculate the percentage bandwidth limit based on the physical interface port speed, even if you include the **logical-bandwidth-policer** statement in the bandwidth policer configuration.

- If you reference a bandwidth policer from a stateless firewall filter term, you must include the **interface-specific** statement in the firewall filter configuration.

Guidelines for Applying a Bandwidth Policer

The following guidelines pertain to applying a bandwidth policer to traffic:

- You can use a bandwidth policer to rate-limit protocol-specific traffic (not **family any**) at the input or output of a logical interface.
- You can apply a bandwidth policer directly to protocol-specific input or output traffic at a logical interface.
- To send only selected packets to a bandwidth policer, you can reference the bandwidth policer from a stateless firewall filter term and then apply the filter to logical interface traffic for a specific protocol family.
 - To reference a *logical bandwidth policer* from a firewall filter, you must include the **interface-specific** statement in the firewall filter configuration.
 - You cannot use a bandwidth policer for forwarding-table filters.
- You cannot apply a bandwidth policer to an aggregate interface, a tunnel interface, or a software interface.

Example: Configuring a Logical Bandwidth Policer

This example shows how to configure a logical bandwidth policer.

- [Requirements on page 4856](#)
- [Overview on page 4856](#)
- [Configuration on page 4857](#)
- [Verification on page 4861](#)

Requirements

Before you begin, make sure that you have two logical units available on a Gigabit Ethernet interface.

Overview

In this example, you configure a single-rate two-color policer that specifies the bandwidth limit as a percentage value rather than as an absolute number of bits per second. This type of policer is called a *bandwidth policer*. By default, a bandwidth policer enforces a bandwidth limit based on the line rate of the underlying physical interface. As an option, you can configure a bandwidth policer to enforce a bandwidth limit based on the configured shaping rate of the logical interface. To configure this type of bandwidth policer, called a *logical bandwidth policer*, you include the [logical-bandwidth-policer](#) statement in the policer configuration.

To configure a logical interface shaping rate, include the **shaping-rate bps** statement at the **[edit class-of-service interfaces interface *interface-name* unit *logical-unit-number*]** hierarchy level. This class-of-service (CoS) configuration statement causes the specified amount of bandwidth to be allocated to the logical interface.



NOTE: If you configure a policer bandwidth limit as a percentage but a shaping rate is not configured for the target logical interface, the policer bandwidth limit is calculated as a percentage of the physical interface media rate, even if you enable the logical-bandwidth policing feature.

To apply a logical bandwidth policer to a logical interface, you can apply the policer directly to the logical interface at the protocol family level or (if you only need to rate-limit filtered packets) you can reference the policer from a stateless firewall filter configured to operate in *interface-specific* mode.

Topology

In this example, you configure two logical interfaces on a single Gigabit Ethernet interface and configure a shaping rate on each logical interface. On logical interface **ge-1/3/0.0**, you allocate 4 Mbps of bandwidth. On logical interface **ge-1/3/0.1**, you allocate 2 Mbps of bandwidth.

You also configure a logical bandwidth policer with a bandwidth limit of 50 percent and a maximum burst size of 125,000 bytes, and then you apply the policer to input and output traffic at the logical units configured on **ge-1/3/0.0**. For logical interface **ge-1/3/0.0**,

the policer rate-limits to a bandwidth limit of 2 Mbps (50 percent of the 4 Mbps shaping rate configured for the logical interface). For logical interface **ge-1/3/0.1**, the policer rate-limits traffic to a bandwidth limit of 1 Mbps (50 percent of the 2 Mbps shaping rate configured for the logical interface).

If no shaping rate is configured for a target logical interface, the policer rate-limits to a bandwidth limit calculated as 50 percent of the physical interface media rate. For example, if you apply a 50 percent bandwidth policer to input or output traffic at a Gigabit Ethernet logical interface without rate shaping, the policer applies a bandwidth limit of 500 Mbps (50 percent of 1000 Mbps).

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configuring the Logical Interfaces on page 4857](#)
- [Configuring Traffic Rate-Shaping by Specifying the Amount of Bandwidth to be Allocated to the Logical Interface on page 4858](#)
- [Configuring the Logical Bandwidth Policar on page 4859](#)
- [Applying the Logical Bandwidth Policers to the Logical Interfaces on page 4860](#)

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces ge-1/3/0 per-unit-scheduler
set interfaces ge-1/3/0 vlan-tagging
set interfaces ge-1/3/0 unit 0 vlan-id 100
set interfaces ge-1/3/0 unit 0 family inet address 172.1.1.1/30
set interfaces ge-1/3/0 unit 1 vlan-id 200
set interfaces ge-1/3/0 unit 1 family inet address 172.2.1.1/30
set class-of-service interfaces ge-1/3/0 unit 0 shaping-rate 4m
set class-of-service interfaces ge-1/3/0 unit 1 shaping-rate 2m
set firewall policer LB-policer logical-bandwidth-policer
set firewall policer LB-policer if-exceeding bandwidth-percent 50
set firewall policer LB-policer if-exceeding burst-size-limit 125k
set firewall policer LB-policer then discard
set interfaces ge-1/3/0 unit 0 family inet policer input LB-policer
set interfaces ge-1/3/0 unit 0 family inet policer output LB-policer
set interfaces ge-1/3/0 unit 1 family inet policer input LB-policer
set interfaces ge-1/3/0 unit 1 family inet policer output LB-policer
```

Configuring the Logical Interfaces

Step-by-Step Procedure

To configure the logical interfaces:

1. Enable configuration of the physical interface.

```
[edit]
user@host# edit interfaces ge-1/3/0
```

```
[edit interfaces ge-1/3/0]
user@host# set per-unit-scheduler
user@host# set vlan-tagging
```

2. Configure the first logical interface.

```
[edit interfaces ge-1/3/0]
user@host# set unit 0 vlan-id 100
user@host# set unit 0 family inet address 172.1.1.1/30
```

3. Configure the second logical interface.

```
[edit interfaces ge-1/3/0]
user@host# set unit 1 vlan-id 200
user@host# set unit 1 family inet address 172.2.1.1/30
```

Results Confirm the configuration of the interfaces by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
ge-1/3/0 {
  per-unit-scheduler;
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    family inet {
      address 172.1.1.1/30;
    }
  }
  unit 1 {
    vlan-id 200;
    family inet {
      address 172.2.1.1/30;
    }
  }
}
```

Configuring Traffic Rate-Shaping by Specifying the Amount of Bandwidth to be Allocated to the Logical Interface

Step-by-Step Procedure To configure rate shaping by specifying the bandwidth to be allocated to the logical interface:

1. Enable CoS configuration on the physical interface.

```
[edit]
user@host# edit class-of-service interfaces ge-1/3/0
```

2. Configure rate shaping for the logical interfaces.

```
[edit]
user@host# set unit 0 shaping-rate 4m
user@host# set unit 1 shaping-rate 2m
```


These statements allocate 4 Mbps of bandwidth to logical unit **ge-1/3/0.0** and 2 Mbps of bandwidth to logical unit **ge-1/3/0.1**.

Results Confirm the configuration of the rate shaping by entering the **show class-of-service** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show class-of-service
interfaces {
  ge-1/3/0 {
    unit 0 {
      shaping-rate 4m;
    }
    unit 1 {
      shaping-rate 2m;
    }
  }
}
```

Configuring the Logical Bandwidth Policer

Step-by-Step Procedure To configure the logical bandwidth policer:

1. Enable configuration of a single-rate two-color policer.

```
[edit]
user@host# edit firewall policer LB-policer
```

2. Configure the policer as a logical-bandwidth policer.

```
[edit firewall policer LB-policer]
user@host# set logical-bandwidth-policer
```

This applies the rate-limiting to logical interfaces.

3. Configure the policer traffic limits and actions.

```
[edit firewall policer LB-policer]
user@host# set if-exceeding bandwidth-percent 50
user@host# set if-exceeding burst-size-limit 125k
user@host# set then discard
```

Results Confirm the configuration of the policer by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
policer LB-policer {
  logical-bandwidth-policer;
  if-exceeding {
    bandwidth-percent 50;
    burst-size-limit 125k;
  }
  then discard;
```

```
}
```

Applying the Logical Bandwidth Policers to the Logical Interfaces

Step-by-Step Procedure

To configure the logical bandwidth policers to the logical interfaces:

1. Enable configuration of the interface.

```
[edit]
user@host# edit interfaces ge-1/3/0
```

2. Apply the logical bandwidth policer to the first logical interface.

```
[edit interfaces ge-1/3/0]
user@host# set unit 0 family inet policer input LB-policer
user@host# set unit 0 family inet policer output LB-policer
```

3. Apply the policing to the second logical interface.

```
[edit interfaces ge-1/3/0]
user@host# set unit 1 family inet policer input LB-policer
user@host# set unit 1 family inet policer output LB-policer
```

Results Confirm the configuration of the interfaces by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
ge-1/3/0 {
  per-unit-scheduler;
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    family inet {
      policer {
        input LB-policer;
        output LB-policer;
      }
    }
    address 172.1.1.1/30;
  }
}
unit 1 {
  vlan-id 200;
  family inet {
    policer {
      input LB-policer;
      output LB-policer;
    }
  }
  address 172.2.1.1/30;
}
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Displaying Traffic Statistics and Policers for the Logical Interface on page 4861](#)
- [Displaying Statistics for the Policer on page 4862](#)

Displaying Traffic Statistics and Policers for the Logical Interface

Purpose Verify the traffic flow through the logical interface and that the policer is evaluated when packets are received on the logical interface.

Action Use the **show interfaces** operational mode command for logical interfaces **ge-1/3/0.0** and **ge-1/3/0.1**, and include the **detail** or **extensive** option. The command output section for **Traffic statistics** lists the number of bytes and packets received and transmitted on the logical interface, and the **Protocol inet** section contains a **Policer** field that lists the policer **LB-policer** as an input or output policer as follows:

- **Input:** LB-policer-ge-1/3/0.0-inet-i
- **Output:** LB-policer-ge-1/3/0.0-inet-o

In this example, the policer is applied to logical interface traffic in both the input and output directions.

```
user@host> show interfaces ge-1/3/0.0 detail
Logical interface ge-1/3/0.0 (Index 80) (SNMP ifIndex 154) (Generation 150)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ] Encapsulation: ENET2
  Traffic statistics:
    Input bytes : 0
    Output bytes : 46
    Input packets: 0
    Output packets: 1
  Local statistics:
    Input bytes : 0
    Output bytes : 46
    Input packets: 0
    Output packets: 1
  Transit statistics:
    Input bytes : 0 0 bps
    Output bytes : 0 0 bps
    Input packets: 0 0 pps
    Output packets: 0 0 pps
  Protocol inet, MTU: 1500, Generation: 174, Route table: 0
    Flags: Sendbcst-pkt-to-re
    Policer: Input: LB-policer-ge-1/3/0.0-inet-i, Output:
LB-policer-ge-1/3/0.0-inet-o
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 172.1.1.0/30, Local: 172.1.1.1, Broadcast: 172.1.1.3,
  Generation: 165
```

```
user@host> show interfaces ge-1/3/0.1 detail
Logical interface ge-1/3/0.1 (Index 81) (SNMP ifIndex 543) (Generation 151)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.200 ] Encapsulation: ENET2
  Traffic statistics:
    Input bytes : 0
    Output bytes : 46
```

```

      Input packets:          0
      Output packets:        1
Local statistics:
      Input bytes  :          0
      Output bytes :         46
      Input packets:          0
      Output packets:         1
Transit statistics:
      Input bytes  :          0          0 bps
      Output bytes :          0          0 bps
      Input packets:          0          0 pps
      Output packets:         0          0 pps
Protocol inet, MTU: 1500, Generation: 175, Route table: 0
Flags: Sendbroadcast-pkt-to-re
Policer: Input: LB-policer-ge-1/3/0.1-inet-i, Output:
LB-policer-ge-1/3/0.1-inet-o
Addresses, Flags: Is-Preferred Is-Primary
Destination: 172.2.1.0/30, Local: 172.2.1.1, Broadcast: 172.2.1.3,
Generation: 167

```

Displaying Statistics for the Policer

Purpose Verify the number of packets evaluated by the policer.

Action Use the [show policer](#) operational mode command and optionally specify the name of the policer. The command output displays the number of packets evaluated by each configured policer (or the specified policer), in each direction. For the policer **LB-policer**, the input and output policer names are displayed as follows:

- **LB-policer-ge-1/3/0.0-inet-i**
- **LB-policer-ge-1/3/0.0-inet-o**
- **LB-policer-ge-1/3/0.1-inet-i**
- **LB-policer-ge-1/3/0.1-inet-o**

The **-inet-i** suffix denotes a policer applied to logical interface input traffic, while the **-inet-o** suffix denotes a policer applied to logical interface output traffic. In this example, the policer is applied to both input and output traffic on logical interface **ge-1/3/0.0** and logical interface **ge-1/3/0.1**.

```

user@host> show policer
Policers:
Name                                     Packets
__default_arp_policer__                 0
LB-policer-ge-1/3/0.0-inet-i            0
LB-policer-ge-1/3/0.0-inet-o            0
LB-policer-ge-1/3/0.1-inet-i            0
LB-policer-ge-1/3/0.1-inet-o            0

```

- Related Documentation**
- [Statement Hierarchy for Configuring Policers on page 4811](#)
 - [Two-Color Policer Configuration Overview on page 4813](#)
 - [Guidelines for Applying Traffic Policers on page 4821](#)
 - [bandwidth-percent on page 4973](#)

- [interface-specific \(Firewall Filters\) on page 4793](#)
- [logical-bandwidth-policer on page 4993](#)
- [shaping-rate \(Applying to an Interface\) on page 1648](#)

Filter-Specific Counters and Policers

- [Filter-Specific Policer Overview on page 4863](#)
- [Example: Configuring a Stateless Firewall Filter to Protect Against TCP and ICMP Floods on page 4863](#)

Filter-Specific Policer Overview

By default, a policer operates in *term-specific* mode so that, for a given firewall filter, the Junos OS creates a separate policer instance for every filter term that references the policer. As an option, you can configure a policer to operate in *filter-specific* mode so that a single policer instance is used by all terms (within the same firewall filter) that reference the policer.

For an IPv4 firewall filter with multiple terms that reference the same policer, configuring the policer to operate in filter-specific mode enables you to count and monitor the activity of the policer at the firewall filter level.



NOTE: Term-specific mode and filter-specific mode also apply to prefix-specific policer sets.

To enable a single-rate two-color policer to operate in filter-specific mode, you can include the **filter-specific** statement at the following hierarchy levels:

- **[edit firewall policer *policer-name*]**
- **[edit logical-systems *logical-system-name* firewall policer *policer-name*]**

You can reference filter-specific policers from IPv4 (**family inet**) firewall filters only.

Example: Configuring a Stateless Firewall Filter to Protect Against TCP and ICMP Floods

This example shows how to create a stateless firewall filter that protects against TCP and ICMP denial-of-service attacks.

- [Requirements on page 4863](#)
- [Overview on page 4864](#)
- [Configuration on page 4865](#)
- [Verification on page 4869](#)

Requirements

No special configuration beyond device initialization is required before configuring stateless firewall filters.

Overview

In this example, you create a stateless firewall filter called **protect-RE** that polices TCP and ICMP packets. This example includes the following policers:

- **tcp-connection-policer**—Limits the traffic rate of the TCP packets to 500,000 bps and the burst size to 15,000 bytes. Packets that exceed the traffic rate are discarded.
- **icmp-policer**—Limits the traffic rate of the ICMP packets to 1,000,000 bps and the burst size to 15,000 bytes. Packets that exceed the traffic rate are discarded.

When specifying limits, the bandwidth limit can be from 32,000 bps to 32,000,000,000 bps and the burst-size limit can be from 1,500 bytes through 100,000,000 bytes. Use the following abbreviations when specifying limits: k (1,000), m (1,000,000), and g (1,000,000,000).

Each policer is incorporated into the action of a filter term. This example includes the following terms:

- **tcp-connection-term**—Policies certain TCP packets with a source address of 192.168.0.0/24 or 10.0.0.0/24. These addresses are defined in the **trusted-addresses** prefix list.

Filtered packets include **tcp-established** packets. The **tcp-established** match condition is an alias for the bit-field match condition **tcp-flags "(ack | rst)"**, which indicates an established TCP session, but not the first packet of a TCP connection.

- **icmp-term**—Policies ICMP packets. All ICMP packets are counted in the **icmp-counter** counter.




NOTE: You can move terms within the firewall filter by using the **insert** command. See *insert* in the *CLI User Guide*.

You can apply a stateless firewall to the input or output sides, or both, of an interface. To filter packets transiting the device, apply the firewall filter to any non-Routing Engine interface. To filter packets originating from, or destined for, the Routing Engine, apply the firewall filter to the loopback (lo0) interface.

Figure 95 on page 4601 shows the sample network.

Figure 109: Firewall Filter to Protect Against TCP and ICMP Floods



Because this firewall filter limits Routing Engine traffic to TCP packets, routing protocols that use other transport protocols for Layer 4 cannot successfully establish sessions when this filter is active. To demonstrate, this example sets up OSPF between Device R1 and Device R2.

“CLI Quick Configuration” on page 4602 shows the configuration for all of the devices in Figure 95 on page 4601.

The section “Step-by-Step Procedure” on page 4603 describes the steps on Device R2.

Configuration

CLI Quick Configuration To quickly configure the stateless firewall filter, copy the following commands to a text file, remove any line breaks, and then paste the commands into the CLI.

Device R1

```

set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces lo0 unit 0 family inet address 192.168.0.1/32 primary
set interfaces lo0 unit 0 family inet address 172.16.0.1/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext peer-as 200
set protocols bgp group ext neighbor 10.0.0.2
set protocols ospf area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.1
set routing-options autonomous-system 100

```

Device R2

```

set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.2/30
set interfaces lo0 unit 0 family inet filter input protect-RE
set interfaces lo0 unit 0 family inet address 192.168.0.2/32 primary
set interfaces lo0 unit 0 family inet address 172.16.0.2/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext neighbor 10.0.0.1 peer-as 100
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.0
set policy-options prefix-list trusted-addresses 10.0.0.0/24
set policy-options prefix-list trusted-addresses 192.168.0.0/24
set policy-options policy-statement send-direct term 1 from protocol direct

```

```
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.2
set routing-options autonomous-system 200
set firewall family inet filter protect-RE term tcp-connection-term from source-prefix-list
trusted-addresses
set firewall family inet filter protect-RE term tcp-connection-term from protocol tcp
set firewall family inet filter protect-RE term tcp-connection-term from tcp-established
set firewall family inet filter protect-RE term tcp-connection-term then policer
tcp-connection-policer
set firewall family inet filter protect-RE term tcp-connection-term then accept
set firewall family inet filter protect-RE term icmp-term from source-prefix-list
trusted-addresses
set firewall family inet filter protect-RE term icmp-term from protocol icmp
set firewall family inet filter protect-RE term icmp-term then policer icmp-policer
set firewall family inet filter protect-RE term icmp-term then count icmp-counter
set firewall family inet filter protect-RE term icmp-term then accept
set firewall policer tcp-connection-policer filter-specific
set firewall policer tcp-connection-policer if-exceeding bandwidth-limit 1m
set firewall policer tcp-connection-policer if-exceeding burst-size-limit 15k
set firewall policer tcp-connection-policer then discard
set firewall policer icmp-policer filter-specific
set firewall policer icmp-policer if-exceeding bandwidth-limit 1m
set firewall policer icmp-policer if-exceeding burst-size-limit 15k
set firewall policer icmp-policer then discard
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure stateless firewall filter policers:

1. Configure the device interfaces.

```
[edit interfaces fe-1/2/0 unit 0 family inet ]
user@R2# set address 10.0.0.2/30
```

```
[edit interfaces lo0 unit 0 family inet]
user@R2# set address 192.168.0.2/32 primary
user@R2# set address 172.16.0.2/32
```

2. Configure the BGP peering session.

```
[edit protocols bgp group ext]
user@R2# set type external
user@R2# set export send-direct
user@R2# set neighbor 10.0.0.1 peer-as 100
```

3. Configure the autonomous system (AS) number and router ID.

```
[edit routing-options]
user@R2# set autonomous-system 200
user@R2# set router-id 192.168.0.2
```

4. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R2# set interface lo0.0 passive
```



```
user@R2# set interface fe-1/2/0.0
```

5. Define the list of trusted addresses.

```
[edit policy-options prefix-list trusted-addresses]
user@R2# set 10.0.0.0/24
user@R2# set 192.168.0.0/24
```

6. Configure a policy to advertise direct routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R2# set from protocol direct
user@R2# set then accept
```

7. Configure the TCP policer.

```
[edit firewall policer tcp-connection-policer]
user@R2# set filter-specific
user@R2# set if-exceeding bandwidth-limit 1m
user@R2# set if-exceeding burst-size-limit 15k
user@R2# set then discard
```

8. Create the ICMP policer.

```
[edit firewall policer icmp-policer]
user@R2# set filter-specific
user@R2# set if-exceeding bandwidth-limit 1m
user@R2# set if-exceeding burst-size-limit 15k
user@R2# set then discard
```

9. Configure the TCP filter rules.

```
[edit firewall family inet filter protect-RE term tcp-connection-term]
user@R2# set from source-prefix-list trusted-addresses
user@R2# set from protocol tcp
user@R2# set from tcp-established
user@R2# set then policer tcp-connection-policer
user@R2# set then accept
```

10. Configure the ICMP filter rules.

```
[edit firewall family inet filter protect-RE term icmp-term]
user@R2# set from source-prefix-list trusted-addresses
user@R2# set from protocol icmp
user@R2# set then policer icmp-policer
user@R2# set then count icmp-counter
user@R2# set then accept
```

11. Apply the filter to the loopback interface.

```
[edit interfaces lo0 unit 0]
user@R2# set family inet filter input protect-RE
```

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

```
user@R2# show interfaces
```

```
fe-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.2/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      filter {
        input protect-RE;
      }
      address 192.168.0.2/32 {
        primary;
      }
      address 172.16.0.2/32;
    }
  }
}

user@R2# show protocols
bgp {
  group ext {
    type external;
    export send-direct;
    neighbor 10.0.0.1 {
      peer-as 100;
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.0 {
      passive;
    }
    interface fe-1/2/0.0;
  }
}

user@R2# show policy-options
prefix-list trusted-addresses {
  10.0.0.0/24;
  192.168.0.0/24;
}
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R2# show routing-options
router-id 192.168.0.2;
autonomous-system 200;

user@R2# show firewall
```

```

family inet {
  filter protect-RE {
    term tcp-connection-term {
      from {
        source-prefix-list {
          trusted-addresses;
        }
        protocol tcp;
        tcp-established;
      }
      then {
        policer tcp-connection-policer;
        accept;
      }
    }
    term icmp-term {
      from {
        source-prefix-list {
          trusted-addresses;
        }
        protocol icmp;
      }
      then {
        policer icmp-policer;
        count icmp-counter;
        accept;
      }
    }
  }
}
policer tcp-connection-policer {
  filter-specific;
  if-exceeding {
    bandwidth-limit 1m;
    burst-size-limit 15k;
  }
  then discard;
}
policer icmp-policer {
  filter-specific;
  if-exceeding {
    bandwidth-limit 1m;
    burst-size-limit 15k;
  }
  then discard;
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.



NOTE: To verify the TCP policer, you can use a packet generation tool. This task is not shown here.

- [Displaying Stateless Firewall Filter That Are in Effect on page 4870](#)
- [Using telnet to Verify the tcp-established Condition in the TCP Firewall Filter on page 4870](#)
- [Using telnet to Verify the Trusted Prefixes Condition in the TCP Firewall Filter on page 4871](#)
- [Using OSPF to Verify the TCP Firewall Filter on page 4872](#)
- [Verifying the ICMP Firewall Filter on page 4873](#)

Displaying Stateless Firewall Filter That Are in Effect

Purpose Verify the configuration of the firewall filter.

Action From operational mode, enter the **show firewall** command.

```
user@R2> show firewall
Filter: protect-RE
Counters:
Name                               Bytes      Packets
icmp-counter                        0           0
Policers:
Name                               Bytes      Packets
icmp-policer                       0           0
tcp-connection-policer             0           0
```

Meaning The output shows the filter, the counter, and the policers that are in effect on Device R2.

Using telnet to Verify the tcp-established Condition in the TCP Firewall Filter

Purpose Make sure that telnet traffic works as expected.

Action Verify that the device can establish only TCP sessions with hosts that meet the **from tcp-established** condition..

1. From Device R2, make sure that the BGP session with Device R1 is established.

```
user@R2> show bgp summary | match down
Groups: 1 Peers: 1 Down peers: 0
```

2. From Device R2, telnet to Device R1.

```
user@R2> telnet 192.168.0.1
Trying 192.168.0.1...
Connected to R1.acme.net.
Escape character is '^['.
```

```
R1 (ttyp4)
```

```
login:
```

3. From Device R1, telnet to Device R2.

```
user@R1> telnet 192.168.0.2
```

```
Trying 192.168.0.2...
telnet: connect to address 192.168.0.2: Operation timed out
telnet: Unable to connect to remote host
```

4. On Device R2, deactivate the **from tcp-established** match condition.

```
[edit firewall family inet filter protect-RE term tcp-connection-term]
user@R2# deactivate from tcp-established
user@R2# commit
```

5. From Device R1, try again to telnet to Device R2.

```
user@R1> telnet 192.168.0.1
Trying 192.168.0.2...
Connected to R2.acme.net.
Escape character is '^['.
```

```
R2 (ttyp4)
```

```
login:
```

Meaning Verify the following information:

- As expected, the BGP session is established. The **from tcp-established** match condition is not expected to block BGP session establishment.
- From Device R2, you can telnet to Device R1. Device R1 has no firewall filter configured, so this is the expected behavior.
- From Device R1, you cannot telnet to Device R2. Telnet uses TCP as the transport protocol, so this result might be surprising. The cause for the lack of telnet connectivity is the **from tcp-established** match condition. This match condition limits the type of TCP traffic that is accepted of Device R2. After this match condition is deactivated, the telnet session is successful.

Using telnet to Verify the Trusted Prefixes Condition in the TCP Firewall Filter

Purpose Make sure that telnet traffic works as expected.

Action Verify that the device can establish only telnet sessions with a host at an IP address that matches one of the trusted source addresses. For example, log in to the device with the **telnet** command from another host with one of the trusted address prefixes. Also, verify that telnet sessions with untrusted source addresses are blocked.

1. From Device R1, telnet to Device R2 from an untrusted source address.

```
user@R1> telnet 172.16.0.2 source 172.16.0.1
Trying 172.16.0.2...
^C
```

2. From Device R2, add 172.16/16 to the list of trusted prefixes.

```
[edit policy-options prefix-list trusted-addresses]
user@R2# set 172.16.0.0/16
user@R2# commit
```

3. From Device R1, try again to telnet to Device R2.

```
user@R1> telnet 172.16.0.2 source 172.16.0.1
```

```
Trying 172.16.0.2...
Connected to R2.acme.net.
Escape character is '^['.
```

```
R2 (ttyp4)
```

```
login:
```

Meaning Verify the following information:

- From Device R1, you cannot telnet to Device R2 with an untrusted source address. After the 172.16/16 prefix is added to the list of trusted prefixes, the telnet request from source address 172.16.0.1 is accepted.
- OSPF session establishment is blocked. OSPF does not use TCP as its transport protocol. After the **from protocol tcp** match condition is deactivated, OSPF session establishment is not blocked.

Using OSPF to Verify the TCP Firewall Filter

Purpose Make sure that OSPF traffic works as expected.

Action Verify that the device cannot establish OSPF connectivity.

1. From Device R1, check the OSPF sessions.

```
user@R1> show ospf neighbor
Address      Interface      State    ID          Pri  Dead
10.0.0.2     fe-1/2/0.0    Init    192.168.0.2 128   34
```

2. From Device R2, check the OSPF sessions.

```
user@R2> show ospf neighbor
```

3. From Device R2, remove the **from protocol tcp** match condition.

```
[edit firewall family inet filter protect-RE term tcp-connection-term]
user@R2# deactivate from protocol
user@R2# commit
```

4. From Device R1, recheck the OSPF sessions.

```
user@R1> show ospf neighbor
Address      Interface      State    ID          Pri  Dead
10.0.0.2     fe-1/2/0.0    Full    192.168.0.2 128   36
```

5. From Device R2, recheck the OSPF sessions.

```
user@R2> show ospf neighbor
Address      Interface      State    ID          Pri  Dead
10.0.0.1     fe-1/2/0.0    Full    192.168.0.1 128   39
```

Meaning Verify the following information:

- ## Verifying the ICMP Firewall Filter

Action 1. Undo the configuration changes made in previous verification steps.

2. From Device R1, ping the loopback interface on Device R2.

```
pinground-trip min/avg/max/stddev = 2.976/3.405/42.380/2.293 ms
```

- 0

- 14 packets transmitted, 0 packets received, 100% packet loss

4873

- The ICMP packet counter is incrementing, and the icmp-policer is incrementing.
- Device R2 does not send ICMP responses to the **ping 172.16.0.2 source 172.16.0.1** command.

Related Documentation

- [Statement Hierarchy for Configuring Policers on page 4811](#)
- [Two-Color Policer Configuration Overview on page 4813](#)
- [Guidelines for Applying Traffic Policers on page 4821](#)
- [Prefix-Specific Counting and Policing Actions on page 4874](#)

Prefix-Specific Counting and Policing Actions

- [Prefix-Specific Counting and Policing Overview on page 4874](#)
- [Filter-Specific Counter and Policer Set Overview on page 4876](#)
- [Example: Configuring Prefix-Specific Counting and Policing on page 4877](#)
- [Prefix-Specific Counting and Policing Configuration Scenarios on page 4883](#)

Prefix-Specific Counting and Policing Overview

This topic covers the following information:

- [Separate Counting and Policing for Each IPv4 Address Range on page 4874](#)
- [Prefix-Specific Action Configuration on page 4875](#)
- [Counter and Policer Set Size and Indexing on page 4876](#)

Separate Counting and Policing for Each IPv4 Address Range

Prefix-specific counting and policing enables you to configure an IPv4 firewall filter term that matches on a source or destination address, applies a single-rate two-color policer as the term action, but associates the matched packet with a specific counter and policer instance based on the source or destination in the packet header. You can implicitly create a separate counter or policer instance for a single address or for a group of addresses.



NOTE: J Series Services Routers do not support prefix-specific counting and policing.

Prefix-specific counting and policing uses a *prefix-specific action* configuration that specifies the name of the policer you want to apply, whether prefix-specific counting is to be enabled, and a source or destination address prefix range.

The prefix range specifies between 1 and 16 sequential set bits of an IPv4 address mask. The length of the prefix range determines the size of the counter and policer set, which consists of as few as 2 or as many as 65,536 counter and policer instances. The position of the bits of the prefix range determines the indexing of filter-matched packets into the set of instances.



NOTE: A prefix-specific action is specific to a source or destination *prefix range*, but it is not specific to a particular source or destination *address range*, and it is not specific to a particular interface.

To apply a prefix-specific action to the traffic at an interface, you configure a firewall filter term that matches on source or destination addresses, and then you apply the firewall filter to the interface. The flow of filtered traffic is rate-limited using prefix-specific counter and policer instances that are selected per packet based on the source or destination address in the header of the filtered packet.

Prefix-Specific Action Configuration

To configure a prefix-specific action, you specify the following information:

- Prefix-specific action name—Name that can be referenced as the action of an IPv4 standard firewall filter term that matches packets on source or destination addresses.
- Policer name—Name of a single-rate two-color policer for which you want to implicitly create prefix-specific instances.



NOTE: For aggregated Ethernet interfaces, you can configure a prefix-specific action that references a logical interface policer (also called an aggregate policer). You can reference this type of prefix-specific action from an IPv4 standard firewall filter and then apply the filter at the aggregate level of the interface.

- Counting option—Option to include if you want to enable prefix-specific counters.
- Filter-specific option—Option to include if you want a single counter and policer set to be shared across all terms in the firewall filter. A prefix-specific action that operates in this way is said to operate in *filter-specific* mode. If you do not enable this option, the prefix-specific action operates in *term-specific* mode, meaning that a separate counter and policer set is created for each filter term that references the prefix-specific action.
- Source address prefix length—Length of the address prefix, from 0 through 32, to be used with a packet matched on the source address.
- Destination address prefix length—Length of the address prefix, from 0 through 32, to be used with a packet matched on the destination address.
- Subnet prefix length—Length of the subnet prefix, from 0 through 32, to be used with a packet matched on either the source or destination address.

You must configure source and destination address prefix lengths to be from 1 to 16 bits longer than the subnet prefix length. If you configure source or destination address prefix lengths to be more than 16 bits beyond the configured subnet prefix length, an error occurs when you try to commit the configuration.

Counter and Policer Set Size and Indexing

The number of prefix-specific actions (counters or policers) implicitly created for a prefix-specific action is determined by the length of the address prefix and the length of the subnet prefix:

$$\text{Size of Counter and Policer Set} = 2^{(\text{source-or-destination-prefix-length} - \text{subnet-prefix-length})}$$

Table 367 on page 4876 shows examples of counter and policer set size and indexing.

Table 367: Examples of Counter and Policer Set Size and Indexing

Example Prefix Lengths Specified in the Prefix-Specific Action	Calculation of Counter or Policer Set Size	Indexing of Instances	
$\text{source-prefix-length} = 32$ $\text{subnet-prefix-length} = 16$	Size = $2^{(32-16)} = 2^{16} = 65,536$ instances NOTE: This calculation shows the largest counter or policer set size supported.	Instance 0:	x.x.0.0
		Instance 1:	x.x.0.1
		Instance 65535:	x.x.255.255
$\text{source-prefix-length} = 32$ $\text{subnet-prefix-length} = 24$	Size = $2^{(32-24)} = 2^8 = 256$ instances	Instance 0:	x.x.x.0
		Instance 1:	x.x.x.1
		Instance 255:	x.x.x.255
$\text{source-prefix-length} = 32$ $\text{subnet-prefix-length} = 25$	Size = $2^{(32-25)} = 2^7 = 128$ instances	Instance 0:	x.x.x.0
		Instance 1:	x.x.x.1
		Instance 127:	x.x.x.127
$\text{source-prefix-length} = 24$ $\text{subnet-prefix-length} = 20$	Size = $2^{(24-20)} = 2^4 = 16$ instances	Instance 0:	x.x.0.x
		Instance 1:	x.x.1.x
		Instance 15:	x.x.15.x

Filter-Specific Counter and Policer Set Overview

By default, a prefix-specific policer set operates in *term-specific* mode so that, for a given firewall filter, the Junos OS creates a separate counter and policer set for every filter term that references the prefix-specific action. As an option, you can configure a prefix-specific policer set to operate in *filter-specific* mode so that a single prefix-specific policer set is used by all terms (within the same firewall filter) that reference the policer.

For an IPv4 firewall filter with multiple terms that reference the same prefix-specific policer set, configuring the policer set to operate in filter-specific mode enables you to count and monitor the activity of the policer set at the firewall filter level.



NOTE: Term-specific mode and filter-specific mode also apply to policers. See [“Filter-Specific Policer Overview” on page 4863](#).

To enable a prefix-specific policer set to operate in filter-specific mode, you can include the **filter-specific** statement at following the hierarchy levels:

- [edit firewall family inet prefix-action *prefix-action-name*]
- [edit logical-systems *logical-system-name* firewall family inet prefix-action *prefix-action-name*]

You can reference filter-specific, prefix-specific policer sets from IPv4 (**family inet**) firewall filters only.

Example: Configuring Prefix-Specific Counting and Policing

This example shows how to configure prefix-specific counting and policing.

- [Requirements on page 4877](#)
- [Overview on page 4877](#)
- [Configuration on page 4878](#)
- [Verification on page 4882](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, you configure prefix-specific counting and policing based on the last octet of the source address field in packets matched by an IPv4 firewall filter.

The single-rate two-color policer named **1Mbps-policer** rate-limits traffic to a bandwidth of 1,000,000 bps and a burst-size limit of 63,000 bytes, discarding any packets in a traffic flow that exceeds the traffic limits.

Independent of the IPv4 addresses contained in any packets passed from a firewall filter, the prefix-specific action named **psa-1Mbps-per-source-24-32-256** specifies a set of 256 counters and policers, numbered from 0 through 255. For each packet, the last octet of the source address field is used to index into the associated prefix-specific counter and policer in the set:

- Packets with a source address ending with the octet 0x0000 0000 index the first counter and policer in the set.
- Packets with a source address ending with the octet 0x0000 0001 index the second counter and policer in the set.
- Packets with a source address ending with the octet 0x1111 1111 index the last counter and policer in the set.

The **limit-source-one-24** firewall filter contains a single term that matches all packets from the /24 subnet of source address **10.10.10.0**, passing these packets to the prefix-specific action **psa-1Mbps-per-source-24-32-256**.

Topology

In this example, because the filter term matches the /24 subnet of a single source address, each counting and policing instance in the prefix-specific set is used for only one source address.

- Packets with a source address **10.10.10.0** index the first counter and policer in the set.
- Packets with a source address **10.10.10.1** index the second counter and policer in the set.
- Packets with a source address **10.10.10.255** index the last counter and policer in the set.

This example shows the simplest case of prefix-specific actions, in which the filter term matches on one address with a prefix length that is the same as the prefix length specified in the prefix-specific action for indexing into the set of prefix-specific counters and policers.

For descriptions of other configurations for prefix-specific counting and policing, see [“Prefix-Specific Counting and Policing Configuration Scenarios” on page 4883](#).

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configuring a Policer for Prefix-Specific Counting and Policing on page 4879](#)
- [Configuring a Prefix-Specific Action Based on the Policer on page 4879](#)
- [Configuring an IPv4 Filter That References the Prefix-Specific Action on page 4880](#)
- [Applying the Firewall Filter to IPv4 Input Traffic at a Logical Interface on page 4881](#)

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set firewall policer 1Mbps-policer if-exceeding bandwidth-limit 1m
set firewall policer 1Mbps-policer if-exceeding burst-size-limit 63k
set firewall policer 1Mbps-policer then discard
set firewall family inet prefix-action psa-1Mbps-per-source-24-32-256 policer
  1Mbps-policer
set firewall family inet prefix-action psa-1Mbps-per-source-24-32-256 count
set firewall family inet prefix-action psa-1Mbps-per-source-24-32-256
  subnet-prefix-length 24
set firewall family inet prefix-action psa-1Mbps-per-source-24-32-256 source-prefix-length
  32
set firewall family inet filter limit-source-one-24 term one from source-address
  10.10.10.0/24
```

```

set firewall family inet filter limit-source-one-24 term one then prefix-action
  psa-1Mbps-per-source-24-32-256
set interfaces so-0/0/2 unit 0 family inet filter input limit-source-one-24
set interfaces so-0/0/2 unit 0 family inet address 10.39.1.1/16

```

Configuring a Policer for Prefix-Specific Counting and Policing

Step-by-Step Procedure

To configure a policer to be used for prefix-specific counting and policing:

1. Enable configuration of a single-rate two-color policer.

```

[edit]
user@host# edit firewall policer 1Mbps-policer

```

2. Define the traffic limit.

```

[edit firewall policer 1Mbps-policer]
user@host# set if-exceeding bandwidth-limit 1m
user@host# set if-exceeding burst-size-limit 63k

```

Packets in a traffic flow that conforms to this limit are passed with the PLP set to **low**.

3. Define the actions for nonconforming traffic.

```

[edit firewall policer 1Mbps-policer]
user@host# set then discard

```

Packets in a traffic flow that exceeds this limit are discarded. Other configurable actions for a single-rate two-color policer are to set the forwarding class and to set the PLP level.

Results Confirm the configuration of the policer by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```

[edit]
user@host# show firewall
policer 1Mbps-policer {
  if-exceeding {
    bandwidth-limit 1m;
    burst-size-limit 63k;
  }
  then discard;
}

```

Configuring a Prefix-Specific Action Based on the Policer

Step-by-Step Procedure

To configure a prefix-specific action that references the policer and specifies a portion of a source address prefix:

1. Enable configuration of a prefix-specific action.

```

[edit]
user@host# edit firewall family inet prefix-action psa-1Mbps-per-source-24-32-256

```

Prefix-specific counting and policing can be defined for IPv4 traffic only.

2. Reference the policer for which a prefix-specific set is to be created.

```
[edit firewall family inet prefix-action psa-1Mbps-per-source-24-32-256]
user@host# set policer 1Mbps-policer
user@host# set count
```



NOTE: For aggregated Ethernet interfaces, you can configure a prefix-specific action that references a logical interface policer (also called an aggregate policer). You can reference this type of prefix-specific action from an IPv4 standard firewall filter and then apply the filter at the aggregate level of the interface.

3. Specify the prefix range on which IPv4 addresses are to be indexed to the counter and policer set.

```
[edit firewall family inet prefix-action psa-1Mbps-per-source-24-32-256]
user@host# set source-prefix-length 32
user@host# set subnet-prefix-length 24
```

Results Confirm the configuration of the prefix-specific action by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
policer 1Mbps-policer {
  if-exceeding {
    bandwidth-limit 1m;
    burst-size-limit 63k;
  }
  then discard;
}
family inet {
  prefix-action psa-1Mbps-per-source-24-32-256 {
    policer 1Mbps-policer;
    subnet-prefix-length 24;
    source-prefix-length 32;
  }
}
```

Configuring an IPv4 Filter That References the Prefix-Specific Action

Step-by-Step Procedure To configure an IPv4 standard firewall filter that references the prefix-specific action:

1. Enable configuration of the IPv4 standard firewall filter.

```
[edit]
user@host# edit firewall family inet filter limit-source-one-24
```

Prefix-specific counting and policing can be defined for IPv4 traffic only.

2. Configure the filter term to match on the packet source address or destination address.

```
[edit firewall family inet filter limit-source-one-24]
user@host# set term one from source-address 10.10.10.0/24
```

3. Configure the filter term to reference the prefix-specific action.

```
[edit firewall family inet filter limit-source-one-24]
user@host# set term one then prefix-action psa-1Mbps-per-source-24-32-256
```

You could also use the **next term** action to configure all Hypertext Transfer Protocol (HTTP) traffic to each host to transmit at 500 Kbps and have the total HTTP traffic limited to 1 Mbps.

Results Confirm the configuration of the prefix-specific action by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
policer 1Mbps-policer {
  if-exceeding {
    bandwidth-limit 1m;
    burst-size-limit 63k;
  }
  then discard;
}
family inet {
  prefix-action psa-1Mbps-per-source-24-32-256 {
    policer 1Mbps-policer;
    subnet-prefix-length 24;
    source-prefix-length 32;
  }
  filter limit-source-one-24 {
    term one {
      from {
        source-address {
          10.10.10.0/24;
        }
      }
      then prefix-action psa-1Mbps-per-source-24-32-256;
    }
  }
}
```

Applying the Firewall Filter to IPv4 Input Traffic at a Logical Interface

Step-by-Step Procedure To apply the firewall filter to IPv4 input traffic at a logical interface:

1. Enable configuration of IPv4 on the logical interface.

```
[edit]
user@host# edit interfaces so-0/0/2 unit 0 family inet
```

2. Configure an IP address.

```
[edit interfaces so-0/0/2 unit 0 family inet]
user@host# set address 10.39.1.1/16
```

3. Apply the IPv4 standard stateless firewall filter.

```
[edit interfaces so-0/0/2 unit 0 family inet]
user@host# set filter input limit-source-one-24
```

Results Confirm the configuration of the prefix-specific action by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
so-0/0/2 {
  unit 0 {
    family inet {
      filter {
        input limit-source-one-24;
      }
      address 10.39.1.1/16;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Displaying the Firewall Filters Applied to an Interface on page 4882](#)
- [Displaying Prefix-Specific Actions Statistics for the Firewall Filter on page 4883](#)

Displaying the Firewall Filters Applied to an Interface

Purpose Verify that the firewall filter **limit-source-one-24** is applied to the IPv4 input traffic at logical interface **so-0/0/2.0**.

Action Use the **show interfaces statistics** operational mode command for logical interface **so-0/0/2.0**, and include the **detail** option. In the command output section for **Protocol inet**, the **Input Filters** field displays **limit-source-one-24**, indicating that the filter is applied to IPv4 traffic in the input direction:

```
user@host> show interfaces statistics so-0/0/2.0 detail
Logical interface so-0/0/2.0 (Index 79) (SNMP ifIndex 510) (Generation 149)
Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
Protocol inet, MTU: 4470, Generation: 173, Route table: 0
Flags: Sendbcast-pkt-to-re, Protocol-Down
Input Filters: limit-source-one-24
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 10.39/16, Local: 10.39.1.1, Broadcast: 10.39.255.255,
Generation: 163
```


Displaying Prefix-Specific Actions Statistics for the Firewall Filter

Purpose Verify the number of packets evaluated by the policer.

Action Use the `show firewall prefix-action-stats filter filter-name prefix-action name` operational mode command to display statistics about a prefix-specific action configured on a firewall filter.

As an option, you can use the `from set-index to set-index` command option to specify the starting and ending counter or policer to be displayed. A policer set is indexed from 0 through 65535.

The command output displays the specified filter name followed by a listing of the number of bytes and packets processed by each policer in the policer set.

For a term-specific policer, each policer in the set is identified as follows:

prefix-specific-action-name-term-name-set-index

For a filter-specific policer, each policer is identified in the command output as follows:

prefix-specific-action-name-set-index

Because the example prefix-specific action `psa-1Mbps-per-source-24-32-256` is referenced by only one term of the example filter `limit-source-one-24`, the example policer `1Mbps-policer` is configured as term-specific. In the `show firewall prefix-action-stats` command output, the policer statistics are displayed as `psa-1Mbps-per-source-24-32-256-one-0`, `psa-1Mbps-per-source-24-32-256-one-1`, and so on through `psa-1Mbps-per-source-24-32-256-one-255`.

```
user@host> show firewall prefix-action-stats filter limit-source-one-24 prefix-action
psa-1Mbps-per-source-24-32-256 from 0 to 9
```

```
Filter: limit-source-one-24
```

```
Counters:
```

Name	Bytes	Packets
psa-1Mbps-per-source-24-32-256-one-0	0	0
psa-1Mbps-per-source-24-32-256-one-1	0	0
psa-1Mbps-per-source-24-32-256-one-2	0	0
psa-1Mbps-per-source-24-32-256-one-3	0	0
psa-1Mbps-per-source-24-32-256-one-4	0	0
psa-1Mbps-per-source-24-32-256-one-5	0	0
psa-1Mbps-per-source-24-32-256-one-6	0	0
psa-1Mbps-per-source-24-32-256-one-7	0	0
psa-1Mbps-per-source-24-32-256-one-8	0	0
psa-1Mbps-per-source-24-32-256-one-9	0	0

Prefix-Specific Counting and Policing Configuration Scenarios

This topic covers the following information:

- [Prefix Length of the Action and Prefix Length of Addresses in Filtered Packets on page 4884](#)
- [Scenario 1: Firewall Filter Term Matches on Multiple Addresses on page 4885](#)

- [Scenario 2: Subnet Prefix Is Longer Than the Prefix in the Filter Match Condition on page 4886](#)
- [Scenario 3: Subnet Prefix Is Shorter Than the Prefix in the Firewall Filter Match Condition on page 4888](#)

Prefix Length of the Action and Prefix Length of Addresses in Filtered Packets

[Table 368 on page 4884](#) describes the relationship between the prefix length specified in the prefix-specific action and the prefix length of the addresses matched by the firewall filter term that references the prefix-specific action.

Table 368: Summary of Prefix-Specific Action Scenarios

Counter and Policer Set	Packet-Filtering Criteria	Indexing of Instances	
Prefix-specific action scenario: "Example: Configuring Prefix-Specific Counting and Policing" on page 4877			
<i>source-prefix-length</i> = 32 <i>subnet-prefix-length</i> = 24 Set size: 2^8 = 256 Instance numbers: 0 - 255	<i>source-address</i> = 10.10.10.0/24	Instance 0	10.10.10.0
		Instance 1:	10.10.10.1
	
		Instance 255:	10.10.10.255
Prefix-specific action scenario: "Scenario 1: Firewall Filter Term Matches on Multiple Addresses" on page 4885			
<i>source-prefix-length</i> = 32 <i>subnet-prefix-length</i> = 24 Set size: 2^8 = 256 Instance numbers: 0 - 255	<i>source-address</i> = 10.10.10.0/24	Instance 0	10.10.10.0, 10.11.x.0
	<i>source-address</i> = 10.11.0.0/16	Instance 1:	10.10.10.1, 10.11.x.1
	
		Instance 255:	10.10.10.255, 10.11.x.255
For addresses in the /16 subnet, x ranges from 0 through 255.			
Prefix-specific action scenario: "Scenario 2: Subnet Prefix Is Longer Than the Prefix in the Filter Match Condition" on page 4886			
<i>source-prefix-length</i> = 32 <i>subnet-prefix-length</i> = 25 Set size: 2^7 = 128 Instance numbers: 0 - 127	<i>source-address</i> = 10.10.10.0/24	Instance 0	10.10.10.0, 10.10.10.128
		Instance 1:	10.10.10.1, 10.10.10.120
	

Table 368: Summary of Prefix-Specific Action Scenarios (*continued*)

Counter and Policer Set	Packet-Filtering Criteria	Indexing of Instances	
		Instance 127:	10.10.10.255, 10.10.10.127
Prefix-specific action scenario: "Scenario 3: Subnet Prefix Is Shorter Than the Prefix in the Firewall Filter Match Condition" on page 4888			
<i>source-prefix-length</i> = 32 <i>subnet-prefix-length</i> = 24 Set size: $2^8 = 256$ Instance numbers: 0 - 255	<i>source-address</i> = 10.10.10.0/25 NOTE: Only packets with source addresses ranging from 10.10.10.0 through 10.10.10.127 are passed to the prefix-specific action.	Instance 0	10.10.10.0
		Instance 1:	10.10.10.1
	
		Instance 127:	10.10.10.127
		Instances 128 – 255: unused	

Scenario 1: Firewall Filter Term Matches on Multiple Addresses

The complete example, "Example: Configuring Prefix-Specific Counting and Policing" on page 4877, shows the simplest case of prefix-specific actions, in which a single-term firewall filter matches on one address with a prefix length that is the same as the subnet prefix length specified in the prefix-specific action. Unlike the example, this scenario describes a configuration in which a single-term firewall filter matches on two IPv4 source addresses. In addition, the additional condition matches on a source address with a prefix length that is different from the subnet prefix length defined in the prefix-specific action. In this case, the additional condition matches on the /16 subnet of the source address 10.11.0.0.



NOTE: Unlike packets that match the source address 10.10.10.0/24, packets that match the source address 10.11.0.0/16 are in a many-to-one correspondence with the instances in the counter and policer set.

The filter-matched packets that are passed to the prefix-specific action index into the counter and policer set in such a way that the counting and policing instances are shared by packets that contain source addresses across the 10.10.10.0/24 and 10.11.0.0/16 subnets as follows:

- The first counter and policer in the set are indexed by packets with source addresses 10.10.10.0 and 10.11.x.0, where x ranges from 0 through 255.
- The second counter and policer in the set are indexed by packets with source addresses 10.10.10.1 and 10.11.x.1, where x ranges from 0 through 255.
- The 256th (last) counter and policer in the set are indexed by packets with source addresses 10.10.10.255 and 10.11.x.255, where x ranges from 0 through 255.

The following configuration shows the statements for configuring the single-rate two-color policer, the prefix-specific action that references the policer, and the IPv4 standard stateless firewall filter that references the prefix-specific action:

```
[edit]
firewall {
  policer 1Mbps-policer {
    if-exceeding {
      bandwidth-limit 1m;
      burst-size-limit 63k;
    }
    then discard;
  }
  family inet {
    prefix-action psa-1Mbps-per-source-24-32-256 {
      policer 1Mbps-policer;
      subnet-prefix-length 24;
      source-prefix-length 32;
    }
    filter limit-source-two-24-16 {
      term one {
        from {
          source-address {
            10.10.10.0/24;
            10.11.0.0/16;
          }
        }
        then prefix-action psa-1Mbps-per-source-24-32-256;
      }
    }
  }
}
interfaces {
  so-0/0/2 {
    unit 0 {
      family inet {
        filter {
          input limit-source-two-24-16;
        }
        address 10.39.1.1/16;
      }
    }
  }
}
```

Scenario 2: Subnet Prefix Is Longer Than the Prefix in the Filter Match Condition

The complete example, [“Example: Configuring Prefix-Specific Counting and Policing” on page 4877](#), shows the simplest case of prefix-specific actions, in which the single-term firewall filter matches on one address with a prefix length that is the same as the subnet prefix length specified in the prefix-specific action. Unlike the example, this scenario describes a configuration in which the prefix-specific action defines a subnet prefix length that is longer than the prefix of the source address matched by the firewall filter. In this case, the prefix-specific action defines a subnet-prefix value of **25**, while the firewall filter matches on a source address in the **/24** subnet.



NOTE: The firewall filter passes the prefix-specific action packets with source addresses that range from 10.10.10.0 through 10.10.10.255, while the prefix-specific action specifies a set of only 128 counters and policers, numbered from 0 through 127.

The filter-matched packets that are passed to the prefix-specific action index into the counter and policer set in such a way that the counting and policing instances are shared by packets that contain either of two source addresses within the **10.10.10.0/24** subnet:

- The first counter and policer in the set are indexed by packets with source addresses **10.10.10.0** and **10.10.10.128**.
- The second counter and policer in the set are indexed by packets with source addresses **10.10.10.1** and **10.10.10.129**.
- The 128th (last) counter and policer in the set are indexed by packets with source addresses **10.10.10.127** and **10.10.10.255**.

The following configuration shows the statements for configuring the single-rate two-color policer, the prefix-specific action that references the policer, and the IPv4 standard stateless firewall filter that references the prefix-specific action:

```
[edit]
firewall {
  policer 1Mbps-policer {
    if-exceeding {
      bandwidth-limit 1m;
      burst-size-limit 63k;
    }
    then discard;
  }
  family inet {
    prefix-action psa-1Mbps-per-source-25-32-128 {
      policer 1Mbps-policer;
      subnet-prefix-length 25;
      source-prefix-length 32;
    }
    filter limit-source-one-24 {
      term one {
        from {
          source-address {
            10.10.10.0/24;
          }
        }
        then prefix-action psa-1Mbps-per-source-25-32-128;
      }
    }
  }
}
interfaces {
  so-0/0/2 {
    unit 0 {
      family inet {
```

```

    filter {
        input limit-source-one-24;
    }
    address 10.39.1.1/16;
}
}
}
}
}

```

Scenario 3: Subnet Prefix Is Shorter Than the Prefix in the Firewall Filter Match Condition

The complete example, “[Example: Configuring Prefix-Specific Counting and Policing](#)” on [page 4877](#), shows the simplest case of prefix-specific actions, in which the single-term firewall filter matches on one address with a prefix length that is the same as the subnet prefix length specified in the prefix-specific action. Unlike the example, this scenario describes a configuration in which the prefix-specific action defines a subnet prefix length that is shorter than the prefix of the source address matched by the firewall filter. In this case, the filter term matches on the /25 subnet of the source address 10.10.10.0.



NOTE: The firewall filter passes the prefix-specific action only packets with source addresses that range from 10.10.10.0 through 10.10.10.127, while the prefix-specific action specifies a set of 256 counters and policers, numbered from 0 through 255.

The matched packets that are passed to the prefix-specific action index into the lower half of the counter and policer set only:

- The first counter and policer in the set are indexed by packets with source address 10.10.10.0.
- The second counter and policer in the set are indexed by packets with source address 10.10.10.1 and 10.10.10.129.
- The 128th counter and policer in the set are indexed by packets with source address 10.10.10.127.
- The upper half of the set (instances numbered from 128 through 255) are not indexed by packets passed to the prefix-specific action from this particular firewall filter.

The following configuration shows the statements for configuring the single-rate two-color policer, the prefix-specific action that references the policer, and the IPv4 standard stateless firewall filter that references the prefix-specific action:

```

[edit]
firewall {
    policer 1Mbps-policer {
        if-exceeding {
            bandwidth-limit 1m;
            burst-size-limit 63k;
        }
        then discard;
    }
    family inet {

```

```

prefix-action psa-1Mbps-per-source-24-32-256 {
  policer 1Mbps-policer;
  subnet-prefix-length 24;
  source-prefix-length 32;
}
filter limit-source-one-25 {
  term one {
    from {
      source-address {
        10.10.10.0/25;
      }
    }
    then prefix-action psa-1Mbps-per-source-24-32-256;
  }
}
}
}
interfaces {
  so-0/0/2 {
    unit 0 {
      family inet {
        filter {
          input limit-source-one-25;
        }
        address 10.39.1.1/16;
      }
    }
  }
}
}

```

Related Documentation

- [Statement Hierarchy for Configuring Policers on page 4811](#)
- [Two-Color Policer Configuration Overview on page 4813](#)
- [Guidelines for Applying Traffic Policers on page 4821](#)

Multifield Classification

- [Multifield Classification Overview on page 4889](#)
- [Multifield Classification Requirements and Restrictions on page 4891](#)
- [Multifield Classification Limitations on M Series Routers on page 4892](#)
- [Example: Configuring Multifield Classification on page 4894](#)
- [Example: Configuring and Applying a Firewall Filter for a Multifield Classifier on page 4900](#)

Multifield Classification Overview

This topic covers the following information:

- [Forwarding Classes and PLP Levels on page 4890](#)
- [Multifield Classification and BA Classification on page 4890](#)
- [Multifield Classification Used In Conjunction with Policers on page 4890](#)

Forwarding Classes and PLP Levels

You can configure the Junos OS class of service (CoS) features to classify incoming traffic by associating each packet with a forwarding class, a packet loss priority (PLP) level, or both:

- Based on the associated forwarding class, each packet is assigned to an output queue, and the router services the output queues according to the associated scheduling you configure.
- Based on the associated PLP, each packet carries a lower or higher likelihood of being dropped if congestion occurs. The CoS random early detection (RED) process uses the drop probability configuration, output queue fullness percentage, and packet PLP to drop packet as needed to control congestion at the output stage.

Multifield Classification and BA Classification

The Junos OS supports two general types of packet classification: behavior aggregate (BA) classification and multifield classification:

- BA classification, or CoS value traffic classification, refers to a method of packet classification that uses a CoS configuration to set the forwarding class or PLP of a packet based on the *CoS value* in the IP packet header. The CoS value examined for BA classification purposes can be the Differentiated Services code point (DSCP) value, DSCP IPv6 value, IP precedence value, MPLS EXP bits, and IEEE 802.1p value. The default classifier is based on the IP precedence value.
- Multifield classification refers to a method of packet classification that uses a standard stateless firewall filter configuration to set the forwarding class or PLP for each packet entering or exiting the interface based on *multiple fields* in the IP packet header, including the DSCP value (for IPv4 only), the IP precedence value, the MPLS EXP bits, and the IEEE 802.1p bits. Multifield classification commonly matches on IP address fields, the IP protocol type field, or the port number in the UDP or TCP pseudoheader field. Multifield classification is used instead of BA classification when you need to classify packets based on information in the packet information other than the CoS values only.

With multifield classification, a firewall filter term can specify the packet classification actions for matching packets through the use of the **forwarding-class *class-name*** or **loss-priority (*high* | *medium-high* | *medium-low* | *low*)** nonterminating actions in the term's **then** clause. For more information about these actions, see .



NOTE: BA classification of a packet can be overridden by the stateless firewall filter actions **forwarding-class** and **loss-priority**.

Multifield Classification Used In Conjunction with Policers

To configure multifield classification in conjunction with rate limiting, a firewall filter term can specify the packet classification actions for matching packets through the use of a **policer** nonterminating action that references a single-rate two-color policer.

When multifield classification is configured to perform classification through a policer, the filter-matched packets in the traffic flow are rate-limited to the policer-specified traffic limits. Packets in a conforming flow of filter-matched packets are implicitly set to a **low** PLP. Packets in a nonconforming traffic flow can be discarded, or the packets can be set to a specified forwarding class, set to a specified PLP level, or both, depending on the type of policer and how the policer is configured to handle nonconforming traffic.



NOTE: Before you apply a firewall filter that performs multifield classification and also a policer to the same logical interface and for the same traffic direction, make sure that you consider the order of policer and firewall filter operations.

As an example, consider the following scenario:

- You configure a firewall filter that performs multifield classification (acts on matched packets by setting the forwarding class, the PLP, or both) based on the packet's existing forwarding class or PLP. You apply the firewall filter at the input of a logical interface.
- You also configure a single-rate two-color policer that acts on a red traffic flow by re-marking (setting the forwarding class, the PLP, or both) rather than discarding those packets. You apply the policer as an interface policer at the input of the same logical interface to which you apply the firewall filter.

Because of the order of policer and firewall operations, the input policer is executed before the input firewall filter. This means that the multifield classification specified by the firewall filter is performed on input packets that have already been re-marked once by policing actions. Consequently, any input packet that matches the conditions specified in a firewall filter term is then subject to a second re-marking according to the forwarding-class or loss-priority nonterminating actions also specified in that term.

Multifield Classification Requirements and Restrictions

This topic covers the following information:

- [Supported Platforms on page 4891](#)
- [CoS Tricolor Marking Requirement on page 4892](#)
- [Restrictions on page 4892](#)

Supported Platforms

The **loss-priority** firewall filter action is supported on the following routing and switching platforms only:

- EX Series switches
- M7i and M10i routers with the Enhanced CFEB (CFEB-E)
- M120 and M320 routers

- MX Series routers
- T Series routers with Enhanced II Flexible PIC Concentrators (FPCs)

CoS Tricolor Marking Requirement

The **loss-priority** firewall filter action has platform-specific requirements dependencies on the CoS tricolor marking feature, as defined in RFC 2698:

- On an M320 router, you cannot commit a configuration that includes the **loss-priority** firewall filter action unless you enable the CoS tricolor marking feature.
- On all routing platforms that support the **loss-priority** firewall filter action, you cannot set the **loss-priority** firewall filter action to **medium-low** or **medium-high** unless you enable the CoS tricolor marking feature. .

To enable the CoS tricolor marking feature, include the **tri-color** statement at the **[edit class-of-service]** hierarchy level.

Restrictions

You cannot configure the **loss-priority** and **three-color-policer** nonterminating actions for the same firewall filter term. These two nonterminating actions are mutually exclusive.

Multifield Classification Limitations on M Series Routers

This topic covers the following information:

- [Problem: Output-Filter Matching on Input-Filter Classification on page 4892](#)
- [Workaround: Configure All Actions in the Ingress Filter on page 4893](#)

Problem: Output-Filter Matching on Input-Filter Classification

On M Series routers (except M120 routers), you cannot classify packets with an output filter match based on the ingress classification that is set with an input filter applied to the same IPv4 logical interface.

For example, in the following configuration, the filter called **ingress** assigns all incoming IPv4 packets to the **expedited-forwarding** class. The filter called **egress** counts all packets that were assigned to the **expedited-forwarding** class in the **ingress** filter. This configuration does not work on most M Series routers. It works on all other routing platforms, including M120 routers, MX Series routers, and T Series routers, and EX Series switches.

```
[edit]
user@host # show firewall
family inet {
  filter ingress {
    term 1 {
      then {
        forwarding-class expedited-forwarding;
        accept;
      }
    }
    term 2 {
      then accept;
    }
  }
}
```

```

}
filter egress {
  term 1 {
    from {
      forwarding-class expedited-forwarding;
    }
    then count ef;
  }
  term 2 {
    then accept;
  }
}
}

```

```

[edit]
user@host# show interfaces
ge-1/2/0 {
  unit 0 {
    family inet {
      filter {
        input ingress;
        output egress;
      }
    }
  }
}

```

Workaround: Configure All Actions in the Ingress Filter

As a workaround, you can configure all of the actions in the ingress filter.

```

user@host # show firewall
family inet {
  filter ingress {
    term 1 {
      then {
        forwarding-class expedited-forwarding;
        accept;
        count ef;
      }
    }
    term 2 {
      then accept;
    }
  }
}

```

```

[edit]
user@host# show interfaces
ge-1/2/0 {
  unit 0 {
    family inet {
      filter {
        input ingress;
      }
    }
  }
}

```

```
}
```

Example: Configuring Multifield Classification

This example shows how to configure multifield classification of IPv4 traffic by using firewall filter actions and two firewall filter policers.

- [Requirements on page 4894](#)
- [Overview on page 4895](#)
- [Configuration on page 4896](#)
- [Verification on page 4900](#)

Requirements

Before you begin, make sure that your environment supports the features shown in this example:

1. The **loss-priority** firewall filter action must be supported on the router and configurable to all four values.
 - a. To be able to set a **loss-priority** firewall filter action, configure this example on logical interface **ge-1/2/0.0** on one of the following routing or switching platforms:
 - EX Series switch
 - MX Series router
 - M120 or M320 router
 - M7i or M10i router with the Enhanced CFEB (CFEB-E)
 - T Series router with Enhanced II Flexible PIC Concentrator (FPC)
 - b. To be able to set a **loss-priority** firewall filter action to **medium-low** or **medium-high**, make sure that the CoS tricolor marking feature is enabled. To enable the CoS tricolor marking feature, include the **tri-color** statement at the **[edit class-of-service]** hierarchy level.
2. The **expedited-forwarding** and **assured-forwarding** forwarding classes must be scheduled on the underlying physical interface **ge-1/2/0**.
 - a. Make sure that the following forwarding classes are assigned to output queues:
 - **expedited-forwarding**
 - **assured-forwarding**



NOTE: You cannot commit a configuration that assigns the same forwarding class to two different queues.

- b. Make sure that the output queues to which the forwarding classes are assigned are associated with schedulers. A scheduler defines the amount of interface bandwidth assigned to the queue, the size of the memory buffer allocated for storing packets, the priority of the queue, and the random early detection (RED) drop profiles associated with the queue.
 - You configure output queue schedulers at the **[edit class-of-service schedulers]** hierarchy level.
 - You associate output queue schedulers with forwarding classes by means of a scheduler map that you configure at the **[edit class-of-service scheduler-maps map-name]** hierarchy level.
- c. Make sure that output-queue scheduling is applied to the physical interface **ge-1/2/0**.

You apply a scheduler map to a physical interface at the **[edit class-of-service interfaces ge-1/2/0 scheduler-map map-name]** hierarchy level.

Overview

In this example, you apply multifield classification to the input IPv4 traffic at a logical interface by using stateless firewall filter actions and two firewall filter policers that are referenced from the firewall filter. Based on the source address field, packets are either set to the **low** loss priority or else policed. Neither of the policers discards nonconforming traffic. Packets in nonconforming flows are marked for a specific forwarding class (**expedited-forwarding** or **assured-forwarding**), set to a specific loss priority, and then transmitted.



NOTE: Single-rate two-color policers always transmit packets in a conforming traffic flow after implicitly setting a low loss priority.

Topology

In this example, you apply multifield classification to the IPv4 traffic on logical interface **ge-1/2/0.0**. The classification rules are specified in the IPv4 stateless firewall filter **mfc-filter** and two single-rate two-color policers, **ef-policer** and **af-policer**.

The IPv4 standard stateless firewall filter **mfc-filter** defines three filter terms:

- **isp1-customers**—The first filter term matches packets with the source address 10.1.1.0/24 or 10.1.2.0/24. Matched packets are assigned to the **expedited-forwarding** forwarding class and set to the **low** loss priority.
- **isp2-customers**—The second filter term matches packets with the source address 10.1.3.0/24 or 10.1.4.0/24. Matched packets are passed to **ef-policer**, a policer that rate-limits traffic to a bandwidth limit of 300 Kbps with a burst-size limit of 50 KB. This policer specifies that packets in a nonconforming flow are marked for the **expedited-forwarding** forwarding class and set to the **high** loss priority.
- **other-customers**—The third and final filter term passes all other packets to **af-policer**, a policer that rate-limits traffic to a bandwidth limit of 300 Kbps and a burst-size limit

of 50 KB (the same traffic limits as defined by **ef-policer**). This policer specifies that packets in a nonconforming flow are marked for the **assured-forwarding** forwarding class and set to the **medium-high** loss priority.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configuring Policers to Rate-Limit Expedited-Forwarding and Assured-Forwarding Traffic on page 4896](#)
- [Configuring a Multifield Classification Filter That Also Applies Policing on page 4897](#)
- [Applying Multifield Classification Filtering and Policing to the Logical Interface on page 4899](#)

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set firewall policer ef-policer if-exceeding bandwidth-limit 300k
set firewall policer ef-policer if-exceeding burst-size-limit 50k
set firewall policer ef-policer then loss-priority high
set firewall policer ef-policer then forwarding-class expedited-forwarding
set firewall policer af-policer if-exceeding bandwidth-limit 300k
set firewall policer af-policer if-exceeding burst-size-limit 50k
set firewall policer af-policer then loss-priority high
set firewall policer af-policer then forwarding-class assured-forwarding
set firewall family inet filter mfc-filter term isp1-customers from source-address 10.1.1.0/24
set firewall family inet filter mfc-filter term isp1-customers from source-address 10.1.2.0/24
set firewall family inet filter mfc-filter term isp1-customers then loss-priority low
set firewall family inet filter mfc-filter term isp1-customers then forwarding-class
  expedited-forwarding
set firewall family inet filter mfc-filter term isp2-customers from source-address
  10.1.3.0/24
set firewall family inet filter mfc-filter term isp2-customers from source-address
  10.1.4.0/24
set firewall family inet filter mfc-filter term isp2-customers then policer ef-policer
set firewall family inet filter mfc-filter term other-customers then policer af-policer
set interfaces ge-1/2/0 unit 0 family inet address 192.168.1.1/24
set interfaces ge-1/2/0 unit 0 family inet filter input mfc-filter
```

Configuring Policers to Rate-Limit Expedited-Forwarding and Assured-Forwarding Traffic

Step-by-Step Procedure

To configure policers to rate-limit expedited-forwarding and assured-forwarding traffic:

1. Define traffic limits for expedited-forwarding traffic.

```
[edit]
user@host# edit firewall policer ef-policer
[edit firewall policer ef-policer]
user@host# set if-exceeding bandwidth-limit 300k
user@host# set if-exceeding burst-size-limit 50k
```

```

user@host# set then loss-priority high
user@host# set then forwarding-class expedited-forwarding

```

2. Configure a policer for assured-forwarding traffic.

```

[edit firewall policer ef-policer]
user@host# up

[edit firewall]
user@host# edit policer af-policer

[edit firewall policer af-policer]
user@host# set if-exceeding bandwidth-limit 300k
user@host# set if-exceeding burst-size-limit 50k
user@host# set then loss-priority high
user@host# set then forwarding-class assured-forwarding

```

Results Confirm the configuration of the policer by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```

[edit]
user@host# show firewall
policer af-policer {
  if-exceeding {
    bandwidth-limit 300k;
    burst-size-limit 50k;
  }
  then {
    loss-priority high;
    forwarding-class assured-forwarding;
  }
}
policer ef-policer {
  if-exceeding {
    bandwidth-limit 300k;
    burst-size-limit 50k;
  }
  then {
    loss-priority high;
    forwarding-class expedited-forwarding;
  }
}

```

Configuring a Multifield Classification Filter That Also Applies Policing

Step-by-Step Procedure To configure a multifield classification filter that additionally applies policing:

1. Enable configuration of a firewall filter term for IPv4 traffic.

```

[edit]
user@host# edit firewall family inet filter mfc-filter

```

2. Configure the first term to match on source addresses and then classify the matched packets.

```

[edit firewall family inet filter mfc-filter]

```

```
user@host# set term isp1-customers from source-address 10.1.1.0/24
user@host# set term isp1-customers from source-address 10.1.2.0/24
user@host# set term isp1-customers then loss-priority low
user@host# set term isp1-customers then forwarding-class expedited-forwarding
```

3. Configure the second term to match on different source addresses and then police the matched packets.

```
[edit firewall family inet filter mfc-filter]
user@host# set term isp2-customers from source-address 10.1.3.0/24
user@host# set term isp2-customers from source-address 10.1.4.0/24
user@host# set term isp2-customers then policer ef-policer
```

4. Configure the third term to police all other packets to a different set of traffic limits and actions.

```
[edit firewall family inet filter mfc-filter]
user@host# set term other-customers then policer af-policer
```

Results Confirm the configuration of the filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
  filter mfc-filter {
    term isp1-customers {
      from {
        source-address 10.1.1.0/24;
        source-address 10.1.2.0/24;
      }
      then {
        loss-priority low;
        forwarding-class expedited-forwarding;
      }
    }
    term isp2-customers {
      from {
        source-address 10.1.3.0/24;
        source-address 10.1.4.0/24;
      }
      then {
        policer ef-policer;
      }
    }
    term other-customers {
      then {
        policer af-policer;
      }
    }
  }
}
policer af-policer {
  if-exceeding {
    bandwidth-limit 300k;
```



```

        burst-size-limit 50k;
    }
    then discard;
}
policer ef-policer {
    if-exceeding {
        bandwidth-limit 200k;
        burst-size-limit 50k;
    }
    then {
        loss-priority high;
        forwarding-class expedited-forwarding;
    }
}
}

```

Applying Multifield Classification Filtering and Policing to the Logical Interface

Step-by-Step Procedure To apply multifield classification filtering and policing to the logical interface:

1. Enable configuration of IPv4 on the logical interface.

```

[edit]
user@host# edit interfaces ge-1/2/0 unit 0 family inet

```

2. Configure an IP address for the logical interface.

```

[edit interfaces ge-1/2/0 unit 0 family inet ]
user@host# set address 192.168.1.1/24

```

3. Apply the firewall filter to the logical interface input.

```

[edit interfaces ge-1/2/0 unit 0 family inet ]
user@host# set filter input mfc-filter

```



NOTE: Because the policer is executed before the filter, if an input policer is also configured on the logical interface, it cannot use the forwarding class and PLP of a multifield classifier associated with the interface.

Results Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```

[edit]
user@host# show interfaces
ge-1/2/0 {
    unit 0 {
        family inet {
            filter {
                input mfc-filter;
            }
            address 192.168.1.1/24;
        }
    }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Displaying the Number of Packets Processed by the Policer at the Logical Interface

Purpose Verify the traffic flow through the logical interface and that the policer is evaluated when packets are received on the logical interface.

Action Use the **show firewall** operational mode command for the filter you applied to the logical interface.

```
user@host> show firewall filter rate-limit-in
Filter: rate-limit-in
Policers:
Name                                     Packets
ef-policer-isp2-customers                32863
af-policer-other-customers                3870
```

The command output lists the policers applied by the firewall filter **rate-limit-in**, and the number of packets that matched the filter term.



NOTE: The packet count includes the number of out-of-specification (out-of-spec) packet counts, not all packets policed by the policer.

The policer name is displayed concatenated with the name of the firewall filter term in which the policer is referenced as an action.

Example: Configuring and Applying a Firewall Filter for a Multifield Classifier

This example shows how to configure a firewall filter to classify traffic using a multifield classifier. The classifier detects packets of interest to CoS as they arrive on an interface.

- [Requirements on page 4900](#)
- [Overview on page 4901](#)
- [Configuration on page 4901](#)
- [Verification on page 4904](#)

Requirements

Before you begin, review how to create and configure a firewall. See [“Guidelines for Configuring Standard Firewall Filters” on page 4478](#).



NOTE: On T4000 Type 5 FPCs, a filter attached at the Layer 2 application point (that is, at the logical interface level) is unable to match with the forwarding class of a packet that is set by a Layer 3 classifier such as DSCP, DSCP V6, inet-precedence, and mpls-exp.

Overview

One common way to detect packets of CoS interest is by source or destination address. The destination address is used in this example, but many other matching criteria for packet detection are available to firewall filters.

In this example, you configure the firewall filter `mf-classifier`. You create and name the assured forwarding traffic class, set the match condition, and specify the destination address as `192.168.44.55`. You create the forwarding class for assured forwarding DiffServ traffic as `af-class` and set the loss priority to low.

Then you create and name the expedited forwarding traffic class, set the match condition, for the expedited forwarding traffic class, and specify the destination address as `192.168.66.77`. You then create the forwarding class for expedited forwarding DiffServ traffic as `ef-class` and set the policer to `ef-policer`. Then you create and name the network-control traffic class and set the match condition.

You then create and name the forwarding class for the network control traffic class as `nc-class`. You create and name the forwarding class for the best-effort traffic class as `be-class`. Finally, you apply the multifield classifier firewall filter as an input filter on each customer-facing or host-facing that needs the filter. In this example, the interface is `ge-0/0/0`.

Configuration

CLI Quick Configuration

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

```
set firewall filter mf-classifier interface-specific
set firewall filter mf-classifier term assured-forwarding from destination-address
  192.168.44.55
set firewall filter mf-classifier term assured-forwarding then forwarding-class af-class
set firewall filter mf-classifier term assured-forwarding then loss-priority low
set firewall filter mf-classifier term expedited-forwarding from destination-address
  192.168.66.77
set firewall filter mf-classifier term expedited-forwarding then forwarding-class ef-class
set firewall filter mf-classifier term expedited-forwarding then policer ef-policer
set firewall filter mf-classifier term network-control from precedence net-control
set firewall filter mf-classifier term network-control then forwarding-class nc-class
set firewall filter mf-classifier term best-effort then forwarding-class be-class
set interfaces ge-0/0/0 unit 0 family inet filter input mf-classifier
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure a firewall filter for a multifield classifier for a device:

1. Create and name the multifield classifier filter.

```
[edit]
user@host# edit firewall filter mf-classifier
```

user@host# set interface-specific

2. Create and name the term for the assured forwarding traffic class.

[edit firewall filter mf-classifier]

user@host# edit term assured-forwarding

3. Specify the destination address for assured forwarding traffic.

[edit firewall filter mf-classifier term assured-forwarding]

user@host# set from destination-address 192.168.44.55

4. Create the forwarding class and set the loss priority for the assured forwarding traffic class.

[edit firewall filter mf-classifier term assured-forwarding]

user@host# set then forwarding-class af-class

user@host# set then loss-priority low

5. Create and name the term for the expedited forwarding traffic class.

[edit]

user@host# edit firewall filter mf-classifier

user@host# edit term expedited-forwarding

6. Specify the destination address for the expedited forwarding traffic.

[edit firewall filter mf-classifier term expedited-forwarding]

user@host# set from destination-address 192.168.66.77

7. Create the forwarding class and apply the policer for the expedited forwarding traffic class.

[edit firewall filter mf-classifier term expedited-forwarding]

user@host# set then forwarding-class ef-class

user@host# set then policer ef-policer

8. Create and name the term for the network control traffic class.

[edit]

user@host# edit firewall filter mf-classifier

user@host# edit term network-control

9. Create the match condition for the network control traffic class.

[edit firewall filter mf-classifier term network-control]

user@host# set from precedence net-control

10. Create and name the forwarding class for the network control traffic class.

[edit firewall filter mf-classifier term network-control]

user@host# set then forwarding-class nc-class

11. Create and name the term for the best-effort traffic class.

[edit]

user@host# edit firewall filter mf-classifier

user@host# edit term best-effort

12. Create and name the forwarding class for the best-effort traffic class.

[edit firewall filter mf-classifier term best-effort]

user@host# set then forwarding-class be-class



NOTE: Because this is the last term in the filter, it has no match condition.

13. Apply the multifield classifier firewall filter as an input filter.

[edit]

```
user@host# set interfaces ge-0/0/0 unit 0 family inet filter input mf-classifier
```

Results From configuration mode, confirm your configuration by entering the **show firewall filter mf-classifier** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

[edit]

```
user@host# show firewall filter mf-classifier
interface-specific;
term assured-forwarding {
  from {
    destination-address {
      192.168.44.55/32;
    }
  }
  then {
    loss-priority low;
    forwarding-class af-class;
  }
}
term expedited-forwarding {
  from {
    destination-address {
      192.168.66.77/32;
    }
  }
  then {
    policer ef-policer;
    forwarding-class ef-class;
  }
}
term network-control {
  from {
    precedence net-control;
  }
  then forwarding-class nc-class;
}
term best-effort {
  then forwarding-class be-class;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying a Firewall Filter for a Multifield Classifier Configuration on page 4904](#)

Verifying a Firewall Filter for a Multifield Classifier Configuration

Purpose Verify that a firewall filter for a multifield classifier is configured properly on a device.

Action From configuration mode, enter the **show firewall filter mf-classifier** command.

Related Documentation

- [Standard Firewall Filter Nonterminating Actions on page 4744](#)
- [Order of Policer and Firewall Filter Operations on page 4810](#)
- [Statement Hierarchy for Configuring Policers on page 4811](#)
- [Two-Color Policer Configuration Overview on page 4813](#)
- [Guidelines for Applying Traffic Policers on page 4821](#)
- [Junos CoS Components on page 1266](#)
- [BA Classifier Overview on page 1320](#)
- [Overview of Forwarding Classes on page 1526](#)
- [Default Forwarding Classes](#)
- [RED Drop Profiles Overview on page 1674](#)
- [tri-color on page 1483](#) statement

Policer Overhead to Account for Rate Shaping in the Traffic Manager

- [Policer Overhead to Account for Rate Shaping Overview on page 4904](#)
- [Example: Configuring Policer Overhead to Account for Rate Shaping on page 4905](#)

Policer Overhead to Account for Rate Shaping Overview

If you configure ingress or egress traffic-shaping overhead values for an interface, the traffic manager cannot apply these values to any rate-limiting also applied to the interface. To enable the router to account for the additional Ethernet frame length when policing actions are being determined, you must configure the ingress or egress overhead values for policers separately.



NOTE: When a policer overhead value is changed, the PIC or DPC goes offline and then comes back online.

For Gigabit Ethernet Intelligent Queuing 2 (IQ2) and Enhanced IQ2 (IQ2E) PICs or interfaces on Dense Port Concentrators (DPCs) in MX Series routers, you can control the rate of traffic that passes through all interfaces on the PIC or DPC by configuring a *policer overhead*. You can configure a policer ingress overhead and a policer egress overhead,

each with values from 0 through 255 bytes. The policer overhead values are added to the length of the final Ethernet frame when determining ingress and egress policer actions.

Example: Configuring Policer Overhead to Account for Rate Shaping

This example shows how to configure overhead values for policers when rate-shaping overhead is configured.

- [Requirements on page 4905](#)
- [Overview on page 4905](#)
- [Configuration on page 4906](#)
- [Verification on page 4911](#)

Requirements

Before you begin, make sure that interface for which you are applying ingress or egress policer overhead is hosted on one of the following:

- Gigabit Ethernet IQ2 PIC
- IQ2E PIC
- DPCs in MX Series routers

Overview

This example shows how to configure policer overhead values for all physical interfaces on a supported PIC or MPC so that the rate shaping value configured on a logical interface is accounted for in any policing on that logical interface.

Topology

The router hosts a Gigabit Ethernet IQ2 PIC, installed in PIC location 3 of the Flexible PIC Concentrator (FPC) in slot number 1. The physical interface on port 1 on that PIC is configured to receive traffic on logical interface 0 and send it back out on logical interface 1. Class-of-service scheduling includes 100 Mbps of traffic rate-shaping overhead for the output traffic. A policer egress overhead of 100 bytes is configured on the entire PIC so that, for any policers applied to the output traffic, 100 bytes are added to the final Ethernet frame length when determining ingress and egress policer actions.



NOTE:

Traffic rate-shaping and corresponding policer overhead are configured separately:

- You configure rate shaping at the `[edit class-of-service interfaces interface-name unit unit-number]` hierarchy level.
- You configure policer overhead at the `[edit chassis fpc slot-number pic pic-number]` hierarchy level.

When a policer overhead value is changed, the PIC or DPC goes offline and then comes back online.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configuring the Logical Interfaces on page 4906](#)
- [Configuring Traffic Rate-Shaping on the Logical Interface That Carries Output Traffic on page 4908](#)
- [Configuring Policer Overhead on the PIC or MPC That Hosts the Rate-Shaped Logical Interface on page 4909](#)
- [Applying a Policer to the Logical Interface That Carries Input Traffic on page 4910](#)

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces ge-1/3/1 per-unit-scheduler
set interfaces ge-1/3/1 vlan-tagging
set interfaces ge-1/3/1 unit 0 vlan-id 100
set interfaces ge-1/3/1 unit 0 family inet address 10.10.10.1/30
set interfaces ge-1/3/1 unit 1 vlan-id 101
set interfaces ge-1/3/1 unit 1 family inet address 20.20.20.1/30 arp 20.20.20.2 mac
  00:00:11:22:33:44
set class-of-service schedulers be transmit-rate percent 5
set class-of-service schedulers ef transmit-rate percent 30
set class-of-service schedulers af transmit-rate percent 30
set class-of-service schedulers nc transmit-rate percent 35
set class-of-service scheduler-maps my-map forwarding-class best-effort scheduler be
set class-of-service scheduler-maps my-map forwarding-class expedited-forwarding
  scheduler ef
set class-of-service scheduler-maps my-map forwarding-class network-control scheduler
  nc
set class-of-service scheduler-maps my-map forwarding-class assured-forwarding
  scheduler af
set class-of-service interfaces ge-1/3/1 unit 1 scheduler-map my-map
set class-of-service interfaces ge-1/3/1 unit 1 shaping-rate 100m
set firewall policer 500Kbps logical-interface-policer
set firewall policer 500Kbps if-exceeding bandwidth-limit 500k
set firewall policer 500Kbps if-exceeding burst-size-limit 625k
set firewall policer 500Kbps then discard
set chassis fpc 1 pic 3 ingress-policer-overhead 100
set chassis fpc 1 pic 3 egress-policer-overhead 100
set interfaces ge-1/3/1 unit 0 family inet policer input 500Kbps
```

Configuring the Logical Interfaces

Step-by-Step Procedure

To configure the logical interfaces:

1. Enable configuration of the interface

[edit]


```
user@host# edit interfaces ge-1/3/1
```

2. Enable multiple queues for each logical interface (so that you can associate an output scheduler with each logical interface).

```
[edit interfaces ge-1/3/1]
user@host# set per-unit scheduler
user@host# set vlan-tagging
```



NOTE: For Gigabit Ethernet IQ2 PICs only, use the **shared-scheduler** statement to enable shared schedulers and shapers on a physical interface.

3. Configure logical interface **ge-1/3/1.0**.

```
[edit interfaces ge-1/3/1]
user@host# set unit 0 vlan-id 100
user@host# set unit 0 family inet address 10.10.10.1/30
```

4. Configure logical interface **ge-1/3/1.1**.

```
[edit interfaces ge-1/3/1]
user@host# set unit 1 vlan-id 101
user@host# set unit 1 family inet address 20.20.20.1/30 arp 20.20.20.2 mac
00:00:11:22:33:44
```

Results Confirm the configuration of the interfaces by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
ge-1/3/1 {
  per-unit-scheduler;
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    family inet {
      address 10.10.10.1/30;
    }
  }
  unit 1 {
    vlan-id 101;
    family inet {
      address 20.20.20.1/30 {
        arp 20.20.20.2 mac 00:00:11:22:33:44;
      }
    }
  }
}
```

*Configuring Traffic Rate-Shaping on the Logical Interface That Carries Output Traffic***Step-by-Step Procedure**

To configure traffic rate-shaping on the logical interface that carries output traffic:

1. Enable configuration of class-of-service features.

```
[edit]
user@host# edit class-of-service
```

2. Configure packet scheduling on logical interface **ge-1/3/1.0**.

- a. Configure schedulers that specify the percentage of transmission capacity.

```
[edit class-of-service]
user@host# edit schedulers
```

```
[edit class-of-service schedulers]
user@host# set be transmit-rate percent 5
user@host# set ef transmit-rate percent 30
user@host# set af transmit-rate percent 30
user@host# set nc transmit-rate percent 35
```

A percentage of zero drops all packets in the queue. When the **rate-limit** option is specified, the transmission rate is limited to the rate-controlled amount. In contrast with the **exact** option, a scheduler with the **rate-limit** option shares unused bandwidth above the rate-controlled amount.

- b. Configure a scheduler map to associate each scheduler with a forwarding class.

```
[edit class-of-service]
user@host# edit scheduler-maps my-map

[edit class-of-service scheduler-maps my-map]
user@host# set forwarding-class best-effort scheduler be
user@host# set forwarding-class expedited-forwarding scheduler ef
user@host# set forwarding-class network-control scheduler nc
user@host# set forwarding-class assured-forwarding scheduler af
```

- c. Associate the scheduler map with logical interface **ge-1/3/1.0**.

```
[edit class-of-service]
user@host# edit interfaces ge-1/3/1 unit 1
```

```
[edit class-of-service interfaces ge-1/3/1 unit 1]
user@host# set scheduler-map my-map
```

3. Configure 100 Mbps of traffic rate-shaping overhead on logical interface **ge-1/3/1.1**.

```
[edit class-of-service interfaces ge-1/3/1 unit 1]
user@host# set shaping-rate 100
```

Alternatively, you can configure a shaping rate for a logical interface and oversubscribe the physical interface by including the **shaping-rate** statement at the **[edit class-of-service traffic-control-profiles]** hierarchy level. With this configuration approach, you can independently control the delay-buffer rate.

Results Confirm the configuration of the class-of-service features (including the 100 Mbp of shaping of the egress traffic) by entering the **show class-of-service** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show class-of-service
interfaces {
  ge-1/3/1 {
    unit 1 {
      scheduler-map my-map;
      shaping-rate 100m;
    }
  }
}
scheduler-maps {
  my-map {
    forwarding-class best-effort scheduler be;
    forwarding-class expedited-forwarding scheduler ef;
    forwarding-class network-control scheduler nc;
    forwarding-class assured-forwarding scheduler af;
  }
}
schedulers {
  be {
    transmit-rate percent 5;
  }
  ef {
    transmit-rate percent 30;
  }
  af {
    transmit-rate percent 30;
  }
  nc {
    transmit-rate percent 35;
  }
}
```

Configuring Policer Overhead on the PIC or MPC That Hosts the Rate-Shaped Logical Interface

Step-by-Step Procedure To configure policer overhead on the PIC or MPC that hosts the rate-shaped logical interface:

1. Enable configuration of the supported PIC or MPC.

```
[edit]
user@host# set chassis fpc 1 pic 3
```

2. Configure 100 bytes of policer overhead on the supported PIC or MPC.

```
[edit]
user@host# set ingress-policer-overhead 100
user@host# set egress-policer-overhead 100
```



NOTE: These values are added to the length of the final Ethernet frame when determining ingress and egress policer actions for all physical interfaces on the PIC or MPC.

You can specify policer overhead with values from 0 through 255 bytes.

Results Confirm the configuration of the policer overhead on the physical interface to account for rate-shaping by entering the **show chassis** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show chassis
chassis {
  fpc 1 {
    pic 3 {
      egress-policer-overhead 100;
      ingress-policer-overhead 100;
    }
  }
}
```

Applying a Policer to the Logical Interface That Carries Input Traffic

Step-by-Step Procedure To apply a policer to the logical interface that carries input traffic:

1. Configure the logical interface (aggregate) policer.

```
[edit]
user@host# edit firewall policer 500Kbps

[edit firewall policer 500Kbps]
user@host# set logical-interface-policer
user@host# set if-exceeding bandwidth-limit 500k
user@host# set if-exceeding burst-size-limit 625k
user@host# set then discard
```

2. Apply the policer to Layer 3 input on the IPv4 logical interface.

```
[edit]
user@host# set interfaces ge-1/3/1 unit 0 family inet policer input 500Kbps
```



NOTE: The 100 Mbps policer overhead is added to the length of the final Ethernet frame when determining ingress and egress policer actions,

Results Confirm the configuration of the policer with rate-shaping overhead by entering the **show firewall** and **show interfaces** configuration mode commands. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
policer 500Kbps {
  logical-interface-policer;
  if-exceeding {
    bandwidth-limit 500k;
    burst-size-limit 625k;
  }
  then discard;
}
[edit]
user@host# show interfaces
ge-1/3/1 {
  per-unit-scheduler;
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    layer2-policer {
      input-policer 500Kbps;
    }
    family inet {
      address 10.10.10.1/30;
    }
  }
  unit 0 {
    vlan-id 101;
    family inet {
      address 20.20.20.1/30 {
        arp 20.20.20.2 mac 00:00:11:22:33:44;
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Displaying Traffic Statistics and Policers for the Logical Interface on page 4911](#)
- [Displaying Statistics for the Policer on page 4912](#)

Displaying Traffic Statistics and Policers for the Logical Interface

Purpose Verify the traffic flow through the logical interface and that the policer is evaluated when packets are received on the logical interface.

Action Use the **show interfaces** operational mode command for logical interface **ge-1/3/1.0**, and include the **detail** or **extensive** option. The command output section for **Traffic statistics** lists the number of bytes and packets received and transmitted on the logical interface, and the **Protocol inet** section contains a **Policer** field that would list the policer **500Kbps** as an input or output policer as follows:

- **Input: 500Kbps-ge-1/3/1.0-log_int-i**
- **Output: 500Kbps-ge-1/3/1.0-log_int-o**

The **log_int-i** suffix denotes a logical interface policer applied to input traffic, while the **log_int-o** suffix denotes a logical interface policer applied to output traffic. In this example, the logical interface policer is applied to Input traffic only.

Displaying Statistics for the Policer

Purpose Verify the number of packets evaluated by the policer.

Action Use the **show policer** operational mode command and optionally specify the name of the policer. The command output displays the number of packets evaluated by each configured policer (or the specified policer), in each direction. For the policer **500Kbps**, the input and output policer names are displayed as follows:

- **500Kbps-ge-1/3/1.0-log_int-i**
- **500Kbps-ge-1/3/1.0-log_int-o**

The **log_int-i** suffix denotes a logical interface policer applied to input traffic, while the **log_int-o** suffix denotes a logical interface policer applied to output traffic. In this example, the logical interface policer is applied to input traffic only.

- Related Documentation**
- [Statement Hierarchy for Configuring Policers on page 4811](#)
 - [Two-Color Policer Configuration Overview on page 4813](#)
 - [Guidelines for Applying Traffic Policers on page 4821](#)
 - *“Configuring a Policer Overhead” in the Junos OS Operational Mode Commands*

Configuring Three-Color Policers

- [Three-Color Policer Configuration Guidelines on page 4912](#)
- [Basic Single-Rate Three-Color Policers on page 4915](#)
- [Basic Two-Rate Three-Color Policers on page 4921](#)

Three-Color Policer Configuration Guidelines

- [Platforms Supported for Three-Color Policers on page 4913](#)
- [Color Modes for Three-Color Policers on page 4913](#)
- [Naming Conventions for Three-Color Policers on page 4914](#)

Platforms Supported for Three-Color Policers

Three-color policers are supported on the following Juniper Networks routers and switch:

- EX Series switches
- M120 Multiservice Edge Routers
- M320 Multiservice Edge Routers and T Series Core Routers with Enhanced II Flexible PIC Concentrators (FPCs)
- MX Series 3D Universal Edge Routers
- T640 Core Routers with Enhanced Scaling FPC4
- T4000 Core Routers with FPC5

On MX Series and M120 routers and EX Series switches, you can apply three-color policers to aggregated interfaces.

The **discard** action for a tricolor marking policer for a firewall filter is supported on the M120 routers, M320 routers with Enhanced-III FPCs, M7i and M10i routers with the Enhanced CFEB (CFEB-E), and MX Series routers with Trio MPCs and EX Series switches, so it is not necessary to include the **logical-interface-policer** statement for them.

Color Modes for Three-Color Policers

Three-color policers—both single-rate and two-rate three-color policer schemes—can operate in either of two modes:

- [Color-Blind Mode on page 4913](#)
- [Color-Aware Mode on page 4913](#)

Color-Blind Mode

In *color-blind* mode, the three-color policer assumes that all packets examined have not been previously marked or metered. If you configure a three-color policer to be color-blind instead of color-aware, the policer ignores preexisting color markings that might have been set for a packet by another traffic policer configured at a previous network node.

Color-Aware Mode

In *color-aware* mode, the three-color policer assumes that all packets examined have been previously marked or metered. In other words, the three-color policer takes into account any coloring markings that might have been set for a packet by another traffic policer configured at a previous network node. At the node where color-aware policing is configured, any preexisting color markings are used in determining the appropriate policing action for the packet.

In color-aware mode, the three-color policer can increase the packet loss priority (PLP) level of a packet, but never decrease it. For example, if a color-aware three-color policer meters a packet with a medium PLP marking, it can raise the PLP level to high, but cannot reduce the PLP level to low.

For two-rate, three-color policing, the Junos OS uses two token buckets to manage bandwidth based on the two rates of traffic. For example, two-rate policing might be

configured on a node upstream in the network. The two-rate policer has marked a packet as yellow (loss priority medium-low). The color-aware policer takes this yellow marking into account when determining the appropriate policing action. In color-aware policing, the yellow packet would never receive the action associated with either the green packets or red packets. This way, tokens for violating packets are never taken from the metering token buckets at the color-aware policing node.



NOTE: For a three-color policer operating in color-aware mode and when the PLP of the input packet is medium-low, the color of the input packet to the policer is mapped to the color yellow.

In such a scenario, if the color of the input packet remains unchanged, the policer operates in the following way:

- On a T1600 Enhanced Scaling Type 4 FPC (T1600-FPC4-ES), the PLP of the output packet remains medium-low.
- On a T4000 Type 5 FPC (T4000-FPC5-3D), the PLP of the output packet is marked as medium-high.

Because of this difference, for any applications (such as rewrite and WRED selection on egress interface) that use PLP, the packets are treated differently for the same flow depending on the FPC type (T1600 Enhanced Scaling FPC4 (T1600-FPC4-ES) or T4000 FPC5 (T4000-FPC5-3D)) on which the policer is applied.

Naming Conventions for Three-Color Policers

Because policers can be numerous and must be applied correctly to work, a simple naming convention makes it easier to apply the policers properly.

We recommend that you name your policer using a convention that identifies the basic components of the policer:

- Three-color policer type—Where **srTCM** identifies a *single-rate* three-color policer and **trTCM** identifies a *two-rate* three-color policer.
- Three-color policer color mode—Where **ca** identifies a *color-aware* three-color policer and **cb** identifies a *color-blind three-color policer*.



NOTE:

TCM stands for tricolor marking.

Table 369 on page 4915 describes a recommended naming convention for policers.

Table 369: Recommended Naming Convention for Policers

Three-Color Policer Type	Naming Convention	Example Names
Single-rate three-color, color-aware	<i>srTCMnumber-ca</i>	srTCM1-ca, srTCM2-ca, srTCM3-ca, ...
Single-rate three-color, color-blind	<i>srTCMnumber-cb</i>	srTCM1-cb, srTCM2-cb, srTCM3-cb, ...
Two-rate three-color, color-aware	<i>trTCMnumber-ca</i>	trTCM1-ca, trTCM2-ca, trTCM3-ca, ...
Two-rate three-color, color-blind	<i>trTCMnumber-cb</i>	trTCM1-cb, trTCM2-cb, trTCM3-cb, ...

**Related
Documentation**

- [Statement Hierarchy for Configuring Policers on page 4811](#)
- [Three-Color Policer Configuration Overview on page 4817](#)
- [Guidelines for Applying Traffic Policers on page 4821](#)

Basic Single-Rate Three-Color Policers

- [Single-Rate Three-Color Policer Overview on page 4915](#)
- [Example: Configuring a Single-Rate Three-Color Policer on page 4916](#)

Single-Rate Three-Color Policer Overview

A single-rate three-color policer defines a bandwidth limit and a maximum burst size for guaranteed traffic and a second burst size for peak traffic. A single-rate three-color policer is most useful when a service is structured according to packet length and not peak arrival rate.

Single-rate three-color policing meters a traffic stream based on the following configured traffic criteria:

- Committed information rate (CIR)—Bandwidth limit for guaranteed traffic.
- Committed burst size (CBS)—Maximum packet size permitted for bursts of data that exceed the CIR.
- Excess burst size (EBS)—Maximum packet size permitted for peak traffic.

Single-rate tricolor marking (single-rate TCM) classifies traffic as belonging to one of three color categories and performs congestion-control actions on the packets based on the color marking:

- Green—Traffic that conforms to *either* the bandwidth limit *or* the burst size for guaranteed traffic (CIR and CBS). For a green traffic flow, single-rate marks the packets with an implicit loss priority of **low** and transmits the packets.
- Yellow—Traffic that exceeds *both* the bandwidth limit *and* the burst size for guaranteed traffic (CIR or CBS) but not the burst size for peak traffic (EBS). For a yellow traffic flow, single-rate marks the packets with an implicit loss priority of **medium-high** and transmits the packets.
- Red—Traffic that exceeds the burst size for peak traffic (EBS), single-rate marks packets with an implicit loss priority of **high** and, optionally, discards the packets.

If congestion occurs downstream, the packets with higher loss priority are more likely to be discarded.



NOTE: For both single-rate and two-rate three-color policers, the only *configurable* action is to discard packets in a red traffic flow.

The **discard** action for a tricolor marking policer for a firewall filter is supported on the M120 routers, M320 routers with Enhanced-III FPCs, M7i and M10i routers with the Enhanced CFEB (CFEB-E), and MX Series routers with Trio MPCs and EX Series switches, so it is not necessary to include the **logical-interface-policer** statement for them.

Example: Configuring a Single-Rate Three-Color Policer

This example shows how to configure a single-rate three-color policer.

- [Requirements on page 4916](#)
- [Overview on page 4916](#)
- [Configuration on page 4917](#)
- [Verification on page 4920](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

A single-rate three-color policer meters a traffic flow against a bandwidth limit and burst-size limit for guaranteed traffic, plus a second burst-size limit for excess traffic. Traffic that conforms to the limits for guaranteed traffic is categorized as green, and nonconforming traffic falls into one of two categories:

- Nonconforming traffic that does not exceed the burst size for excess traffic is categorized as yellow.

- Nonconforming traffic that exceeds the burst size for excess traffic is categorized as red.

Each category is associated with an action. For green traffic, packets are implicitly set with a loss-priority value of **low** and then transmitted. For yellow traffic, packets are implicitly set with a loss-priority value of **medium-high** and then transmitted. For red traffic, packets are implicitly set with a loss-priority value of **high** and then transmitted. If the policer configuration includes the optional **action** statement (**action loss-priority high then discard**), then packets in a red flow are discarded instead.

You can apply a three-color policer to Layer 3 traffic as a firewall filter policer only. You reference the policer from a stateless firewall filter term, and then you apply the filter to the input or output of a logical interface at the protocol level.

Topology

In this example, you apply a color-aware, single-rate three-color policer to the input IPv4 traffic at logical interface **ge-2/0/5.0**. The IPv4 firewall filter term that references the policer does not apply any packet-filtering. The filter is used only to apply the three-color policer to the interface.

You configure the policer to rate-limit traffic to a bandwidth limit of 40 Mbps and a burst-size limit of 100 KB for green traffic but also allow an excess burst-size limit of 200 KB for yellow traffic. Only nonconforming traffic that exceeds the peak burst-size limit is categorized as red. In this example, you configure the three-color policer action **loss-priority high then discard**, which overrides the implicit marking of red traffic to a **high** loss priority.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configuring a Single-Rate Three-Color Policer on page 4918](#)
- [Configuring an IPv4 Stateless Firewall Filter That References the Policer on page 4919](#)
- [Applying the Filter to the Logical Interface on page 4919](#)

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set firewall three-color-policer srTCM1-ca single-rate color-aware
set firewall three-color-policer srTCM1-ca single-rate committed-information-rate 40m
set firewall three-color-policer srTCM1-ca single-rate committed-burst-size 100k
set firewall three-color-policer srTCM1-ca single-rate excess-burst-size 200k
set firewall three-color-policer srTCM1-ca action loss-priority high then discard
set firewall family inet filter filter-srTCM1ca-all term 1 then three-color-policer single-rate
srTCM1-ca
set class-of-service interfaces ge-2/0/5 unit 0 forwarding-class af
set interfaces ge-2/0/5 unit 0 family inet address 10.20.130.1/24
set interfaces ge-2/0/5 unit 0 family inet filter input filter-srTCM1ca-all
```

Configuring a Single-Rate Three-Color Policer

Step-by-Step Procedure

To configure a single-rate three-color policer:

1. Enable configuration of a three-color policer.

```
[edit]  
user@host# edit firewall three-color-policer srTCM1-ca
```
2. Configure the color mode of the single-rate three-color policer.

```
[edit firewall three-color-policer srTCM1-ca]  
user@host# set single-rate color-aware
```
3. Configure the single-rate guaranteed traffic limits.

```
[edit firewall three-color-policer srTCM1-ca]  
user@host# set single-rate committed-information-rate 40m  
user@host# set single-rate committed-burst-size 100k
```
4. Configure the single-rate burst-size limit that is used to classify nonconforming traffic.

```
[edit firewall three-color-policer srTCM1-ca]  
user@host# set single-rate excess-burst-size 200k
```
5. (Optional) Configure the action for nonconforming traffic.

```
[edit firewall three-color-policer srTCM1-ca]  
user@host# set action loss-priority high then discard
```

For three-color policers, the only configurable action is to discard packets in a red traffic flow. In this example, packets in a red traffic flow have been implicitly marked with a **high** packet loss priority (PLP) level because the traffic flow exceeded the rate-limiting defined by the single rate-limit (specified by the **committed-information-rate 40m** statement) and the larger burst-size limit (specified by the **excess-burst-size 200k** statement). Because the optional **action** statement is included, this example takes the more severe action of discarding packets in a red traffic flow.

Results Confirm the configuration of the hierarchical policer by entering the **show firewall** configuration command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
three-color-policer srTCM1-ca {  
  action {  
    loss-priority high then discard;  
  }  
  single-rate {  
    color-aware;  
    committed-information-rate 40m;  
    committed-burst-size 100k;  
    excess-burst-size 200k;  
  }  
}
```

*Configuring an IPv4 Stateless Firewall Filter That References the Policier***Step-by-Step Procedure**

To configure a standard stateless firewall filter that references the policier:

1. Enable configuration of an IPv4 standard stateless firewall filter.

```
[edit]
user@host# edit firewall family inet filter filter-srtcm1ca-all
```

2. Specify the filter term that references the policier.

```
[edit firewall family inet filter filter-srtcm1ca-all]
user@host# set term 1 then three-color-policer single-rate srTCM1-ca
```

Note that the term does not specify any match conditions. The firewall filter passes all packets to the policier.

Results

Confirm the configuration of the firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
  filter filter-srtcm1ca-all {
    term 1 {
      then {
        three-color-policer {
          single-rate srTCM1-ca;
        }
      }
    }
  }
}
three-color-policer srTCM1-ca {
  action {
    loss-priority high then discard;
  }
  single-rate {
    color-aware;
    committed-information-rate 40m;
    committed-burst-size 100k;
    excess-burst-size 200k;
  }
}
```

*Applying the Filter to the Logical Interface***Step-by-Step Procedure**

To apply the filter to the logical interface:

1. (MX Series routers and EX Series switches only) (Optional) Reclassify all incoming packets on the logical interface **ge-2/0/5.0** to assured forwarding, regardless of any preexisting classification.

```
[edit]
user@host# set class-of-service interfaces ge-2/0/5 unit 0 forwarding-class af
```

The classifier name can be a configured classifier or one of the default classifiers.

2. Enable configuration of the logical interface.

```
[edit]
user@host# edit interfaces ge-2/0/5 unit 0 family inet
```

3. Configure an IP address.

```
[edit interfaces ge-2/0/5 unit 0 family inet]
user@host# set address 10.20.130.1/24
```

4. Reference the filter as an input filter.

```
[edit interfaces ge-2/0/5 unit 0 family inet]
user@host# set filter input filter-srtcm1ca-all
```

Results Confirm the configuration of the interface by entering the **show class-of-service** and **show interfaces** configuration mode commands. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show class-of-service
interfaces {
  ge-2/0/5 {
    unit 0 {
      forwarding-class af;
    }
  }
}
[edit]
user@host# show interfaces
ge-2/0/5 {
  unit 0 {
    family inet {
      filter {
        input filter-srtcm1ca-all;
      }
      address 10.20.130.1/24;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Displaying the Firewall Filters Applied to the Logical Interface

Purpose Verify that the firewall filter is applied to IPv4 input traffic at the logical interface.

Action Use the **show interfaces** operational mode command for the logical interface **ge-2/0/5.0**, and specify **detail** mode. The **Protocol inet** section of the command output displays IPv4

information for the logical interface. Within that section, the **Input Filters** field displays the name of the firewall filter applied to IPv4 input traffic at the logical interface.

```
user@host> show interfaces ge-2/0/5.0 detail
Logical interface ge-2/0/5.0 (Index 105) (SNMP ifIndex 556) (Generation 170)
Flags: Device-Down SNMP-Traps 0x4004000 Encapsulation: ENET2
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
Protocol inet, MTU: 1500, Generation: 242, Route table: 0
Flags: Sendbcst-pkt-to-re
Input Filters: filter-srtcm1ca-all
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
  Destination: 10.20.130/24, Local: 10.20.130.1, Broadcast: 10.20.130.255,
  Generation: 171
Protocol multiservice, MTU: Unlimited, Generation: 243, Route table: 0
Policer: Input: __default_arp_policer__
```

- Related Documentation**
- [Statement Hierarchy for Configuring Policers on page 4811](#)
 - [Three-Color Policer Configuration Overview on page 4817](#)
 - [Three-Color Policer Configuration Guidelines on page 4912](#)

Basic Two-Rate Three-Color Policers

- [Two-Rate Three-Color Policer Overview on page 4921](#)
- [Example: Configuring a Two-Rate Three-Color Policer on page 4922](#)

Two-Rate Three-Color Policer Overview

A two-rate three-color policer defines two bandwidth limits (one for guaranteed traffic and one for peak traffic) and two burst sizes (one for each of the bandwidth limits). A two-rate three-color policer is most useful when a service is structured according to arrival rates and not necessarily packet length.

Two-rate three-color policing meters a traffic stream based on the following configured traffic criteria:

- Committed information rate (CIR)—Bandwidth limit for guaranteed traffic.
- Committed burst size (CBS)—Maximum packet size permitted for bursts of data that exceed the CIR.

- Peak information rate (PIR)—Bandwidth limit for peak traffic.
- Peak burst size (PBS)—Maximum packet size permitted for bursts of data that exceed the PIR.

Two-rate tricolor marking (two-rate TCM) classifies traffic as belonging to one of three color categories and performs congestion-control actions on the packets based on the color marking:

- Green—Traffic that conforms to the bandwidth limit and burst size for guaranteed traffic (CIR and CBS). For a green traffic flow, two-rate TCM marks the packets with an implicit loss priority of **low** and transmits the packets.
- Yellow—Traffic that exceeds the bandwidth limit or burst size for guaranteed traffic (CIR or CBS) but not the bandwidth limit and burst size for peak traffic (PIR and PBS). For a yellow traffic flow, two-rate TCM marks packets with an implicit loss priority of **medium-high** and transmits the packets.
- Red—Traffic that exceeds the bandwidth limit and burst size for peak traffic (PIR and PBS). For a red traffic flow, two-rate TCM marks packets with an implicit loss priority of **high** and, optionally, discards the packets.

If congestion occurs downstream, the packets with higher loss priority are more likely to be discarded.



NOTE: For both single-rate and two-rate three-color policers, the only *configurable* action is to discard packets in a red traffic flow.

For a tricolor marking policer referenced by a firewall filter term, the **discard** policing action is supported on the following routing and switching platforms:

- EX Series switches
- M7i and M10i routers with the Enhanced CFEB (CFEB-E)
- M120 and M320 routers with Enhanced-III FPCs
- MX Series routers with Trio MPCs

To apply a tricolor marking policer on these routing platforms, it is not necessary to include the **logical-interface-policer** statement.

Example: Configuring a Two-Rate Three-Color Policer

This example shows how to configure a two-rate three-color policer.

- [Requirements on page 4923](#)
- [Overview on page 4923](#)
- [Configuration on page 4923](#)
- [Verification on page 4926](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

A two-rate three-color policer meters a traffic flow against a bandwidth limit and burst-size limit for guaranteed traffic, plus a bandwidth limit and burst-size limit for peak traffic. Traffic that conforms to the limits for guaranteed traffic is categorized as green, and nonconforming traffic falls into one of two categories:

- Nonconforming traffic that does not exceed peak traffic limits is categorized as yellow.
- Nonconforming traffic that exceeds peak traffic limits is categorized as red.

Each category is associated with an action. For green traffic, packets are implicitly set with a loss-priority value of **low** and then transmitted. For yellow traffic, packets are implicitly set with a loss-priority value of **medium-high** and then transmitted. For red traffic, packets are implicitly set with a loss-priority value of **high** and then transmitted. If the policer configuration includes the optional **action** statement (**action loss-priority high then discard**), then packets in a red flow are discarded instead.

You can apply a three-color policer to Layer 3 traffic as a firewall filter policer only. You reference the policer from a stateless firewall filter term, and then you apply the filter to the input or output of a logical interface at the protocol level.

Topology

In this example, you apply a color-aware, two-rate three-color policer to the input IPv4 traffic at logical interface **fe-0/1/1.0**. The IPv4 firewall filter term that references the policer does not apply any packet-filtering. The filter is used only to apply the three-color policer to the interface.

You configure the policer to rate-limit traffic to a bandwidth limit of 40 Mbps and a burst-size limit of 100 KB for green traffic, and you configure the policer to also allow a peak bandwidth limit of 60 Mbps and a peak burst-size limit of 200 KB for yellow traffic. Only nonconforming traffic that exceeds the peak traffic limits is categorized as red. In this example, you configure the three-color policer action **loss-priority high then discard**, which overrides the implicit marking of red traffic to a **high** loss priority.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configuring a Two-Rate Three-Color Policer on page 4924](#)
- [Configuring an IPv4 Stateless Firewall Filter That References the Policer on page 4925](#)
- [Applying the Filter to a Logical Interface at the Protocol Family Level on page 4926](#)

CLI Quick Configuration To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the [edit] hierarchy level.

```
set firewall three-color-policer trTCM1-ca two-rate color-aware
set firewall three-color-policer trTCM1-ca two-rate committed-information-rate 40m
set firewall three-color-policer trTCM1-ca two-rate committed-burst-size 100k
set firewall three-color-policer trTCM1-ca two-rate peak-information-rate 60m
set firewall three-color-policer trTCM1-ca two-rate peak-burst-size 200k
set firewall three-color-policer trTCM1-ca action loss-priority high then discard
set firewall family inet filter filter-trtcm1ca-all term 1 then three-color-policer two-rate
trTCM1-ca
set interfaces ge-2/0/5 unit 0 family inet address 10.10.10.1/30
set interfaces ge-2/0/5 unit 0 family inet filter input filter-trtcm1ca-all
set class-of-service interfaces ge-2/0/5 forwarding-class af
```

Configuring a Two-Rate Three-Color Policer

Step-by-Step Procedure To configure a two-rate three-color policer:

1. Enable configuration of a three-color policer.

```
[edit]
user@host# set firewall three-color-policer trTCM1-ca
```

2. Configure the color mode of the two-rate three-color policer.

```
[edit firewall three-color-policer trTCM1-ca]
user@host# set two-rate color-aware
```

3. Configure the two-rate guaranteed traffic limits.

```
[edit firewall three-color-policer trTCM1-ca]
user@host# set two-rate committed-information-rate 40m
user@host# set two-rate committed-burst-size 100k
```

Traffic that does not exceed both of these limits is categorized as green. Packets in a green flow are implicitly set to **low** loss priority and then transmitted.

4. Configure the two-rate peak traffic limits.

```
[edit firewall three-color-policer trTCM1-ca]
user@host# set two-rate peak-information-rate 60m
user@host# set two-rate peak-burst-size 200k
```

Nonconforming traffic that does not exceed both of these limits is categorized as yellow. Packets in a yellow flow are implicitly set to **medium-high** loss priority and then transmitted. Nonconforming traffic that exceeds both of these limits is categorized as red. Packets in a red flow are implicitly set to **high** loss priority.

5. (Optional) Configure the policer action for red traffic.

```
[edit firewall three-color-policer trTCM1-ca]
user@host# set action loss-priority high then discard
```

For three-color policers, the only configurable action is to discard red packets. Red packets are packets that have been assigned high loss priority because they exceeded the peak information rate (PIR) and the peak burst size (PBS).

Results Confirm the configuration of the policer by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
three-color-policer trTCM1-ca {
  action {
    loss-priority high then discard;
  }
  two-rate {
    color-aware;
    committed-information-rate 40m;
    committed-burst-size 100k;
    peak-information-rate 60m;
    peak-burst-size 200k;
  }
}
```

Configuring an IPv4 Stateless Firewall Filter That References the Policer

Step-by-Step Procedure To configure an IPv4 stateless firewall filter that references the policer:

1. Enable configuration of an IPv4 standard stateless firewall filter.

```
[edit]
user@host# set firewall family inet filter filter-trtcm1ca-all
```

2. Specify the filter term that references the policer.

```
[edit firewall family inet filter filter-trtcm1ca-all]
user@host# set term 1 then three-color-policer two-rate trTCM1-ca
```

Note that the term does not specify any match conditions. The firewall filter passes all packets to the policer.

Results Confirm the configuration of the firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
  filter filter-trtcm1ca-all {
    term 1 {
      then {
        three-color-policer {
          two-rate trTCM1-ca;
        }
      }
    }
  }
}
three-color-policer trTCM1-ca {
  action {
    loss-priority high then discard;
  }
}
```

```
two-rate {  
  color-aware;  
  committed-information-rate 40m;  
  committed-burst-size 100k;  
  peak-information-rate 60m;  
  peak-burst-size 200k;  
}  
}
```

Applying the Filter to a Logical Interface at the Protocol Family Level

Step-by-Step Procedure

To apply the filter to the logical interface at the protocol family level:

1. Enable configuration of an IPv4 firewall filter.

```
[edit]  
user@host# edit interfaces ge-2/0/5 unit 0 family inet
```

2. Apply the policer to the logical interface at the protocol family level.

```
[edit interfaces ge-2/0/5 unit 0 family inet]  
user@host# set address 10.10.10.1/30  
user@host# set filter input filter-trtcm1ca-all
```

3. (MX Series routers and EX Series switches only) (Optional) For input policers, you can configure a fixed classifier. A fixed classifier reclassifies all incoming packets, regardless of any preexisting classification.

```
[edit]  
user@host# set class-of-service interfaces ge-2/0/5 forwarding-class af
```

The classifier name can be a configured classifier or one of the default classifiers.

Results Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]  
user@host# show interfaces  
ge-2/0/5 {  
  unit 0 {  
    family inet {  
      address 10.10.10.1/30;  
      filter {  
        input filter-trtcm1ca-all;  
      }  
    }  
  }  
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Displaying the Firewall Filters Applied to the Logical Interface on page 4927](#)

Displaying the Firewall Filters Applied to the Logical Interface

Purpose Verify that the firewall filter is applied to IPv4 input traffic at the logical interface.

Action Use the **show interfaces** operational mode command for the logical interface **ge-2/0/5.0**, and specify **detail** mode. The **Protocol inet** section of the command output displays IPv4 information for the logical interface. Within that section, the **Input Filters** field displays the name of IPv4 firewall filters associated with the logical interface.

```
user@host> show interfaces ge-2/0/5.0 detail
Logical interface ge-2/0/5.0 (Index 105) (SNMP ifIndex 556) (Generation 170)
  Flags: Device-Down SNMP-Traps 0x4004000 Encapsulation: ENET2
  Traffic statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Local statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Transit statistics:
    Input bytes : 0 0 bps
    Output bytes : 0 0 bps
    Input packets: 0 0 pps
    Output packets: 0 0 pps
  Protocol inet, MTU: 1500, Generation: 242, Route table: 0
    Flags: Sendbcst-pkt-to-re
    Input Filters: filter-trtcm1ca-all
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.20.130/24, Local: 10.20.130.1, Broadcast: 10.20.130.255,

    Generation: 171
  Protocol multiservice, MTU: Unlimited, Generation: 243, Route table: 0
    Policers: Input: __default_arp_policer__
```

- Related Documentation**
- [Statement Hierarchy for Configuring Policers on page 4811](#)
 - [Three-Color Policer Configuration Overview on page 4817](#)
 - [Three-Color Policer Configuration Guidelines on page 4912](#)

Configuring Logical and Physical Interface Policers

- [Two-Color and Three-Color Logical Interface Policers on page 4927](#)
- [Two-Color and Three-Color Physical Interface Policers on page 4940](#)

Two-Color and Three-Color Logical Interface Policers

- [Logical Interface \(Aggregate\) Policer Overview on page 4928](#)
- [Example: Configuring a Two-Color Logical Interface \(Aggregate\) Policer on page 4928](#)
- [Example: Configuring a Three-Color Logical Interface \(Aggregate\) Policer on page 4934](#)

Logical Interface (Aggregate) Policer Overview

A *logical interface policer*—also called an *aggregate policer*—is a two-color or three-color policer that defines traffic rate limiting that you can apply to input or output traffic for multiple protocol families on the same logical interface without creating multiple instances of the policer.

To configure a single-rate two-color logical interface policer, include the **logical-interface-policer** statement at one of the following hierarchy levels:

- [edit **firewall policer** *policer-name*]
- [edit logical-systems *logical-system-name* **firewall policer** *policer-name*]

To configure a single-rate or two-rate three-color logical interface policer, include the **logical-interface-policer** statement at one of the following hierarchy levels:

- [edit **firewall three-color-policer** *name*]
- [edit logical-systems *logical-system-name* **firewall three-color-policer** *name*]



NOTE: A three-color policer can be applied to Layer 2 traffic as a logical interface policer only. You cannot apply a three-color policer to Layer 2 traffic as a physical interface policer (through a firewall filter).

You apply a logical interface policer to Layer 3 traffic directly to the interface configuration at the logical unit level (to rate-limit all traffic types, regardless of the protocol family) or at the protocol family level (to rate-limit traffic of a specific protocol family). You cannot reference a logical interface policer from a stateless firewall filter term and then apply the filter to a logical interface.

You can apply a logical interface policer to unicast traffic only. For information about configuring a stateless firewall filter for flooded traffic, see “*Applying Filters to Forwarding Tables*” in the “Traffic Sampling, Forwarding, and Monitoring” section of the *Routing Policy Configuration Guide*.

To display a logical interface policer on a particular interface, issue the **show interfaces policers** operational mode command.

Example: Configuring a Two-Color Logical Interface (Aggregate) Policer

This example shows how to configure a single-rate two-color policer as a logical interface policer and apply it to incoming IPv4 traffic on a logical interface.

- [Requirements on page 4929](#)
- [Overview on page 4929](#)
- [Configuration on page 4929](#)
- [Verification on page 4933](#)

Requirements

Before you begin, make sure that the logical interface to which you apply the two-color logical interface policer is hosted on a Gigabit Ethernet interface (**ge-**) or a 10-Gigabit Ethernet interface (**xe-**).

Overview

In this example, you configure the single-rate two-color policer **policer_IFL** as a logical interface policer and apply it to incoming IPv4 traffic at logical interface **ge-1/3/1.0**.

Topology

If the input IPv4 traffic on the physical interface **ge-1/3/1** exceeds the bandwidth limit equal to 90 percent of the media rate with a 300 KB burst-size limit, then the logical interface policer **policer_IFL** rate-limits the input IPv4 traffic on the logical interface **ge-1/3/1.0**. Configure the policer to mark nonconforming traffic by setting packet loss priority (PLP) levels to **high** and classifying packets as **best-effort**.

As the incoming IPv4 traffic rate on the physical interface slows and conforms to the configured limits, Junos OS stops marking the incoming IPv4 packets at the logical interface.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configuring the Logical Interfaces on page 4930](#)
- [Configuring the Single-Rate Two-Color Policer as a Logical Interface Policer on page 4931](#)
- [Applying the Logical Interface Policer to Input IPv4 Traffic at a Logical Interface on page 4932](#)

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces ge-1/3/1 vlan-tagging
set interfaces ge-1/3/1 unit 0 vlan-id 100
set interfaces ge-1/3/1 unit 0 family inet address 10.10.10.1/30
set interfaces ge-1/3/1 unit 1 vlan-id 101
set interfaces ge-1/3/1 unit 1 family inet address 20.20.20.1/30 arp 20.20.20.2 mac
00:00:11:22:33:44
set firewall policer policer_IFL logical-interface-policer
set firewall policer policer_IFL if-exceeding bandwidth-percent 90
set firewall policer policer_IFL if-exceeding burst-size-limit 300k
set firewall policer policer_IFL then loss-priority high
set firewall policer policer_IFL then forwarding-class best-effort
set interfaces ge-1/3/1 unit 0 family inet policer input policer_IFL
```

Configuring the Logical Interfaces

Step-by-Step Procedure

To configure the logical interfaces:

1. Enable configuration of the interface.

[edit]
user@host# edit interfaces ge-1/3/1
2. Configure single tagging.

[edit interfaces ge-1/3/1]
user@host# set vlan-tagging
3. Configure logical interface **ge-1/3/1.0**.

[edit interfaces ge-1/3/1]
user@host# set unit 0 vlan-id 100
user@host# set unit 0 family inet address 10.10.10.1/30
4. Configure logical interface **ge-1/3/1.0**.

[edit interfaces ge-1/3/1]
user@host# set unit 1 vlan-id 101
user@host# set unit 1 family inet address 20.20.20.1/30 arp 20.20.20.2 mac 00:00:11:22:33:44

Results Confirm the configuration of the logical interfaces by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
ge-1/3/1 {
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    family inet {
      address 10.10.10.1/30;
    }
  }
  unit 1 {
    vlan-id 101;
    family inet {
      address 20.20.20.1/30 {
        arp 20.20.20.2 mac 00:00:11:22:33:44;
      }
    }
  }
}
```


*Configuring the Single-Rate Two-Color Policer as a Logical Interface Policer***Step-by-Step Procedure**

To configure a single-rate two-color policer as a logical interface policer:

1. Enable configuration of a single-rate two-color policer.

```
[edit]
user@host# edit firewall policer policer_IFL
```

2. Specify that the policer is a logical interface (aggregate) policer.

```
[edit firewall policer policer_IFL]
user@host# set logical-interface-policer
```

A logical interface policer rate-limits traffic based on a percentage of the media rate of the physical interface underlying the logical interface to which the policer is applied. The policer is applied directly to the interface rather than referenced by a firewall filter.

3. Specify the policer traffic limits.

- a. Specify the bandwidth limit.

- To specify the bandwidth limit as an absolute rate, from 8,000 bits per second through 50,000,000,000 bits per second, include the **bandwidth-limit *bps*** statement.
- To specify the bandwidth limit as a percentage of the physical port speed on the interface, include the **bandwidth-percent *percent*** statement.

In this example, the CLI commands and output are based on a bandwidth limit specified as a percentage rather than as an absolute rate.

```
[edit firewall policer policer_IFL]
user@host# set if-exceeding bandwidth-percent 90
```

- b. Specify the burst-size limit, from 1,500 bytes through 100,000,000,000 bytes, which is the maximum packet size to be permitted for bursts of data that exceed the specified bandwidth limit.

```
[edit firewall policer policer_IFL]
user@host# set if-exceeding burst-size-limit 300k
```

4. Specify the policer actions to be taken on traffic that exceeds the configured rate limits.

- To discard the packet, include the **discard** statement.
- To set the loss-priority value of the packet, include the **loss-priority (*low* | *medium-low* | *medium-high* | *high*)** statement.
- To classify the packet to a forwarding class, include the **forwarding-class (*forwarding-class* | *assured-forwarding* | *best-effort* | *expedited-forwarding* | *network-control*)** statement.

In this example, the CLI commands and output are based on both setting the packet loss priority level and classifying the packet.

```
[edit firewall policer policer_IFL]
user@host# set then loss-priority high
user@host# set then forwarding-class best-effort
```

Results Confirm the configuration of the policer by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
policer policer_IFL {
  logical-interface-policer;
  if-exceeding {
    bandwidth-percent 90;
    burst-size-limit 300k;
  }
  then {
    loss-priority high;
    forwarding-class best-effort;
  }
}
```

Applying the Logical Interface Policer to Input IPv4 Traffic at a Logical Interface

Step-by-Step Procedure To apply the two-color logical interface policer to input IPv4 traffic a logical interface:

1. Enable configuration of the logical interface.

```
[edit]
user@host# edit interfaces ge-1/3/1 unit 0
```
2. Apply the policer to all traffic types or to a specific traffic type on the logical interface.
 - To apply the policer to all traffic types, regardless of the protocol family, include the **policer (input | output) *policer-name*** statement at the **[edit interfaces *interface-name* unit *number*]** hierarchy level.
 - To apply the policer to traffic of a specific protocol family, include the **policer (input | output) *policer-name*** statement at the **[edit interfaces *interface-name* unit *unit-number* family *family-name*]** hierarchy level.

To apply the logical interface policer to incoming packets, use the **policer input *policer-name*** statement. To apply the logical interface policer to outgoing packets, use the **policer output *policer-name*** statement.

In this example, the CLI commands and output are based on rate-limiting the IPv4 input traffic at logical interface **ge-1/3/1.0**.

```
[edit interfaces ge-1/3/1 unit 0]
user@host# set family inet policer input policer_IFL
```

Results Confirm the configuration of the interface by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
ge-1/3/1 {
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    family inet {
      policer input policer_IFL;
      address 10.10.10.1/30;
    }
  }
  unit 1 {
    vlan-id 101;
    family inet {
      address 20.20.20.1/30 {
        arp 20.20.20.2 mac 00:00:11:22:33:44;
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Displaying Traffic Statistics and Policers for the Logical Interface on page 4933](#)
- [Displaying Statistics for the Policer on page 4934](#)

Displaying Traffic Statistics and Policers for the Logical Interface

Purpose Verify the traffic flow through the logical interface and that the policer is evaluated when packets are received on the logical interface.

Action Use the **show interfaces** operational mode command for logical interface **ge-1/3/1.0**, and include the **detail** or **extensive** option. The command output section for **Traffic statistics** lists the number of bytes and packets received and transmitted on the logical interface. The **Protocol inet** subsection contains a **Policer** field that would list the policer **policer_IFL** as an input or output logical interface policer as follows:

- **Input:** policer_IFL-ge-1/3/1.0-log_int-i
- **Output:** policer_IFL-ge-1/3/1.0-log_int-o

The **log_int-i** suffix denotes a logical interface policer applied to input traffic, while the **log_int-o** suffix denotes a logical interface policer applied to output traffic. In this example, the logical interface policer is applied to input traffic only.

Displaying Statistics for the Policer

Purpose Verify the number of packets evaluated by the policer.

Action Use the **show policer** operational mode command and optionally specify the name of the policer. The command output displays the number of packets evaluated by each configured policer (or the specified policer), in each direction. For the policer **policer_IFL**, the input and output policer names are displayed as follows:

- **policer_IFL-ge-1/3/1.0-log_int-i**
- **policer_IFL-ge-1/3/1.0-log_int-o**

The **log_int-i** suffix denotes a logical interface policer applied to input traffic, while the **log_int-o** suffix denotes a logical interface policer applied to output traffic. In this example, the logical interface policer is applied to input traffic only.

Example: Configuring a Three-Color Logical Interface (Aggregate) Policer

This example shows how to configure a two-rate three-color color-blind policer as a logical interface (aggregate) policer and apply the policer directly to Layer 2 input traffic at a supported logical interface.

- [Requirements on page 4934](#)
- [Overview on page 4934](#)
- [Configuration on page 4935](#)
- [Verification on page 4939](#)

Requirements

Before you begin, make sure that the logical interface to which you apply the three-color logical interface policer is hosted on a Gigabit Ethernet interface (**ge-**) or a 10-Gigabit Ethernet interface (**xe-**) on an MX Series router and EX Series switch.

Overview

A two-rate three-color policer meters a traffic flow against a bandwidth limit and burst-size limit for guaranteed traffic, plus a second set of bandwidth and burst-size limits for peak traffic. Traffic that conforms to the limits for guaranteed traffic is categorized as green, and nonconforming traffic falls into one of two categories:

- Nonconforming traffic that does not exceed the bandwidth and burst-size limits for peak traffic is categorized as yellow.
- Nonconforming traffic that exceeds the bandwidth and burst-size limits for peak traffic is categorized as red.

A logical interface policer defines traffic rate-limiting rules that you can apply to multiple protocol families on the same logical interface without creating multiple instances of the policer.



NOTE: You apply a logical interface policer directly to a logical interface at the logical unit level, and not by referencing the policer in a stateless firewall filter and then applying the filter to the logical interface at the protocol family level.

Topology

In this example, you configure the two-rate three-color policer **trTCM2-cb** as a color-blind logical interface policer and apply the policer to incoming Layer 2 traffic on logical interface **ge-1/3/1.0**.



NOTE: When using a three-color policer to rate-limit Layer 2 traffic, color-aware policing can be applied to egress traffic only.

The policer defines guaranteed traffic rate limits such that traffic that conforms to the bandwidth limit of 40 Mbps with a 100 KB allowance for traffic bursting (based on the token-bucket formula) is categorized as green. As with any policed traffic, the packets in a green flow are implicitly set to a **low** loss priority and then transmitted.

Nonconforming traffic that falls within the peak traffic limits of a 60 Mbps bandwidth limit and a 200 KB allowance for traffic bursting (based on the token-bucket formula) is categorized as yellow. The packets in a yellow traffic flow are implicitly set to a **medium-high** loss priority and then transmitted.

Nonconforming traffic that exceeds the peak traffic limits are categorized as red. The packets in a red traffic flow are implicitly set to a **high** loss priority. In this example, the optional policer action for red traffic (**loss-priority high then discard**) is configured, so packets in a red traffic flow are discarded instead of transmitted.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configuring the Logical Interfaces on page 4936](#)
- [Configuring the Two-Rate Three-Color Policer as a Logical Interface Policer on page 4937](#)
- [Applying the Three-Color Policer to the Layer 2 Input at the Logical Interface on page 4938](#)

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces ge-1/3/1 vlan-tagging
set interfaces ge-1/3/1 unit 0 vlan-id 100
set interfaces ge-1/3/1 unit 0 family inet address 10.10.10.1/30
set interfaces ge-1/3/1 unit 1 vlan-id 101
```

```
set interfaces ge-1/3/1 unit 1 family inet address 20.20.20.1/30 arp 20.20.20.2 mac
00:00:11:22:33:44
set firewall three-color-policer trTCM2-cb logical-interface-policer
set firewall three-color-policer trTCM2-cb two-rate color-blind
set firewall three-color-policer trTCM2-cb two-rate committed-information-rate 40m
set firewall three-color-policer trTCM2-cb two-rate committed-burst-size 100k
set firewall three-color-policer trTCM2-cb two-rate peak-information-rate 60m
set firewall three-color-policer trTCM2-cb two-rate peak-burst-size 200k
set firewall three-color-policer trTCM2-cb action loss-priority high then discard
set interfaces ge-1/3/1 unit 0 layer2-policer input-three-color trTCM2-cb
```

Configuring the Logical Interfaces

Step-by-Step Procedure

To configure the logical interfaces:

1. Enable configuration of the interface.

```
[edit]
user@host# edit interfaces ge-1/3/1
```

2. Configure single tagging.

```
[edit interfaces ge-1/3/1]
user@host# set vlan-tagging
```

3. Configure logical interface **ge-1/3/1.0**.

```
[edit interfaces ge-1/3/1]
user@host# set unit 0 vlan-id 100
user@host# set unit 0 family inet address 10.10.10.1/30
```

4. Configure logical interface **ge-1/3/1.0**.

```
[edit interfaces ge-1/3/1]
user@host# set unit 1 vlan-id 101
user@host# set unit 1 family inet address 20.20.20.1/30 arp 20.20.20.2 mac
00:00:11:22:33:44
```

Results Confirm the configuration of the logical interfaces by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
ge-1/3/1 {
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    family inet {
      address 10.10.10.1/30;
    }
  }
  unit 1 {
    vlan-id 101;
    family inet {
      address 20.20.20.1/30 {
        arp 20.20.20.2 mac 00:00:11:22:33:44;
      }
    }
  }
}
```

```

    }
  }
}

```

Configuring the Two-Rate Three-Color Policer as a Logical Interface Policer

Step-by-Step Procedure

To configure the two-rate three-color policer as a logical interface policer:

1. Enable configuration of a three-color policer.

```

[edit]
user@host# edit firewall three-color-policer trTCM2-cb

```

2. Specify that the policer is a logical interface (aggregate) policer.

```

[edit firewall three-color-policer trTCM2-cb]
user@host# set logical-interface-policer

```

A logical interface policer rate-limits traffic based on a percentage of the media rate of the physical interface underlying the logical interface to which the policer is applied, and the policer is applied directly to the interface rather than referenced by a firewall filter.

3. Specify that the policer is two-rate and color-blind.

```

[edit firewall three-color-policer trTCM2-cb]
user@host# set two-rate color-blind

```

A color-aware three-color policer takes into account any coloring markings that might have been set for a packet by another traffic policer configured at a previous network node, and any preexisting color markings are used in determining the appropriate policing action for the packet.

Because you are applying this three-color policer applied to input at Layer 2, you must configure the policer to be color-blind.

4. Specify the policer traffic limits used to classify a green traffic flow.

```

[edit firewall three-color-policer trTCM2-cb]
user@host# set two-rate committed-information-rate 40m
user@host# set two-rate committed-burst-size 100k

```

5. Specify the additional policer traffic limits used to classify a yellow or red traffic flow.

```

[edit firewall three-color-policer trTCM2-cb]
user@host# set two-rate peak-information-rate 60m
user@host# set two-rate peak-burst-size 200k

```

6. (Optional) Specify the configured policer action for packets in a red traffic flow.

```

[edit firewall three-color-policer trTCM2-cb]
user@host# set action loss-priority high then discard

```

In color-aware mode, the three-color policer configured action can increase the packet loss priority (PLP) level of a packet, but never decrease it. For example, if a color-aware three-color policer meters a packet with a medium PLP marking, it can raise the PLP level to high, but cannot reduce the PLP level to low.

Results Confirm the configuration of the three-color policer by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
three-color-policer trTCM2-cb {
  logical-interface-policer;
  action {
    loss-priority high then discard;
  }
  two-rate {
    color-blind;
    committed-information-rate 40m;
    committed-burst-size 100k;
    peak-information-rate 60m;
    peak-burst-size 200k;
  }
}
```

Applying the Three-Color Policer to the Layer 2 Input at the Logical Interface

Step-by-Step Procedure To apply the three-color policer to the Layer 2 input at the logical interface:

1. Enable application of Layer 2 logical interface policers.

```
[edit]
user@host# edit interfaces ge-1/3/1 unit 0
```

2. Apply the three-color logical interface policer to a logical interface input.

```
[edit interfaces ge-1/3/1 unit 0]
user@host# set layer2-policerinput-three-color trTCM2-cb
```

Results Confirm the configuration of the logical interfaces by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
ge-1/3/1 {
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    layer2-policer {
      input-three-color trTCM2-cb;
    }
    family inet {
      address 10.10.10.1/30;
    }
  }
  unit 1 {
    vlan-id 101;
    family inet {
      address 20.20.20.1/30 {
        arp 20.20.20.2 mac 00:00:11:22:33:44;
      }
    }
  }
}
```



```

    }
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Displaying Traffic Statistics and Policers for the Logical Interface on page 4939](#)
- [Displaying Statistics for the Policer on page 4939](#)

Displaying Traffic Statistics and Policers for the Logical Interface

Purpose Verify the traffic flow through the logical interface and that the policer is evaluated when packets are received on the logical interface.

Action Use the **show interfaces** operational mode command for logical interface **ge-1/3/1.0**, and include the **detail** or **extensive** option. The command output section for **Traffic statistics** lists the number of bytes and packets received and transmitted on the logical interface, and the **Protocol inet** section contains a **Policer** field that would list the policer **trTCM2-cb** as an input or output policer as follows:

- Input: trTCM2-cb-ge-1/3/1.0-log_int-i
- Output: trTCM2-cb-ge-1/3/1.0-log_int-o

The **log_int-i** suffix denotes a logical interface policer applied to input traffic, while the **log_int-o** suffix denotes a logical interface policer applied to output traffic. In this example, the logical interface policer is applied to in the input direction only.

Displaying Statistics for the Policer

Purpose Verify the number of packets evaluated by the policer.

Action Use the **show policer** operational mode command and optionally specify the name of the policer. The command output displays the number of packets evaluated by each configured policer (or the specified policer), in each direction. For the policer **trTCM2-cb**, the input and output policer names are displayed as follows:

- trTCM2-cb-ge-1/3/1.0-log_int-i
- trTCM2-cb-e-1/3/1.0-log_int-o

The **log_int-i** suffix denotes a logical interface policer applied to input traffic, while the **log_int-o** suffix denotes a logical interface policer applied to output traffic. In this example, the logical interface policer is applied to input traffic only.

Related Documentation

- [Statement Hierarchy for Configuring Policers on page 4811](#)
- [Two-Color Policer Configuration Overview on page 4813](#)

- [Three-Color Policer Configuration Overview on page 4817](#)
- [Guidelines for Applying Traffic Policers on page 4821](#)

Two-Color and Three-Color Physical Interface Policers

- [Physical Interface Policer Overview on page 4940](#)
- [Example: Configuring a Physical Interface Policer for Aggregate Traffic at a Physical Interface on page 4941](#)

Physical Interface Policer Overview

A *physical interface policer* is a two-color or three-color policer that defines traffic rate limiting that you can apply to input or output traffic for all the logical interfaces and protocol families configured on a physical interface, even if the logical interfaces belong to different routing instances. This feature is useful when you want to perform aggregate policing for different protocol families and different logical interfaces on the same physical interface.

For example, suppose that a provider edge (PE) router has numerous logical interfaces, each corresponding to a different customer, configured on the same link to a customer edge (CE) device. Now suppose that a customer wants to apply one set of rate limits aggregately for certain types of traffic on a single physical interface. To accomplish this, you could apply a single physical interface policer to the physical interface, which rate-limits all the logical interfaces configured on the interface and all the routing instances to which those interfaces belong.

To configure a single-rate two-color physical interface policer, include the **physical-interface-policer** statement at one of the following hierarchy levels:

- [edit **firewall policer** *policer-name*]
- [edit logical-system *logical-system-name* **firewall policer** *policer-name*]
- [edit routing-instances *routing-instance-name* **firewall policer** *policer-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* **firewall policer** *policer-name*]

To configure a single-rate or two-rate three-color physical interface policer, include the **physical-interface-policer** statement at one of the following hierarchy levels:

- [edit **firewall three-color-policer** *policer-name*]
- [edit logical-system *logical-system-name* **firewall three-color-policer** *policer-name*]
- [edit routing-instances *routing-instance-name* **firewall three-color-policer** *policer-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* **firewall three-color-policer** *policer-name*]

You apply a physical interface policer to Layer 3 traffic by referencing the policer from a stateless firewall filter term and then applying the filter to a logical interface. You cannot apply a physical interface to Layer 3 traffic directly to the interface configuration.

To reference a single-rate two-color policer from a stateless firewall filter term, use the **policer** nonterminating action. To reference a single-rate or two-rate three-color policer from a stateless firewall filter term, use the **three-color-policer** nonterminating action.

The following requirements apply to a stateless firewall filter that references a physical interface policer:

- You must configure the firewall filter for a specific, supported protocol family: **ipv4**, **ipv6**, **mpls**, **vpls**, or circuit cross-connect (**ccc**), but not for **family any**.
- You must configure the firewall filter as a *physical interface filter* by including the **physical-interface-filter** statement at the **[edit firewall family family-name filter filter-name]** hierarchy level.
- A firewall filter that is defined as a physical interface filter can reference a physical interface policer only.
- A firewall filter that is defined as a physical interface filter cannot reference a policer configured with the **interface-specific** statement.
- You cannot configure a firewall filter as both a physical interface filter and as a logical interface filter that also includes the **interface-specific** statement.

Example: Configuring a Physical Interface Policer for Aggregate Traffic at a Physical Interface

This example shows how to configure a single-rate two-color policer as a physical interface policer.

- [Requirements on page 4941](#)
- [Overview on page 4941](#)
- [Configuration on page 4942](#)
- [Verification on page 4946](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

A *physical interface policer* specifies rate-limiting for aggregate traffic, which encompasses all protocol families and logical interfaces configured on a physical interface, even if the interfaces belong to different routing instances.

You can apply a physical interface policer to Layer 3 input or output traffic only by referencing the policer from a stateless firewall filter that is configured for specific a specific protocol family (not for **family any**) and configured as a physical interface filter. You configure the filter terms with match conditions that select the types of packets you want to rate-limit, and you specify the physical interface policer as the action to apply to matched packets.

Topology

The physical interface policer in this example, **shared-policer-A**, rate-limits to 10,000,000 bps and permits a maximum burst of traffic of 500,000 bytes. You configure the policer to discard packets in nonconforming flows, but you could instead configure the policer to re-mark nonconforming traffic with a forwarding class, a packet loss priority (PLP) level, or both.

To be able to use the policer to rate-limit IPv4 traffic, you reference the policer from an IPv4 physical interface filter. For this example, you configure the filter to pass the policer IPv4 packets that meet either of the following match terms:

- Packets received through TCP and with the IP precedence fields **critical-ecp** (0xa0), **immediate** (0x40), or **priority** (0x20)
- Packets received through TCP and with the IP precedence fields **internet-control** (0xc0) or **routine** (0x00)

You could also reference the policer from physical interface filters for other protocol families.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configuring the Logical Interfaces on the Physical Interface on page 4943](#)
- [Configuring a Physical Interface Policer on page 4943](#)
- [Configuring an IPv4 Physical Interface Filter on page 4944](#)
- [Applying the IPv4 Physical interface Filter to a Physical Interface on page 4945](#)

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the [edit] hierarchy level.

```
set interfaces so-1/0/0 unit 0 family inet address 192.168.1.1/24
set interfaces so-1/0/0 unit 0 family vpls
set interfaces so-1/0/0 unit 1 family mpls
set firewall policer shared-policer-A physical-interface-policer
set firewall policer shared-policer-A if-exceeding bandwidth-limit 100m burst-size-limit 500k
set firewall policer shared-policer-A then discard
set firewall family inet filter ipv4-filter physical-interface-filter
set firewall family inet filter ipv4-filter term tcp-police-1 from precedence [ critical-ecp immediate priority ]
set firewall family inet filter ipv4-filter term tcp-police-1 from protocol tcp
set firewall family inet filter ipv4-filter term tcp-police-1 then policer shared-policer-A
set firewall family inet filter ipv4-filter term tcp-police-2 from precedence [ internet-control routine ]
set firewall family inet filter ipv4-filter term tcp-police-2 from protocol tcp
set firewall family inet filter ipv4-filter term tcp-police-2 then policer shared-policer-A
```

```
set interfaces so-1/0/0 unit 0 family inet filter input ipv4-filter
```

Configuring the Logical Interfaces on the Physical Interface

Step-by-Step Procedure

To configure the logical interfaces on the physical interface:

1. Enable configuration of logical interfaces.

[edit]
user@host# edit interfaces so-1/0/0
2. Configure protocol families on logical unit 0.

[edit interfaces so-1/0/0]
user@host# set unit 0 family inet address 192.168.1.1/24
user@host# set unit 0 family vpls
3. Configure protocol families on logical unit 1.

[edit interfaces so-1/0/0]
user@host# set unit 1 family mpls

Results

Confirm the configuration of the firewall filter by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
so-1/0/0 {
  unit 0 {
    family inet {
      address 192.168.1.1/24;
    }
    family vpls;
  }
  unit 1 {
    family mpls;
  }
}
```

Configuring a Physical Interface Policer

Step-by-Step Procedure

To configure a physical interface policer:

1. Enable configuration of the two-color policer.

[edit]
user@host# edit firewall policer shared-policer-A
2. Configure the type of two-color policer.

[edit firewall policer shared-policer-A]
user@host# set physical-interface-policer
3. Configure the traffic limits and the action for packets in a nonconforming traffic flow.

[edit firewall policer shared-policer-A]

```
user@host# set if-exceeding bandwidth-limit 100m burst-size-limit 500k
user@host# set then discard
```

For a physical interface filter, the actions you can configure for packets in a nonconforming traffic flow are to discard the packets, assign a forwarding class, assign a PLP value, or assign both a forwarding class and a PLP value.

Results Confirm the configuration of the policer by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
policer shared-policer-A {
  physical-interface-policer;
  if-exceeding {
    bandwidth-limit 100m;
    burst-size-limit 500k;
  }
  then discard;
}
```

Configuring an IPv4 Physical Interface Filter

Step-by-Step Procedure To configure a physical interface policer as the action for terms in an IPv4 physical interface policer:

1. Configure a standard stateless firewall filter under a specific protocol family.

```
[edit]
user@host# edit firewall family inet filter ipv4-filter
```

You cannot configure a physical interface firewall filter for **family any**.

2. Configure the filter as a physical interface filter so that you can apply the physical interface policer as an action.

```
[edit firewall family inet filter ipv4-filter]
user@host# set physical-interface-filter
```

3. Configure the first term to match IPv4 packets received through TCP with the IP precedence fields **critical-ecp**, **immediate**, or **priority** and to apply the physical interface policer as a filter action.

```
[edit firewall family inet filter ipv4-filter]
user@host# set term tcp-police-1 from precedence [ critical-ecp immediate priority ]
user@host# set term tcp-police-1 from protocol tcp
user@host# set term tcp-police-1 then policer shared-policer-A
```

4. Configure the first term to match IPv4 packets received through TCP with the IP precedence fields **internet-control** or **routine** and to apply the physical interface policer as a filter action.

```
[edit firewall family inet filter ipv4-filter]
user@host# set term tcp-police-2 from precedence [ internet-control routine ]
user@host# set term tcp-police-2 from protocol tcp
user@host# set term tcp-police-2 then policer shared-policer-A
```

Results Confirm the configuration of the firewall filter by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
family inet {
  filter ipv4-filter {
    physical-interface-filter;
    term tcp-police-1 {
      from {
        precedence [ critical-ecp immediate priority ];
        protocol tcp;
      }
      then policer shared-policer-A;
    }
    term tcp-police-2 {
      from {
        precedence [ internet-control routine ];
        protocol tcp;
      }
      then policer shared-policer-A;
    }
  }
}
policer shared-policer-A {
  physical-interface-policer;
  if-exceeding {
    bandwidth-limit 100m;
    burst-size-limit 500k;
  }
  then discard;
}
```

Applying the IPv4 Physical interface Filter to a Physical Interface

Step-by-Step Procedure To apply the physical interface filter to a physical interface:

1. Enable configuration of IPv4 on the logical interface.

```
[edit]
user@host# edit interfaces so-1/0/0 unit 0 family inet
```

2. Apply the IPv4 physical interface filter in the input direction.

```
[edit interfaces so-1/0/0 unit 0 family inet]
user@host# set filter input ipv4-filter
```

Results Confirm the configuration of the firewall filter by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
so-1/0/0 {
  unit 0 {
```

```

family inet {
  filter {
    input ipv4-filter;
  }
  address 192.168.1.1/24;
}
family vpls;
}
unit 1 {
  family mpls;
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Displaying the Firewall Filters Applied to an Interface on page 4946](#)
- [Displaying the Number of Packets Processed by the Policer at the Logical Interface on page 4946](#)

Displaying the Firewall Filters Applied to an Interface

Purpose Verify that the firewall filter **ipv4-filter** is applied to the IPv4 input traffic at logical interface **so-1/0/0.0**.

Action Use the **show interfaces statistics** operational mode command for logical interface **so-1/0/0.0**, and include the **detail** option. In the **Protocol inet** section of the command output, the **Input Filters** field shows that the firewall filter **ipv4-filter** is applied in the input direction.

```

user@host> show interfaces statistics so-1/0/0 detail
Logical interface so-1/0/0.0 (Index 79) (SNMP ifIndex 510) (Generation 149)
Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
Protocol inet, MTU: 4470, Generation: 173, Route table: 0
Flags: Sendbcast-pkt-to-re, Protocol-Down
Input Filters: ipv4-filter
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 10.39/16, Local: 10.39.1.1, Broadcast: 10.39.255.255,
Generation: 163

```

Displaying the Number of Packets Processed by the Policer at the Logical Interface

Purpose Verify the traffic flow through the logical interface and that the policer is evaluated when packets are received on the logical interface.

Action Use the **show firewall** operational mode command for the filter you applied to the logical interface.

```

user@host> show firewall filter ipv4-filter
Filter: ipv4-filter
Policers:
Name                                     Packets

```


shared-policer-A-tcp-police-1	32863
shared-policer-A-tcp-police-2	3870

The command output displays the name of policer (**shared-policer-A**), the name of the filter term (**police-1**) under which the policer action is specified, and the number of packets that matched the filter term. This is only the number of out-of-specification (out-of-spec) packet counts, not all packets policed by the policer.

Related Documentation

- [Firewall Filter Match Conditions Based on Numbers or Text Aliases on page 4488](#)
- [Firewall Filter Match Conditions Based on Bit-Field Values on page 4489](#)
- [Firewall Filter Match Conditions Based on Address Fields on page 4493](#)
- [Firewall Filter Match Conditions Based on Address Classes on page 4502](#)
- [Statement Hierarchy for Configuring Policers on page 4811](#)
- [Two-Color Policer Configuration Overview on page 4813](#)
- [Three-Color Policer Configuration Overview on page 4817](#)
- [Guidelines for Applying Traffic Policers on page 4821](#)
- [physical-interface-filter on page 5002](#)
- [physical-interface-policer on page 5003](#)

Configuring Layer 2 Policers

- [Two-Color and Three-Color Policers at Layer 2 on page 4947](#)

Two-Color and Three-Color Policers at Layer 2

- [Two-Color Policing at Layer 2 Overview on page 4947](#)
- [Three-Color Policing at Layer 2 Overview on page 4949](#)
- [Example: Configuring a Three-Color Logical Interface \(Aggregate\) Policer on page 4950](#)

Two-Color Policing at Layer 2 Overview

This topic covers the following information:

- [Guidelines for Configuring Two-Color Policing of Layer 2 Traffic on page 4947](#)
- [Statement Hierarchy for Configuring a Two-Color Policer for Layer 2 Traffic on page 4948](#)
- [Statement Hierarchy for Applying a Two-Color Policer to Layer 2 Traffic on page 4948](#)

Guidelines for Configuring Two-Color Policing of Layer 2 Traffic

The following guidelines apply to two-color policing of Layer 2 traffic:

- You can apply a two-color policer to ingress or egress Layer 2 traffic at a logical interface hosted on a Gigabit Ethernet interface (**ge-**) or a 10-Gigabit Ethernet interface (**xe-**) only.
- A single logical interface supports Layer 2 policing in both directions.

- You can apply a two-color policer to Layer 2 traffic as a logical interface policer only. You cannot apply a two-color policer to Layer 2 traffic as a stateless firewall filter action.
- You can apply a two-color policer to Layer 2 traffic by referencing the policer in the interface configuration at the logical unit level, and not at the protocol level.

For information about configuring three-color policing of Layer 2 traffic, see [“Three-Color Policing at Layer 2 Overview” on page 4949](#).

Statement Hierarchy for Configuring a Two-Color Policer for Layer 2 Traffic

To enable a single-rate two-color policer to rate-limit Layer 2 traffic, include the **logical-interface-policer** statement in the **policer** configuration.

```
firewall {
  policer policer-name {
    logical-interface-policer;
    if-exceeding {
      (bandwidth-limit bps | bandwidth-percent percentage);
      burst-size-limit bytes;
    }
    then {
      discard;
      forwarding-class class-name;
      loss-priority (high | low | medium-high | medium-low);
    }
  }
}
```

You can include the configuration at the following hierarchy levels:

- [edit]
- [edit logical-systems *logical-system-name*]

Statement Hierarchy for Applying a Two-Color Policer to Layer 2 Traffic

To apply a logical interface policer to Layer 2 traffic, include the **layer2-policer input-policer *policer-name*** statement or the **layer2-policer output-policer *policer-name*** statement to a supported logical interface. Use the **input-policer** or **output-policer** statements to apply a two-color policer at Layer 2.

```
interfaces {
  (ge-fpc/pic/port | xe-fpc/pic/port) {
    unit unit-number {
      layer2-policer {
        input-policer policer-name;
        output-policer policer-name;
      }
    }
  }
}
```

You can include the configuration at the following hierarchy levels:

- [edit]

- `[edit logical-systems logical-system-name]`

Three-Color Policing at Layer 2 Overview

This topic covers the following information:

- [Guidelines for Configuring Three-Color Policing of Layer 2 Traffic on page 4949](#)
- [Statement Hierarchy for Configuring a Three-Color Policer for Layer 2 Traffic on page 4949](#)
- [Statement Hierarchy for Applying a Three-Color Policer to Layer 2 Traffic on page 4950](#)

Guidelines for Configuring Three-Color Policing of Layer 2 Traffic

The following guidelines apply to three-color policing of Layer 2 traffic:

- You can apply a three-color policer to Layer 2 traffic at a logical interface hosted on a Gigabit Ethernet interface (**ge-**) or a 10-Gigabit Ethernet interface (**xe-**) only.
- A single logical interface supports Layer 2 policing in both directions.
- You can apply a three-color policer to Layer 2 traffic as a logical interface policer only. You cannot apply a two-color policer to Layer 2 traffic as a stateless firewall filter action.
- You can apply a three-color policer to Layer 2 traffic by referencing the policer in the interface configuration at the logical unit level, and not at the protocol level.
- You can apply a color-aware three-color policer to Layer 2 traffic in the egress direction only, but you apply a color-blind three-color policer to Layer 2 traffic in either direction.

For information about configuring two-color policing of Layer 2 traffic, see [“Two-Color Policing at Layer 2 Overview” on page 4947](#).

Statement Hierarchy for Configuring a Three-Color Policer for Layer 2 Traffic

To enable a single-rate or two-rate three-color policer to rate-limit Layer 2 traffic, include the **logical-interface-policer** statement in the **three-color-policer** configuration.

```
firewall {
  three-color-policer policer-name {
    action {
      loss-priority high then discard;
    }
    logical-interface-policer;
    single-rate {
      (color-aware | color-blind);
      committed-burst-size bytes;
      committed-information-rate bps;
      excess-burst-size bytes;
    }
    two-rate {
      (color-aware | color-blind);
      committed-burst-size bytes;
      committed-information-rate bps;
      peak-burst-size bytes;
      peak-information-rate bps;
    }
  }
}
```

```
}
```

You can include the configuration at the following hierarchy levels:

- **[edit]**
- **[edit logical-systems *logical-system-name*]**

Statement Hierarchy for Applying a Three-Color Policer to Layer 2 Traffic

To apply a logical interface policer to Layer 2 traffic, include the **layer2-policer** statement for a supported logical interface at the logical unit level. Use the **input-three-color *policer-name*** statement or **output-three-color *policer-name*** statement to specify the direction of the traffic to be policed.

```
interfaces {  
  (ge-fpc/pic/port | xe-fpc/pic/port) {  
    unit unit-number {  
      layer2-policer {  
        input-three-color policer-name;  
        output-three-color policer-name;  
      }  
    }  
  }  
}
```

You can include the configuration at the following hierarchy levels:

- **[edit]**
- **[edit logical-systems *logical-system-name*]**

Example: Configuring a Three-Color Logical Interface (Aggregate) Policer

This example shows how to configure a two-rate three-color color-blind policer as a logical interface (aggregate) policer and apply the policer directly to Layer 2 input traffic at a supported logical interface.

- [Requirements on page 4950](#)
- [Overview on page 4951](#)
- [Configuration on page 4952](#)
- [Verification on page 4955](#)

Requirements

Before you begin, make sure that the logical interface to which you apply the three-color logical interface policer is hosted on a Gigabit Ethernet interface (**ge-**) or a 10-Gigabit Ethernet interface (**xe-**) on an MX Series router and EX Series switch.

Overview

A two-rate three-color policer meters a traffic flow against a bandwidth limit and burst-size limit for guaranteed traffic, plus a second set of bandwidth and burst-size limits for peak traffic. Traffic that conforms to the limits for guaranteed traffic is categorized as green, and nonconforming traffic falls into one of two categories:

- Nonconforming traffic that does not exceed the bandwidth and burst-size limits for peak traffic is categorized as yellow.
- Nonconforming traffic that exceeds the bandwidth and burst-size limits for peak traffic is categorized as red.

A logical interface policer defines traffic rate-limiting rules that you can apply to multiple protocol families on the same logical interface without creating multiple instances of the policer.



NOTE: You apply a logical interface policer directly to a logical interface at the logical unit level, and not by referencing the policer in a stateless firewall filter and then applying the filter to the logical interface at the protocol family level.

Topology

In this example, you configure the two-rate three-color policer **trTCM2-cb** as a color-blind logical interface policer and apply the policer to incoming Layer 2 traffic on logical interface **ge-1/3/1.0**.



NOTE: When using a three-color policer to rate-limit Layer 2 traffic, color-aware policing can be applied to egress traffic only.

The policer defines guaranteed traffic rate limits such that traffic that conforms to the bandwidth limit of 40 Mbps with a 100 KB allowance for traffic bursting (based on the token-bucket formula) is categorized as green. As with any policed traffic, the packets in a green flow are implicitly set to a **low** loss priority and then transmitted.

Nonconforming traffic that falls within the peak traffic limits of a 60 Mbps bandwidth limit and a 200 KB allowance for traffic bursting (based on the token-bucket formula) is categorized as yellow. The packets in a yellow traffic flow are implicitly set to a **medium-high** loss priority and then transmitted.

Nonconforming traffic that exceeds the peak traffic limits are categorized as red. The packets in a red traffic flow are implicitly set to a **high** loss priority. In this example, the optional policer action for red traffic (**loss-priority high then discard**) is configured, so packets in a red traffic flow are discarded instead of transmitted.

Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#).

To configure this example, perform the following tasks:

- [Configuring the Logical Interfaces on page 4952](#)
- [Configuring the Two-Rate Three-Color Policer as a Logical Interface Policer on page 4953](#)
- [Applying the Three-Color Policer to the Layer 2 Input at the Logical Interface on page 4954](#)

CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces ge-1/3/1 vlan-tagging
set interfaces ge-1/3/1 unit 0 vlan-id 100
set interfaces ge-1/3/1 unit 0 family inet address 10.10.10.1/30
set interfaces ge-1/3/1 unit 1 vlan-id 101
set interfaces ge-1/3/1 unit 1 family inet address 20.20.20.1/30 arp 20.20.20.2 mac
00:00:11:22:33:44
set firewall three-color-policer trTCM2-cb logical-interface-policer
set firewall three-color-policer trTCM2-cb two-rate color-blind
set firewall three-color-policer trTCM2-cb two-rate committed-information-rate 40m
set firewall three-color-policer trTCM2-cb two-rate committed-burst-size 100k
set firewall three-color-policer trTCM2-cb two-rate peak-information-rate 60m
set firewall three-color-policer trTCM2-cb two-rate peak-burst-size 200k
set firewall three-color-policer trTCM2-cb action loss-priority high then discard
set interfaces ge-1/3/1 unit 0 layer2-policer input-three-color trTCM2-cb
```

Configuring the Logical Interfaces

Step-by-Step Procedure

To configure the logical interfaces:

1. Enable configuration of the interface.

```
[edit]
user@host# edit interfaces ge-1/3/1
```

2. Configure single tagging.

```
[edit interfaces ge-1/3/1]
user@host# set vlan-tagging
```

3. Configure logical interface **ge-1/3/1.0**.

```
[edit interfaces ge-1/3/1]
user@host# set unit 0 vlan-id 100
user@host# set unit 0 family inet address 10.10.10.1/30
```

4. Configure logical interface **ge-1/3/1.0**.

```
[edit interfaces ge-1/3/1]
user@host# set unit 1 vlan-id 101
user@host# set unit 1 family inet address 20.20.20.1/30 arp 20.20.20.2 mac
00:00:11:22:33:44
```

Results Confirm the configuration of the logical interfaces by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
ge-1/3/1 {
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    family inet {
      address 10.10.10.1/30;
    }
  }
  unit 1 {
    vlan-id 101;
    family inet {
      address 20.20.20.1/30 {
        arp 20.20.20.2 mac 00:00:11:22:33:44;
      }
    }
  }
}
```

Configuring the Two-Rate Three-Color Policer as a Logical Interface Policer

Step-by-Step Procedure To configure the two-rate three-color policer as a logical interface policer:

1. Enable configuration of a three-color policer.

```
[edit]
user@host# edit firewall three-color-policer trTCM2-cb
```

2. Specify that the policer is a logical interface (aggregate) policer.

```
[edit firewall three-color-policer trTCM2-cb]
user@host# set logical-interface-policer
```

A logical interface policer rate-limits traffic based on a percentage of the media rate of the physical interface underlying the logical interface to which the policer is applied, and the policer is applied directly to the interface rather than referenced by a firewall filter.

3. Specify that the policer is two-rate and color-blind.

```
[edit firewall three-color-policer trTCM2-cb]
user@host# set two-rate color-blind
```

A color-aware three-color policer takes into account any coloring markings that might have been set for a packet by another traffic policer configured at a previous network node, and any preexisting color markings are used in determining the appropriate policing action for the packet.

Because you are applying this three-color policer applied to input at Layer 2, you must configure the policer to be color-blind.

- Specify the policer traffic limits used to classify a green traffic flow.

```
[edit firewall three-color-policer trTCM2-cb]
user@host# set two-rate committed-information-rate 40m
user@host# set two-rate committed-burst-size 100k
```

- Specify the additional policer traffic limits used to classify a yellow or red traffic flow.

```
[edit firewall three-color-policer trTCM2-cb]
user@host# set two-rate peak-information-rate 60m
user@host# set two-rate peak-burst-size 200k
```

- (Optional) Specify the configured policer action for packets in a red traffic flow.

```
[edit firewall three-color-policer trTCM2-cb]
user@host# set action loss-priority high then discard
```

In color-aware mode, the three-color policer configured action can increase the packet loss priority (PLP) level of a packet, but never decrease it. For example, if a color-aware three-color policer meters a packet with a medium PLP marking, it can raise the PLP level to high, but cannot reduce the PLP level to low.

Results Confirm the configuration of the three-color policer by entering the **show firewall** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show firewall
three-color-policer trTCM2-cb {
  logical-interface-policer;
  action {
    loss-priority high then discard;
  }
  two-rate {
    color-blind;
    committed-information-rate 40m;
    committed-burst-size 100k;
    peak-information-rate 60m;
    peak-burst-size 200k;
  }
}
```

Applying the Three-Color Policer to the Layer 2 Input at the Logical Interface

Step-by-Step Procedure

To apply the three-color policer to the Layer 2 input at the logical interface:

- Enable application of Layer 2 logical interface policers.

```
[edit]
user@host# edit interfaces ge-1/3/1 unit 0
```

- Apply the three-color logical interface policer to a logical interface input.

```
[edit interfaces ge-1/3/1 unit 0]
user@host# set layer2-policerinput-three-color trTCM2-cb
```


Results Confirm the configuration of the logical interfaces by entering the **show interfaces** configuration mode command. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
[edit]
user@host# show interfaces
ge-1/3/1 {
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    layer2-policer {
      input-three-color trTCM2-cb;
    }
    family inet {
      address 10.10.10.1/30;
    }
  }
  unit 1 {
    vlan-id 101;
    family inet {
      address 20.20.20.1/30 {
        arp 20.20.20.2 mac 00:00:11:22:33:44;
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Displaying Traffic Statistics and Policers for the Logical Interface on page 4955](#)
- [Displaying Statistics for the Policer on page 4956](#)

Displaying Traffic Statistics and Policers for the Logical Interface

Purpose Verify the traffic flow through the logical interface and that the policer is evaluated when packets are received on the logical interface.

Action Use the **show interfaces** operational mode command for logical interface **ge-1/3/1.0**, and include the **detail** or **extensive** option. The command output section for **Traffic statistics** lists the number of bytes and packets received and transmitted on the logical interface, and the **Protocol inet** section contains a **Policer** field that would list the policer **trTCM2-cb** as an input or output policer as follows:

- **Input:** trTCM2-cb-ge-1/3/1.0-log_int-i
- **Output:** trTCM2-cb-ge-1/3/1.0-log_int-o

The **log_int-i** suffix denotes a logical interface policer applied to input traffic, while the **log_int-o** suffix denotes a logical interface policer applied to output traffic. In this example, the logical interface policer is applied to in the input direction only.

Displaying Statistics for the Policer

Purpose Verify the number of packets evaluated by the policer.

Action Use the [show policer](#) operational mode command and optionally specify the name of the policer. The command output displays the number of packets evaluated by each configured policer (or the specified policer), in each direction. For the policer **trTCM2-cb**, the input and output policer names are displayed as follows:

- **trTCM2-cb-ge-1/3/1.0-log_int-i**
- **trTCM2-cb-e-1/3/1.0-log_int-o**

The **log_int-i** suffix denotes a logical interface policer applied to input traffic, while the **log_int-o** suffix denotes a logical interface policer applied to output traffic. In this example, the logical interface policer is applied to input traffic only.

- Related Documentation**
- [Statement Hierarchy for Configuring Policers on page 4811](#)
 - [Guidelines for Applying Traffic Policers on page 4821](#)
 - [layer2-policer on page 1523](#)
 - [logical-interface-policer on page 1500](#)
 - [policer \(Configuring\) on page 1502](#)
 - [three-color-policer \(Configuring\) on page 1508](#)

Configuration Statements

- [\[edit firewall\] Hierarchy Level on page 4956](#)

[edit firewall] Hierarchy Level

Several statements in the **[edit firewall]** hierarchy are valid at numerous locations within the hierarchy. To make the complete hierarchy easier to read, the repeated statements are listed in the following sections, which are referenced at the appropriate locations in “Complete [edit firewall] Hierarchy” on page 325.

- [Common Firewall Actions on page 4957](#)
- [Common IP Firewall Actions on page 4957](#)
- [Common IPv4 Firewall Actions on page 4957](#)
- [Common IP Firewall Match Conditions on page 4958](#)
- [Common IPv4 Firewall Match Conditions on page 4959](#)
- [Common Layer 2 Firewall Match Conditions on page 4959](#)
- [Complete \[edit firewall\] Hierarchy on page 4961](#)

Common Firewall Actions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- [edit firewall family (any | ccc | ethernet-switching | inet | inet6 | mpls | vpls) filter *filter-name* term *term-name* then]
- [edit firewall filter *filter-name* term *term-name* then]

The common firewall actions are as follows:

```
count counter-name;
forwarding-class class-name;
loss-priority (high | low | medium-high | medium-low);
next term;
policer policer-name;
three-color-policer policer-name {
    (single-rate single-rate-policer-name | two-rate two-rate-policer-name);
}
```

Common IP Firewall Actions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- [edit firewall family inet filter *filter-name* term *term-name* then]
- [edit firewall family inet6 filter *filter-name* term *term-name* then]
- [edit firewall filter *filter-name* term *term-name* then]

The common IP firewall actions are as follows:

```
log;
logical-system logical-system-name <routing-instance routing-instance-name>
    <topology topology-name>;
port-mirror;
port-mirror-instance instance-name;
routing-instance routing-instance-name <topology topology-name>;
sample;
service-filter-hit;
syslog;
topology topology-name;
```

Common IPv4 Firewall Actions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- [edit firewall family inet filter *filter-name* term *term-name* then]
- [edit firewall filter *filter-name* term *term-name* then]

The common IP version 4 (IPv4) firewall actions are as follows:

```
(accept | discard <accounting collector-name> | reject <administratively-prohibited |
bad-host-tos | bad-network-tos | fragmentation-needed | host-prohibited |
host-unknown | host-unreachable | network-prohibited | network-unknown |
network-unreachable | port-unreachable | precedence-cutoff | precedence-violation |
protocol-unreachable | source-host-isolated | source-route-failed | tcp-reset>);
ipsec-sa sa-name;
load-balance sa-name;
next-hop-group group-name;
prefix-action action-name;
```

Common IP Firewall Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit firewall\] Hierarchy](#)” on page 325 instead of the statements being repeated.

- **[edit firewall family inet dialer-filter *filter-name* term *term-name* from]** (with the exceptions noted at this level in “[Complete \[edit firewall\] Hierarchy](#)” on page 325)
- **[edit firewall family inet filter *filter-name* term *term-name* from]**
- **[edit firewall family inet6 dialer-filter *filter-name* term *term-name* from]** (with the exceptions noted at this level in “[Complete \[edit firewall\] Hierarchy](#)” on page 325)
- **[edit firewall family inet6 filter *filter-name* term *term-name* from]**
- **[edit firewall filter *filter-name* term *term-name* from]**

The common IP firewall match conditions are as follows:

```
address {
  ip-prefix </prefix-length> <except>;
}
destination-address {
  ip-prefix </prefix-length> <except>;
}
destination-class [ class-names ] | destination-class-except [ class-names ];
(destination-port [ port-names ] | destination-port-except [ port-names ]);
destination-prefix-list {
  list-name <except>;
}
(forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
icmp-code [ codes ] | icmp-code-except [ codes ];
icmp-type [ types ] | icmp-type-except [ types ];
interface interface-name;
(interface-group [ group-names ] | interface-group-except [ group-names ]);
interface-set set-name;
(loss-priority [ priorities ] | loss-priority-except [ priorities ]);
(packet-length [ values ] | packet-length-except [ values ]);
(port [ port-names ] | port-except [ port-names ]);
prefix-list {
  list-name <except>;
}
service-filter-hit;
source-address {
```

```

    ip-prefix </prefix-length> <except>;
  }
  (source-class [ class-names ] | source-class-except [ class-names ]);
  (source-port [ port-names ] | source-port-except [ port-names ]);
  source-prefix-list {
    list-name <except>;
  }
  tcp-established;
  tcp-flags flag;
  tcp-initial;

```

Common IPv4 Firewall Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- **[edit firewall family inet dialer-filter *filter-name* term *term-name* from]** (with the exceptions noted at this level in [“Complete \[edit firewall\] Hierarchy” on page 325](#))
- **[edit firewall family inet filter *filter-name* term *term-name* from]**
- **[edit firewall filter *filter-name* term *term-name* from]**

The common IPv4 firewall match conditions are as follows:

```

  (ah-spi [ values ] | ah-spi-except [ values ]);
  (dscp [ code-point-values ] | dscp-except [ code-point-values ]);
  (esp-spi [ values ] | esp-spi-except [ values ]);
  first-fragment;
  fragment-flags flag;
  (fragment-offset [ offsets ] | fragment-offset-except [ offsets ]);
  (ip-options [ option-names ] | ip-options-except [ option-names ]);
  is-fragment;
  (precedence [ precedence-names ] | precedence-except [ precedence-names ]);
  (protocol [ protocol-names ] | protocol-except [ protocol-names ]);
  (ttl [ ttl-values ] | ttl-except [ ttl-values ]);

```

Common Layer 2 Firewall Match Conditions

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in [“Complete \[edit firewall\] Hierarchy” on page 325](#) instead of the statements being repeated.

- **[edit firewall family ethernet-switching filter *filter-name* term *term-name* from]**
- **[edit firewall family vpls filter *filter-name* term *term-name* from]**

The common Layer 2 firewall match conditions are as follows:

```

  destination-mac-address {
    mac-address <except>;
  }
  (destination-port [ port-names ] | destination-port-except [ port-names ]);
  (dscp [ code-point-values ] | dscp-except [ code-point-values ]);
  (ether-type [ protocol-types ] | ether-type-except [ protocol-types ]);
  (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);

```

```
(icmp-code [ codes ] | icmp-code-except [ codes ]);
icmp-type [ types ] | icmp-type-except [ types ];
(interface-group [ group-names ] | interface-group-except [ group-names ]);
ip-address {
    ip-prefix </prefix-length> <except>;
}
ip-destination-address {
    ip-prefix </prefix-length> <except>;
}
(ip-precedence [ precedence-names ] | ip-precedence-except [ precedence-names ]);
(ip-protocol [ protocol-names ] | ip-protocol-except [ protocol-names ]);
ip-source-address ip-prefix </prefix-length>;
(learn-vlan-lp-priority [ priorities ] | learn-vlan-lp-priority [ priorities ]);
(learn-vlan-id [ vlan-ids ] | learn-vlan-id-except [ vlan-ids ]);
(loss-priority [ priorities ] | loss-priority-except [ priorities ]);
(port [ port-names ] | port-except [ port-names ]);
source-mac-address {
    mac-address <except>;
}
(source-port [ port-names ] | source-port-except [ port-names ]);
tcp-flags flag;
(traffic-type [ broadcast known-unicast multicast unknown-unicast ] |
    traffic-type-except [ broadcast known-unicast multicast unknown-unicast ]);
(user-vlan-lp-priority [ priorities ] | user-vlan-lp-priority [ priorities ]);
(user-vlan-id [ vlan-ids ] | user-vlan-id-except [ vlan-ids ]);
(vlan-ether-type [ protocol-types ] | vlan-ether-type-except [ protocol-types ]);
```

Complete [edit firewall] Hierarchy

```

firewall {
  family (any | ccc | ethernet-switching | inet | inet6 | mpls | vpls) {
    ... the family subhierarchies appear after the main [edit firewall] hierarchy ...
  }
  filter filter-name {
    accounting-profile [ profile-names ];
    enhanced-mode;
    interface-shared-with;
    interface-specific;
    physical-interface-policer;
    term term-name {
      filter filter-name;
      from {
        ... statements in Common IP Firewall Match Conditions on page 322 AND
        statements in Common IPv4 Firewall Match Conditions on page 323 ...
      }
      then {
        ... statements in Common Firewall Actions on page 320 AND
        statements in Common IP Firewall Actions on page 321 AND
        statements in Common IPv4 Firewall Actions on page 321 ...
      }
    }
  }
  hierarchical-policer policer-name {
    aggregate {
      if-exceeding {
        bandwidth-limit bps;
        burst-size-limit bytes;
      }
      then {
        discard;
        forwarding-class class-name;
        loss-priority (high | low | medium-high | medium-low);
      }
    }
    logical-interface-policer;
    physical-interface-policer;
    premium {
      if-exceeding {
        bandwidth-limit bps;
        burst-size-limit bytes;
      }
      then {
        discard;
      }
    }
  }
  shared-bandwidth-policer;
  interface-set interface-set-name {
    interface-name;
  }
  load-balance-group group-name {
    next-hop-group [ group-names ];
  }
}

```

```
}
policer policer-name {
  filter-specific;
  if-exceeding {
    (bandwidth-limit bps | bandwidth-percent percentage);
    burst-size-limit bytes;
  }
  logical-bandwidth-policer;
  logical-interface-policer;
  physical-interface-policer;
  then {
    discard;
    forwarding-class class-name;
    loss-priority (high | low | medium-high | medium-low);
  }
}
three-color-policer policer-name {
  action {
    loss-priority high then discard;
  }
  filter-specific;
  logical-interface-policer;
  physical-interface-policer;
  shared-bandwidth-policer;
  single-rate {
    (color-aware | color-blind);
    committed-burst-size bytes;
    committed-information-rate bps;
    excess-burst-size bytes;
  }
  two-rate {
    (color-aware | color-blind);
    committed-burst-size bytes;
    committed-information-rate bps;
    peak-burst-size bytes;
    peak-information-rate bps;
  }
}
}

firewall {
  family any {
    filter filter-name {
      interface-shared;
      term term-name {
        from {
          (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
          interface interface-name;
          interface-set set-name;
          (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
          (packet-length [ values ] | packet-length-except [ values ]);
        }
        then {
          ... statements in Common Firewall Actions on page 320 PLUS ...
          (accept | discard);
        }
      }
    }
  }
}
```



```

    }
  }
}

firewall {
  family ccc {
    filter filter-name {
      accounting-profile [ profile-names ];
      physical-interface-filter;
      interface-specific;
      term term-name {
        filter filter-name;
        from {
          (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
          (interface-group [ group-names ] | interface-group-except [ group-names ]);
          (learn-vlan-1p-priority [ priorities ] | learn-vlan-1p-priority [ priorities ]);
          (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
          (user-vlan-1p-priority [ priorities ] | user-vlan-1p-priority [ priorities ]);
        }
        then {
          ... statements in Common Firewall Actions on page 320 PLUS ...
          (accept | discard);
          port-mirror-instance instance-name;
        }
      }
    }
  }
}

firewall {
  family ethernet-switching {
    filter filter-name {
      interface-specific;
      term term-name {
        from {
          destination-address {
            ip-prefix</prefix-length>;
          }
          destination-mac-address {
            mac-address;
          }
          destination-port [ port-names ];
          destination-prefix-list {
            list-name;
          }
          dot1q-tag [ tag-values ];
          dot1q-user-priority [ priority-values ];
          dscp [ code-point-values ];
          ether-type [ protocol-names ];
          fragment-flags flag;
          icmp-code [ codes ];
          icmp-type [ types ];
          interface interface-name;
          is-fragment;

```

```

precedence [ precedence-names ];
protocol [ protocol-names ];
source-address {
    ip-prefix < / prefix-length >;
}
source-mac-address {
    mac-address;
}
source-port [ port-names ];
source-prefix-list {
    list-name;
}
tcp-established;
tcp-flags flag;
tcp-initial;
vlan [ vlan-names ];
}
then {
    (accept | discard);
    analyzer analyzer-name;
    count counter-name;
    forwarding-class class-name;
    interface interface-name;
    log;
    loss-priority (high | low);
    policer policer-name;
    syslog;
    vlan vlan-name;
}
}
}
}
}

firewall {
    family inet {
        dialer-filter filter-name {
            accounting-profile [ profile-names ];
            term term-name {
                from {
                    ... statements in Common IP Firewall Match Conditions on page 322 AND
                    statements in Common IPv4 Firewall Match Conditions on page 323 EXCEPT
                    FOR ...
                    (ah-spi [ values ] | ah-spi-except [ values ]); # NOT valid at this level
                    (destination-class [ class-names ] |
                     destination-class-except [ class-names ]); # NOT valid at this level
                    interface interface-name; # NOT valid at this level
                    (loss-priority [ priorities ] | loss-priority-except [ priorities ]); # NOT valid at
                     this level
                    service-filter-hit; # NOT valid at this level
                    (source-class [ class-names ] | source-class-except [ class-names ]); # NOT
                     valid at this level
                }
            }
        }
    }
}

```

```

        sample;
        syslog;
    }
}
filter filter-name {
    accounting-profile [ profile-names ];
    interface-specific;
    term term-name {
        filter filter-name;
        from {
            ... statements in Common IP Firewall Match Conditions on page 322 AND
               statements in Common IPv4 Firewall Match Conditions on page 323 ...
        }
        then {
            ... statements in Common Firewall Actions on page 320 AND
               statements in Common IP Firewall Actions on page 321 AND
               statements in Common IPv4 Firewall Actions on page 321 ...
        }
    }
}
prefix-action name {
    count;
    destination-prefix-length prefix-length;
    filter-specific;
    policer policer-name;
    source-prefix-length prefix-length;
    subnet-prefix-length prefix-length;
}
service-filter filter-name {
    term term-name {
        from {
            address {
                ip-prefix</prefix-length>;
            }
            (ah-spi [ values ] | ah-spi-except [ values ]);
            destination-address {
                ip-prefix</prefix-length>;
            }
            (destination-port [ port-names ] | destination-port-except [ port-names ]);
            destination-prefix-list {
                list-name;
            }
            (esp-spi [ values ] | esp-spi-except [ values ]);
            first-fragment;
            fragment-flags flag;
            (fragment-offset [ offsets ] | fragment-offset-except [ offsets ]);
            (interface-group [ group-names ] | interface-group-except [ group-names ]);
            (ip-options [ option-names ] | ip-options-except [ option-names ]);
            is-fragment;
            (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
            (port [ port-names ] | port-except [ port-names ]);
            prefix-list {
                list-name;
            }
            (protocol [ protocol-names ] | protocol-except [ protocol-names ]);

```

```

        source-address {
            ip-prefix </prefix-length>;
        }
        (source-port [ port-names ] | source-port-except [ port-names ]);
        source-prefix-list {
            list-name;
        }
        tcp-flags flag-name;
    }
    then {
        count counter-name;
        log;
        port-mirror;
        sample;
        (service | skip);
    }
}
}
simple-filter filter-name {
    term term-name {
        from {
            destination-address ip-prefix </prefix-length>;
            destination-port port-name;
            forwarding-class [ class-names ];
            protocol protocol-name;
            source-address ip-prefix </prefix-length>;
            source-port port-name;
        }
        then {
            forwarding-class class-name;
            loss-priority (high | low | medium-high | medium-low);
            policer policer-name;
        }
    }
}
}
}
}
firewall {
    family inet6 {
        dialer-filter filter-name {
            accounting-profile [ profile-names ];
            term term-name {
                from {
                    ... statements in Common IP Firewall Match Conditions on page 322 PLUS ...
                    (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
                    ... BUT NOT ...
                    (destination-class [ class-names ] |
                     destination-class-except [ class-names ]); # NOT valid at this level
                    (forwarding-class [ class-names ] |
                     forwarding-class-except [ class-names ]); # NOT valid at this level
                    interface interface-name; # NOT valid at this level
                    (interface-group [ group-names ] | interface-group-except [ group-names ]); #
                     NOT valid at this level
                    (loss-priority [ priorities ] | loss-priority-except [ priorities ]); # NOT valid at
                     this level
                }
            }
        }
    }
}

```

```

        service-filter-hit; # NOT valid at this level
        (source-class [ class-names ] | source-class-except [ class-names ]); # NOT
            valid at this level
        tcp-established; # NOT valid at this level
        tcp-flags flag; # NOT valid at this level
        tcp-initial; # NOT valid at this level
    }
    then {
        (ignore | note);
        log;
        sample;
        syslog;
    }
}
}
filter filter-name {
    accounting-profile [ profile-names ];
    interface-specific;
    term term-name {
        filter filter-name;
        from {
            ... statements in Common IP Firewall Match Conditions on page 322 PLUS ...
            (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
            (traffic-class [ code-point-values ] | traffic-class-except [ code-point-values ]);
        }
        then {
            ... statements in Common Firewall Actions on page 320 AND
            statements in Common IP Firewall Actions on page 321 PLUS ...
            (accept | discard | reject <address-unreachable | administratively-prohibited |
                beyond-scope | fragmentation-needed | no-route | port-unreachable |
                tcp-reset>);
        }
    }
}
}
service-filter filter-name {
    term term-name {
        from {
            address {
                ip-prefix</prefix-length>;
            }
            (ah-spi [ values ] | ah-spi-except [ values ]);
            destination-address {
                ip-prefix</prefix-length>;
            }
            (destination-port [ port-names ] | destination-port-except [ port-names ]);
            destination-prefix-list {
                list-name;
            }
            (esp-spi [ values ] | esp-spi-except [ values ]);
            (interface-group [ group-names ] | interface-group-except [ group-names ]);
            (next-header [ protocol-types ] | next-header-except [ protocol-types ]);
            (port [ port-names ] | port-except [ port-names ]);
            prefix-list {
                list-name;
            }
            source-address {

```

```

        ip-prefix </prefix-length>;
    }
    (source-port [ port-names ] | source-port-except [ port-names ]);
    source-prefix-list {
        list-name;
    }
    tcp-flags flag-name;
}
then {
    count counter-name;
    log;
    port-mirror;
    sample;
    (service | skip);
}
}
}
}
}

firewall {
    family mpls {
        filter filter-name {
            accounting-profile [ profile-names ];
            interface-specific;
            physical-interface-filter;
            term term-name {
                from {
                    (exp [ exp-bits ] | exp-except [ exp-bits ]);
                }
                then {
                    (ignore | note);
                    log;
                    sample;
                    syslog;
                }
            }
        }
    }
    filter filter-name {
        accounting-profile [ profile-names ];
        interface-specific;
        physical-interface-filter;
        term term-name {
            filter filter-name;
            from {
                (exp [ exp-bits ] | exp-except [ exp-bits ]);
                (forwarding-class [ class-names ] | forwarding-class-except [ class-names ]);
                interface interface-name;
                interface-set set-name;
                (loss-priority [ priorities ] | loss-priority-except [ priorities ]);
            }
            then {
                ... statements in Common Firewall Actions on page 320 PLUS ...
                (accept | discard);
                sample;
            }
        }
    }
}

```

```

    }
  }
}


firewall {
  family vpls {
    filter filter-name {
      accounting-profile [ profile-names ];
      interface-specific;
      term term-name {
        filter filter-name;
        from {
          ... statements in Common Layer 2 Firewall Match Conditions on page 323 ...
        }
        then {
          ... statements in Common Firewall Actions on page 320 PLUS ...
          (accept | discard);
          port-mirror;
          port-mirror-instance instance-name;
        }
      }
    }
  }
}

```


Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*

action

Syntax	<pre>action { loss-priority high then discard; }</pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall three-color-policer <i>name</i>], [edit firewall three-color-policer <i>name</i>], [edit logical-systems <i>logical-system-name</i> firewall three-color-policer <i>name</i>]
Release Information	<p>Statement introduced in Junos OS Release 8.2.</p> <p>Logical systems support introduced in Junos OS Release 9.3.</p> <p>Support at the [edit dynamic-profiles ... three-color-policer] hierarchy level introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Discard traffic on a logical interface using tricolor marking policing.
	<div>  <p>NOTE: This statement is supported only on IQ2 interfaces.</p> </div> <p>The remaining statement is explained separately.</p>
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Three-Color Policer Configuration Overview on page 4817 • Basic Single-Rate Three-Color Policers on page 4915 • Basic Two-Rate Three-Color Policers on page 4921 • Two-Color and Three-Color Logical Interface Policers on page 4927 • Two-Color and Three-Color Physical Interface Policers on page 4940 • Two-Color and Three-Color Policers at Layer 2 on page 4947 • loss-priority high then discard on page 4996

bandwidth-limit (Policer)

Syntax	<code>bandwidth-limit <i>bps</i>;</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall policer <i>policer-name</i> if-exceeding], [edit firewall policer <i>policer-name</i> if-exceeding], [edit logical-systems <i>logical-system-name</i> policer <i>policer-name</i> if-exceeding]
Release Information	Statement introduced before Junos OS Release 7.4. Support at the [edit dynamic-profiles ... if-exceeding] hierarchy level introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>For a single-rate two-color policer, configure the bandwidth limit as a number of bits per second. Single-rate two-color policing uses the single token bucket algorithm to measure traffic-flow conformance to a two-color policer rate limit.</p> <p>Traffic at the interface that conforms to the bandwidth limit is categorized green. Traffic that exceeds the specified rate is also categorized as green provided that sufficient tokens remain in the single token bucket. Packets in a green flow are implicitly marked with low packet loss priority (PLP) and then passed through the interface.</p> <p>Traffic that exceeds the specified rate when insufficient tokens remain in the single token bucket is categorized red. Depending on the configuration of the two-color policer, packets in a red traffic flow might be implicitly discarded; or the packets might be re-marked with a specified forwarding class, a specified PLP, or both, and then passed through the interface.</p> <div style="margin-top: 20px;">  <p>NOTE: This statement specifies the bandwidth limit as an absolute number of bits per second. Alternatively, for single-rate two-color policers only, you can use the bandwidth-percent <i>percentage</i> statement to specify the bandwidth limit as a percentage of either the physical interface port speed or the configured logical interface shaping rate.</p> </div> <p>Single-rate two-color policing allows bursts of traffic for short periods, whereas single-rate and two-rate three-color policing allows more sustained bursts of traffic.</p> <p>Hierarchical policing is a form of two-color policing that applies different policing actions based on whether the packets are classified for expedited forwarding (EF) or for a lower priority. You apply a hierarchical policer to ingress Layer 2 traffic to allow bursts of EF traffic for short period and bursts of non-EF traffic for short periods, with EF traffic always taking precedence over non-EF traffic.</p>
Options	<i>bps</i> —You can specify the number of bits per second either as a decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).

Range: (M Series, MX Series, T Series routers, and EX Series switches)
8000 through 100,000,000,000

Default: None.

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

- Related Documentation**
- [Two-Color Policer Configuration Overview on page 4813](#)
 - [Policer Bandwidth and Burst-Size Limits](#)
 - [Policer Color-Marking and Actions on page 4822](#)
 - [Single Token Bucket Algorithm on page 4824](#)
 - [Determining Proper Burst Size for Traffic Policers on page 4834](#)
 - [bandwidth-percent on page 4973](#)
 - [burst-size-limit \(Policer\) on page 4975](#)

bandwidth-percent

Syntax	<code>bandwidth-percent <i>percentage</i>;</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall policer <i>policer-name</i> if-exceeding], [edit firewall policer <i>policer-name</i> if-exceeding], [edit logical-systems <i>logical-system-name</i> policer <i>policer-name</i> if-exceeding]
Release Information	Statement introduced before Junos OS Release 7.4. Support at the [edit dynamic-profiles ... if-exceeding] hierarchy level introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	For a single-rate two-color policer, configure the bandwidth limit as a percentage value. Single-rate two-color policing uses the <i>single token bucket algorithm</i> to measure traffic-flow conformance to a two-color policer rate limit. Traffic at the interface that conforms to the bandwidth limit is categorized green. Traffic that exceeds the specified rate is also categorized as green provided that sufficient tokens remain in the single token bucket. Packets in a green flow are implicitly marked with low packet loss priority and then passed through the interface. Traffic that exceeds the specified rate when insufficient tokens remain in the single token bucket is categorized red. Depending on the configuration of the two-color policer, packets in a red traffic flow might be implicitly discarded; or the packets might be re-marked with a specified forwarding class, a specified PLP, or both, and then passed through the interface.



NOTE: This statement specifies the bandwidth limit as a percentage of either the physical interface port speed or the configured logical interface shaping rate. Alternatively, you can use the **bandwidth-limit *bps*** statement to specify the bandwidth limit as an absolute number of bits per second.

The function of the bandwidth limit is extended by the burst size (configured using the **burst-size-limit *bytes*** statement) to allow bursts of traffic up to a limit based on the overall traffic load:

- When a single-rate two-color policer is applied to the input or output traffic at an interface, the initial capacity for traffic bursting is equal to the number of bytes specified by this statement.
- During periods of relatively low traffic (traffic that arrives at or departs from the interface at overall rates below the token arrival rate), unused tokens accumulate in the bucket, but only up to the configured token bucket depth.

Single-rate two-color policing allows bursts of traffic for short periods, whereas single-rate and two-rate three-color policing allows more sustained bursts of traffic.

Hierarchical policing is a form of two-color policing that applies different policing actions based on whether the packets are classified for expedited forwarding (EF) or for a lower priority. You apply a hierarchical policer to ingress Layer 2 traffic to allow bursts of EF traffic for short periods and bursts of non-EF traffic for short periods, with EF traffic always taking precedence over non-EF traffic.

Options *percentage*—Traffic rate as a percentage of either the physical interface media rate or the logical interface configured shaping rate. You can configure a shaping rate on a logical interface by using class-of-service statement.



NOTE: You cannot rate-limit based on bandwidth percentage for aggregate, tunnel, and software interfaces. The bandwidth percentage policer cannot be used for forwarding table filters. Bandwidth percentage policers can only be used for interface-specific filters.

Range: 0 through 100

Default: None.

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

Related Documentation

- [Two-Color Policer Configuration Overview on page 4813](#)
- [Policer Bandwidth and Burst-Size Limits](#)
- [Policer Color-Marking and Actions on page 4822](#)
- [Single Token Bucket Algorithm on page 4824](#)
- [Determining Proper Burst Size for Traffic Policers on page 4834](#)
- [Bandwidth Policers on page 4854](#)
- [bandwidth-limit \(Policer\) on page 4971](#)
- [burst-size-limit \(Policer\) on page 4975](#)

burst-size-limit (Policer)

Syntax	<code>burst-size-limit bytes;</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall policer <i>policer-name</i> if-exceeding], [edit firewall policer <i>policer-name</i> if-exceeding], [edit logical-systems <i>logical-system-name</i> policer <i>policer-name</i> if-exceeding]
Release Information	Statement introduced before Junos OS Release 7.4. Support at the [edit dynamic-profiles ... if-exceeding] hierarchy level introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>For a single-rate two-color policer, configure the burst size as a number of bytes. The burst size allows for short periods of traffic bursting (back-to-back traffic at average rates that exceed the configured bandwidth limit). Single-rate two-color policing uses the <i>single token bucket algorithm</i> to measure traffic-flow conformance to a two-color policer rate limit.</p> <p>Traffic at the interface that conforms to the bandwidth limit is categorized green. Traffic that exceeds the specified rate is also categorized as green provided that sufficient tokens remain in the single token bucket. Packets in a green flow are implicitly marked with low packet loss priority and then passed through the interface.</p> <p>Traffic that exceeds the specified rate when insufficient tokens remain in the single token bucket is categorized red. Depending on the configuration of the two-color policer, packets in a red traffic flow might be implicitly discarded; or the packets might be re-marked with a specified forwarding class, a specified PLP, or both, and then passed through the interface.</p> <p>The burst size extends the function of the bandwidth limit (configured using either the bandwidth-limit <i>bps</i> statement or the bandwidth-percent <i>percentage</i> statement) to allow bursts of traffic up to a limit based on the overall traffic load:</p> <ul style="list-style-type: none"> • When a single-rate two-color policer is applied to the input or output traffic at an interface, the initial capacity for traffic bursting is equal to the number of bytes specified by this statement. • During periods of relatively low traffic (traffic that arrives at or departs from the interface at overall rates below the token arrival rate), unused tokens accumulate in the bucket, but only up to the configured token bucket depth. <p>Single-rate two-color policing allows bursts of traffic for short periods, whereas single-rate and two-rate three-color policing allows more sustained bursts of traffic.</p> <p>Hierarchical policing is a form of two-color policing that applies different policing actions based on whether the packets are classified for expedited forwarding (EF) or for a lower priority. You apply a hierarchical policer to ingress Layer 2 traffic to allow bursts of EF traffic for short period and bursts of non-EF traffic for short periods, with EF traffic always taking precedence over non-EF traffic.</p>

Table 370 on page 4976 summarizes the relationship between the **bandwidth-limit** and the token arrival rate. This information is useful in calculating the minimum **burst-size-limit**.

Table 370: Bandwidth Limits and Token Rates

Bandwidth Limit	Token Rate
0-333 Mbps	low (262 μ s)
334-666 Mbps	high (8.2 μ s)
667-1333 Mbps	low
1334 Mbps and above	high

The burst-size limit enforced is based on the burst-size limit you configure. For a rate-limited logical interface, the Packet Forwarding Engine calculates the optimum burst-size-limit values and then applies the value closest to the burst-size-limit value specified in the policer configuration.

On MX Series routers and EX Series switches, the burst-size limit is not as freely configurable as it is on other platforms. Junos OS does not support an unlimited combination of policer bandwidth and burst-size limits on MX Series routers and EX Series switches. For a single-rate two-color policer on an MX Series router and on an EX Series switch, the minimum supported burst-size limit is equivalent to the amount of traffic allowed by the policer bandwidth limit in a time span of 1 millisecond. For example, for a policer configured with a **bandwidth-limit** value of 1 Gbps, the minimum supported value for **burst-size-limit** on an MX Series router and on an EX Series switch is 125 KB. If you configure a value that is smaller than the minimum, Junos OS overrides the configuration and applies the actual minimum.

Options **bytes**—Burst-size limit in bytes. The minimum recommended value is the maximum transmission unit (MTU) of the IP packets being policed. You can specify the value either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 1500 through 100,000,000,000


Default: None

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


**Related
Documentation**

- [Two-Color Policer Configuration Overview on page 4813](#)
- [*Policer Bandwidth and Burst-Size Limits*](#)
- [Policer Color-Marking and Actions on page 4822](#)
- [Single Token Bucket Algorithm on page 4824](#)
- [Determining Proper Burst Size for Traffic Policers on page 4834](#)
- [bandwidth-limit \(Policer\) on page 4971](#)
- [bandwidth-percent on page 4973](#)

color-aware

Syntax	color-aware;
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall three-color-policer <i>name</i> single-rate], [edit dynamic-profiles <i>profile-name</i> firewall three-color-policer <i>name</i> two-rate], [edit firewall three-color-policer <i>policer-name</i> single-rate], [edit firewall three-color-policer <i>policer-name</i> two-rate]
Release Information	Statement introduced in Junos OS Release 7.4. Support at the [edit dynamic-profiles ... single-rate] and [edit dynamic-profiles ... two-rate] hierarchy levels introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>For a three-color policer, configure the way preclassified packets are metered. In color-aware mode, the local router can assign a higher packet loss priority, but cannot assign a lower packet loss priority.</p> <p>For example, suppose an upstream router assigned medium-high packet loss priority to a packet because the packet exceeded the committed information rate on the upstream router interface.</p> <ul style="list-style-type: none"> • If the local router applies color-aware policing to the packet, the router <i>cannot</i> change the packet loss priority to low, even if the packet conforms to the configured committed information route on the local router interface. • If the local router applies color-blind policing to the packet, the router <i>can</i> change the packet loss priority to low if the packet conforms to the configured committed information route on the local router interface.
<div>  <p>NOTE: A color-aware policer cannot be applied to Layer 2 traffic.</p> </div>	
Default	If you omit the color-aware statement, the default behavior is color-aware mode.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Three-Color Policer Configuration Overview on page 4817 • Color Modes for Three-Color Policers on page 4913 • color-blind on page 4979

color-blind

Syntax	color-blind;
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall three-color-policer <i>name</i> single-rate], [edit dynamic-profiles <i>profile-name</i> firewall three-color-policer <i>name</i> two-rate], [edit firewall three-color-policer <i>policer-name</i> single-rate], [edit firewall three-color-policer <i>policer-name</i> two-rate]
Release Information	Statement introduced in Junos OS Release 7.4. Support at the [edit dynamic-profiles ... single-rate] and [edit dynamic-profiles ... two-rate] hierarchy levels introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>For a three-color policer, configure the way preclassified packets are metered. In color-blind mode, the local router ignores the preclassification of packets and can assign a higher or lower packet loss priority.</p> <p>For example, suppose an upstream router assigned medium-high packet loss priority to a packet because the packet exceeded the committed information rate on the upstream router interface.</p> <ul style="list-style-type: none"> • If the local router applies color-aware policing to the packet, the router <i>cannot</i> change the packet loss priority to low, even if the packet conforms to the configured committed information route on the local router interface. <div style="margin-top: 10px;">  <p>NOTE: A color-aware policer cannot be applied to Layer 2 traffic.</p> </div> <ul style="list-style-type: none"> • If the local router applies color-blind policing to the packet, the router <i>can</i> change the packet loss priority to low if the packet conforms to the configured committed information route on the local router interface.
Default	If you omit the color-blind statement, the default behavior is color-aware mode.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Three-Color Policer Configuration Overview on page 4817 • Color Modes for Three-Color Policers on page 4913 • color-aware on page 4978

committed-burst-size

Syntax	<code>committed-burst-size bytes;</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall three-color-policer <i>name</i> single-rate], [edit dynamic-profiles <i>profile-name</i> firewall three-color-policer <i>name</i> two-rate], [edit firewall three-color-policer <i>policer-name</i> single-rate], [edit firewall three-color-policer <i>policer-name</i> two-rate]
Release Information	Statement introduced in Junos OS Release 7.4. Support at the [edit dynamic-profiles ... single-rate] and [edit dynamic-profiles ... two-rate] hierarchy levels introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	For a three-color policer, configure the committed burst size (CBS) as a number of bytes.



NOTE: When you include the **committed-burst-size** statement in the configuration, you must also include the **committed-information-rate** statement at the same hierarchy level.

In three-color policing, a committed information rate (CIR) defines the guaranteed bandwidth for traffic arriving at or departing from the interface under normal line conditions. A flow of traffic at an average rate that conforms to the CIR is categorized green.

During periods of average traffic rates below the CIR, any unused bandwidth capacity accumulates up to a maximum amount defined by the CBS. Short periods of bursting traffic (back-to-back traffic at averages rates that exceed the CIR) are also categorized as green provided that unused bandwidth capacity is available.

Traffic that exceeds both the CIR and the CBS is considered nonconforming.

Single-rate three-color policers use a *dual token bucket algorithm* to measure traffic against a single rate limit. Nonconforming traffic is categorized as yellow or red, based on the **excess-burst-size** statement included in the policer configuration.

Two-rate three-color policers use a *dual-rate dual token bucket algorithm* to measure traffic against two rate limits. Nonconforming traffic is categorized as yellow or red based on the **peak-information-rate** and **peak-burst-rate** statements included in the policer configuration.

Options **bytes**—Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 1500 through 100,000,000,000 bytes

Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Three-Color Policer Configuration Overview on page 4817• Policer Bandwidth and Burst-Size Limits• Policer Color-Marking and Actions on page 4822• Dual Token Bucket Algorithms on page 4826• Determining Proper Burst Size for Traffic Policers on page 4834• committed-information-rate on page 4982• excess-burst-size on page 4984• peak-burst-size on page 4999• peak-information-rate on page 5001

committed-information-rate

Syntax	<code>committed-information-rate <i>bps</i>;</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall three-color-policer <i>name</i> single-rate], [edit dynamic-profiles <i>profile-name</i> firewall three-color-policer <i>name</i> two-rate], [edit firewall three-color-policer <i>policer-name</i> single-rate], [edit firewall three-color-policer <i>policer-name</i> two-rate]
Release Information	Statement introduced in Junos OS Release 7.4. Support at the [edit dynamic-profiles ... single-rate] and [edit dynamic-profiles ... two-rate] hierarchy levels introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	For a three-color policer, configure the committed information rate as a number of bits per second. The committed information rate (CIR) is the guaranteed bandwidth for traffic arriving at or departing from the interface under normal line conditions.



NOTE: When you include the **committed-information-rate** statement in the configuration, you must also include the **committed-burst-size** statement at the same hierarchy level.

In three-color policing, a CIR defines the guaranteed bandwidth for traffic arriving at or departing from the interface under normal line conditions. A flow of traffic at an average rate that conforms to the CIR is categorized green.

During periods of average traffic rates below the CIR, any unused bandwidth capacity accumulates up to a maximum amount defined by the committed burst size (CBS). Short periods of bursting traffic (back-to-back traffic at averages rates that exceed the CIR) are also categorized as green provided that unused bandwidth capacity is available.

Traffic that exceeds both the CIR and the CBS is considered nonconforming.

Single-rate three-color policers use a *dual token bucket algorithm* to measure traffic against a single rate limit. Nonconforming traffic is categorized as yellow or red, based on the **excess-burst-size** statement included in the policer configuration.

Two-rate three-color policers use a *dual-rate dual token bucket algorithm* to measure traffic against two rate limits. Nonconforming traffic is categorized as yellow or red based on the **peak-information-rate** and **peak-burst-rate** statements included in the policer configuration.

Options ***bps***—Number of bits per second. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Range: 1500 through 100,000,000,000 bps

Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Three-Color Policer Configuration Overview on page 4817• Policer Bandwidth and Burst-Size Limits• Policer Color-Marking and Actions on page 4822• Dual Token Bucket Algorithms on page 4826• Determining Proper Burst Size for Traffic Policers on page 4834• committed-burst-size on page 4980• excess-burst-size on page 4984• peak-burst-size on page 4999• peak-information-rate on page 5001

excess-burst-size

Syntax	<code>excess-burst-size bytes;</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall three-color-policer name single-rate], [edit firewall three-color-policer policer-name single-rate]
Release Information	Statement introduced in Junos OS Release 7.4. Support at the [edit dynamic-profiles ... single-rate] hierarchy level introduced in Junos Release OS 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	For a single-rate three-color policer, configure the excess burst size (EBS) as a number of bytes. The EBS allows for moderate periods of bursting traffic that exceeds both the committed information rate (CIR) and the committed burst size (CBS).



NOTE: When you include the **excess-burst-size** statement in the configuration, you must also include the **committed-burst-size** and **committed-information-rate** statements at the same hierarchy level.

Traffic that exceeds both the CIR and the CBS is considered nonconforming.

Single-rate three-color policing uses a *dual token bucket algorithm* to measure traffic against a single rate limit. Nonconforming traffic is categorized as yellow or red based on the **excess-burst-size** statement included in the policer configuration.

During periods of traffic that conforms to the CIR, any unused portion of the guaranteed bandwidth capacity accumulates in the first token bucket, up to the maximum number of bytes defined by the CBS. If any accumulated bandwidth capacity overflows the first bucket, the excess accumulates in a second token bucket, up to the maximum number of bytes defined by the EBS.

A nonconforming traffic flow is categorized yellow if its size conforms to bandwidth capacity accumulated in the first token bucket. Packets in a yellow flow are marked with **medium-high** packet loss priority (PLP) and then passed through the interface.

A nonconforming traffic flow is categorized red if its size exceeds the bandwidth capacity accumulated in the second token bucket. Packets in a red traffic flow are marked with **high** PLP and then either passed through the interface or optionally discarded.

Options	bytes —Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). Range: 1500 through 100,000,000,000 bytes
Required Privilege Level	firewall —To view this statement in the configuration. firewall-control —To add this statement to the configuration.

- Related Documentation**
- [Three-Color Policer Configuration Overview on page 4817](#)
 - [Policer Bandwidth and Burst-Size Limits](#)
 - [Policer Color-Marking and Actions on page 4822](#)
 - [Dual Token Bucket Algorithms on page 4826](#)
 - [Determining Proper Burst Size for Traffic Policers on page 4834](#)
 - [committed-burst-size on page 4980](#)
 - [committed-information-rate on page 4982](#)

filter-specific

Syntax	filter-specific;
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall policer <i>policer-name</i>], [edit firewall family inet prefix-action <i>name</i>], [edit firewall policer <i>policer-name</i>], [edit logical-systems <i>logical-system-name</i> firewall policer <i>policer-name</i>], [edit logical-systems <i>logical-system-name</i> firewall family inet prefix-action <i>name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Logical systems support introduced in Junos OS Release 9.3. Support at the [edit dynamic-profiles ... policer <i>policer-name</i>] hierarchy level introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Set the prefix-specific action or policer to operate in <i>filter-specific</i> mode, meaning that a single policer and counter are shared by all filter terms that reference the prefix-specific action or policer. By default, the prefix-specific action or policer operates in <i>term-specific</i> mode, meaning that a separate policer and counter are used for each filter term that references the prefix-specific action or policer.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Filter-Specific Policer Overview on page 4863 • Prefix-Specific Counting and Policing Overview on page 4874 • Filter-Specific Counter and Policer Set Overview on page 4876

forwarding-class (Firewall Filter Action)

Syntax	<code>forwarding-class class-name;</code>
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> then], [edit logical-systems <i>logical-system-name</i> firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> then]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Set the forwarding class of incoming packets.
Options	<i>class-name</i> —Name of the forwarding class.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Standard Firewall Filter Nonterminating Actions on page 4744• Policer Color-Marking and Actions on page 4822• Multifield Classification Overview on page 4889

hierarchical-policer

Syntax	<pre> hierarchical-policer <i>policer-name</i> { aggregate { if-exceeding { bandwidth-limit <i>bps</i>; burst-size-limit <i>bytes</i>; } then { discard; } } premium { if-exceeding { bandwidth-limit <i>bps</i>; burst-size-limit <i>bytes</i>; } then { discard; } } } </pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall], [edit firewall]
Release Information	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Support at the [edit dynamic-profiles ... firewall] hierarchy level introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	On M40e, M120, and M320 edge routers with Flexible PIC Concentrator (FPC) input as FFPC and FPC output as SFPC, and on MX Series, T320, T640, and T1600 edge routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, specify a hierarchical policer.
Options	<p><i>policer-name</i>—Name that identifies the policer. The name can contain letters, numbers, and hyphens (-), and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Hierarchical Policer Configuration Overview</i> • <i>Hierarchical Policers</i> • <i>aggregate (Hierarchical Policer)</i> • <i>bandwidth-limit (Hierarchical Policer)</i>

- *burst-size-limit* (*Hierarchical Policer*)
- *if-exceeding* (*Hierarchical Policer*)
- *premium* (*Hierarchical Policer*)

if-exceeding (Policer)

Syntax	<pre>if-exceeding { (bandwidth-limit <i>bps</i> bandwidth-percent <i>number</i>); burst-size-limit <i>bytes</i>; }</pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall policer <i>policer-name</i>], [edit firewall policer <i>policer-name</i>], [edit logical-systems <i>logical-system-name</i> firewall policer <i>policer-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Logical systems support introduced in Junos OS Release 9.3. Support at the [edit dynamic-profiles ... policer <i>policer-name</i>] hierarchy level introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure rate limits for a single-rate two-color policer. The remaining statements are explained separately.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Two-Color Policer Configuration Overview on page 4813• Hierarchical Policer Configuration Overview• Basic Single-Rate Two-Color Policers on page 4838• Bandwidth Policers on page 4854• Filter-Specific Counters and Policers on page 4863• Prefix-Specific Counting and Policing Actions on page 4874• Multifield Classification on page 4889• Policer Overhead to Account for Rate Shaping in the Traffic Manager on page 4904• Hierarchical Policers

input-hierarchical-policer

Syntax	input-hierarchical-policer <i>policer-name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> layer2-policer], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer],
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Apply a hierarchical policer to the Layer 2 input traffic for all protocol families at the physical or logical interface.
Options	<i>policer-name</i> —Name of the hierarchical policer.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Hierarchical Policers</i>• <i>layer2-policer (Hierarchical Policer)</i>

input-policer

Syntax	<code>input-policer <i>policer-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer] [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Apply a single-rate two-color policer to the Layer 2 input traffic at the logical interface. The input-policer and input-three-color statements are mutually exclusive.
Options	<i>policer-name</i> —Name of the single-rate two-color policer that you define at the [edit firewall] hierarchy level.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Two-Color and Three-Color Policers at Layer 2 on page 4947• Applying Layer 2 Policers to Gigabit Ethernet Interfaces on page 1459• Configuring a Gigabit Ethernet Policer• input-three-color on page 1522• layer2-policer on page 1523• logical-interface-policer on page 1500• output-policer on page 1524• output-three-color on page 1525

input-three-color

Syntax	<code>input-three-color <i>policer-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer] [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Apply a single-rate or two-rate three-color policer to the Layer 2 input traffic at the logical interface. The input-three-color and input-policer statements are mutually exclusive.
Options	<i>policer-name</i> —Name of the single-rate or two-rate three-color policer.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Two-Color and Three-Color Policers at Layer 2 on page 4947 • Applying Layer 2 Policers to Gigabit Ethernet Interfaces on page 1459 • Configuring a Gigabit Ethernet Policer • input-policer on page 1521 • layer2-policer on page 1523 • logical-interface-policer on page 1500 • output-policer on page 1524 • output-three-color on page 1525

layer2-policer

Syntax	<pre>layer2-policer { input-policer policer-name; input-three-color policer-name; output-policer policer-name; output-three-color policer-name; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>For 1-Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series, MX Series, and T Series routers, and for aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces on EX Series switches, apply Layer 2 logical interface policers. The following policers are supported:</p> <ul style="list-style-type: none">• Two-color• Single-rate tricolor marking (srTCM)• Two-rate tricolor marking (trTCM) <p>Two-color and tricolor policers are configured at the [edit firewall] hierarchy level.</p>
Options	<p>input-policer <i>policer-name</i>—Two-color input policer to associate with the interface. This statement is mutually exclusive with the input-three-color statement.</p> <p>input-three-color <i>policer-name</i>—Tricolor input policer to associate with the interface. This statement is mutually exclusive with the input-policer statement.</p> <p>output-policer <i>policer-name</i>—Two-color output policer to associate with the interface. This statement is mutually exclusive with the output-three-color statement.</p> <p>output-three-color <i>policer-name</i>—Tricolor output policer to associate with the interface. This statement is mutually exclusive with the output-policer statement.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Applying Layer 2 Policers to Gigabit Ethernet Interfaces on page 1459• <i>Configuring Gigabit Ethernet Two-Color and Tricolor Policers</i>


load-balance-group

Syntax	<code>load-balance-group <i>group-name</i> { next-hop-group [<i>group-names</i>]; }</code>
Hierarchy Level	[edit firewall]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure a load-balance group.
Options	<p><i>group-name</i>—Name of load-balance group.</p> <p><i>group-names</i>—Name of next-hop groups to include in the load-balance group set.</p>
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Load-Balance Groups</i> in the <i>Routing Policy Configuration Guide</i>

logical-bandwidth-policer

Syntax	<code>logical-bandwidth-policer;</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall policer <i>policer-name</i>], [edit firewall policer <i>policer-name</i>], [edit logical-systems <i>logical-system-name</i> firewall policer <i>policer-name</i>]
Release Information	<p>Statement introduced in Junos OS Release 8.2.</p> <p>Logical systems support introduced in Junos OS Release 9.3.</p> <p>Support at the [edit dynamic-profiles ... policer <i>policer-name</i>] hierarchy level introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	For a policer with a bandwidth limit configured as a percentage (using the bandwidth-percent statement), specify that the percentage be based on the shaping rate defined on the logical interface, rather than on the media rate of the physical interface.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Bandwidth Policers on page 4854 • Configuring Logical Bandwidth Policers on page 1391 • bandwidth-percent on page 4973 statement • interface-specific on page 4793 statement

logical-interface-policer

Syntax	logical-interface-policer;
Hierarchy Level	<p>[edit dynamic-profiles <i>profile-name</i> firewall policer <i>policer-name</i>], [edit dynamic-profiles <i>profile-name</i> firewall three-color-policer <i>name</i>], [edit firewall atm-policer <i>atm-policer-name</i>] [edit firewall policer <i>policer-name</i>], [edit firewall policer <i>policer-template-name</i>], [edit firewall three-color-policer <i>policer-name</i>], [edit logical-systems <i>logical-system-name</i> firewall policer <i>policer-name</i>], [edit logical-systems <i>logical-system-name</i> firewall three-color-policer <i>name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support at the [edit firewall three-color-policer <i>policer-name</i>] hierarchy level introduced in Junos OS Release 8.2.</p> <p>Logical systems support introduced in Junos OS Release 9.3.</p> <p>Support at the [edit dynamic-profiles ... policer <i>policer-name</i>] and [edit dynamic-profiles ... three-color-policer <i>name</i>] hierarchy levels introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Configure a logical interface policer.
	<div><p>NOTE: Starting in Junos OS Release 12.2R2, on T Series Core Routers only, you can configure an MPLS LSP policer for a specific LSP to be shared across different protocol family types. You must include the logical-interface-policer statement to do so.</p></div>
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Two-Color and Three-Color Logical Interface Policers on page 4927• Traffic Policer Types• Configuring Tricolor Marking Policers on page 1455• action on page 1497• Configuring Gigabit Ethernet Two-Color and Tricolor Policers• action

loss-priority (Firewall Filter Action)

Syntax	loss-priority (high low);
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> then], [edit logical-systems <i>logical-system-name</i> firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> then]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Set the loss priority of incoming packets.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Standard Firewall Filter Nonterminating Actions on page 4744• Policer Color-Marking and Actions on page 4822• Multifield Classification Overview on page 4889

loss-priority high then discard (Three-Color Policer)

Syntax	loss-priority high then discard;
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall three-color-policer <i>name</i> action], [edit firewall three-color-policer <i>policer-name</i> action], [edit logical-systems <i>logical-system-name</i> firewall three-color-policer <i>policer-name</i> action]
Release Information	Statement introduced before Junos OS Release 8.2. Logical systems support introduced in Junos OS Release 9.3. Support at the [edit dynamic-profiles ... action] hierarchy level introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>For packets with high loss priority, discard the packets. The loss priority setting is implicit and is not configurable. Include this statement if you do not want the local router to forward packets that have high packet loss priority.</p> <p>For single-rate three-color policers, the Junos OS assigns high loss priority to packets that exceed the committed information rate and the excess burst size.</p> <p>For two-rate three-color policers, the Junos OS assigns high loss priority to packets that exceed the peak information rate and the peak burst size.</p>
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Three-Color Policer Configuration Overview on page 4817• Basic Single-Rate Three-Color Policers on page 4915• Basic Two-Rate Three-Color Policers on page 4921• Two-Color and Three-Color Logical Interface Policers on page 4927• Two-Color and Three-Color Physical Interface Policers on page 4940• Two-Color and Three-Color Policers at Layer 2 on page 4947• action on page 1497


output-policer

Syntax	<code>output-policer <i>policer-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Apply a single-rate two-color policer to the Layer 2 output traffic at the logical interface. The output-policer and output-three-color statements are mutually exclusive.
Options	<i>policer-name</i> —Name of the single-rate two-color policer that you define at the [edit firewall] hierarchy level.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Two-Color and Three-Color Policers at Layer 2 on page 4947 • Applying Layer 2 Policers to Gigabit Ethernet Interfaces on page 1459 • Configuring a Gigabit Ethernet Policer • input-policer on page 1521 • input-three-color on page 1522 • layer2-policer on page 1523 • logical-interface-policer on page 1500 • output-three-color on page 1525

output-three-color


Syntax	<code>output-three-color <i>policer-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer] [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Apply a single-rate or two-rate three-color policer to the Layer 2 output traffic at the logical interface. The output-three-color and output-policer statements are mutually exclusive.
Options	<i>policer-name</i> —Name of the single-rate or two-rate three-color policer.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Two-Color and Three-Color Policers at Layer 2 on page 4947• Applying Layer 2 Policers to Gigabit Ethernet Interfaces on page 1459• Configuring a Gigabit Ethernet Policer• input-three-color on page 1522• input-policer on page 1521• layer2-policer on page 1523• logical-interface-policer on page 1500• output-policer on page 1524

peak-burst-size

Syntax	<code>peak-burst-size bytes;</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall three-color-policer name two-rate], [edit firewall three-color-policer <i>policer-name</i> two-rate]
Release Information	Statement introduced in Junos OS Release 7.4. Support at the [edit dynamic-profiles ... two-rate] hierarchy level introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	For a two-rate three-color policer, configure the peak burst size (PBS) as a number of bytes. The PBS defines the maximum number of bytes of unused peak bandwidth capacity that can be accumulated. The accumulated bandwidth allows for moderate periods of bursting traffic that exceeds the peak information rate (PIR) and the committed burst size (CBS).
<div>  <p>NOTE: When you include the peak-burst-size statement in the configuration, you must also include the committed-burst-size and peak-information-rate statements at the same hierarchy level.</p> </div>	
<p>Two-rate three-color policers use a <i>dual-rate dual token bucket algorithm</i> to measure traffic against two rate limits.</p> <ul style="list-style-type: none"> • A traffic flow is categorized green if it conforms to both the committed information rate (CIR) and the CBS-bounded accumulation of available committed bandwidth capacity. • A traffic flow is categorized yellow if exceeds the CIR and CBS but conforms to the PIR. Packets in a yellow flow are marked with medium-high packet loss priority (PLP) and then passed through the interface. • A traffic flow is categorized red if exceeds the PIR and the PBS-bounded accumulation of available peak bandwidth capacity. Packets in a red traffic flow are marked with high PLP and then either passed through the interface or optionally discarded. 	
Options	bytes —Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). Range: 1500 through 100,000,000,000 bytes
Required Privilege Level	firewall —To view this statement in the configuration. firewall-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Three-Color Policer Configuration Overview on page 4817 • Policer Bandwidth and Burst-Size Limits

- [Policer Color-Marking and Actions on page 4822](#)
- [Dual Token Bucket Algorithms on page 4826](#)
- [Determining Proper Burst Size for Traffic Policers on page 4834](#)
- [committed-burst-size on page 4980](#)
- [committed-information-rate on page 4982](#)
- [excess-burst-size on page 4984](#)
- [peak-information-rate on page 5001](#)

peak-information-rate

Syntax	<code>peak-information-rate bps;</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall three-color-policer name two-rate], [edit firewall three-color-policer <i>policer-name</i> two-rate]
Release Information	Statement introduced in Junos OS Release 7.4. Support at the [edit dynamic-profiles ... two-rate] hierarchy level introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	For a two-rate three-color policer, configure the peak information rate (PIR) as a number of bits per second. The PIR is the maximum rate for traffic arriving at or departing from the interface under peak line conditions. Traffic that exceeds the committed information rate (CIR) and the committed burst size (CBS) is metered to the PIR.
<div>  <p>NOTE: When you include the peak-information-rate statement in the configuration, you must also include the committed-information-rate and peak-burst-size statements at the same hierarchy level.</p> </div>	
<p>Two-rate three-color policers use a <i>dual-rate dual token bucket algorithm</i> to measure traffic against two rate limits.</p> <ul style="list-style-type: none"> • A traffic flow is categorized green if it conforms to both the CIR and the CBS-bounded accumulation of available committed bandwidth capacity. • A traffic flow is categorized yellow if exceeds the CIR and CBS but conforms to the PIR. Packets in a yellow flow are marked with medium-high packet loss priority (PLP) and then passed through the interface. • A traffic flow is categorized red if exceeds the PIR and the PBS-bounded accumulation of available peak bandwidth capacity. Packets in a red traffic flow are marked with high PLP and then either passed through the interface or optionally discarded. 	
Options	bps —Number of bits per second. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). Range: 1500 through 100,000,000,000 bps
Required Privilege Level	firewall —To view this statement in the configuration. firewall-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Three-Color Policer Configuration Overview on page 4817 • Policer Bandwidth and Burst-Size Limits • Policer Color-Marking and Actions on page 4822

- [Dual Token Bucket Algorithms on page 4826](#)
- [Determining Proper Burst Size for Traffic Policers on page 4834](#)
- [committed-burst-size on page 4980](#)
- [committed-information-rate on page 4982](#)
- [excess-burst-size on page 4984](#)
- [peak-burst-size on page 4999](#)


physical-interface-filter

Syntax	physical-interface-filter;
Hierarchy Level	[edit firewall family family-name filter filter-name], [edit logical-systems <i>logical-system-name</i> firewall family family-name filter filter-name], [edit routing-instances <i>routing-instance-name</i> firewall family family-name filter filter-name], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> firewall family family-name filter filter-name]
Release Information	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Configure a physical interface filter. Use this statement to reference a physical interface policer for the specified protocol family.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Two-Color and Three-Color Physical Interface Policers on page 4940• physical-interface-policer on page 5003• policer (Configuring) on page 1502

physical-interface-policer

Syntax	physical-interface-policer;
Hierarchy Level	<p>[edit dynamic-profiles <i>profile-name</i> firewall policer <i>policer-name</i>], [edit firewall policer <i>policer-name</i>], [edit firewall three-color-policer <i>policer-name</i>], [edit logical-system <i>logical-system-name</i> firewall policer <i>policer-name</i>], [edit logical-system <i>logical-system-name</i> three-color-policer <i>policer-name</i>], [edit routing-instances <i>routing-instance-name</i> firewall policer <i>policer-name</i>], [edit routing-instances <i>routing-instance-name</i> firewall three-color-policer <i>policer-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> firewall policer <i>policer-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> firewall three-color-policer <i>policer-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.6.</p> <p>Support at the [edit dynamic-profiles ... policer <i>policer-name</i>] hierarchy level introduced in Junos Release OS 11.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure an aggregate policer for a physical interface.</p> <p>A physical interface policer can be a two-color or three-color policer. When you apply physical interface policer, to different protocol families on the same logical interface, the protocol families share the same policer instance. This means that rate limiting is performed aggregately for the protocol families for which the policer is applied. This feature enables you to use a single policer instance to perform aggregate policing for different protocol families on the same physical interface. If you want a policer instance to be associated with a protocol family, the corresponding physical interface filter needs to be applied to that protocol family. The policer is not automatically applied to all protocol families configured on the physical interface.</p> <p>In contrast, with logical interface policers there are multiple separate policer instances.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Two-Color and Three-Color Physical Interface Policers on page 4940 • physical-interface-filter on page 5002

policer (Applying to a Logical Interface)

Syntax	<pre>policer { input <i>policer-name</i>; output <i>policer-name</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>unit-number</i>], [edit interfaces <i>interface-name</i> unit <i>unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>unit-number</i> family <i>family</i>]
Release Information	Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>Apply a single-rate two-color policer—except for a physical interface policer—to Layer 3 input or output traffic at a logical interface.</p> <ul style="list-style-type: none">• To rate-limit all traffic types, regardless of the protocol family, you can apply a logical interface policer at the logical unit level of a supported interface.• To rate-limit traffic of a specific protocol family, you can apply a basic two-color policer, a bandwidth policer, or a logical interface policer at the protocol family level of a supported interface.
	<div> NOTE: You cannot apply a physical interface policer as part of the interface configuration. You can apply a physical interface policer by referencing the policer from a physical interface filter term.</div>
Options	<p>input <i>policer-name</i>—Name of one policer to evaluate packets received on the interface.</p> <p>output <i>policer-name</i>—Name of one policer to evaluate packets transmitted on the interface.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Single-Rate Two-Color Policer Overview on page 4838• Bandwidth Policer Overview on page 4854• Logical Interface (Aggregate) Policer Overview on page 4928

policer (Configuring)

Syntax	<pre> policer <i>policer-name</i> { filter-specific; if-exceeding { bandwidth-limit <i>bps</i>; bandwidth-percent <i>number</i>; burst-size-limit <i>bytes</i>; } logical-bandwidth-policer; logical-interface-policer; physical-interface-policer; shared-bandwidth-policer; then { <i>policer-action</i>; } } </pre>
Hierarchy Level	<p>[edit dynamic-profiles <i>profile-name</i> firewall], [edit firewall], [edit logical-systems <i>logical-system-name</i> firewall]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>The out-of-profile policer action added in Junos OS Release 8.1.</p> <p>The logical-bandwidth-policer statement added in Junos OS Release 8.2.</p> <p>Logical systems support introduced in Junos OS Release 9.3.</p> <p>The physical-interface-policer statement introduced in Junos OS Release 9.6.</p> <p>The shared-bandwidth-policer statement added in Junos OS Release 11.2.</p> <p>Support at the [edit dynamic-profiles ... firewall] hierarchy level introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	<p>Configure policer rate limits and actions. When included at the [edit firewall] hierarchy level, the policer statement creates a template, and you do not have to configure a policer individually for every firewall filter or interface. To activate a policer, you must include the policer-action modifier in the then statement in a firewall filter term or on an interface.</p>
Options	<p><i>policer-action</i>—One or more actions to take:</p> <ul style="list-style-type: none"> • discard—Discard traffic that exceeds the rate limits. • forwarding-class <i>class-name</i>—Specify the particular forwarding class. • loss-priority—Set the packet loss priority (PLP) to low, medium-low, medium-high, or high. • out-of-profile—On J Series routers with strict priority queuing, prevent starvation of other queues by rate limiting the data stream entering the strict priority queue, marking the packets that exceed the rate limit as out-of-profile, and dropping the out-of-profile packets if the physical interface is congested.

policer-name—Name that identifies the policer. The name can contain letters, numbers, and hyphens (-), and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" "). Policer names cannot begin with an underscore in the form `_.*`.

then—Actions to take on matching packets.

The remaining statements are explained separately.

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

Related Documentation

- [Bandwidth Policer Overview on page 4854](#)
- [Configuring Multifield Classifiers on page 1390](#)
- [Logical Interface \(Aggregate\) Policer Overview on page 4928](#)
- [Physical Interface Policer Overview on page 4940](#)
- [Statement Hierarchy for Configuring Policers on page 4811](#)
- [Single-Rate Two-Color Policer Overview on page 4838](#)
- [Using Multifield Classifiers to Set PLP on page 1461](#)
- [filter \(Configuring\) on page 1422](#)
- [priority \(Schedulers\) on page 1641](#)

policer (Firewall Filter Action)

Syntax	<code>policer <i>policer-name</i>;</code>
Hierarchy Level	[edit <code>firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i></code> then], [edit <code>logical-systems <i>logical-system-name</i> firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i></code> then]
Release Information	Statement introduced before Junos OS Release 7.4. Logical systems support introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	For T Series routers and M320 routers with Enhanced II Flexible PIC Concentrators (FPCs) and the T640 Core Router with Enhanced Scaling FPC4, apply a tricolor marking policer.
Options	<i>policer-name</i> —Name of a single-rate two-color policer to use to rate-limit traffic.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Standard Firewall Filter Nonterminating Actions on page 4744• Two-Color Policer Configuration Overview on page 4813

prefix-action (Configuring)

Syntax	<pre>prefix-action <i>prefix-action-name</i> { count; destination-prefix-length <i>prefix-length</i>; filter-specific; policer <i>policer-name</i>; source-prefix-length <i>prefix-length</i>; subnet-prefix-length <i>prefix-length</i>; }</pre>
Hierarchy Level	<pre>[edit firewall family inet], [edit logical-systems <i>logical-system-name</i> firewall family inet]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Logical systems support introduced in Junos OS Release 9.3.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Configure a prefix-specific action.
Options	<p>count—Enable counter.</p> <p>destination-prefix-length <i>prefix-length</i>—Destination prefix length. Range: 0 through 32</p> <p>filter-specific—Create the prefix-specific set of policers and counters as a filter-specific set. If this option is not specified, the prefix-specific set of policers and counters are created as term-specific.</p> <p>policer <i>policer-name</i>—Policer name.</p> <p>source-prefix-length <i>prefix-length</i>—Source prefix length. Range: 0 through 32</p> <p>subnet-prefix-length <i>prefix-length</i>—Subnet prefix length. Range: 0 through 32</p>
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Prefix-Specific Counting and Policing Actions on page 4874

prefix-action (Firewall Filter Action)

Syntax	<code>prefix-action <i>prefix-action-name</i>;</code>
Hierarchy Level	[edit firewall family inet filter <i>filter-name</i> term <i>term-name</i> then], [edit logical-systems <i>logical-system-name</i> firewall family inet filter <i>filter-name</i> term <i>term-name</i> then]
Release Information	Statement introduced before Junos OS Release 7.4. Logical systems support introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Reference a prefix-specific action.
Options	<i>prefix-action-name</i> —Name of a prefix-specific action to use to rate-limit traffic.
Related Documentation	<ul style="list-style-type: none">• Standard Firewall Filter Nonterminating Actions on page 4744• Prefix-Specific Counting and Policing Actions on page 4874

single-rate

Syntax	<pre>single-rate { (color-aware color-blind); committed-information-rate <i>bps</i>; committed-burst-size <i>bytes</i>; excess-burst-size <i>bytes</i>; }</pre>
Hierarchy Level	<pre>[edit dynamic-profiles <i>profile-name</i> firewall three-color-policer <i>name</i>], [edit firewall three-color-policer <i>policer-name</i>], [edit logical-systems <i>logical-system-name</i> firewall three-color-policer <i>policer-name</i>]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Logical systems support introduced in Junos OS Release 9.3.</p> <p>Support at the <code>[edit dynamic-profiles ... three-color-policer <i>name</i>]</code> hierarchy level introduced in Junos OS Release 11.4.</p>
Description	<p>Configure a single-rate three-color policer in which marking is based on the committed information rate (CIR), committed burst size (CBS), and excess burst size (EBS).</p> <p>Packets that conform to the CIR or the CBS are assigned low loss priority (green). Packets that exceed the CIR and the CBS but are within the EBS are assigned medium-high loss priority (yellow). Packets that exceed the EBS are assigned high loss priority (red).</p> <p>Green and yellow packets are always forwarded; this action is not configurable. You can configure red packets to be discarded. By default, red packets are forwarded.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Three-Color Policer Configuration Overview on page 4817 • color-aware on page 4978 • color-blind on page 4979 • two-rate on page 5012

three-color-policer (Applying)

Syntax	<pre>three-color-policer { (single-rate two-rate) <i>policer-name</i>; }</pre>
Hierarchy Level	[edit firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> then] [edit logical-systems <i>logical-system-name</i> firewall family <i>family-name</i> filter <i>filter-name</i> term <i>term-name</i> then]
Release Information	Statement introduced in Junos OS Release 7.4. single-rate statement added in Junos OS Release 8.2. Logical systems support introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	For M320 and T Series routers with Enhanced II Flexible PIC Concentrators (FPCs) and the T640 router with Enhanced Scaling FPC4, apply a tricolor marking policer.
Options	single-rate —Named tricolor policer is a single-rate policer. two-rate —Named tricolor policer is a two-rate policer. <i>policer-name</i> —Name of a tricolor policer.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Applying Tricolor Marking Policers to Firewall Filters on page 1456• Standard Firewall Filter Nonterminating Actions on page 4744• Three-Color Policer Configuration Overview on page 4817

three-color-policer (Configuring)

Syntax	<pre> three-color-policer <i>policer-name</i> { action { loss-priority high then discard; } filter-specific; logical-interface-policer; physical-interface-policer; shared-bandwidth-policer; single-rate { (color-aware color-blind); committed-burst-size <i>bytes</i>; committed-information-rate <i>bps</i>; excess-burst-size <i>bytes</i>; } two-rate { (color-aware color-blind); committed-burst-size <i>bytes</i>; committed-information-rate <i>bps</i>; peak-burst-size <i>bytes</i>; peak-information-rate <i>bps</i>; } } </pre>
Hierarchy Level	<p>[edit dynamic-profiles <i>profile-name</i> firewall], [edit firewall], [edit logical-systems <i>logical-system-name</i> firewall]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. The action and single-rate statements added in Junos OS Release 8.2. Logical systems support introduced in Junos OS Release 9.3. Support at the [edit dynamic-profiles ... firewall] hierarchy level introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Configure a three-color policer.
Options	<p><i>policer-name</i>—Name of the three-color policer. Reference this name when you apply the policer to an interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Statement Hierarchy for Configuring Policers on page 4811 • Configuring Tricolor Marking Policers on page 1455 • Three-Color Policer Configuration Guidelines on page 4912 • Basic Single-Rate Three-Color Policers on page 4915

- [Basic Two-Rate Three-Color Policers on page 4921](#)
- [Two-Color and Three-Color Logical Interface Policers on page 4927](#)
- [Two-Color and Three-Color Physical Interface Policers on page 4940](#)
- [Two-Color and Three-Color Policers at Layer 2 on page 4947](#)

two-rate

Syntax	<pre>two-rate { (color-aware color-blind); committed-information-rate <i>bps</i>; committed-burst-size <i>bytes</i>; peak-information-rate <i>bps</i>; peak-burst-size <i>bytes</i>; }</pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall three-color-policer <i>name</i>], [edit firewall three-color-policer <i>policer-name</i>], [edit logical-systems <i>logical-system-name</i> firewall three-color-policer <i>policer-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Logical systems support introduced in Junos OS Release 9.3. Support at the [edit dynamic-profiles ... three-color-policer <i>name</i>] hierarchy levels introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	<p>Configure a two-rate three-color policer in which marking is based on the committed information rate (CIR), committed burst size (CBS), peak information rate (PIR), and peak burst size (PBS).</p> <p>Packets that conform to the CIR or the CBS are assigned low loss priority (green). Packets that exceed the CIR and the CBS but are within the PIR or the PBS are assigned medium-high loss priority (yellow). Packets that exceed the PIR and the PBS are assigned high loss priority (red).</p> <p>Green and yellow packets are always forwarded; this action is not configurable. You can configure red packets to be discarded. By default, red packets are forwarded.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Three-Color Policer Configuration Overview on page 4817• color-aware on page 4978• color-blind on page 4979• single-rate on page 5009

Administration

- [Traffic Policing Standards on page 5013](#)
- [Firewall Filter and Policer Operational Mode Commands on page 5013](#)

Traffic Policing Standards

- [Supported Standards for Policing on page 5013](#)

Supported Standards for Policing

Three-color policers are part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment, which is described and defined in the following RFCs:

- RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*
- RFC 2475, *An Architecture for Differentiated Service*
- RFC 2597, *Assured Forwarding PHB Group*
- RFC 2598, *An Expedited Forwarding PHB*
- RFC 2698, *A Two Rate Three Color Marker*

In a DiffServ environment, the most significant 6 bits of the type-of-service (ToS) octet in the IP header contain a value called the *Differentiated Services code point* (DSCP). Within the DSCP field, the most significant 3 bits are interpreted as the *IP precedence* field, which can be used to select different per-hop forwarding treatments for the packet.

Firewall Filter and Policer Operational Mode Commands

clear firewall

Syntax	clear firewall (all counter <i>counter-name</i> filter <i>filter-name</i> log (all <i>logical-system-name</i>) logical-system <i>logical-system-name</i>)
Syntax (EX Series Switches)	clear firewall (all counter <i>counter-name</i> filter <i>filter-name</i> log (all <i>logical-system-name</i>) policer counter (all counter-id <i>counter-index</i>))
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. logical-system option introduced in Junos OS Release 9.3. log option introduced before Junos OS Release 11.4.
Description	Clear statistics about configured firewall filters. When you clear the counters of a filter, this impacts not only the counters shown by the CLI, but also the ones tracked by SNMP2.



NOTE: The clear firewall command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for graceful Routing Engine switchover (GRES).

If you clear statistics for firewall filters that are applied to Trio-based DPCs and that also use the **prefix-action** action on matched packets, wait at least 5 seconds before you enter the **show firewall prefix-action-stats** command. A 5-second pause between issuing the **clear firewall** and **show firewall prefix-action-stats** commands avoids a possible timeout of the **show firewall prefix-action-stats** command.

- Options**
- all**—Clear the packet and byte counts for all filters. On EX Series switches, this option also clears the packet counts for all policer counters.
 - counter *counter-name***—Clear the packet and byte counts for a filter counter that has been configured with the counter firewall filter action.
 - filter *filter-name***—Clear the packet and byte counts for the specified firewall filter.
 - log (all | *logical-system-name*)**—Clear log entries for IPv4 firewall filters that have **then log** as an action. Use **log all** to clear all log entries or **log *logical-system-name*** to clear log entries for the specified logical system.
 - logical-system *logical-system-name***—Clear the packet and byte counts for the specified logical system.
 - policer counter (all | counter-id *counter-index*)**—(EX8200 switches only) Clear all policer counters using the **policer counter all** command, or clear a specific policer counter using the **policer counter counter-id *counter-index*** command. The value of *counter-index* can be 0, 1, or 2.

Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show firewall on page 4450
List of Sample Output	clear firewall all on page 5015 clear firewall (counter counter-name) on page 5015 clear firewall (filter filter-name) on page 5015 clear firewall (policer counter all) (EX8200 Switch) on page 5015 clear firewall (policer counter counter-id counter-index) (EX8200 Switch) on page 5015

Sample Output

clear firewall all

```
user@host> clear firewall all
```

clear firewall (counter counter-name)

```
user@host> clear firewall counter port-filter-counter
```

clear firewall (filter filter-name)

```
user@host> clear firewall filter ingress-port-filter
```

clear firewall (policer counter all) (EX8200 Switch)

```
user@switch> clear firewall policer counter all
```

clear firewall (policer counter counter-id counter-index) (EX8200 Switch)

```
user@switch> clear firewall policer counter counter-id 0
```

show firewall

Syntax	<code>show firewall</code> <code><counter <i>counter-name</i>></code> <code><filter <i>filter-name</i>></code> <code><log></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><terse></code>
Syntax (EX Series Switches)	<code>show firewall</code> <code><counter <i>counter-name</i>></code> <code><detail></code> <code><filter <i>filter-name</i>></code> <code><log <(detail interface <i>interface-name</i>)>></code> <code><policer counters <(detail counter-id <i>counter-index</i> <detail>)>></code> <code><terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. logical-system option introduced in Junos OS Release 9.3. terse option introduced in Junos OS Release 9.4. policer counters option introduced in Junos OS Release 12.2 for EX Series switches. detail option introduced in Junos OS Release 12.3.
Description	Display statistics about configured firewall filters.
Options	<p>none—(Optional) Display statistics about all configured firewall filters and counters. For EX Series switches, this command also displays statistics about all configured policers.</p> <p>counter <i>counter-name</i>—(Optional) Name of a filter counter.</p> <p>detail—(EX Series switches only) (Optional) Display firewall filter statistics with enhanced policer.</p> <p>filter <i>filter-name</i>—(Optional) Name of a configured filter.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>log—(Optional) Display log entries for firewall filters.</p> <p>log <(detail interface <i>interface-name</i>)>—(EX Series switches only) (Optional) Display detailed log entries of firewall activity or log information about a specific interface.</p> <p>policer counters <(detail counter-id <i>counter-index</i> <detail>)>—(EX8200 switches only) (Optional) Display policer counter statistics in brief or in detail.</p> <p>terse—(Optional) Display firewall filter names only.</p>
Required Privilege Level	view

- Related Documentation**
- [clear firewall on page 4448](#)
 - [show firewall log on page 4457](#)
 - *Verifying That Firewall Filters Are Operational*
 - *Verifying That Policers Are Operational*

- List of Sample Output**
- [show firewall filter \(MX Series Router and EX Series Switch\) on page 5019](#)
 - [show firewall filter \(non MX Series Router and EX Series Switch\) on page 5019](#)
 - [show firewall filter \(Hierarchical Policers, MX Series with MPC\) on page 5019](#)
 - [show firewall filter \(Dynamic Input Filter\) on page 5019](#)
 - [show firewall \(Logical Systems\) on page 5019](#)
 - [show firewall \(counter counter-name\) on page 5020](#)
 - [show firewall log on page 5020](#)
 - [show firewall policer counters \(EX8200 Switch\) on page 5020](#)
 - [show firewall policer counters \(detail\) \(EX8200 Switch\) on page 5021](#)
 - [show firewall policer counters \(counter-id counter-index\) \(EX8200 Switch\) on page 5021](#)
 - [show firewall policer counters \(counter-id counter-index detail\) \(EX8200 Switch\) on page 5021](#)
 - [show firewall detail on page 5022](#)

- Output Fields** [Table 323 on page 4451](#) lists the output fields for the **show firewall** command. Output fields are listed in the approximate order in which they appear.

Table 371: show firewall Output Fields

Field Name	Field Description
Filter	<p>Name of a filter that has been configured with the filter statement at the [edit firewall] hierarchy level.</p> <p>Except on EX Series switches:</p> <ul style="list-style-type: none"> • When an interface-specific filter is displayed, the name of the filter is followed by the full interface name and by either -i for an input filter or -o for an output filter. • When dynamic filters are displayed, the name of the filter is followed by the full interface name and by either -in for an input filter or -out for an output filter. When a logical system-specific filter is displayed, the name of the filter is prefixed with two underscore (__) characters and the name of the logical system (for example, __ls1/filter1).
Counters	<p>Display filter counter information:</p> <ul style="list-style-type: none"> • Name—Name of a filter counter that has been configured with the counter firewall filter action. • Bytes—Number of bytes that match the filter term under which the counter action is specified. • Packets—Number of packets that matched the filter term under which the counter action is specified. <p>NOTE: On M and T series routers, firewall filters cannot count ip-options packets on a per option type and per interface basis. A limited work around is to use the show pfe statistics ip options command to see ip-options statistics on a per Packet Forwarding Engine (PFE) basis. See <i>show pfe statistics ip</i> for sample output.</p>

Table 371: show firewall Output Fields (*continued*)

Field Name	Field Description
Policers	<p>Display policer information:</p> <ul style="list-style-type: none"> • Name—Name of policer. • Bytes—(For two-color policers on MX Series routers and EX Series switches, and for hierarchical policers on interfaces hosted on MICs and MPCs in MX Series routers) Number of bytes that match the filter term under which the policer action is specified. This is only the number out-of-specification (out-of-spec) byte counts, not all the bytes in all packets policed by the policer. For other platforms, this field is blank. • Packets—Number of packets that matched the filter term under which the policer action is specified. This is only the number of out-of-specification (out-of-spec) packet counts, not all packets policed by the policer.
Policer Counter Index	(EX8200 switch only) Global management counter ID. The counter ID value (<i>counter-index</i>) can be 0, 1, or 2.
Green	(EX8200 switch only) Number of packets within the limits. The number of packets is smaller than the committed information rate (CIR).
Yellow	(EX8200 switch only) Number of packets partially within the limits. The number of packets is greater than the CIR, but the burst size is within the excess burst size (EBS) limit.
Discard	(EX8200 switch only) Number of discarded packets.
Bytes	(EX8200 switch only) Number of green, yellow, red, or discarded packets in bytes.
Packets	(EX8200 switch only) Number of green, yellow, red, or discarded packets.
Filter name	(EX8200 switch only) Name of the filter with a term associated to a policer.
Term name	(EX8200 switch only) Name of the term associated with a policer.
Policer name	(EX8200 switch only) Name of the policer that is associated with a global management counter.

Sample Output

show firewall filter (MX Series Router and EX Series Switch)

```
user@host> show firewall filter test
Filter: test
Counters:
Name                               Bytes          Packets
Counter-1                          0               0
Counter-2                          0               0
Policers:
Name                               Bytes          Packets
Policer-1                         2770            70
```

show firewall filter (non MX Series Router and EX Series Switch)

```
user@host> show firewall filter test
Filter: test
Counters:
Name                               Bytes          Packets
Counter-1                          0               0
Counter-2                          0               0
Policers:
Name                               Bytes          Packets
Policer-1                         70
```

show firewall filter (Hierarchical Policier, MX Series with MPC)

```
user@host> show firewall filter
FL_V4_PHY-HP-EF-AWARE-Gold=400k-MCAST=200k-Total=1M-ds-10/0/0:2:1-i

Filter: FL_V4_PHY-HP-EF-AWARE-Gold=400k-MCAST=200k-Total=1M-ds-10/0/0:2:1-i
Counters:
Name                               Bytes          Packets
AF1x_counter-ds-10/0/0:2:1-i      0               0
AF2x_counter-ds-10/0/0:2:1-i      25529445976    24500428
AF3x_counter-ds-10/0/0:2:1-i      2182022        39482
AF4x_counter-ds-10/0/0:2:1-i      0               0
BE_counter-ds-10/0/0:2:1-i        0               0
EF_counter-ds-10/0/0:2:1-i        14817044120    12265765
STD_counter-ds-10/0/0:2:1-i       0               0
Policers:
Name                               Bytes          Packets
POL_CE-PE_M=200k-filter-ds-10/0/0:2:1-i  5948099658    5708349
POL_CE-PE_G=400K_R=1M-filter-ds-10/0/0:2:1-i  ??????????    3572794
?????????????????????????????????????????  ??????????    ????????
```

show firewall filter (Dynamic Input Filter)

```
user@host> show firewall filter dfwd-ge-5/0/0.1-in
Filter: dfwd-ge-5/0/0.1-in
Counters:
Name                               Bytes          Packets
c1-ge-5/0/0.1-in                  0               0
```

show firewall (Logical Systems)

```
user@host> show firewall
```

```

Filter: __lr1/test
Counters:
Name                               Bytes          Packets
icmp                               420            5
Filter: __default_bpdu_filter__
Filter: __lr1/inet_filter1
Counters:
Name                               Bytes          Packets
inet_tcp_count                     0              0
inet_udp_count                     0              0
Filter: __lr1/inet_filter2
Counters:
Name                               Bytes          Packets
inet_icmp_count                    0              0
inet_pim_count                     0              0
Filter: __lr2/inet_filter1
Counters:
Name                               Bytes          Packets
inet_tcp_count                     0              0
inet_udp_count                     0              0

```

show firewall (counter counter-name)

```

user@host> show firewall counter icmp-counter
Filter: ingress-port-voip-class-filter
Counters:
Name                               Bytes          Packets
icmp-counter                       0              0

```

show firewall log

```

user@host> show firewall log
Log :

Time      Filter  Action Interface  Protocol  Src Addr
      Dest Addr
08:00:53 pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:52 pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:51 pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:50 pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:49 pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:48 pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:47 pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4

```

show firewall policer counters (EX8200 Switch)

```

user@switch> show firewall policer counters
Policer Counter Index 0:
          Bytes          Packets
Green:           73       15914
Yellow:           9       1962
Discard:        119      25942

```

```

Policer Counter Index 1:
      Bytes      Packets
Green:         0         0
Yellow:        0         0
Discard:       0         0

Policer Counter Index 2:
      Bytes      Packets
Green:         0         0
Yellow:        0         0
Discard:       0         0

```

show firewall policer counters (detail) (EX8200 Switch)

```

user@switch> show firewall policer counters detail
Policer Counter Index 0:
      Bytes      Packets
Green:         73      15914
Yellow:         9      1962
Discard:       119     25942

Filter name      Term name      Policer name
myfilter         polcr-term-1   myfilter-polcr-1
inet-filter-ae   ae-snmp        policer-1
inet-filter-ae   ae-ssh         policer-2

Policer Counter Index 1:
      Bytes      Packets
Green:         0         0
Yellow:        0         0
Discard:       0         0

Filter name      Term name      Policer name

Policer Counter Index 2:
      Bytes      Packets
Green:         0         0
Yellow:        0         0
Discard:       0         0

Filter name      Term name      Policer name

```

show firewall policer counters (counter-id counter-index) (EX8200 Switch)

```

user@switch> show firewall policer counters counter-id 0
Policer Counter Index 0:
      Bytes      Packets
Green:         73      15914
Yellow:         9      1962
Discard:       119     25942

```

show firewall policer counters (counter-id counter-index detail) (EX8200 Switch)

```

user@switch> show firewall policer counters counter-id 0 detail
Policer Counter Index 0:
      Bytes      Packets
Green:         73      15914
Yellow:         9      1962
Discard:       119     25942

Filter name      Term name      Policer name

```

```

myfilter          polcr-term-1      myfilter-polcr-1
inet-filter-ae    ae-snmp           policer-1
inet-filter-ae    ae-ssh            policer-2
    
```

show firewall detail

```

user@host> show firewall detail
Filter: __default_bpdu_filter__
    
```

Filter: foo

Counters:

Name

c1

Bytes	Packets
17652140	160474

Policers:

Name

P1-t1

Bytes	Packets
0	18286
0 18446744073709376546	
0 18446744073709358260	

OOS

Offered

Transmitted

show firewall filter version

Syntax	show firewall filter version <filter-name>
Release Information	Command introduced in Junos OS Release 10.2R2. Command introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Display the version number of the installed firewall filter in the Routing Engine.
Options	none—(Optional) Display the version number of all installed firewall filters. filter-name—(Optional) Name of a configured filter. If you specify the name of a filter, only the version number of that filter is displayed.
Additional Information	The initial version number is 1. This number increments by one when you modify the firewall filter settings or an associated prefix action. The maximum version number is 4,294,967,295. When the version number reaches 4,294,967,295, this number is reset to 1.
Required Privilege Level	view
List of Sample Output	show firewall filter version on page 5023
Output Fields	Table 372 on page 5023 lists the output fields for the show firewall filter version command. Output fields are listed in the approximate order in which they appear.

Table 372: show firewall filter version Output Fields

Field Name	Field Description
Filter	Name of a filter that has been configured with the filter statement at the [edit firewall] hierarchy level.
Version	Display the version number of the firewall filter.

Sample Output

show firewall filter version

```

user@host> show firewall filter version
Filter version information :
Filter                                     Version
test                                     10

```

show firewall log

Syntax	show firewall log <detail> <interface <i>interface-name</i> > <logical-system (<i>logical-system-name</i> all)>
Syntax (EX Series Switches)	show firewall log <detail> <interface <i>interface-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. logical-system option introduced in Junos OS Release 9.3.
Description	Display log information about firewall filters.
Options	none —Display log information about firewall filters. detail —(Optional) Display detailed information. interface <i>interface-name</i> —(Optional) Display log information about a specific interface. logical-system (<i>logical-system-name</i> all) —(Optional) Perform this operation on all logical systems or on a particular system.
Required Privilege Level	view
List of Sample Output	show firewall log on page 5025 show firewall log detail on page 5025
Output Fields	Table 324 on page 4457 lists the output fields for the show firewall log command. Output fields are listed in the approximate order in which they appear.

Table 373: show firewall log Output Fields

Field Name	Field Description
Time of Log	Time that the event occurred.
Filter	<p>Name of a filter that has been configured with the filter statement at the [edit firewall] hierarchy level.</p> <ul style="list-style-type: none"> A hyphen (-) indicates that the packet was handled by the Packet Forwarding Engine. A space (no hyphen) indicates the packet was handled by the Routing Engine. The notation pfe indicates packets logged by the Packet Forwarding Engine hardware filters.

Table 373: show firewall log Output Fields (*continued*)

Field Name	Field Description
Filter Action	Filter action: <ul style="list-style-type: none"> • A—Accept • D—Discard • R—Reject
Name of Interface	Ingress interface for the packet.
Name of protocol	Packet's protocol name: egp, gre, icmp, ipip, ospf, pim, rsvp, tcp, or udp .
Packet length	Length of the packet.
Source address	Packet's source address.
Destination address	Packet's destination address and port.

Sample Output

show firewall log

```

user@host>show firewall log
Time      Filter  Action Interface    Protocol  Src Addr    Dest Addr
13:10:12  pfe       D      rlsq0.902     ICMP      180.1.177.2 180.1.177.1
13:10:11  pfe       D      rlsq0.902     ICMP      180.1.177.2 180.1.177.1

```

show firewall log detail

```

user@host> show firewall log detail
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0Name of protocol: TCP, Packet Length: 50824, Source address:
172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 1020, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0

```

Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
....

show firewall prefix-action-stats

Syntax	show firewall prefix-action-stats filter <i>filter-name</i> prefix-action <i>prefix-action-name</i> <from <i>number</i> to <i>number</i> > <logical-system (<i>logical-system-name</i> all)>
Release Information	Command introduced before Junos OS Release 7.4. logical-system option introduced in Junos OS Release 9.3. Command introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Display prefix action statistics about configured firewall filters. If you clear statistics for firewall filters that are applied to Trio-based DPCs and that also use the prefix-action action on matched packets, wait at least 5 seconds before you enter the show firewall prefix-action-stats command. A 5-second pause between issuing the clear firewall and show firewall prefix-action-stats commands avoids a possible timeout of the show firewall prefix-action-stats command.
Options	filter <i>filter-name</i> —Name of a filter. prefix-action <i>prefix-action-name</i> —Name of a prefix action. from <i>number</i> to <i>number</i> —(Optional) Starting and ending counter or policer. logical-system (<i>logical-system-name</i> all) —(Optional) Perform this operation on all logical systems or on a particular system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear firewall on page 4448
List of Sample Output	show firewall prefix-action-stats on page 5028
Output Fields	Table 374 on page 5027 lists the output fields for the show firewall prefix-action-stats command. Output fields are listed in the approximate order in which they appear.

Table 374: show firewall prefix-action-stats Output Fields

Field Name	Field Description
Filter	Filter name. Filters configured for logical systems include the name of the filter prefixed with the two underscore characters (__) and the name of the logical system (for example, __ls1/filter1).

Sample Output

show firewall prefix-action-stats

```
user@host> show firewall prefix-action-stats filter test prefix-action act1
Filter: __ls2/test
```

show policer

Syntax	show policer <detail> <policer-name>
Release Information	Command introduced before Junos OS Release 7.4. Option detail introduced in Junos OS Release 12.3. Command introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Display the number of policed packets for a given policer or an aggregate policer. An aggregate policer is an aggregate of different policers on the same logical interface.
Options	none —Display the number of policed packets for all configured policers. detail —(Optional) Display enhanced statistics for policers. policer-name —(Optional) Display the number of policed packets for the specified policer.
Required Privilege Level	view
List of Sample Output	show policer (MX Series Router and EX Series Switch) on page 5029 show policer (non MX Series Router) on page 5030 show policer (Aggregate Policer, non MX Series Router) on page 5030 show policer detail on page 5030
Output Fields	Table 375 on page 5029 lists the output fields for the show policer command. Output fields are listed in the approximate order in which they appear.

Table 375: show policer Output Fields

Field Name	Field Description
Name	Name of the policer.
Bytes	(For two-color policers on MX Series routers and EX Series switches, and for hierarchical policers on interfaces hosted on MICs and MPCs in MX Series routers) Total number of bytes policed by the specified policer. For other platforms, this field is blank.
Packets	Total number of packets policed by the specified policer.

Sample Output

show policer (MX Series Router and EX Series Switch)

```

user@host> show policer
Policers:
Name                               Bytes          Packets
__default_arp_policer__           314520         5242

```

pol-2M-ge-1/2/0.1-inet-i	10372300	103723
pol-2M-ge-1/2/0.1-inet6-i	7727800	77278
pol-2M-ge-1/2/0.1-mp1s-i	7070336	67984
pol-2M-ge-1/2/0.1001-vp1s-i	65153700	651537
pol-2M-ge-1/2/0.2001-vp1s-i	65180900	651809
pol-2M-ge-1/2/0.3001-ccc-i	62202144	647939

show policer (non MX Series Router)

```
user@host> show policer
Policers:
Name                               Bytes          Packets
__default_arp_policer__
pol-2M-ge-1/2/0.1-inet-i          103723
pol-2M-ge-1/2/0.1-inet6-i         77278
pol-2M-ge-1/2/0.1-mp1s-i          67984
pol-2M-ge-1/2/0.1001-vp1s-i       651537
pol-2M-ge-1/2/0.2001-vp1s-i       651809
pol-2M-ge-1/2/0.3001-ccc-i        647939
```

show policer (Aggregate Policar, non MX Series Router)

```
user@host> show policer
Policers:
Name                               Bytes          Packets
__default_arp_policer__
P1-ae0.0-log_int-o                0
P2-ge-7/0/2.0-inet-o              0
P2-ge-7/0/2.0-inet6-o             0
__policer_tmpl__-term              0
__policer_tmpl__-fc0              0
__policer_tmpl__-fc0              0
__policer_tmpl__-fc1              0
__policer_tmpl__-fc0              0
__policer_tmpl__-fc1              0
__policer_tmpl__-fc2              0
__policer_tmpl__-fc0              0
__policer_tmpl__-fc1              0
__policer_tmpl__-fc2              0
__policer_tmpl__-fc3              0
```

show policer detail

```
user@host> show policer detail
Policers:
Name                               Bytes          Packets
__default_arp_policer__
  OOS                             0              0
  Offered                         0             496
  Transmitted                     0             496
P1-xe-1/0/0.0-inet-i
  OOS                             0             11329
  Offered                         0            111188
  Transmitted                     0            99859
```

CHAPTER 19

Spanning-Tree Protocols

- [Overview on page 5031](#)
- [Configuration on page 5052](#)
- [Administration on page 5125](#)

Overview

- [Spanning-Tree on page 5031](#)

Spanning-Tree

- [Spanning-Tree Protocols Supported on page 5032](#)
- [BPDU Overview on page 5033](#)
- [Loop Protection for Spanning-Tree Instance Interfaces Overview on page 5034](#)
- [Root Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [BPDU Protection for Spanning-Tree Instance Interfaces Overview on page 5036](#)
- [VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Overview on page 5036](#)
- [VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Topology on page 5037](#)
- [Layer 2 Protocol Tunneling Through a Network Overview on page 5039](#)
- [Spanning-Tree Protocols in Logical Systems on page 5041](#)
- [Bridge Priority for Election of Root Bridge and Designated Bridge on page 5041](#)
- [Maximum Age for Awaiting Arrival of Hello BPDUs on page 5042](#)
- [Hello Time for Root Bridge to Transmit Hello BPDUs on page 5042](#)
- [Forward Delay Before Ports Transition to Forwarding State on page 5042](#)
- [Spanning-Tree Instance Interface on page 5042](#)
- [Spanning-Tree Instance Interface Priority on page 5043](#)
- [Spanning-Tree Instance Interface Cost on page 5044](#)
- [Spanning-Tree Instance Interface Point-to-Point Link Mode on page 5044](#)
- [Spanning-Tree Instance Interface Configured as an Edge Port on page 5045](#)
- [Spanning-Tree Protocol Trace Options on page 5045](#)
- [MAC Address Rewriting Enabled for Layer 2 Protocol Tunneling on page 5046](#)

- [Layer 2 Protocol Tunnel Interface on page 5046](#)
- [Layer 2 Protocol to be Tunneled on page 5047](#)
- [Loop Protection for a Spanning-Tree Instance Interface on page 5047](#)
- [Root Protect for a Spanning-Tree Instance Interface on page 5048](#)
- [BPDU Protection for Individual Spanning-Tree Instance Interfaces on page 5048](#)
- [BPDU Protection on All Edge Ports of the Bridge on page 5049](#)
- [VPLS Multihoming: Priority of the Backup Bridge on page 5049](#)
- [VPLS Multihoming: Hold Time Before Switching to Primary Priority on page 5050](#)
- [VPLS Multihoming: System Identifier for Bridges in the Ring on page 5050](#)
- [VPLS Multihoming: Bridge Flush of MAC Cache on Topology Change on page 5051](#)

Spanning-Tree Protocols Supported

In a Layer 2 environment, you can configure various spanning-tree protocol versions to create a loop-free topology in Layer 2 networks.

A spanning-tree protocol is a Layer 2 control protocol (L2CP) that calculates the best path through a switched network containing redundant paths. A spanning-tree protocol uses bridge protocol data unit (BPDU) data frames to exchange information with other switches. A spanning-tree protocol uses the information provided by the BPDUs to elect a root bridge, identify root ports for each switch, identify designated ports for each physical LAN segment, and prune specific redundant links to create a loop-free tree topology. The resulting tree topology provides a single active Layer 2 data path between any two end stations.



NOTE: In discussions of spanning-tree protocols, the terms *bridge* and *switch* are often used interchangeably.

The Juniper Networks MX Series 3D Universal Edge Routers and EX Series switches support STP, RSTP, MSTP, and VSTP.

- The original Spanning-Tree Protocol (STP) is defined in the IEEE 802.1D 1998 specification. A newer version called Rapid Spanning-Tree Protocol (RSTP) was originally defined in the IEEE 802.1w draft specification and later incorporated into the IEEE 802.1D-2004 specification. A recent version called Multiple Spanning-Tree Protocol (MSTP) was originally defined in the IEEE 802.1s draft specification and later incorporated into the IEEE 802.1Q-2003 specification. The VLAN Spanning-Tree Protocol (VSTP) is compatible with the Per-VLAN Spanning Tree Plus (PVST+) and Rapid-PVST+ protocols supported on Cisco Systems routers and switches.
- RSTP provides faster reconvergence time than the original STP by identifying certain links as point to point and by using protocol handshake messages rather than fixed timeouts. When a point-to-point link fails, the alternate link can transition to the forwarding state without waiting for any protocol timers to expire.
- MSTP provides the capability to logically divide a Layer 2 network into regions. Every region has a unique identifier and can contain multiple instances of spanning trees. All

regions are bound together using a Common Instance Spanning Tree (CIST), which is responsible for creating a loop-free topology *across* regions, whereas the Multiple Spanning-Tree Instance (MSTI) controls topology *within* regions. MSTP uses RSTP as a converging algorithm and is fully interoperable with earlier versions of STP.

- VSTP maintains a separate spanning-tree instance for each VLAN. Different VLANs can use different spanning-tree paths. When different VLANs use different spanning-tree paths, the CPU processing resources being consumed increase as more VLANs are configured. VSTP BPDU packets are tagged with the corresponding VLAN identifier and are transmitted to the multicast destination media access control (MAC) address **01-00-0c-cc-cc-cd** with a protocol type of **0x010b**. VSTP BPDUs are tunneled by pure IEEE 802.1q bridges.



NOTE: All virtual switch routing instances configured on an MX Series router are supported using only one spanning-tree process. The Layer 2 control protocol process is named l2cpd.

For more information about the various versions of spanning-tree protocols, see the appropriate IEEE specification.

Related Documentation

- [Loop Protection for Spanning-Tree Instance Interfaces Overview on page 5034](#)
- [Root Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [BPDU Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Overview on page 5036](#)

BPDU Overview

In a Layer 2 bridge environment, spanning-tree protocols use data frames called Bridge Protocol Data Units (BPDUs) to exchange information among bridges.

Spanning-tree protocols on peer systems exchange BPDUs, which contain information about port roles, bridge IDs, and root path costs. On each MX Series router or EX Series switch, the spanning-tree protocol uses this information to elect a root bridge, identify root ports for each switch, identify designated ports for each physical LAN segment, and prune specific redundant links to create a loop-free tree topology. The resulting tree topology provides a single active Layer 2 data path between any two end stations.



NOTE: In discussions of spanning-tree protocols, the terms *bridge* and *switch* are often used interchangeably.

The transmission of BPDUs is controlled by the Layer 2 Control Protocol process (l2cpd) on MX Series 3D Universal Edge Routers.

The transmission of periodic packets on behalf of the l2cpd process is carried out by periodic packet management (PPM), which, by default, is configured to run on the Packet Forwarding Engine. The ppm process on the Packet Forwarding Engine ensures that

the BPDUs are transmitted even when the l2cpd process control plane is unavailable, and keeps the remote adjacencies alive during a unified in-service software upgrade (unified ISSU). However, if you want the distributed PPM (ppmd) process to run on the Routing Engine instead of the Packet Forwarding Engine, you can disable the ppm process on the Packet Forwarding Engine. For more information, see the *Junos OS High Availability Configuration Guide*.

On MX Series routers or EX Series switches with redundant Routing Engines (two Routing Engines that are installed in the same router), you can configure nonstop bridging. Nonstop bridging enables the router to switch from a primary Routing Engine to a backup Routing Engine without losing Layer 2 Control Protocol (L2CP) information. Nonstop bridging uses the same infrastructure as graceful Routing Engine switchover (GRES) to preserve interface and kernel information. However, nonstop bridging also saves L2CP information by running the l2cpd process on the backup Routing Engine.



NOTE: To use nonstop bridging, you must first enable GRES.

Nonstop bridging is supported for the following Layer 2 control protocols:

- Spanning-Tree Protocol (STP)
- Rapid Spanning-Tree Protocol (RSTP)
- Multiple Spanning-Tree Protocol (MSTP)

For more information about GRES and nonstop bridging, see the *Junos OS High Availability Configuration Guide*.

**Related
Documentation**

- [Spanning-Tree Protocols Supported on page 5032](#)
- [Loop Protection for Spanning-Tree Instance Interfaces Overview on page 5034](#)
- [Root Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [BPDU Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)

Loop Protection for Spanning-Tree Instance Interfaces Overview

Spanning-tree protocol loop protection enhances the normal checks that spanning-tree protocols perform on interfaces. Loop protection performs a specified action when BPDUs are not received on a nondesignated port interface. You can choose to block the interface or issue an alarm when bridge protocol data units (BPDUs) are not received on the port.

The spanning-tree protocol family is responsible for breaking loops in a network of bridges with redundant links. However, hardware failures can create forwarding loops (STP loops) and cause major network outages. Spanning-tree protocols break loops by blocking ports (interfaces). However, errors occur when a blocked port transitions erroneously to a forwarding state.

Ideally, a spanning-tree protocol bridge port remains blocked as long as a superior alternate path to the root bridge exists for a connected LAN segment. This designated port is determined by receiving superior BPDUs from a peer on that port. When other

ports no longer receive BPDUs, the spanning-tree protocol considers the topology to be loop free. However, if a blocked or alternate port moves into a forwarding state, this creates a loop.

By default (that is, without spanning-tree protocol loop protection configured), an interface that stops receiving BPDUs will assume the designated port role and possibly result in a spanning-tree protocol loop.

You can configure spanning-tree protocol loop protection to improve the stability of Layer 2 networks.

You configure spanning-tree protocol loop protection to prevent selected interfaces from interpreting the lack of received BPDUs as a “false positive” condition for making the interface the designated port.

**Related
Documentation**

- [Loop Protection for a Spanning-Tree Instance Interface on page 5047](#)
- [Configuring Loop Protection for a Spanning-Tree Instance Interface on page 5082](#)
- [Example: Enabling Loop Protection for Spanning-Tree Protocols on page 5052](#)

Root Protection for Spanning-Tree Instance Interfaces Overview

Root protect helps to enforce the root bridge placement in a Layer 2 switched network. Enable root protect on interfaces that should not receive superior bridge protocol data units (BPDUs) from the root bridge. Typically, these ports are Spanning-Tree-Protocol-designated ports on an administrative boundary. Enabling root protect ensures the port remains a spanning-tree designated port.

If the bridge receives superior BPDUs on a port that has root protect enabled, that port transitions to a root-prevented STP state and the interface is blocked. This prevents a bridge that should not be the root bridge from being elected the root bridge.

After the bridge stops receiving superior BPDUs on the port with root protect enabled and the received BPDUs time out, that port transitions back to the STP-designated port state.

**Related
Documentation**

- [Root Protect for a Spanning-Tree Instance Interface on page 5048](#)
- [Enabling Root Protect for a Spanning-Tree Instance Interface on page 5083](#)

BPDU Protection for Spanning-Tree Instance Interfaces Overview

By default, if a Bridge Protocol Data Unit (BPDU) data frame is received on a blocked interface, the system will disable the interface and stop forwarding frames out the interface until the interface is explicitly cleared.

The Spanning-Tree Protocol (STP) family is designed to break possible loops in a Layer 2 bridged network. Loop prevention avoids damaging broadcast storms that can potentially render the network useless. STP processes on bridges exchange BPDUs to determine the LAN topology, decide the root bridge, stop forwarding on some ports, and so on.

However, a misbehaving user application or device can interfere with the operation of the STP protocols and cause network problems.

On the MX Series routers or EX Series switches only, you can configure BPDU protection to ignore BPDUs received on interfaces where none should be expected (for example, a LAN interface on a network edge with no other bridges present). If a BPDU is received on a blocked interface, the interface is disabled and stops forwarding frames. By default, all BPDUs are accepted and processed on all interfaces.

You can configure BPDU protection on interfaces with the following encapsulation types:

- **ethernet-bridge**
- **ethernet-vpls**
- **extended-vlan-bridge**
- **vlan-vpls**
- **extended-vlan-vpls**

You can configure BPDU protection on individual interfaces or on all the edge ports of the bridge.

**Related
Documentation**

- [Configuring BPDU Protection on Individual Interfaces on page 5084](#)
- [Configuring BPDU Protection on All Edge Ports on page 5085](#)
- [Spanning-Tree Protocols Supported on page 5032](#)
- [Root Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Overview on page 5036](#)

VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Overview

Redundancy is built into many networks through the use of alternate links and paths, which often take the shape of rings.

In the case of multiple hosts attached to customer edge (CE) routers and provider edge (PE) routers to secure virtual private LAN service (VPLS), this practice is often called *multihoming*:

- Multiple hosts attach to CE routers, which are attached to each other as well as to the PE routers that access the VPLS network cloud. Any single link between the edge routers can fail without impacting the hosts' access to the VPLS services.
- This Layer 2 ring connects to the multiprotocol link switching (MPLS) infrastructure through two PE routers. Link breaks on the ring are protected by running a version of the spanning-tree protocol with the root-protect option enabled.

The virtual private network (VPN) protocols at Layer 3, however, are not aware of the blocking state that results from this root protection setup (rings or loops are not permitted at Layer 2 because the Layer 2 protocols will not function properly).

However, to keep the Layer 2 ring functioning in a multihomed environment with link failures, the spanning-tree protocol running on the MX Series routers requires the following additional configuration:

- The VPN protocols have to act on the blocking and unblocking of interfaces by the spanning-tree protocol. Specifically, media access control (MAC) flush messages need to be sent by the blocking PE router to LDP peers in order to flush the MAC addresses learned when other interface ports were blocked.
- Also, if an active PE router with VPLS root protection bridging enabled loses VPLS connectivity, root protection requires that the bridge switch to the other PE router to maintain connectivity. The spanning-tree protocol needs to be aware of the status of the VPLS connectivity on the PE router. If the MAC address cache is not flushed when the topology changes, frames could be sent to the wrong device.

You can control the actions taken by the MX Series router when the topology changes in a multihomed Layer 2 ring VPLS environment using *VPLS root protection*.


**Related
Documentation**

- [VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Topology on page 5037](#)
- [Configuring VPLS Root Protection Topology Change Actions to Control Global Spanning Tree Behavior on page 5080](#)
- [Configuring VPLS Root Protection Topology Change Actions to Control VLAN Spanning Tree Behavior on page 5085](#)
- [Example: Configuring VPLS Root Topology Change Actions on page 5053](#)

VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Topology

[Figure 110 on page 5038](#) shows hosts connected to CE routers and to a VPLS network through two PE routers. The CE routers are also connected, forming a kind of ring structure.

Figure 110: VPLS Multihoming Configuration



The two PE routers have their own links to a VPLS network service, but are not directly connected to each other. All four edge routers run some type of spanning-tree protocol with root protection enabled, and only one PE interface will be in the forwarding state, the other being blocked.

Assume this forwarding interface is through PE1. If the link between CE1 and CE2 fails, then the blocking PE2 interface must detect a root protection switch and move to the forwarding state. All of the MAC addresses learned by CE2 that connect to the VPLS network service through PE1 need to be flushed. In the same way, when the link between CE1 and CE2 is restored, PE2 again detects the root protection switch and begins blocking again. Now all of the MAC addresses learned by CE2 that connect through PE2 need to be flushed. All of this is controlled by configuring VPLS root protection topology change actions on the CE routers.

Also, at a global level, each type of spanning-tree protocol will have a priority hold time associated with it. This is the number of seconds in the range from 1 through 255 seconds that the system waits to switch to the primary priority when the first core domain comes up. The default is 2 seconds. This allows the maximum number of core domains to come up, and some might be slower than others.

Related Documentation

- [VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Overview on page 5036](#)
- [Configuring VPLS Root Protection Topology Change Actions to Control Global Spanning Tree Behavior on page 5080](#)
- [Configuring VPLS Root Protection Topology Change Actions to Control VLAN Spanning Tree Behavior on page 5085](#)
- [Example: Configuring VPLS Root Topology Change Actions on page 5053](#)

Layer 2 Protocol Tunneling Through a Network Overview

Layer 2 protocol tunneling allows Layer 2 protocol data units (PDUs) to be tunneled through a network. This is useful to provide a single spanning-tree protocol domain for subscribers across a service provider network. It is also useful for tunneling Cisco Discovery Protocol (CDP) or VLAN Trunk Protocol (VTP) PDUs across a network.

Layer 2 protocol tunneling is supported on MX Series routers with Enhanced (Dense Port Concentrators) DPCs and Enhanced Queuing (DPCs), see [Table 377 on page 5040](#) for a list of the DPCs supported. Layer 2 protocol tunneling is supported on all Modular Port Concentrators (MPCs),



NOTE: Layer 2 protocol tunneling is not supported on Rev-A DPCs on MX Series routers because of microcode space limitations.

When a control packet for STP, CDP, or VTP is received on a service provider edge port configured for Layer 2 protocol tunneling, the multicast destination MAC address is rewritten with the predefined multicast tunnel MAC address of **01:00:0c:cd:cd:d0**. The packet is transported across the provider network transparently to the other end of the tunnel and the original multicast destination MAC address is restored when the packet is transmitted.

If a packet is received on a tunnel interface that already has a destination multicast MAC address of **01:00:0c:cd:cd:d0**, the port enters an error state and is shut down. To clear the error condition, the administrator must enter the **clear error mac-rewrite interface *interface-name*** command.

Layer 2 protocol tunneling and MAC rewrite are supported in VPLS, but only certain hardware configurations are supported.

[Table 376 on page 5039](#) shows the MPCs and Enhanced DPCs supported when configuring Layer 2 protocol tunneling and VPLS.

Table 376: MAC Rewrite and VPLS Configurations

CE-Facing Interface	PE-Core Facing Interface	Layer 2 Protocol Tunneling
MPC	MPC	Yes
MPC	Enhanced DPC	Yes
Enhanced DPC	MPC	Yes
Enhanced DPC	Enhanced DPC	No

[Table 377 on page 5040](#) lists the DPCs that support the Layer 2 tunneling protocol.

Table 377: DPCs Supported for Layer 2 Protocol Tunneling

DPC Name	DPC Model Number
Gigabit Ethernet	
<i>Gigabit Ethernet Enhanced DPC with SFP</i>	DPCE-R-40GE-SFP
<i>Gigabit Ethernet Enhanced Ethernet Services DPC with SFP</i>	DPCE-X-40GE-SFP
<i>Gigabit Ethernet Enhanced Queuing Ethernet Services DPC with SFP</i>	DPCE-X-Q-40GE-SFP
<i>Gigabit Ethernet Enhanced Queuing IP Services DPCs with SFP</i>	DPCE-R-Q-20GE-SFP
<i>Gigabit Ethernet Enhanced Queuing IP Services DPCs with SFP</i>	DPCE-R-Q-40GE-SFP
10-Gigabit Ethernet	
<i>10-Gigabit Ethernet Enhanced DPCs with XFP</i>	DPCE-R-2XGE-XFP
<i>10-Gigabit Ethernet Enhanced DPCs with XFP</i>	DPCE-R-4XGE-XFP
<i>10-Gigabit Ethernet Enhanced Ethernet Services DPC with XFP</i>	DPCE-X-4XGE-XFP
<i>10-Gigabit Ethernet Enhanced Queuing Ethernet Services DPC with XFP</i>	DPCE-X-Q-4XGE-XFP
<i>10-Gigabit Ethernet Enhanced Queuing IP Services DPC with XFP</i>	DPCE-R-Q-4XGE-XFP
Multi-Rate Ethernet	
<i>Multi-Rate Ethernet Enhanced DPC with SFP and XFP</i>	DPCE-R-20GE-2XGE
<i>Multi-Rate Ethernet Enhanced Ethernet Services DPC with SFP and XFP</i>	DPCE-X-20GE-2XGE
<i>Multi-Rate Ethernet Enhanced Queuing IP Services DPC with SFP and XFP</i>	DPCE-R-Q-20GE-2XGE
Tri-Rate Ethernet	
<i>Tri-Rate Enhanced DPC</i>	DPCE-R-40GE-TX
<i>Tri-Rate Enhanced Ethernet Services DPC</i>	DPCE-X-40GE-TX



NOTE: When an MX Series router or EX Series switch sends a RADIUS access request, the Chargeable-User-Identity parameter is sent with an empty field. For more information about configuring RADIUS, see the *Junos Subscriber Access Configuration Guide*.

Related Documentation

- [Configuring Layer 2 Protocol Tunneling on page 5086](#)
- [Checking for a MAC Rewrite Error Condition Blocking a Spanning-Tree Instance Interface on page 5127](#)
- [Clearing a MAC Rewrite Error Condition Blocking a Spanning-Tree Instance Interface on page 5127](#)

Spanning-Tree Protocols in Logical Systems

On MX Series routers and EX Series switches only, you can configure spanning-tree protocols in logical systems for bridge domains and other virtual-switch routing instances.

When configuring spanning-tree protocols in logical systems for bridge domains and other virtual-switch routing instances, the following guidelines apply:

- You can only configure 16 logical systems.
- Logging is performed for the entire device and not per logical system.
- You cannot restart Layer 2 learning for an individual logical system.

Related Documentation

- [Spanning-Tree Protocols Supported on page 5032](#)

Bridge Priority for Election of Root Bridge and Designated Bridge

Use the bridge priority to control which bridge is elected as the root bridge and also to control which bridge is elected the root bridge when the initial root bridge fails.

The root bridge for each spanning-tree protocol instance is determined by the bridge ID. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge. The bridge with the lowest bridge ID is elected as the root bridge. If the bridge priorities are equal or if the bridge priority is not configured, the bridge with the lowest MAC address is elected the root bridge.

The bridge priority can also be used to determine which bridge becomes the designated bridge for a LAN segment. If two bridges have the same path cost to the root bridge, the bridge with the lowest bridge ID becomes the designated bridge.

The bridge priority can be set only in increments of 4096.

Related Documentation

- [Configuring Rapid Spanning-Tree Protocol on page 5073](#)
- [Configuring Multiple Spanning-Tree Protocol on page 5067](#)
- [Configuring VLAN Spanning-Tree Protocol on page 5077](#)

- [bridge-priority on page 1876](#)

Maximum Age for Awaiting Arrival of Hello BPDUs

The maximum age timer specifies the maximum expected arrival time of hello BPDUs. If the maximum age timer expires, the bridge detects that the link to the root bridge has failed and initiates a topology reconvergence. The maximum age timer should be longer than the configured hello timer.

Related Documentation

- [Configuring Rapid Spanning-Tree Protocol on page 5073](#)
- [Configuring Multiple Spanning-Tree Protocol on page 5067](#)
- [Configuring VLAN Spanning-Tree Protocol on page 5077](#)
- [max-age on page 5101](#)

Hello Time for Root Bridge to Transmit Hello BPDUs

The hello timer specifies the time interval at which the root bridge transmits configuration BPDUs.

Related Documentation

- [Configuring Rapid Spanning-Tree Protocol on page 5073](#)
- [Configuring Multiple Spanning-Tree Protocol on page 5067](#)
- [Configuring VLAN Spanning-Tree Protocol on page 5077](#)
- [hello-time on page 5099](#)

Forward Delay Before Ports Transition to Forwarding State

The forwarding delay timer specifies the length of time a spanning-tree protocol bridge port remains in the listening and learning states before transitioning to the forwarding state. Setting the interval too short could cause unnecessary spanning-tree reconvergence. Before changing this parameter, you should have a thorough understanding of spanning-tree protocols.

Related Documentation

- [Configuring Rapid Spanning-Tree Protocol on page 5073](#)
- [Configuring Multiple Spanning-Tree Protocol on page 5067](#)
- [Configuring VLAN Spanning-Tree Protocol on page 5077](#)
- [forward-delay on page 5098](#)

Spanning-Tree Instance Interface

STP and RSTP are limited to a single instance on any physical interface. Use the **interface** statement to configure which interfaces participate in the STP or RSTP instance.

MSTP supports multiple instances on a single physical interface. Use the **interface** statement to configure which logical interfaces participate in MSTP.

For VSTP, interfaces can be configured at the global level or at the VLAN level. Interfaces configured at the global VSTP level will be enabled for all the configured VLANs. If an interface is configured at both the global and VLAN levels, the configuration at the VLAN level overrides the global configuration.

**Related
Documentation**

- [Configuring Rapid Spanning-Tree Protocol on page 5073](#)
- [Configuring Multiple Spanning-Tree Protocol on page 5067](#)
- [Configuring VLAN Spanning-Tree Protocol on page 5077](#)
- [cost on page 5094](#)
- [edge on page 5096](#)
- [interface \(Spanning Tree\) on page 1896](#)
- [mode on page 5103](#)
- [priority on page 5106](#)

Spanning-Tree Instance Interface Priority

The root port is the interface on the nonroot bridge with the lowest path cost to the root bridge. When multiple interfaces have the same path cost to the root bridge, the interface with the lowest interface priority is selected as the root port.

If the interface priority is not configured and multiple interfaces have the same path cost to the root bridge, the interface with the lowest interface identifier is selected as the root port.

If the interface priority is configured under the MSTP protocol, this becomes the default value for all interfaces. If the interface priority is configured under the MSTI interface, the value overrides the default for that interface.

If the interface priority is configured at both the VSTP global and VLAN levels, the configuration at the VLAN level overrides the global configuration.

**Related
Documentation**

- [Configuring Rapid Spanning-Tree Protocol on page 5073](#)
- [Configuring Multiple Spanning-Tree Protocol on page 5067](#)
- [Configuring VLAN Spanning-Tree Protocol on page 5077](#)
- [interface \(Spanning Tree\) on page 1896](#)
- [priority on page 5106](#)

Spanning-Tree Instance Interface Cost

The path cost used to calculate the root path cost from any given LAN segment is determined by the total cost of each link in the path. By default, the link cost is determined by the speed of the link. The interface cost can be configured to override the default cost and control which bridge is the designated bridge and which port is the designated port. In MSTP the CIST external path cost is determined by the link speed and the number of hops.

If the interface cost is not configured, the cost is determined by the speed of the interface. For example, a 100-Mbps link has a default path cost of 19, a 1000-Mbps link has a default path cost of 4, and a 10-Gbps link has a default path cost of 2.

If the interface cost is configured under MSTP, this becomes the default value for all interfaces. If the interface cost is configured under the MSTI interface, the value overrides the default for that interface.

If the interface cost is configured at both the VSTP global and VLAN levels, the configuration at the VLAN level overrides the global configuration.

The interface cost should be set the same for all interfaces connected to the same LAN segment.

Related Documentation

- [Configuring Rapid Spanning-Tree Protocol on page 5073](#)
- [Configuring Multiple Spanning-Tree Protocol on page 5067](#)
- [Configuring VLAN Spanning-Tree Protocol on page 5077](#)
- [cost on page 5094](#)
- [interface \(Spanning Tree\) on page 1896](#)

Spanning-Tree Instance Interface Point-to-Point Link Mode

The interface mode allows RSTP, MSTP, and VSTP to converge faster than the original STP on point-to-point links. The protocol does not need to wait for timers on point-to-point links. Configure interfaces that have a point-to-point link to another Layer 2 bridge as **p2p**. This parameter is ignored if the STP is configured to run the original spanning-tree version.

If the interface mode is configured at both the VSTP global and VLAN levels, the configuration at the VLAN level overrides the global configuration.

Related Documentation

- [Configuring Rapid Spanning-Tree Protocol on page 5073](#)
- [Configuring Multiple Spanning-Tree Protocol on page 5067](#)
- [Configuring VLAN Spanning-Tree Protocol on page 5077](#)
- [mode on page 5103](#)
- [interface \(Spanning Tree\) on page 1896](#)

Spanning-Tree Instance Interface Configured as an Edge Port

RSTP, MSTP, and VSTP instance interfaces configured as *edge ports* enable the protocol to converge faster than the original IEEE 802.1D STP version. Edge ports transition directly to the forwarding state, and so the protocol does not need to wait for BPDUs to be received on edge ports.

The Junos OS supports automatic detection of edge ports as described in the RSTP standard. Layer 2 bridges do not expect to receive BPDUs for edge ports. If a BPDU is received for an edge port, the port becomes a non-edge port.

Keep the following guidelines in mind when configuring spanning-tree instance interfaces as edge ports:

- Do not configure a spanning-tree instance interface as an edge port if it is connected to any Layer 2 bridge. An instance interface connected to Layer 2 bridges but configured as an edge port can cause physical loops.
- If the spanning-tree protocol is configured to run the original IEEE 802.1D spanning-tree version, the edge-port option (if configured) is ignored.
- If edge ports are configured at both the VSTP global and VLAN levels, the configuration at the VLAN level overrides the global configuration.

Related Documentation

- [Configuring Rapid Spanning-Tree Protocol on page 5073](#)
- [Configuring Multiple Spanning-Tree Protocol on page 5067](#)
- [Configuring VLAN Spanning-Tree Protocol on page 5077](#)
- [edge on page 5096](#)
- [interface \(Spanning Tree\) on page 1896](#)

Spanning-Tree Protocol Trace Options

In order to trace spanning-tree protocol operations, you can set spanning-tree protocol-specific trace options in the spanning-tree protocol configuration.

For general information about tracing and global tracing options, see the statement summary for the global **traceoptions** statement in the *Junos OS Routing Protocols Configuration Guide*.

Related Documentation

- [Configuring Rapid Spanning-Tree Protocol on page 5073](#)
- [Configuring Multiple Spanning-Tree Protocol on page 5067](#)
- [Configuring VLAN Spanning-Tree Protocol on page 5077](#)
- [Example: Tracing Spanning-Tree Protocol Operations on page 5065](#)
- [traceoptions \(Spanning Tree\) on page 5109](#)

MAC Address Rewriting Enabled for Layer 2 Protocol Tunneling

To configure Layer 2 protocol tunneling, you must enable MAC address rewriting by installing the destination multicast tunnel MAC address of **01:00:0c:cd:cd:d0** in the MAC table.

To enable MAC address rewriting, include the **mac-rewrite** statement at the **[edit protocols layer2-control]** hierarchy level.

When enabling MAC address rewriting for Layer 2 protocol tunneling, the following guidelines apply:

- You can enable Layer 2 protocol tunneling for untagged interfaces.
- You can enable Layer 2 protocol tunneling for single-identifier tagged ports.
- You cannot enable Layer 2 protocol tunneling for double identifier tagged interfaces

Related Documentation

- [Layer 2 Protocol Tunneling Through a Network Overview on page 5039](#)
- [Layer 2 Protocol Tunnel Interface on page 5046](#)
- [Layer 2 Protocol to be Tunneled on page 5047](#)
- [Configuring Layer 2 Protocol Tunneling on page 5086](#)

Layer 2 Protocol Tunnel Interface

To configure the interface where Layer 2 protocol tunneling is enabled, include the **interface ge-fpc/pic/port** statement at the **[edit protocols layer2-control]** hierarchy level.

Keep the following guidelines in mind when configuring Layer 2 protocol tunneling:

- Layer 2 protocol tunneling is supported on MX Series routers with enhanced queuing Dense Port Concentrators (DPCs).
- Layer 2 protocol tunneling must be configured on the interfaces at each end of the tunnel.
- You can enable Layer 2 protocol tunneling for untagged interfaces and single-identifier tagged interfaces only.
- For single-identifier tagged ports, configure a logical interface with the native VLAN identifier. This configuration associates the untagged control packets with a logical interface.
- You cannot enable Layer 2 protocol tunneling for double identifier tagged interfaces.

Related Documentation

- [Layer 2 Protocol Tunneling Through a Network Overview on page 5039](#)
- [MAC Address Rewriting Enabled for Layer 2 Protocol Tunneling on page 5046](#)
- [Layer 2 Protocol to be Tunneled on page 5047](#)
- [Configuring Layer 2 Protocol Tunneling on page 5086](#)

Layer 2 Protocol to be Tunneled

To configure Layer 2 protocol tunneling, you must specify the protocol that is to be tunneled using the Layer 2 tunnel:

- **cdp**—Cisco Discovery Protocol.
- **stp**—All versions of the spanning-tree protocol.
- **vtp**—Tunnel the VLAN trunk protocol.

For each protocol specified, a static destination MAC address corresponding to the protocol being tunneled is installed in the MAC table.

To specify the protocol that will be tunneled by the Layer 2 protocol tunneling, you can include the **protocol (cdp | stp | vtp)** statement at the **[edit protocols layer2-control mac-rewrite interface ge-fpc/pic/port]** hierarchy level.



NOTE: When CDP, STP, or VTP is configured for tunneling on a customer-facing port in a provider bridge, the corresponding protocol should not be enabled for operation on that interface.

Related Documentation

- [Layer 2 Protocol Tunneling Through a Network Overview on page 5039](#)
- [MAC Address Rewriting Enabled for Layer 2 Protocol Tunneling on page 5046](#)
- [Layer 2 Protocol Tunnel Interface on page 5046](#)
- [Configuring Layer 2 Protocol Tunneling on page 5086](#)

Loop Protection for a Spanning-Tree Instance Interface

By default, a spanning-tree protocol interface that stops receiving Bridge Protocol Data Unit (BPDU) data frames will transition to the designated port (forwarding) state, creating a potential loop. To prevent a spanning-tree instance interface from interpreting a lack of received BPDUs as a “false positive” condition for assuming the designated port role, you can configure one of the following loop protection options:

- Configure the router to raise an alarm condition if the spanning-tree instance interface has not received BPDUs during the timeout interval.
- Configure the router to block the spanning-tree instance interface if the interface has not received BPDUs during the timeout interval.



NOTE: Spanning-tree instance interface loop protection is enabled for all spanning-tree instances on the interface, but blocks or alarms only those instances that stop receiving BPDUs.

We recommend you configure loop protection only on non-designated interfaces such as the root or alternate interfaces. Otherwise, if you configure loop protection on both sides of a designated link, then certain STP configuration events (such as setting the root bridge priority to an inferior value in a topology with many loops) can cause both interfaces to transition to blocking mode.

**Related
Documentation**

- [Loop Protection for Spanning-Tree Instance Interfaces Overview on page 5034](#)
- [Configuring Loop Protection for a Spanning-Tree Instance Interface on page 5082](#)
- [Example: Enabling Loop Protection for Spanning-Tree Protocols on page 5052](#)
- [bpdudisable-timeout-action](#)
- [interface \(Spanning Tree\) on page 1896](#)

Root Protect for a Spanning-Tree Instance Interface

When root protect is enabled on an interface, it is enabled for all spanning-tree protocol instances on that interface. The interface is blocked only for those instances that receive superior BPDUs.

By default, root protect is disabled.

**Related
Documentation**

- [Root Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [Enabling Root Protect for a Spanning-Tree Instance Interface on page 5083](#)
- [interface \(Spanning Tree\) on page 1896](#)
- [no-root-port on page 5119](#)

BPDU Protection for Individual Spanning-Tree Instance Interfaces

To configure BPDU protection on one or more spanning-tree instance interfaces, include the **bpdublock** statement:

```
bpdublock {  
  interface interface-name;  
  disable-timeout seconds;  
}
```



NOTE: If you also include the optional **disable-timeout *seconds*** statement, *blocked interfaces* are automatically cleared after the specified time interval unless the interval is 0.

**Related
Documentation**

- [Root Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [BPDU Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [Configuring BPDU Protection on Individual Interfaces on page 5084](#)

BPDU Protection on All Edge Ports of the Bridge

To configure edge port blocking for a particular STP family member, include the **bpdu-block-on-edge** statement for **mstp**, **rstp**, or **vstp**:

```
bpdu-block-on-edge;
interface interface-name;
```



NOTE: In contrast to BPDU protection configured on individual spanning-tree instance interfaces, BPDU protection configured on all edge ports of an entire spanning-tree protocol *disables designated edge ports* and does not enable them again.

Related Documentation

- [Root Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [BPDU Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [Configuring BPDU Protection on All Edge Ports on page 5085](#)

VPLS Multihoming: Priority of the Backup Bridge

When an MX Series router in a VPLS multihomed Layer 2 ring is running a spanning-tree protocol with root protection enabled, you can modify the default actions taken by the MX Series router when the topology changes. To do this, configure the VPLS root protection topology change actions.

The default value of the backup bridge is **32,768**. You can set the backup bridge priority to a value from **0** through **61440**, in increments of 4096.

To change the default value, you can use the following statement:

```
backup-bridge-priority vpls-ring-backup-bridge-priority
```

You can include the statement at the **[edit protocols (mstp | rstp | vstp)]** hierarchy level (to control global spanning-tree protocol behavior) or at the **[edit protocols vstp vlan vlan-id]** hierarchy level (to control a particular VLAN).



NOTE: VPLS root topology change actions are configured independently of VPLS, the spanning-tree protocol, or the spanning-tree protocol root protect option.

Related Documentation

- [VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Overview on page 5036](#)
- [VPLS Multihoming: Hold Time Before Switching to Primary Priority on page 5050](#)
- [Configuring VPLS Root Protection Topology Change Actions to Control Global Spanning Tree Behavior on page 5080](#)
- [Example: Configuring VPLS Root Topology Change Actions on page 5053](#)

VPLS Multihoming: Hold Time Before Switching to Primary Priority

When an MX Series router or an EX Series switch in a VPLS multihomed Layer 2 ring is running a spanning-tree protocol with root protection enabled, you can modify the default actions taken by the router or switch when the topology changes. To do this, configure the VPLS root protection topology change actions.

The default number of seconds to hold before switching to the primary priority when the first core domain comes up is 2 seconds.

To change the default value, you can use the following statement:

priority-hold-time *seconds*

You can include the statement at the **[edit protocols (mstp | rstp | vstp)]** hierarchy level (to control global spanning-tree protocol behavior) or at the **[edit protocols vstp vlan *vlan-id*]** hierarchy level (to control a particular VLAN).



NOTE: VPLS root topology change actions are configured independently of VPLS, the spanning-tree protocol, or the spanning-tree protocol root protect option.

Related Documentation

- [VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Overview on page 5036](#)
- [VPLS Multihoming: Priority of the Backup Bridge on page 5049](#)
- [Configuring VPLS Root Protection Topology Change Actions to Control Global Spanning Tree Behavior on page 5080](#)
- [Example: Configuring VPLS Root Topology Change Actions on page 5053](#)

VPLS Multihoming: System Identifier for Bridges in the Ring

When an MX Series router or an EX Series switch in a VPLS multihomed Layer 2 ring is running a spanning-tree protocol with root protection enabled, you can modify the default actions taken by the router or switch when the topology changes. To do this, configure the VPLS root protection topology change actions.

The system identifier for bridges in the ring is not configured by default.

To configure a system identifier for bridges in the ring, you can use the following statement:

system-id *system-id-value* *bridge-host-ip-address(es)*

You can include the statement at the **[edit protocols (mstp | rstp | vstp)]** hierarchy level (to control global spanning-tree protocol behavior) or at the **[edit protocols vstp vlan *vlan-id*]** hierarchy level (to control a particular VLAN).



NOTE: VPLS root topology change actions are configured independently of VPLS, the spanning-tree protocol, or the spanning-tree protocol root protect option.

Related Documentation

- [VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Overview on page 5036](#)
- [VPLS Multihoming: Priority of the Backup Bridge on page 5049](#)
- [VPLS Multihoming: Hold Time Before Switching to Primary Priority on page 5050](#)
- [Configuring VPLS Root Protection Topology Change Actions to Control Global Spanning Tree Behavior on page 5080](#)
- [Example: Configuring VPLS Root Topology Change Actions on page 5053](#)

VPLS Multihoming: Bridge Flush of MAC Cache on Topology Change

When an MX Series router or an EX Series switch in a VPLS multihomed Layer 2 ring is running a spanning-tree protocol with root protection enabled, you can modify the default actions taken by the router or switch when the topology changes. To do this, configure the VPLS root protection topology change actions.

By default, if root protect is enabled and then the topology changes, the bridges do not flush the media access control (MAC) address cache of the MAC addresses learned when other interface ports were blocked.

To change the default behavior, you can use the following statement:

vpls-flush-on-topology-change

You can include the statement at the **[edit protocols (mstp | rstp | vstp)]** hierarchy level (to control global spanning-tree protocol behavior) or at the **[edit protocols vstp vlan *vlan-id*]** hierarchy level (to control a particular VLAN).

Specifically, MAC flush messages are sent from the blocked PE to LDP peers based on the mapping of system identifier to IP addresses as specified using the **system-id** statement.



NOTE: VPLS root topology change actions are configured independently of VPLS, the spanning-tree protocol, or the spanning-tree protocol root protect option.

Related Documentation

- [VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Overview on page 5036](#)
- [Configuring VPLS Root Protection Topology Change Actions to Control Global Spanning Tree Behavior on page 5080](#)
- [Configuring VPLS Root Protection Topology Change Actions to Control VLAN Spanning Tree Behavior on page 5085](#)

- [Example: Configuring VPLS Root Topology Change Actions on page 5053](#)

Configuration

- [Configuration Examples on page 5052](#)
- [Configuration Tasks on page 5066](#)
- [Configuration Statements on page 5087](#)

Configuration Examples

- [Example: Blocking BPDUs on Aggregated Ethernet for 600 Seconds on page 5052](#)
- [Example: Enabling Loop Protection for Spanning-Tree Protocols on page 5052](#)
- [Example: Configuring VPLS Root Topology Change Actions on page 5053](#)
- [Example: Configuring VSTP on a Trunk Port with Tagged Traffic on page 5053](#)
- [Example: Tracing Spanning-Tree Protocol Operations on page 5065](#)

Example: Blocking BPDUs on Aggregated Ethernet for 600 Seconds

The following example, when used with a full bridge configuration with aggregated Ethernet, blocks BPDUs on aggregated interface **ae0** for ten minutes (600 seconds) before enabling the interface again:

```
[edit protocols layer2-control]
bpd-block {
  interface ae0;
  disable-timeout 600;
}
```

Related Documentation

- [Root Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [BPDU Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [BPDU Protection for Individual Spanning-Tree Instance Interfaces on page 5048](#)
- [BPDU Protection on All Edge Ports of the Bridge on page 5049](#)
- [Checking the Status of Spanning-Tree Instance Interfaces on page 5125](#)
- [Clearing the Blocked Status of a Spanning-Tree Instance Interface on page 5126](#)

Example: Enabling Loop Protection for Spanning-Tree Protocols

This example blocks and logs the non-designated RSTP port **ge-1/2/0** after the BPDU timeout interval expires:

```
[edit]
protocols {
  rstp {
    interface ge-1/2/0 {
      bpd-timeout-action block;
    }
  }
}
```



NOTE: This is not a complete configuration. You must also fully configure RSTP, including the ge-1/2/0 interface.

Related Documentation

- [Loop Protection for a Spanning-Tree Instance Interface on page 5047](#)
- [Loop Protection for Spanning-Tree Instance Interfaces Overview on page 5034](#)

Example: Configuring VPLS Root Topology Change Actions

This example configures a bridge priority of **36k**, a backup bridge priority of **44k**, a priority hold time value of **60** seconds, a system identifier of **000203:040506** for IP address **10.1.1.1/32**, and sets the bridge to flush the MAC cache on a topology change for MSTP only.

```
[edit]
protocols {
  mstp {
    bridge-priority 36k;
    backup-bridge-priority 44k;
    priority-hold-time 60;
    system-id 000203:040506 {
      10.1.1.1/32;
    }
    vpls-flush-on-topology-change;
  }
}
```



NOTE: This is not a complete configuration.

Related Documentation

- [VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Overview on page 5036](#)
- [VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Topology on page 5037](#)
- [Configuring VPLS Root Protection Topology Change Actions to Control Global Spanning Tree Behavior on page 5080](#)
- [Configuring VPLS Root Protection Topology Change Actions to Control VLAN Spanning Tree Behavior on page 5085](#)

Example: Configuring VSTP on a Trunk Port with Tagged Traffic

- [VSTP on a Trunk Port with Tagged Traffic Overview on page 5054](#)
- [Example: Configuring VSTP on a Trunk Port with Tagged Traffic on page 5054](#)

VSTP on a Trunk Port with Tagged Traffic Overview

In 802.1ad provider bridge networks (stacked VLANs), single-tagged access ports and double-tagged trunk ports can co-exist in a single spanning tree context. In this mode, the VLAN Spanning Tree Protocol (VSTP) can send and receive untagged Rapid Spanning Tree Protocol (RSTP) bridge protocol data units (BPDUs) on gigabit Ethernet (ge), 10 gigabit Ethernet (xe), and aggregated Ethernet (ae) interfaces. The untagged RSTP BPDUs interoperate with tagged VSTP BPDUs sent over the double-tagged trunk ports.

Double-tagging can be useful for Internet service providers, allowing them to use VLANs internally while mixing traffic from clients that are already VLAN-tagged.

Example: Configuring VSTP on a Trunk Port with Tagged Traffic

This example shows how to configure the VSTP to send and receive standard untagged Rapid Spanning Tree Protocol (RSTP) bridge protocol data units (BPDUs) on access trunks that interoperate with tagged VSTP BPDUs sent over the double-tagged trunk ports.

- [Requirements on page 5054](#)
- [Overview on page 5054](#)
- [Configuration on page 5055](#)
- [Verification on page 5063](#)

Requirements

This example uses the following hardware and software components:

- Two CE devices (MX Series routers with DPCE or MPC cards)
- Two PE devices (MX Series routers with DPCE or MPC cards)
- Junos OS Release 12.3 or later running on the PE devices

Overview


This example shows how to configure VSTP on a trunk port with tagged traffic.

Topology

[Figure 111 on page 5055](#) shows a sample topology in which two customer edge (CE) bridges are dual-homed to two provider edge (PE) devices. All of the PE-CE links are single-tagged trunks using C-VLANs 1-100. The core link between Devices PE1 and PE2 is a double-tagged trunk that carries traffic from both CE devices, using S-VLANs 100 and 200 to distinguish the CE traffic.

Two VSTP instances are created on the PE devices, one for each S-VLAN. The CE devices run the standard RSTP. The PE devices run VSTP on the core link while sending standard untagged RSTP BPDUs toward the CE devices.

Figure 111: Topology for VSTP Configured on a Trunk Port with Tagged Traffic



Configuration

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

Device PE1

```

set interfaces ge-2/0/1 flexible-vlan-tagging
set interfaces ge-2/0/1 encapsulation flexible-ethernet-services
set interfaces ge-2/0/1 unit 1 vlan-id 100
set interfaces ge-2/0/1 unit 1 family bridge interface-mode trunk
set interfaces ge-2/0/1 unit 1 family bridge inner-vlan-id-list 1-100
set interfaces ge-2/0/1 unit 2 vlan-id 200
set interfaces ge-2/0/1 unit 2 family bridge interface-mode trunk
set interfaces ge-2/0/1 unit 2 family bridge inner-vlan-id-list 1-100
set interfaces ge-2/0/4 encapsulation ethernet-vpls
set interfaces ge-2/0/4 unit 0 description to_CE1
set interfaces ge-2/0/4 unit 0 family bridge interface-mode trunk
set interfaces ge-2/0/4 unit 0 family bridge vlan-id-list 1-100
set interfaces ge-2/1/1 unit 0 description to_CE2
set interfaces ge-2/1/1 unit 0 family bridge interface-mode trunk
set interfaces ge-2/1/1 unit 0 family bridge vlan-id-list 1-100
set routing-instances vs1 instance-type virtual-switch
set routing-instances vs1 interface ge-2/0/1
set routing-instances vs1 interface ge-2/0/4.0
set routing-instances vs1 protocols vstp vlan 100 interface ge-2/0/1
set routing-instances vs1 protocols vstp vlan 100 interface ge-2/0/4 access-trunk
set routing-instances vs1 bridge-domains bd vlan-id-list 1-100
set routing-instances vs2 instance-type virtual-switch
set routing-instances vs2 interface ge-2/0/1.2
set routing-instances vs2 interface ge-2/1/1.0
set routing-instances vs2 protocols vstp vlan 200 interface ge-2/0/1
set routing-instances vs2 protocols vstp vlan 200 interface ge-2/1/1 access-trunk
set routing-instances vs2 bridge-domains bd vlan-id-list 1-100

```

Device PE2

```

set interfaces ge-2/0/1 flexible-vlan-tagging
set interfaces ge-2/0/1 encapsulation flexible-ethernet-services
set interfaces ge-2/0/1 unit 1 vlan-id 100
set interfaces ge-2/0/1 unit 1 family bridge interface-mode trunk

```

```
set interfaces ge-2/0/1 unit 1 family bridge inner-vlan-id-list 1-100
set interfaces ge-2/0/1 unit 2 vlan-id 200
set interfaces ge-2/0/1 unit 2 family bridge interface-mode trunk
set interfaces ge-2/0/1 unit 2 family bridge inner-vlan-id-list 1-100
set interfaces ge-2/1/3 description to_CE1
set interfaces ge-2/1/3 unit 0 family bridge interface-mode trunk
set interfaces ge-2/1/3 unit 0 family bridge vlan-id-list 1-100
set interfaces ge-2/0/5 description to_CE2
set interfaces ge-2/0/5 unit 0 family bridge interface-mode trunk
set interfaces ge-2/0/5 unit 0 family bridge vlan-id-list 1-100
set routing-instances vs1 instance-type virtual-switch
set routing-instances vs1 interface ge-2/0/1.1
set routing-instances vs1 interface ge-2/1/3.0
set routing-instances vs1 protocols vstp vlan 100 interface ge-2/0/1
set routing-instances vs1 protocols vstp vlan 100 interface ge-2/1/3 access-trunk
set routing-instances vs1 bridge-domains bd vlan-id-list 1-100
set routing-instances vs2 instance-type virtual-switch
set routing-instances vs2 interface ge-2/0/1.2
set routing-instances vs2 interface ge-2/0/5.0
set routing-instances vs2 protocols vstp vlan 200 interface ge-2/0/1
set routing-instances vs2 protocols vstp vlan 200 interface ge-2/0/5 access-trunk
set routing-instances vs2 bridge-domains bd vlan-id-list 1-100
```

Device CE1

```
set interfaces ge-2/0/7 unit 0 description to_PE1
set interfaces ge-2/0/7 unit 0 family bridge interface-mode trunk
set interfaces ge-2/0/7 unit 0 family bridge vlan-id-list 1-100
set interfaces ge-2/1/4 unit 0 description to_PE2
set interfaces ge-2/1/4 unit 0 family bridge interface-mode trunk
set interfaces ge-2/1/4 unit 0 family bridge vlan-id-list 1-100
set protocols rstp interface ge-2/0/7
set protocols rstp interface ge-2/1/4
set bridge-domains bd vlan-id-list 1-100
```

Device CE2

```
set interfaces ge-2/0/8 unit 0 description to_PE2
set interfaces ge-2/0/8 unit 0 family bridge interface-mode trunk
set interfaces ge-2/0/8 unit 0 family bridge vlan-id-list 1-100
set interfaces ge-2/1/1 unit 0 description to_PE1
set interfaces ge-2/1/1 unit 0 family bridge interface-mode trunk
set interfaces ge-2/1/1 unit 0 family bridge vlan-id-list 1-100
set protocols rstp interface ge-2/0/8
set protocols rstp interface ge-2/1/1
set bridge-domains bd vlan-id-list 1-100
```

Configuring PE1, PE2, CE1, and CE2

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure Device PE1:

1. Configure the network interfaces.

```
[edit interfaces]
user@PE1# set ge-2/0/1 flexible-vlan-tagging
user@PE1# set ge-2/0/1 encapsulation flexible-ethernet-services
```

```

user@PE1# set ge-2/0/1 unit 1 vlan-id 100
user@PE1# set ge-2/0/1 unit 1 family bridge interface-mode trunk
user@PE1# set ge-2/0/1 unit 1 family bridge inner-vlan-id-list 1-100
user@PE1# set ge-2/0/1 unit 2 vlan-id 200
user@PE1# set ge-2/0/1 unit 2 family bridge interface-mode trunk
user@PE1# set ge-2/0/1 unit 2 family bridge inner-vlan-id-list 1-100

```

```

user@PE1# set ge-2/0/4 encapsulation ethernet-vpls
user@PE1# set ge-2/0/4 unit 0 description to_CE1
user@PE1# set ge-2/0/4 unit 0 family bridge interface-mode trunk
user@PE1# set ge-2/0/4 unit 0 family bridge vlan-id-list 1-100

```

```

user@PE1# set ge-2/1/1 unit 0 description to_CE2
user@PE1# set ge-2/1/1 unit 0 family bridge interface-mode trunk
user@PE1# set ge-2/1/1 unit 0 family bridge vlan-id-list 1-100

```

2. Configure the routing instances.

```

[edit routing-instances]
user@PE1# set vs1 instance-type virtual-switch
user@PE1# set vs1 interface ge-2/0/1.1
user@PE1# set vs1 interface ge-2/0/4.0
user@PE1# set vs1 protocols vstp vlan 100 interface ge-2/0/1
user@PE1# set vs1 protocols vstp vlan 100 interface ge-2/0/4 access-trunk
user@PE1# set vs1 bridge-domains bd vlan-id-list 1-100

```

```

user@PE1# set vs2 instance-type virtual-switch
user@PE1# set vs2 interface ge-2/0/1.2
user@PE1# set vs2 interface ge-2/1/1.0
user@PE1# set vs2 protocols vstp vlan 200 interface ge-2/0/1
user@PE1# set vs2 protocols vstp vlan 200 interface ge-2/1/1 access-trunk
user@PE1# set vs2 bridge-domains bd vlan-id-list 1-100

```

Step-by-Step Procedure

To configure Device PE2:

1. Configure the interfaces.

```

[edit interfaces]
user@PE2# set ge-2/0/1 flexible-vlan-tagging
user@PE2# set ge-2/0/1 encapsulation flexible-ethernet-services
user@PE2# set ge-2/0/1 unit 1 vlan-id 100
user@PE2# set ge-2/0/1 unit 1 family bridge interface-mode trunk
user@PE2# set ge-2/0/1 unit 1 family bridge inner-vlan-id-list 1-100
user@PE2# set ge-2/0/1 unit 2 vlan-id 200
user@PE2# set ge-2/0/1 unit 2 family bridge interface-mode trunk
user@PE2# set ge-2/0/1 unit 2 family bridge inner-vlan-id-list 1-100

```

```

user@PE2# set ge-2/1/3 description to_CE1
user@PE2# set ge-2/1/3 unit 0 family bridge interface-mode trunk
user@PE2# set ge-2/1/3 unit 0 family bridge vlan-id-list 1-100

```

```

user@PE2# set ge-2/0/5 description to_CE2
user@PE2# set ge-2/0/5 unit 0 family bridge interface-mode trunk
user@PE2# set ge-2/0/5 unit 0 family bridge vlan-id-list 1-100

```

2. Configure the routing instances.

```
[edit routing-instances]
user@PE2# set vs1 instance-type virtual-switch
user@PE2# set vs1 interface ge-2/0/1.1
user@PE2# set vs1 interface ge-2/1/3.0
user@PE2# set vs1 protocols vstp vlan 100 interface ge-2/0/1
user@PE2# set vs1 protocols vstp vlan 100 interface ge-2/1/3 access-trunk
user@PE2# set vs1 bridge-domains bd vlan-id-list 1-100

user@PE2# set vs2 instance-type virtual-switch
user@PE2# set vs2 interface ge-2/0/1.2
user@PE2# set vs2 interface ge-2/0/5.0
user@PE2# set vs2 protocols vstp vlan 200 interface ge-2/0/1
user@PE2# set vs2 protocols vstp vlan 200 interface ge-2/0/5 access-trunk
user@PE2# set vs2 bridge-domains bd vlan-id-list 1-100
```

Step-by-Step Procedure To configure CE1:

1. Configure the interfaces.

```
[edit interfaces]
user@CE1# set ge-2/0/7 unit 0 description to_PE1
user@CE1# set ge-2/0/7 unit 0 family bridge interface-mode trunk
user@CE1# set ge-2/0/7 unit 0 family bridge vlan-id-list 1-100

user@CE1# set ge-2/1/4 unit 0 description to_PE2
user@CE1# set ge-2/1/4 unit 0 family bridge interface-mode trunk
user@CE1# set ge-2/1/4 unit 0 family bridge vlan-id-list 1-100
```

2. Configure the protocols.

```
[edit protocols]
user@CE1# set rstp interface ge-2/0/7
user@CE1# set rstp interface ge-2/1/4
```

3. Configure the bridge domain.

```
[edit bridge-domains]
user@CE1# set bd vlan-id-list 1-100
```

Step-by-Step Procedure To configure CE2:

1. Configure the interfaces.

```
[edit interfaces]
user@CE2# set ge-2/0/8 unit 0 description to_PE2
user@CE2# set ge-2/0/8 unit 0 family bridge interface-mode trunk
user@CE2# set ge-2/0/8 unit 0 family bridge vlan-id-list 1-100

user@CE2# set ge-2/1/1 unit 0 description to_PE1
user@CE2# set ge-2/1/1 unit 0 family bridge interface-mode trunk
user@CE2# set ge-2/1/1 unit 0 family bridge vlan-id-list 1-100
```

2. Configure the protocols.

```
[edit protocols]
```



```
user@CE2# set rstp interface ge-2/0/8
user@CE2# set rstp interface ge-2/1/1
```

3. Configure the bridge domain.

```
[edit bridge-domains]
user@CE2# set bd vlan-id-list 1-100
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-instances**, **show protocols**, and **show bridge-domains** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
Device PE1 user@PE1# show interfaces
ge-2/0/1 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 1 {
    vlan-id 100;
    family bridge {
      interface-mode trunk;
      inner-vlan-id-list 1-100;
    }
  }
  unit 2 {
    vlan-id 200;
    family bridge {
      interface-mode trunk;
      inner-vlan-id-list 1-100;
    }
  }
}
ge-2/0/4 {
  encapsulation ethernet-vpls;
  unit 0 {
    description to_CE1;
    family bridge {
      interface-mode trunk;
      vlan-id-list 1-100;
    }
  }
}
ge-2/1/1 {
  unit 0 {
    description to_CE2;
    family bridge {
      interface-mode trunk;
      vlan-id-list 1-100;
    }
  }
}

user@PE1# show routing-instances
vs1 {
```

```
instance-type virtual-switch;
interface ge-2/0/1.1;
interface ge-2/0/4.0;
protocols {
  vstp {
    vlan 100 {
      interface ge-2/0/1;
      interface ge-2/0/4 {
        access-trunk;
      }
    }
  }
}
bridge-domains {
  bd {
    vlan-id-list 1-100;
  }
}
}
vs2 {
  instance-type virtual-switch;
  interface ge-2/0/1.2;
  interface ge-2/0/1.0;
  protocols {
    vstp {
      vlan 200 {
        interface ge-2/0/1;
        interface ge-2/1/1 {
          access-trunk;
        }
      }
    }
  }
}
bridge-domains {
  bd {
    vlan-id-list 1-100;
  }
}
}
```

```
Device PE2 user@PE2# show interfaces
ge-2/0/1 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 1 {
    vlan-id 100;
    family bridge {
      interface-mode trunk;
      inner-vlan-id-list 1-100;
    }
  }
  unit 2 {
    vlan-id 200;
    family bridge {
      interface-mode trunk;
      inner-vlan-id-list 1-100;
    }
  }
}
```

```

    }
  }
}
ge-2/0/5 {
  description to_CE2;
  unit 0 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 1-100;
    }
  }
}
ge-2/1/3 {
  description to_CE1;
  unit 0 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 1-100;
    }
  }
}
}

user@PE2# show routing-instances
vs1 {
  instance-type virtual-switch;
  interface ge-2/0/1.1;
  interface ge-2/1/3.0;
  protocols {
    vstp {
      vlan 100 {
        interface ge-2/0/1;
        interface ge-2/1/3 {
          access-trunk;
        }
      }
    }
  }
  bridge-domains {
    bd {
      vlan-id-list 1-100;
    }
  }
}
vs2 {
  instance-type virtual-switch;
  interface ge-2/0/1.2;
  interface ge-2/0/5.0;
  protocols {
    vstp {
      vlan 200 {
        interface ge-2/0/1;
        interface ge-2/0/5 {
          access-trunk;
        }
      }
    }
  }
}

```

```
    }
    bridge-domains {
      bd {
        vlan-id-list 1-100;
      }
    }
  }
}

Device CE1 user@CE1# show interfaces
ge-2/0/7 {
  unit 0 {
    description to_PE1;
    family bridge {
      interface-mode trunk;
      vlan-id-list 1-100;
    }
  }
}
ge-2/1/4 {
  unit 0 {
    description to_PE2;
    family bridge {
      interface-mode trunk;
      vlan-id-list 1-100;
    }
  }
}

user@CE1# show protocols
rstp {
  interface ge-2/0/7;
  interface ge-2/1/4;
}

user@CE1# show bridge-domains
bd {
  vlan-id-list 1-100;
}

Device CE2 user@CE2 show interfaces
ge-2/0/8 {
  unit 0 {
    description to_PE2;
    family bridge {
      interface-mode trunk;
      vlan-id-list 1-100;
    }
  }
}
ge-2/1/1 {
  unit 0 {
    description to_PE1;
    family bridge {
      interface-mode trunk;
      vlan-id-list 1-100;
    }
  }
}
```

```

}

user@CE2# show protocols
rstp {
    interface ge-2/0/8;
    interface ge-2/1/1;
}

user@CE2# show bridge-domains
bd {
    vlan-id-list 1-100;
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying That the Interfaces Are Operational on page 5063](#)
- [Verifying the STP Bridge Parameters of the Routing Instances on page 5063](#)
- [Displaying STP Statistics for the Configured Bridge on page 5064](#)

Verifying That the Interfaces Are Operational

Purpose Verify that the interfaces are operational.

Action From operational mode, enter the **show spanning-tree interface routing-instance** command.

```

user@PE1> show spanning-tree interface routing-instance vs1
Spanning tree interface parameters for VLAN 100

```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ge-2/0/1	128:82	128:82	32868.0021590f37d0	20000	FWD	DESG
ge-2/0/4	128:85	128:85	32868.0021590f37d0	20000	FWD	DESG

Meaning The output shows the status of the interfaces configured for VLAN 100.

Verifying the STP Bridge Parameters of the Routing Instances

Purpose Verify the STP bridge parameters configured for the routing instances.

Action From operational mode, enter the **show spanning-tree bridge routing-instance** command.

```
user@PE1> show spanning-tree bridge routing-instance vs1
STP bridge parameters
Routing instance name      : vs1
Enabled protocol           : RSTP

STP bridge parameters for VLAN 100
Root ID                    : 32868.00:21:59:0f:37:d0
Hello time                  : 2 seconds
Maximum age                 : 20 seconds
Forward delay               : 15 seconds
Message age                 : 0
Number of topology changes  : 2
Time since last topology change : 687 seconds
Local parameters
  Bridge ID                 : 32868.00:21:59:0f:37:d0
  Extended system ID        : 100
```

Meaning The output shows the status of the STP bridge parameters for routing instance vs1.

Displaying STP Statistics for the Configured Bridge

Purpose Display spanning-tree statistics for the configured bridge.

Action From operational mode, enter the **show spanning-tree statistics bridge** command.

```

user@PE1> show spanning-tree statistics bridge
STP Context : default
STP Instance : 0
Number of Root Bridge Changes: 0
Number of Root Port Changes: 0

STP Context : x/default
STP Instance : 0
Number of Root Bridge Changes: 0
Number of Root Port Changes: 0

STP Context : vs1
STP Instance : 0
Number of Root Bridge Changes: 2          Last Changed: Thu Sep 20 15:12:18
2012
Number of Root Port Changes: 1          Last Changed: Thu Sep 20 15:01:13
2012
Recent TC Received: ge-2/0/1.1          Received : Thu Sep 20 15:01:17
2012

STP Context : vs2
STP Instance : 0
Number of Root Bridge Changes: 2          Last Changed: Thu Sep 20 15:10:25
2012
Number of Root Port Changes: 2          Last Changed: Thu Sep 20 15:10:25
2012
Recent TC Received: ge-2/1/1.0          Received : Thu Sep 20 15:10:47
2012

STP Context : CE1/default
STP Instance : 0
Number of Root Bridge Changes: 0
Number of Root Port Changes: 0
Recent TC Received: ge-2/1/4.0          Received : Thu Sep 20 15:12:15
2012

```

Meaning The command output shows spanning-tree statistics for the configured bridge.

Example: Tracing Spanning-Tree Protocol Operations

Trace only unusual or abnormal operations to `/var/log/stp-log`:

```

[edit]
routing-options {
  traceoptions {
    file routing-log size 10m world-readable;
    flag all;
  }
}
protocols {
  rstp {
    traceoptions {
      file rstp-log size 10m world-readable;
      flag all;
    }
  }
}

```

}

**Related
Documentation**

- [Spanning-Tree Protocols Supported on page 5032](#)
- [Tracing Spanning-Tree Operations on page 5128](#)

Configuration Tasks

- [Configuring Multiple Spanning-Tree Protocol on page 5067](#)
- [Configuring MST Instances on a Physical Interface on page 5070](#)
- [Disabling MSTP on page 5071](#)
- [Provider Bridge Participation in RSTP or MSTP Instances on page 5072](#)
- [RSTP or VSTP Forced to Run as IEEE 802.1D STP on page 5072](#)
- [Configuring Rapid Spanning-Tree Protocol on page 5073](#)
- [Reverting RSTP or VSTP Back From Forced IEEE 802.1D STP on page 5075](#)
- [Provider Bridge Participation in RSTP or MSTP Instances on page 5076](#)
- [System Identifier for Bridges in STP or RSTP Instances on page 5076](#)
- [Configuring VLAN Spanning-Tree Protocol on page 5077](#)
- [Configuring VPLS Root Protection Topology Change Actions to Control Global Spanning Tree Behavior on page 5080](#)
- [BPDU Protection for Individual Spanning-Tree Instance Interfaces on page 5081](#)
- [Configuring Loop Protection for a Spanning-Tree Instance Interface on page 5082](#)
- [Enabling Root Protect for a Spanning-Tree Instance Interface on page 5083](#)
- [Configuring BPDU Protection on Individual Interfaces on page 5084](#)
- [Configuring BPDU Protection on All Edge Ports on page 5085](#)
- [Configuring VPLS Root Protection Topology Change Actions to Control VLAN Spanning Tree Behavior on page 5085](#)
- [Configuring Layer 2 Protocol Tunneling on page 5086](#)

Configuring Multiple Spanning-Tree Protocol

You can configure the Multiple Spanning-Tree Protocol (MSTP) under the following hierarchy levels:

- [edit *logical-systems logical-system-name protocols*]
- [edit *logical-systems logical-system-name routing-instances routing-instance-name protocols*]
- [edit *protocols*]
- [edit *routing-instances routing-instance-name protocols*]

The routing instance type can be either **virtual-switch** or **layer2-control**.

To configure the Multiple Spanning-Tree Protocol:

1. Enable RSTP as the version of spanning-tree protocol to be configured:

```
[edit]
user@host@ edit ... protocols (STP Type) mstp
```

2. (Optional) Enable provider bridge participation in the MSTP instance:

```
[edit ... protocols mstp]
user@host# set bpd-destination-mac-address provider-bridge-group
```

3. Configure the interfaces that participate in the MSTP instance.

- a. Enable configuration of the interface:

```
[edit ... protocols mstp]
user@host# edit interface interface-name
```

- b. Configure the interface priority:

```
[edit ... protocols mstp interface interface-name]
user@host# set priority interface-priority
```

- c. (Optional) By default, the interface link cost is determined by the link speed. You can configure the interface link cost to control which bridge is the designated bridge and which port is the designated port:

```
[edit ... protocols mstp interface interface-name]
user@host# set cost interface-link-cost
```

- d. Configure the interface link mode to identify point-to-point links:

```
[edit ... protocols mstp interface interface-name]
user@host# set mode (p2p | shared)
```

Specify **p2p** if the link is point to point. Specify **shared** if the link is a shared media.

- e. (Optional) Configure the interface as an edge port:

```
[edit ... protocols mstp interface interface-name]
user@host# set edge
```

Edge ports do not expect to receive bridge protocol data unit (BPDU) packets. If a BPDU packet is received for an edge port, the port becomes a nonedge port.

You can also enable BPDU root protection for all spanning-tree protocol instances on the interface. BPDU root protect ensures the port is the spanning-tree designated port. If the port receives superior BPDU packets, root protect moves this port to a root-prevented spanning-tree state. For configuration details, see [“Checking the Status of Spanning-Tree Instance Interfaces” on page 5125](#).

4. Configure the bridge priority

```
[edit ... protocols mstp]
user@host# set bridge-priority bridge-priority
```

For more information, see [“Bridge Priority for Election of Root Bridge and Designated Bridge” on page 5041](#).

5. Configure hello BPDU timers.

- a. Configure the maximum expected arrival time of hello BPDUs:

```
[edit ... protocols mstp]
user@host# set max-age seconds
```

- b. Configure the time interval at which the root bridge transmits configuration BPDUs:

```
[edit ... protocols mstp]
user@host# set hello-time seconds
```

6. (Optional) By default, the bridge port remains in the listening and learning states for 15 seconds before transitioning to the forwarding state. You can specify a delay from 4 through 20 seconds instead:

```
[edit ... protocols mstp]
```

```
user@host# set forward-delay seconds
```

7. Configure MSTP-specific options.

a. Configure the MSTP region configuration name:

```
[edit ... protocols mstp]
user@host# set configuration-name configuration-name
```

b. Configure the MSTP revision level:

```
[edit ... protocols mstp]
user@host# set revision-level revision-level
```

c. Configure the maximum number of hops a BPDU can be forwarded in the MSTP region:

```
[edit ... protocols mstp]
user@host# set max-hops hops
```

8. Verify the MSTP configuration:

```
[edit]
... { # Optional logical system and/or routing instance
  protocols (STP Type) {
    mstp {
      bpd-destination-mac-address provider-bridge-group; # Optional
      interface interface-name {
        priority interface-priority;
        cost interface-link-cost; # Optional.
        mode (p2p | shared);
        edge; # Optional.
      }
      bridge-priority bridge-priority;
      max-age seconds;
      hello-time seconds;
      forward-delay seconds; # Optional.
      configuration-name configuration-name; # MST region configuration name.
      revision-level revision-level; # MST revision number.
      max-hops hops; # MST maximum hops.
    }
  }
}
```

**Related
Documentation**

- [Spanning-Tree Protocols Supported on page 5032](#)
- [Configuring MST Instances on a Physical Interface on page 5070](#)
- [Disabling MSTP on page 5071](#)

Configuring MST Instances on a Physical Interface

You can configure a Multiple Spanning-Tree Instance (MSTI) under the following hierarchy levels:

- [edit *logical-systems logical-system-name protocols mstp*]
- [edit *logical-systems logical-system-name routing-instances routing-instance-name protocols mstp*]
- [edit *protocols mstp*]
- [edit *routing-instances routing-instance-name protocols mstp*]

The routing instance type can be either **virtual-switch** or **layer2-control**.

Before you begin, configure Multiple Spanning-Tree Protocol. For configuration details, see “Configuring MSTP” on page 5067.

1. Enable configuration of a MST instance:

```
[edit]
user@host# edit ... protocols mstp msti msti-id
The msti-id value must be from 1 through 64.
```

2. Configure the interfaces that participate in the MST instance.

- a. Enable configuration of the interface:

```
[edit ... protocols mstp msti msti-id]
user@host# edit interface interface-name
```

- b. Configure the interface priority:

```
[edit ... protocols mstp msti msti-id interface interface-name]
user@host# set priority interface-priority
```

- c. (Optional) By default, the interface link cost is determined by the link speed. You can configure the interface link cost to control which bridge is the designated bridge and which port is the designated port:

```
[edit ... protocols mstp msti msti-id interface interface-name]
user@host# set cost interface-link-cost
```

- d. (Optional) Configure the interface as an edge port:

```
[edit ... protocols mstp msti msti-id interface interface-name]
user@host# set edge
```

Edge ports do not expect to receive bridge protocol data unit (BPDU) packets. If a BPDU packet is received for an edge port, the port becomes a nonedge port

3. Configure the bridge priority

```
[edit ... protocols mstp msti msti-id]
user@host# set bridge-priority bridge-priority
```

For more information, see “Bridge Priority for Election of Root Bridge and Designated Bridge” on page 5041.

4. (Optional) An MSTI can map to a range of VLANs just as a logical port can map to a range of VLANs. The MSTP VLAN specifies the VLAN or VLAN range to which this

MSTI is mapped. The vlan-id is configured under the logical interface. Configure the VLAN or VLAN range of the MSTI instance:

```
[edit]
user@host# set vlan (vlan-id | vlan-id-range)
```

5. Verify the MST interface configuration.

```
[edit]
protocols {
  mstp {
    ...basic-mstp-configuration...
    msti msti-id { # Instance identifier 1 – 64.
      bridge-priority priority;
      vlan vlan-id; # Optional
      interface interface-name {
        cost cost;
        edge;
        priority interface-priority;
      }
    }
  }
}
```

Related Documentation

- [Spanning-Tree Protocols Supported on page 5032](#)
- [Configuring Multiple Spanning-Tree Protocol on page 5067](#)
- [Disabling MSTP on page 5071](#)

Disabling MSTP

To disable the entire MSTP instance:

- Include the **disable** statement. You can include this statement at the following hierarchy levels:
 - **[edit logical-systems logical-system-name protocols mstp]**
 - **[edit logical-systems logical-system-name routing-instances routing-instance-name protocols mstp]**
 - **[edit protocols mstp]**
 - **[edit routing-instances routing-instance-name protocols mstp]**

Related Documentation

- [Spanning-Tree Protocols Supported on page 5032](#)
- [Configuring Multiple Spanning-Tree Protocol on page 5067](#)
- [Configuring MST Instances on a Physical Interface on page 5070](#)

Provider Bridge Participation in RSTP or MSTP Instances

A provider network can bridge the customer STP BPDU packets between customer sites by default. At the same time, the provider network can prevent forwarding loops by running a spanning-tree protocol in the provider network. On an MX Series router running Rapid Spanning-Tree Protocol (RSTP) or Multiple Spanning-Tree Protocol (MSTP) in a provider network, you can enable provider bridge participation in the RSTP or MSTP instance.

The IEEE 802.1ad specification reserves the group MAC address value of **01:80:c2:00:00:08** to designate the *provider bridge group*. On an MX Series router for which you have enabled provider bridge participation in the RSTP or MSTP instance, the router exchanges BPDU packets with the provider bridge group as follows:

- Transmitted BPDU packets contain the destination MAC address **01:80:c2:00:00:08**.
- Received BPDU packets with the destination MAC address **01:80:c2:00:00:08** are accepted and passed to the Routing Engine.

Related Documentation

- [Spanning-Tree Protocols Supported on page 5032](#)
- [Configuring Rapid Spanning-Tree Protocol on page 5073](#)
- [Configuring Multiple Spanning-Tree Protocol on page 5067](#)
- [bpdu-destination-mac-address on page 5092](#)

RSTP or VSTP Forced to Run as IEEE 802.1D STP

On MX Series routers or EX Series switches in a Layer 2 environment, you can force the configured Rapid Spanning-Tree Protocol (RSTP) or VLAN Spanning-Tree Protocol (VSTP) to run as the original IEEE 802.1D Spanning-Tree Protocol (STP) version. Configure this option for compatibility with older bridges that do not support RSTP or VSTP.

Keep the following limitations in mind when RSTP or VSTP are run as the original STP version:

- If you configure an instance interface as an edge port, the configuration statement is ignored.
- If you configure point-to-point link mode for an instance interface, the configuration statement is ignored.

Related Documentation

- [Spanning-Tree Protocols Supported on page 5032](#)
- [Configuring Rapid Spanning-Tree Protocol on page 5073](#)
- [Configuring VLAN Spanning-Tree Protocol on page 5077](#)
- [Reverting RSTP or VSTP Back From Forced IEEE 802.1D STP on page 5075](#)
- [force-version on page 5097](#)

Configuring Rapid Spanning-Tree Protocol

You can configure Rapid Spanning-Tree Protocol (RSTP) under the following hierarchy levels:

- [edit *logical-systems logical-system-name protocols*]
- [edit *logical-systems logical-system-name routing-instances routing-instance-name protocols*]
- [edit *protocols*]
- [edit *routing-instances routing-instance-name protocols*]

The routing instance type can be either **virtual-switch** or **layer2-control**.

To configure the Rapid Spanning-Tree Protocol:

1. Enable RSTP as the version of spanning-tree protocol to be configured:

```
[edit]
user@host@ edit ... protocols (STP Type) rstp
```

2. (Optional) For compatibility with older bridges that do not support RSTP, you can run force RSTP to run as the original IEEE 802.1D Spanning-Tree Protocol (STP) version:

```
[edit ... protocols rstp]
user@host# set force-version stp
```



NOTE: If RSTP has been forced to run as the original STP version, you can revert back to RSTP by first removing the *force-version* statement from the configuration and then entering the *clear spanning-tree protocol-migration* configuration mode command.

3. (Optional) Enable provider bridge participation in the RSTP instance:

```
[edit ... protocols rstp]
user@host# set bpd-destination-mac-address provider-bridge-group
```

4. (Optional) Specify the extended system identifier used in identifiers bridges that participate in RSTP:

```
[edit ... protocols rstp]
user@host# set extended-system-id identifier
```

5. Configure the interfaces that participate in the RSTP instance.

- a. Enable configuration of the interface:

```
[edit ... protocols rstp]
user@host# edit interface interface-name
```

- b. Configure the interface priority:

```
[edit ... protocols rstp interface interface-name]
user@host# set priority interface-priority
```

- c. (Optional) By default, the interface link cost is determined by the link speed. You can configure the interface link cost to control which bridge is the designated bridge and which port is the designated port:

```
[edit ... protocols rstp interface interface-name]
user@host# set cost interface-link-cost
```

- d. Configure the interface link mode to identify point-to-point links:

```
[edit ... protocols rstp interface interface-name]
user@host# set mode (p2p | shared)
```

Specify **p2p** if the link is point to point. Specify **shared** if the link is a shared media.

- e. (Optional) Configure the interface as an edge port:

```
[edit ... protocols rstp interface interface-name]
user@host# set edge
```

Edge ports do not expect to receive bridge protocol data unit (BPDU) packets. If a BPDU packet is received for an edge port, the port becomes a nonedge port.

You can also enable BPDU root protection for all spanning-tree protocol instances on the interface. BPDU root protect ensures the port is the spanning-tree designated port. If the port receives superior BPDU packets, root protect moves this port to a root-prevented spanning-tree state. For configuration details, see [“Checking the Status of Spanning-Tree Instance Interfaces” on page 5125](#).

6. Configure the bridge priority

```
[edit ... protocols rstp]
user@host# set bridge-priority bridge-priority
```

For more information, see [“Bridge Priority for Election of Root Bridge and Designated Bridge” on page 5041](#).

7. Configure hello BPDU timers.

- a. Configure the maximum expected arrival time of hello BPDUs:

```
[edit ... protocols rstp]
user@host# set max-age seconds
```

- b. Configure the time interval at which the root bridge transmits configuration BPDUs:

```
[edit ... protocols rstp]
user@host# set hello-time seconds
```

8. (Optional) By default, the bridge port remains in the listening and learning states for 15 seconds before transitioning to the forwarding state. You can specify a delay from 4 through 20 seconds instead:

```
[edit ... protocols rstp]
```



```
user@host# set forward-delay seconds
```

9. Verify the RSTP configuration:

```
[edit]
... { # Optional logical system and/or routing instance
  protocols (STP Type) {
    rstp {
      force-version stp; # Optional.
      bpdudestination-mac-address provider-bridge-group; # Optional
      extended-system-id identifier; # Optional.
      interface interface-name {
        priority interface-priority;
        cost interface-link-cost; # Optional.
        mode (p2p | shared);
        edge; # Optional.
      }
      bridge-priority bridge-priority;
      max-age seconds;
      hello-time seconds;
      forward-delay seconds; # Optional.
    }
  }
}
```

Related Documentation

- [Spanning-Tree Protocols Supported on page 5032](#)
- [RSTP or VSTP Forced to Run as IEEE 802.1D STP on page 5072](#)
- [Reverting RSTP or VSTP Back From Forced IEEE 802.1D STP on page 5075](#)
- [Provider Bridge Participation in RSTP or MSTP Instances on page 5072](#)
- [System Identifier for Bridges in STP or RSTP Instances on page 5076](#)

Reverting RSTP or VSTP Back From Forced IEEE 802.1D STP

On MX Series routers or EX Series switches on which Rapid Spanning-Tree Protocol (RSTP) or VLAN Spanning-Tree Protocol (VSTP) has been forced to run as the original IEEE 802.1D Spanning-Tree Protocol (STP) version, you can revert back to RSTP or VSTP.

To revert from the forced instance of the original IEEE 802.1D STP version back to the configured RSTP or VSTP version:

1. Remove the **force-version** statement from the RSTP or VSTP configuration:

```
user@host# delete force-version
```

Include this statement under the RSTP or VSTP hierarchy level:

- `[edit protocols rstp]`
- `[edit protocols vstp]`
- `[edit routing-instances routing-instance-name protocols rstp]`
- `[edit routing-instances routing-instance-name protocols vstp]`

2. Revert the forced IEEE 802.1D STP to run as the configured RSTP or VSTP:

```
user@host# clear spanning-tree protocol-migration <interface interface-name>
<routing-instance routing-instance-name>
```

To revert the STP protocol for the specified interface only, specify the **interface *interface-name*** option.

To revert the STP protocol for a particular routing instance only, specify the **routing-instance *routing-instance-name*** option.

**Related
Documentation**

- [Spanning-Tree Protocols Supported on page 5032](#)
- [RSTP or VSTP Forced to Run as IEEE 802.1D STP on page 5072](#)
- [Configuring Rapid Spanning-Tree Protocol on page 5073](#)
- [Configuring VLAN Spanning-Tree Protocol on page 5077](#)

Provider Bridge Participation in RSTP or MSTP Instances

A provider network can bridge the customer STP BPDU packets between customer sites by default. At the same time, the provider network can prevent forwarding loops by running a spanning-tree protocol in the provider network. On an MX Series router running Rapid Spanning-Tree Protocol (RSTP) or Multiple Spanning-Tree Protocol (MSTP) in a provider network, you can enable provider bridge participation in the RSTP or MSTP instance.

The IEEE 802.1ad specification reserves the group MAC address value of **01:80:c2:00:00:08** to designate the *provider bridge group*. On an MX Series router for which you have enabled provider bridge participation in the RSTP or MSTP instance, the router exchanges BPDU packets with the provider bridge group as follows:

- Transmitted BPDU packets contain the destination MAC address **01:80:c2:00:00:08**.
- Received BPDU packets with the destination MAC address **01:80:c2:00:00:08** are accepted and passed to the Routing Engine.

**Related
Documentation**

- [Spanning-Tree Protocols Supported on page 5032](#)
- [Configuring Rapid Spanning-Tree Protocol on page 5073](#)
- [Configuring Multiple Spanning-Tree Protocol on page 5067](#)
- [bpd-destination-mac-address on page 5092](#)

System Identifier for Bridges in STP or RSTP Instances

The extended system identifier is used to specify different bridge identifiers for different STP or RSTP routing instances.

**Related
Documentation**

- [Configuring Rapid Spanning-Tree Protocol on page 5073](#)
- [extended-system-id on page 5097](#)

Configuring VLAN Spanning-Tree Protocol

You can configure the VLAN Spanning-Tree Protocol (VSTP) under the following hierarchy levels:

- [edit *logical-systems logical-system-name protocols*]
- [edit *logical-systems logical-system-name routing-instances routing-instance-name protocols*]
- [edit *protocols*]
- [edit *routing-instances routing-instance-name protocols*]

The routing instance type can be either **virtual-switch** or **layer2-control**.

To configure the VLAN Spanning-Tree Protocol:

1. Enable VSTP as the version of spanning-tree protocol to be configured:

```
[edit]
user@host@ edit ... protocols (STP Type) vstp
```

2. (Optional) For compatibility with older bridges that do not support VSTP, you can run force VSTP to run as the original IEEE 802.1D Spanning-Tree Protocol (STP) version:

```
[edit ... protocols vstp]
user@host# set force-version stp
```



NOTE: If VSTP has been forced to run as the original STP version, you can revert back to VSTP by first removing the **force-version** statement from the configuration and then entering the **clear spanning-tree protocol-migration** configuration mode command.

3. Configure the interfaces that participate in the VSTP instance.

- a. Enable configuration of the interface:

```
[edit ... protocols vstp]
user@host# edit interface interface-name
```

- b. Configure the interface priority:

```
[edit ... protocols vstp interface interface-name]
user@host# set priority interface-priority
```

- c. (Optional) By default, the interface link cost is determined by the link speed. You can configure the interface link cost to control which bridge is the designated bridge and which port is the designated port:

```
[edit ... protocols vstp interface interface-name]
user@host# set cost interface-link-cost
```

- d. Configure the interface link mode to identify point-to-point links:

```
[edit ... protocols vstp interface interface-name]
user@host# set mode (p2p | shared)
```

Specify **p2p** if the link is point to point. Specify **shared** if the link is a shared media.

- e. (Optional) Configure the interface as an edge port:

```
[edit ... protocols vstp interface interface-name]
user@host# set edge
```

Edge ports do not expect to receive bridge protocol data unit (BPDU) packets. If a BPDU packet is received for an edge port, the port becomes a nonedge port

You can also enable BPDU root protection for all spanning-tree protocol instances on the interface. BPDU root protect ensures the port is the spanning-tree designated port. If the port receives superior BPDU packets, root protect moves this port to a root-prevented spanning-tree state. For configuration details, see [“Checking the Status of Spanning-Tree Instance Interfaces” on page 5125](#).

4. Enable configuration of a VLAN instance:

```
[edit ... protocols vstp]
user@host# edit vlan vlan-id
```

5. Configure the bridge priority

```
[edit ... protocols vstp vlan vlan-id]
user@host# set bridge-priority bridge-priority
```

For more information, see [“Bridge Priority for Election of Root Bridge and Designated Bridge” on page 5041](#).

6. Configure hello BPDU timers.

- a. Configure the maximum expected arrival time of hello BPDUs:

```
[edit ... protocols vstp vlan vlan-id]
user@host# set max-age seconds
```

- b. Configure the time interval at which the root bridge transmits configuration BPDUs:

```
[edit ... protocols vstp vlan vlan-id]
user@host# set hello-time seconds
```

7. (Optional) By default, the bridge port remains in the listening and learning states for 15 seconds before transitioning to the forwarding state. You can specify a delay from 4 through 20 seconds instead:

```
[edit ... protocols vstp vlan vlan-id]
user@host# set forward-delay seconds
```

8. Configure the interfaces that participate in the VSTP instance.

- a. Enable configuration of the interface:

```
[edit ... protocols vstp vlan vlan-id]
user@host# edit interface interface-name
```

- b. Configure the interface priority:

```
[edit ... protocols vstp vlan vlan-id interface interface-name]
user@host# set priority interface-priority
```

- c. (Optional) By default, the interface link cost is determined by the link speed. You can configure the interface link cost to control which bridge is the designated bridge and which port is the designated port:

```
[edit ... protocols vstp vlan vlan-id interface interface-name]
user@host# set cost interface-link-cost
```

- d. Configure the interface link mode to identify point-to-point links:

```
[edit ... protocols vstp vlan vlan-id interface interface-name]
user@host# set mode (p2p | shared)
```

Specify **p2p** if the link is point to point. Specify **shared** if the link is a shared media.

- e. (Optional) Configure the interface as an edge port:

```
[edit ... protocols vstp vlan vlan-id interface interface-name]
user@host# set edge
```

Edge ports do not expect to receive bridge protocol data unit (BPDU) packets. If a BPDU packet is received for an edge port, the port becomes a nonedge port

You can also enable BPDU root protection for all spanning-tree protocol instances on the interface. BPDU root protect ensures the port is the spanning-tree designated port. If the port receives superior BPDU packets, root protect moves this port to a root-prevented spanning-tree state. For configuration details, see [“Checking the Status of Spanning-Tree Instance Interfaces” on page 5125](#).

9. Verify the VSTP configuration:

```
[edit]
... { # Optional logical system and/or routing instance
  protocols (STP Type) {
    vstp {
```

```
force-version stp; # Optional.
interface interface-name {
    priority interface-priority;
    cost interface-link-cost; # Optional.
    mode (p2p | shared);
    edge; # Optional.
}
vlan vlan-id {
    bridge-priority bridge-priority;
    max-age seconds;
    hello-time seconds;
    forward-delay seconds; # Optional.
    interface interface-name {
        priority interface-priority;
        cost interface-link-cost; # Optional.
        mode (p2p | shared);
        edge; # Optional.
    }
}
}
```

Related Documentation

- [Spanning-Tree Protocols Supported on page 5032](#)
- [RSTP or VSTP Forced to Run as IEEE 802.1D STP on page 5072](#)
- [Reverting RSTP or VSTP Back From Forced IEEE 802.1D STP on page 5075](#)
- [VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Overview on page 5036](#)
- [VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Topology on page 5037](#)

Configuring VPLS Root Protection Topology Change Actions to Control Global Spanning Tree Behavior

To configure VPLS root protection topology change actions to control global spanning tree behavior:

1. Enable configuration of the spanning-tree protocol:

```
[edit]
user@host# edit protocols (STP Type) (mstp | rstp | vstp)
```

2. (Optional) Change the priority of the backup bridge in a VPLS multihomed Layer 2 ring with MPLS infrastructure:

```
[edit protocols (rstp | mstp | vstp)]
user@host# set backup-bridge-priority vpls-ring-backup-bridge-priority
```

3. (Optional) Change number of seconds to hold before switching to the primary priority when the first core domain comes up:

```
[edit protocols (rstp | mstp | vstp)]
user@host# set priority-hold-time seconds
```

4. Configure the system identifier for bridges in the ring:

```
[edit protocols (rstp | mstp | vstp)]
user@host# set system-id system-id-value bridge-host-ip-address(es)
```

The **system-id-value** is configured in the format *nnnnnn:nnnnnn*, where *n* = any digit from 0 to 9.

Each **bridge-host-ip-address** is a valid host IP address with a /32 mask.



NOTE: There are no default values for the system identifier or host IP addresses.

5. Configure bridges to flush the MAC address cache (of the MAC addresses learned when other interfaces ports were blocked) when the spanning-tree topology changes:

```
[edit protocols (rstp | mstp | vstp)]
user@host# set vpls-flush-on-topology-change
```

6. Verify the configuration of VPLS root protection topology change actions to control global spanning tree behavior:

```
[edit]
protocols {
  (mstp | rstp | vstp) {
    backup-bridge-priority priority; # Default is 32,768.
    priority-hold-time seconds; # Default is 2 seconds.
    system-id system-id-value {
      ip-address;
    }
    vpls-flush-on-topology-change;
  }
}
```

Related Documentation

- [VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Overview on page 5036](#)
- [VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Topology on page 5037](#)
- [Configuring VPLS Root Protection Topology Change Actions to Control VLAN Spanning Tree Behavior on page 5085](#)
- [Example: Configuring VPLS Root Topology Change Actions on page 5053](#)

BPDU Protection for Individual Spanning-Tree Instance Interfaces

To configure BPDU protection on one or more spanning-tree instance interfaces, include the **bpdu-block** statement:

```
bpdu-block {
  interface interface-name;
  disable-timeout seconds;
}
```



NOTE: If you also include the optional **disable-timeout seconds** statement, *blocked interfaces* are automatically cleared after the specified time interval unless the interval is 0.

Related Documentation

- [Root Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [BPDU Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [Configuring BPDU Protection on Individual Interfaces on page 5084](#)

Configuring Loop Protection for a Spanning-Tree Instance Interface

Before you begin, you must fully configure the spanning-tree protocol, including instance interfaces. You can configure RSTP, MSTP, or VSTP at the following hierarchy levels:

- **[edit protocols]**
- **[edit routing-instances *routing-instance-name* protocols]**

To configure enhanced loop protection:

1. Include the **bpdu-timeout-action** statement with either the **block** or **log** option for the spanning-tree protocol interface.

- For the STP or RSTP instance on a physical interface:

```
[edit]
protocols {
  rstp {
    interface interface-name {
      bpdu-timeout-action (log | block);
    }
  }
}
```

- For all MSTP instances on a physical interface:

```
[edit]
protocols {
  mstp {
    interface interface-name {
      bpdu-timeout-action (log | block);
    }
  }
}
```

- For all VSTP instances on a physical interface configured at the global level or at the VLAN level:

```
[edit]
protocols {
  vstp {
    interface interface-name {
      bpdu-timeout-action (log | block);
    }
    vlan vlan-id {
      interface interface-name {
        bpdu-timeout-action (log | block);
      }
    }
  }
}
```



```
}
```

2. To display the spanning-tree protocol loop protection characteristics on an interface, use the **show spanning-tree interface** operational command.

Related Documentation

- [Loop Protection for Spanning-Tree Instance Interfaces Overview on page 5034](#)
- [Loop Protection for a Spanning-Tree Instance Interface on page 5047](#)
- [Example: Enabling Loop Protection for Spanning-Tree Protocols on page 5052](#)

Enabling Root Protect for a Spanning-Tree Instance Interface

To enable root protect for a spanning-tree instance interface:

1. Enable configuration of the spanning-tree protocol:

```
[edit]
user@host# edit protocols (mstp | rstp | vstp <vlan vlan-id>)
```

2. Enable configuration of the spanning-tree instance interface:

```
[edit ... protocols (mstp | rstp | vstp <vlan vlan-id>)]
user@host# edit interface interface-name
```

3. Enable root protection on the interface:

```
[edit ... protocols (mstp | rstp | vstp <vlan vlan-id>) interface interface-name]
user@host# set no-root-port
```

4. Verify the configuration of root protect for the spanning-tree instance interface:

```
[edit ... protocols (mstp | rstp | vstp <vlan vlan-id>) interface interface-name]
user@host# top
user@host# show ... protocols
```

```
...
(mstp | rstp | vstp <vlan vlan-id>) {
  interface interface-name {
    no-root-port;
  }
}
```



NOTE: This is not a complete configuration.

Related Documentation

- [Root Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [Root Protect for a Spanning-Tree Instance Interface on page 5048](#)

Configuring BPDU Protection on Individual Interfaces

On MX Series routers and EX Series switches, you can configure BPDU protection to ignore BPDU received on interfaces where none should be expected. If a BPDU is received on a blocked interface, the interface is disabled and stops forwarding frames. By default, all BPDUs are accepted and processed on all interfaces.

To configure BPDU protection for individual spanning-tree instance interfaces:

1. Enable BPDU protection on a specific spanning-tree instance interface:

```
[edit]
user@host# edit protocols layer2-control bpd-block
user@host# set interface interface (aex | (ge-fpc/pic/port | xe-fpc/pic/port))
```

If a BPDU is received on the interface, the system will disable the interface and stop forwarding frames out the interface until the bridging process is restarted.

2. (Optional) Configure the amount of time the system waits before *automatically* unblocking this interface after it has received a BPDU.

```
[edit protocols layer2-control] bpd-block interface interface-name]
user@host# set disable-timeout seconds
```

The range of the *seconds* option value is from 10 through 3600 seconds (one hour). A *seconds* option value of 0 is allowed, but this results in the default behavior (the interface is blocked until the interface is cleared).

3. Verify the configuration of BPDU blocking for individual interfaces:

```
[edit]
interfaces {
  ge-fpc/pic/port { # VLAN encapsulation on Gigabit Ethernet.
    encapsulation (ethernet-bridge | extended-vlan-bridge | extended-vlan-vpls |
      vlan-vpls);
  }
  xe-fpc/pic/port { # VLAN encapsulation on 10-Gigabit Ethernet.
    encapsulation (ethernet-bridge | extended-vlan-bridge | extended-vlan-vpls |
      vlan-vpls);
  }
  ae-X { # VLAN encapsulation
    encapsulation (ethernet-vpls vlan-vpls); # on Aggregated Ethernet.
    ...
  }
  ae-X { # Extended VLAN encapsulation
    vlan-tagging; # on Aggregated Ethernet.
    encapsulation extended-vlan-vpls;
    unit logical-unit-number {
      vlan-id number;
      .....
    }
    .....
  }
}
protocols
  layer2-control {
    bpd-block
      interface interface-name;
```

```

        disable-timeout seconds;
    }
}

```

Related Documentation

- [Root Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [BPDU Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [BPDU Protection for Individual Spanning-Tree Instance Interfaces on page 5048](#)

Configuring BPDU Protection on All Edge Ports

On MX Series routers and EX Series switches, you can configure BPDU protection to ignore BPDU received on interfaces where none should be expected. If a BPDU is received on a blocked interface, the interface is disabled and stops forwarding frames. By default, all BPDUs are accepted and processed on all interfaces.

To configure BPDU protection for all edge ports for a particular spanning-tree protocol::

1. Enable edge port blocking for a particular spanning-tree protocol:

```

[edit]
user@host# set protocols (STP Type) (mstp | rstp | vstp) bpdu-block-on-edge

```

2. Verify BPDU protection for edge ports.

```

[edit]
protocols (STP Type) {
  (mstp | rstp | vstp) {
    bpdu-block-on-edge;
  }
}

```

Related Documentation

- [Root Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [BPDU Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [BPDU Protection on All Edge Ports of the Bridge on page 5049](#)

Configuring VPLS Root Protection Topology Change Actions to Control VLAN Spanning Tree Behavior

To configure VPLS root protection topology change actions to control a particular VLAN:

1. Enable configuration of the spanning-tree protocol VLAN:

```

[edit]
user@host# edit protocols (STP Type) vstp vlan vlan-id

```

2. Optional) Change the priority of the backup bridge in a VPLS multihomed Layer 2 ring with MPLS infrastructure:

```

[edit protocols vstp vlan vlan-id]
user@host# set backup-bridge-priority vpls-ring-backup-bridge-priority

```

3. (Optional) Change the hold time before switching to the primary priority when the first core domain comes up:

```

[edit protocols vstp vlan vlan-id]

```

```
user@host# set priority-hold-time seconds
```

4. Configure the system identifier for bridges in the ring:

```
[edit protocols vstp vlan vlan-id]
```

```
user@host# set system-id system-id-value bridge-host-ip-address(es)
```

The *system-id-value* is configured in the format *nnnnnn:nnnnnn*, where *n* = any digit from 0 to 9.

Each *bridge-host-ip-address* is a valid host IP address with a /32 mask.



NOTE: There are no default values for the system identifier or host IP addresses.

5. Configure bridges to flush the MAC address cache (of the MAC addresses learned when other interfaces ports were blocked) when the spanning-tree topology changes:

```
[edit protocols vstp vlan vlan-id]
```

```
user@host# set vpls-flush-on-topology-change
```

6. Verify the configuration of VPLS root protection topology change actions to control a particular VLAN:

```
[edit]
protocols {
  vstp {
    vlan vlan-id {
      backup-bridge-priority priority; # Default is 32,768.
      priority-hold-time seconds; # Default is 2 seconds.
      system-id system-id-value {
        ip-address;
      }
      vpls-flush-on-topology-change;
    }
  }
}
```

Related Documentation

- [VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Overview on page 5036](#)
- [VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Topology on page 5037](#)
- [Configuring VPLS Root Protection Topology Change Actions to Control Global Spanning Tree Behavior on page 5080](#)
- [Example: Configuring VPLS Root Topology Change Actions on page 5053](#)

Configuring Layer 2 Protocol Tunneling

To configure Layer 2 protocol tunneling:

1. Enable MAC address rewriting for Layer 2 protocol tunneling:

```
[edit]
```

```
user@host# edit protocols layer2-control mac-rewrite
```

2. Configure the Layer 2 protocol tunnel interface:

```
[edit ... protocols layer2-control mac-rewrite]
```

```

user@host# edit interface ge-fpc/pic/port
3. Configure the Layer 2 protocol to be tunneled:

[edit protocols layer2-control mac-rewrite interface ge-fpc/pic/port]
user@host# set protocol (cdp | stp | vtp)
4. Verify the configuration:

user@host# show protocols
layer2-control {
  mac-rewrite {
    interface ge-fpc/pic/port {
      protocol (cdp | stp | vtp);
    }
  }
}

```

- Related Documentation**
- [Layer 2 Protocol Tunneling Through a Network Overview on page 5039](#)
 - [Checking for a MAC Rewrite Error Condition Blocking a Spanning-Tree Instance Interface on page 5127](#)
 - [Clearing a MAC Rewrite Error Condition Blocking a Spanning-Tree Instance Interface on page 5127](#)

Configuration Statements

- [\[edit protocols mstp\] Hierarchy Level on page 5087](#)
- [\[edit protocols rstp\] Hierarchy Level on page 5088](#)
- [\[edit protocols vstp\] Hierarchy Level on page 5089](#)
- [\[edit protocols layer2-control\] Hierarchy Level on page 5115](#)
- [\[edit protocols layer2-control\] Hierarchy Level on page 5122](#)

[edit protocols mstp] Hierarchy Level

The following statement hierarchy can also be included at the `[edit logical-systems logical-system-name]` hierarchy level.

```

protocols {
  mstp {
    disable;
    backup-bridge-priority priority;
    bpdu-block-on-edge;
    bpdu-destination-mac-address provider-bridge-group;
    bridge-priority priority;
    configuration-name configuration-name;
    forward-delay seconds;
    hello-time seconds;
    interface interface-name {
      bpdu-timeout-action {
        alarm;
        block;
      }
      cost cost;
      edge;
    }
  }
}

```

```
    mode (point-to-point | shared);
    no-root-port;
    priority interface-priority;
  }
  max-age seconds;
  max-hops hops;
  msti identifier {
    backup-bridge-priority priority;
    bridge-priority priority;
    interface interface-name {
      cost cost;
      priority interface-priority;
    }
    vlan [ vlan-ids ];
  }
  priority-hold-time seconds;
  revision-level revision-level;
  system-id mac-address {
    ip-address ip-address </prefix-length>;
  }
  traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag <disable>;
  }
  vpls-flush-on-topology-change;
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

[edit protocols rstp] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```
protocols {
  rstp {
    disable;
    backup-bridge-priority priority;
    bpdu-block-on-edge;
    bpdu-destination-mac-address provider-bridge-group;
    bridge-priority priority;
    extended-system-id id;
    force-version stp;
    forward-delay seconds;
    hello-time seconds;
    interface interface-name {
      bpdu-timeout-action {
        alarm;
        block;
      }
    }
    cost cost;
```

```

    edge;
    mode (point-to-point | shared);
    no-root-port;
    priority interface-priority;
  }
  max-age seconds;
  priority-hold-time seconds;
  system-id mac-address {
    ip-address ip-address </prefix-length>;
  }
  traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag <disable>;
  }
  vpls-flush-on-topology-change;
}

```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

[edit protocols vstp] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```

protocols {
  vstp {
    disable;
    bpdu-block-on-edge;
    force-version stp;
    interface interface-name {
      access-trunk
      bpdu-timeout-action {
        alarm;
        block;
      }
      cost cost;
      edge;
      mode (point-to-point | shared);
      no-root-port;
      priority interface-priority;
    }
    priority-hold-time seconds;
    system-id mac-address {
      ip-address ip-address </prefix-length>;
    }
    vlan vlan-id {
      ... the vlan subhierarchy appears after the main [edit protocols vstp] hierarchy level ...
    }
    vpls-flush-on-topology-change;
  }
}

```

```
vstp {  
  vlan vlan-id {  
    backup-bridge-priority priority;  
    bridge-priority priority;  
    forward-delay seconds;  
    hello-time seconds;  
    interface interface-name {  
      ...same statements as at the [edit protocols vstp interface interface-name] hierarchy  
      level ...  
    }  
    max-age seconds;  
    traceoptions {  
      file filename <files number> <size maximum-file-size> <world-readable |  
        no-world-readable>;  
      flag flag <disable>;  
    }  
  }  
}
```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

protocols (STP Type)

Syntax	<pre>protocols { mstp { ... } rstp { ... } vstp { ... } }</pre>
Hierarchy Level	[edit], [edit logical-systems <i>logical-system-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6.
Description	Configure the Spanning Tree Protocol type as MSTP, RSTP, or VSTP.
Options	mstp —Configure the protocol as Multiple Spanning Tree. rstp —Configure the protocol as Rapid Spanning Tree. vstp —Configure the protocol as VLAN Spanning Tree.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Rapid Spanning-Tree Protocol on page 5073 • Configuring Multiple Spanning-Tree Protocol on page 5067 • Configuring VLAN Spanning-Tree Protocol on page 5077

access-trunk

Syntax	access-trunk;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols vstp] [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vstp] [edit protocols vstp vlan <i>vlan-identifier</i> interface <i>interface-name</i>] [edit routing-instances <i>routing-instance-name</i> instance-type (layer2-control virtual-switch)]
Description	Enable untagged RTSP BDPUs to be sent and received on the interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring VSTP on a Trunk Port with Tagged Traffic on page 5053

bpdu-destination-mac-address (Spanning Tree)

Syntax	<code>bpdu-destination-mac-address provider-bridge-group;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols (mstp rstp)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp rstp)], [edit protocols (mstp rstp)], [edit routing-instances <i>routing-instance-name</i> protocols (mstp rstp)]
Release Information	Statement introduced in Junos OS Release 9.2. Support for logical systems added in Junos OS Release 9.6.
Description	Enable MX Series router to participate in the provider Rapid Spanning-Tree Protocol (RSTP) instance or a provider Multiple Spanning-Tree Protocol (MSTP) instance.
Default	If the bpdu-destination-mac-address statement is not configured, the bridge participates in the customer RSTP instance, transmitting and receiving standard RSTP BPDU packets.
Options	provider-bridge-group —The destination MAC address of the BPDU packets transmitted is the provider bridge group address 01:80:c2:00:00:08 . Received BPDU packets with this destination MAC address are accepted and passed to the Routing Engine.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• BPDU Overview on page 5033• Provider Bridge Participation in RSTP or MSTP Instances on page 5072• Configuring Rapid Spanning-Tree Protocol on page 5073• Configuring Multiple Spanning-Tree Protocol on page 5067

configuration-name

Syntax	configuration-name <i>configuration-name</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mstp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mstp], [edit protocols mstp], [edit routing-instances <i>routing-instance-name</i> protocols mstp]
Release Information	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6.
Description	The configuration name is the MSTP region name carried in the MSTP BPDUs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• BPDU Overview on page 5033• Configuring Multiple Spanning-Tree Protocol on page 5067

cost

Syntax	<code>cost cost;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (mstp rstp vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols mstp <i>msti</i> <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols vstp <i>vlan</i> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp rstp vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mstp <i>msti</i> <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vstp <i>vlan</i> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (mstp rstp vstp) interface <i>interface-name</i>],</p> <p>[edit protocols mstp <i>msti</i> <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols vstp <i>vlan</i> <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (mstp rstp vstp) interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols mstp <i>msti</i> <i>msti-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vstp <i>vlan</i> <i>vlan-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p>
Description	Configure link cost to control which bridge is the designated bridge and which port is the designated port. By default, the link cost is determined by the link speed.
Options	<p>cost—(Optional) Link cost associated with the port.</p> <p>Range: 1 through 200,000,000</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Spanning-Tree Instance Interface on page 5042 • Spanning-Tree Instance Interface Cost on page 5044

disable

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mstp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mstp], [edit protocols mstp], [edit routing-instances <i>routing-instance-name</i> protocols mstp]
Release Information	Statement introduced in Junos OS Release 9.1. Support for logical systems added in Junos OS Release 9.6.
Description	Disable the entire MSTP instance.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Multiple Spanning-Tree Protocol on page 5067• Disabling MSTP on page 5071

edge

Syntax	edge;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (mstp rstp vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols mstp msti msti-id interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols vstp vlan vlan-id interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp rstp vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mstp msti msti-id interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vstp vlan vlan-id interface <i>interface-name</i>],</p> <p>[edit protocols (mstp rstp vstp) interface <i>interface-name</i>],</p> <p>[edit protocols mstp msti msti-id interface <i>interface-name</i>],</p> <p>[edit protocols vstp vlan vlan-id interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (mstp rstp vstp) interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols mstp msti msti-id interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vstp vlan vlan-id interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p>
Description	Configure interfaces as edge ports. Edge ports do not expect to receive BPDUs. If a BPDU is received, the port becomes a nonedge port.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Spanning-Tree Instance Interface on page 5042 • Spanning-Tree Instance Interface Configured as an Edge Port on page 5045

extended-system-id

Syntax	<code>extended-system-id <i>identifier</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rstp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rstp], [edit protocols rstp], [edit routing-instances <i>routing-instance-name</i> protocols rstp]
Release Information	Statement introduced in Junos OS Release 8.3. Support for logical systems added in Junos OS Release 9.6.
Description	The extended system ID is used to specify different bridge identifiers for different RSTP or STP routing instances.
Options	<i>identifier</i> —Specify the system identifier to use for the RSTP or STP instance. Range: 0 through 4095
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • System Identifier for Bridges in STP or RSTP Instances on page 5076 • Configuring Rapid Spanning-Tree Protocol on page 5073

force-version (IEEE 802.1D STP)

Syntax	<code>force-version stp;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols (rstp vstp)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (rstp vstp)], [edit protocols (rstp vstp)], [edit routing-instances <i>routing-instance-name</i> protocols (rstp vstp)]
Release Information	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6.
Description	Force the spanning-tree version to be the original IEEE 803.1D STP.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Spanning-Tree Protocols Supported on page 5032 • RSTP or VSTP Forced to Run as IEEE 802.1D STP on page 5072 • Reverting RSTP or VSTP Back From Forced IEEE 802.1D STP on page 5075

forward-delay

Syntax	<code>forward-delay seconds;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols (<i>mstp</i> <i>rstp</i>)],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols <i>vstp</i> vlan <i>vlan-id</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> (<i>mstp</i> <i>rstp</i>)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> <i>vstp</i> vlan <i>vlan-id</i>],</code> <code>[edit protocols (<i>mstp</i> <i>rstp</i>)],</code> <code>[edit protocols <i>vstp</i> vlan <i>vlan-id</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols (<i>mstp</i> <i>rstp</i>)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols <i>vstp</i> vlan <i>vlan-id</i>]</code>
Release Information	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6.
Description	Specify the length of time an STP bridge port remains in the listening and learning states before transitioning to the forwarding state.
Options	seconds —(Optional) Number of seconds the bridge port remains in the listening and learning states. Range: 4 through 30 Default: 15 seconds
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Forward Delay Before Ports Transition to Forwarding State on page 5042

hello-time

Syntax	<code>hello-time seconds;</code>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols (<i>mstp</i> <i>rstp</i>)], [edit logical-systems <i>logical-system-name</i> protocols <i>vstp</i> <i>vlan</i> <i>vlan-id</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<i>mstp</i> <i>rstp</i>)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>vstp</i> <i>vlan</i> <i>vlan-id</i>], [edit protocols (<i>mstp</i> <i>rstp</i>)], [edit protocols <i>vstp</i> <i>vlan</i> <i>vlan-id</i>], [edit routing-instances <i>routing-instance-name</i> protocols (<i>mstp</i> <i>rstp</i>)], [edit routing-instances <i>routing-instance-name</i> protocols <i>vstp</i> <i>vlan</i> <i>vlan-id</i>]</pre>
Release Information	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p>
Description	Specify the number of seconds between transmissions of configuration BPDUs by the root bridge.
Options	<p>seconds—(Optional) Number of seconds between transmissions of configuration BPDUs.</p> <p>Range: 1 through 10</p> <p>Default: 2 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Hello Time for Root Bridge to Transmit Hello BPDUs on page 5042

interface (Spanning Tree)

Syntax	<pre>interface <i>interface-name</i> { bpd-<i>timeout-action</i> { alarm; block; } cost <i>cost</i>; edge; mode (p2p shared); no-root-port; priority <i>interface-priority</i>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols (mstp rstp vstp)], [edit logical-systems <i>logical-system-name</i> protocols vstp vlan <i>vlan-id</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp rstp vstp)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vstp vlan <i>vlan-id</i>], [edit protocols (mstp rstp vstp)], [edit protocols vstp vlan <i>vlan-id</i>], [edit routing-instances <i>routing-instance-name</i> protocols (mstp rstp vstp)], [edit routing-instances <i>routing-instance-name</i> protocols vstp vlan <i>vlan-id</i>]</pre>
Release Information	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6.
Description	Configure the interface to participate in the RSTP or MSTP instance.
Options	<p><i>interface-name</i>—Name of a Gigabit Ethernet or 10-Gigabit Ethernet interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Spanning-Tree Instance Interface on page 5042

max-age

Syntax	<code>max-age seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (mstp rstp)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols vstp vlan <i>vlan-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp rstp)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vstp vlan <i>vlan-id</i>],</p> <p>[edit protocols (mstp rstp)],</p> <p>[edit protocols vstp vlan <i>vlan-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (mstp rstp)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vstp vlan <i>vlan-id</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p>
Description	Specify the maximum expected arrival time of hello BPDUs.
Options	<p>seconds—(Optional) Number of seconds expected between hello BPDUs.</p> <p>Range: 6 through 40</p> <p>Default: 20 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Maximum Age for Awaiting Arrival of Hello BPDUs on page 5042

max-hops

Syntax	<code>max-hops hops;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mstp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mstp], [edit protocols mstp], [edit routing-instances <i>routing-instance-name</i> protocols mstp]
Release Information	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6.
Description	Configure the maximum number of hops a BPDU can be forwarded in the MSTP region.
Options	hops —(Optional) Number of hops the BPDU can be forwarded. Range: 1 through 255 Default: 19 hops
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Multiple Spanning-Tree Protocol on page 5067

mode (Protocols STP)

Syntax	<code>mode (p2p shared);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (mstp rstp vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols vstp vlan <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp rstp vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vstp vlan <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (mstp rstp vstp) interface <i>interface-name</i>],</p> <p>[edit protocols vstp vlan <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (mstp rstp vstp) interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vstp vlan <i>vlan-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p>
Description	Configure link mode to identify point-to-point links.
Default	When the link is configured as full-duplex, the default link mode is p2p . When the link is configured half-duplex, the default link mode is shared .
Options	<p>p2p—The link is point to point.</p> <p>shared—The link is shared media.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Spanning-Tree Instance Interface on page 5042 • Spanning-Tree Instance Interface Point-to-Point Link Mode on page 5044

msti

Syntax	<pre>msti <i>msti-id</i> { bridge-priority <i>priority</i>; vlan <i>vlan-id</i>; interface <i>interface-name</i> { cost <i>cost</i>; edge; priority <i>interface-priority</i>; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mstp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mstp], [edit protocols mstp], [edit routing-instances <i>routing-instance-name</i> protocols mstp]
Release Information	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6.
Description	Configure the Multiple Spanning Tree Protocol (MSTI) instance identifier.
Options	<p>msti-id—MSTI instance identifier.</p> <p>Range: 1 through 64</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Multiple Spanning-Tree Protocol on page 5067• Configuring MST Instances on a Physical Interface on page 5070

mstp

Syntax	<pre> mstp { bpd-block-on-edge; bridge-priority <i>priority</i>; configuration-name <i>configuration-name</i>; disable; forward-delay <i>seconds</i>; hello-time <i>seconds</i>; max-age <i>seconds</i>; max-hops <i>hops</i>; priority-hold-time <i>seconds</i>; revision-level <i>revision-level</i>; interface <i>interface-name</i> { bpd-timeout-action { alarm; block; } cost <i>cost</i>; edge; mode (p2p shared); no-root-port; priority <i>interface-priority</i>; } msti <i>msti-id</i> { bridge-priority <i>priority</i>; interface <i>interface-name</i> { cost <i>cost</i>; edge; priority <i>interface-priority</i>; } vlan <i>vlan-id</i>; } traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],</p> <p>[edit protocols],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.4.</p> <p>bpd-block-on-edge statement added in Junos OS Release 9.4.</p> <p>bpd-timeout-action statement added in Junos OS Release 9.4.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p>
Description	Configure MSTP parameters.
Options	The statements are explained separately.

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

Related Documentation

- [Configuring Multiple Spanning-Tree Protocol on page 5067](#)

priority (Protocols STP)

Syntax *priority interface-priority;*

Hierarchy Level [edit logical-systems *logical-system-name* protocols (**mstp** | **rstp** | **vstp**) **interface** *interface-name*],
[edit logical-systems *logical-system-name* protocols **mstp** **msti** *msti-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* protocols **vstp** **vlan** *vlan-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**mstp** | **rstp** | **vstp**) **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **mstp** **msti** *msti-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **vstp** **vlan** *vlan-id* **interface** *interface-name*],
[edit protocols (**mstp** | **rstp** | **vstp**) **interface** *interface-name*],
[edit protocols **mstp** **msti** *msti-id* **interface** *interface-name*],
[edit protocols **vstp** **vlan** *vlan-id* **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols (**mstp** | **rstp** | **vstp**) **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols **mstp** **msti** *msti-id* **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols **vstp** **vlan** *vlan-id* **interface** *interface-name*]

Release Information Statement introduced in Junos OS Release 8.4.
Support for logical systems added in Junos OS Release 9.6.

Description Use the interface priority to control which interface is elected as the root port. The interface priority must be set in increments of 16.

Options *priority*—(Optional) Interface priority.
Range: 0 through 240

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

Related Documentation

- [Spanning-Tree Instance Interface on page 5042](#)
- [Spanning-Tree Instance Interface Configured as an Edge Port on page 5045](#)
- [Spanning-Tree Instance Interface Priority on page 5043](#)

revision-level

Syntax	<code>revision-level <i>revision-level</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mstp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mstp], [edit protocols mstp], [edit routing-instances <i>routing-instance-name</i> protocols mstp]
Release Information	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6.
Description	Set the revision number of the MSTP configuration.
Options	<i>revision-level</i> —Configure the revision number of the MSTP region configuration. Range: 0 through 65,535
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Multiple Spanning-Tree Protocol on page 5067

rstp

Syntax	<pre> rstp { bpd-block-on-edge; bpd-destination-mac-address provider-bridge-group; bridge-priority priority; extended-system-id; force-version stp; forward-delay seconds; hello-time seconds; max-age seconds; interface interface-name { bpd-timeout-action { alarm; block; } cost cost; edge; mode (p2p shared); no-root-port; priority interface-priority; } priority-hold-time seconds; traceoptions { file filename <files number> <size size> <world-readable no-world-readable>; flag flag <flag-modifier> <disable>; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],</p> <p>[edit protocols],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.4.</p> <p>bpd-block-on-edge statement added in Junos OS Release 9.4.</p> <p>bpd-timeout-action statement added in Junos OS Release 9.4.</p> <p>Support for logic systems added in Junos OS Release 9.6.</p>
Description	Configure RSTP parameters.
Options	The statements are explained separately.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Rapid Spanning-Tree Protocol on page 5073

traceoptions (Spanning Tree)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols (mstp rstp vstp)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp rstp vstp)], [edit protocols (mstp rstp vstp)], [edit routing-instances <i>routing-instance-name</i> protocols (mstp rstp vstp)]
Release Information	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6.
Description	Set STP protocol-level tracing options.
Default	The default STP protocol-level trace options are inherited from the global traceoptions statement.
Options	<p>disable—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place STP tracing output in the file <code>/var/log/stp-log</code>.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files Default: 1 trace file only</p> <p>flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. The following are the STP-specific tracing options:</p> <ul style="list-style-type: none"> • all—Trace all operations. • all-failures—Trace all failure conditions. • bpdud—Trace BPDU reception and transmission. • bridge-detection-state-machine—Trace the bridge detection state machine. • events—Trace events of the protocol state machine.

- **port-information-state-machine**—Trace the port information state machine.
- **port-migration-state-machine**—Trace the port migration state machine.
- **port-receive-state-machine**—Trace the port receive state machine.
- **port-role-transit-state-machine**—Trace the port role transit state machine.
- **port-role-select-state-machine**—Trace the port role selection state machine.
- **port-state-transit-state-machine**—Trace the port state transit state machine.
- **port-transmit-state-machine**—Trace the port transmit state machine.
- **ppmd**—Trace the state and events for the ppmmd process.
- **state-machine-variables**—Trace when the state machine variables change.
- **timers**—Trace protocol timers.
- **topology-change-state-machine**—Trace the topology change state machine.

The following are the global tracing options:

- **all**—All tracing operations.
- **config-internal**—Trace configuration internals.
- **general**—Trace general events.
- **normal**—All normal events.

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **parse**—Trace configuration parsing.
- **policy**—Trace policy operations and actions.
- **regex-parse**—Trace regular-expression parsing.
- **route**—Trace routing table changes.
- **state**—Trace state transitions.
- **task**—Trace protocol task processing.
- **timer**—Trace protocol task timer processing.

no-world-readable—(Optional) Prevent any user from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

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

Related Documentation	<ul style="list-style-type: none"> • Spanning-Tree Protocol Trace Options on page 5045 • Tracing Spanning-Tree Operations on page 5128 • Example: Tracing Spanning-Tree Protocol Operations on page 5065
------------------------------	---

vlan (MSTP)

Syntax	<code>vlan <i>vlan-id</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mstp msti msti-id], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mstp msti msti-id], [edit protocols mstp msti msti-id], [edit routing-instances <i>routing-instance-name</i> protocols mstp msti msti-id]
Release Information	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6.
Description	Configure the VLAN of an MSTI or VSTP instance or configure the VLAN range of an MSTI instance.
Options	<i>vlan-id</i> —The VLAN identifier associated with the MSTI. <i>vlan-id-range</i> —Range of VLAN identifiers associated with the MSTI in the form <i>minimum-vlan-id-maximum-vlan-id</i> . VLAN identifier ranges are not supported for VSTP. Range: 1 through 4096
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Multiple Spanning-Tree Protocol on page 5067

vlan (VSTP)

Syntax	<pre> vlan <i>vlan-id</i> { bridge-priority <i>priority</i>; forward-delay <i>seconds</i>; hello-time <i>seconds</i>; max-age <i>seconds</i>; interface <i>interface-name</i> { cost <i>cost</i>; edge; mode (p2p shared); no-root-port; priority <i>interface-priority</i>; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols vstp], [edit protocols vstp]
Release Information	Statement introduced in Junos OS Release 9.0. Support for logical systems added in Junos OS Release 9.6.
Description	Configure VSTP VLAN parameters.
Options	The statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring VLAN Spanning-Tree Protocol on page 5077

vstp

```
Syntax  vstp {
        bpdu-block-on-edge;
        force-version stp;
        interface interface-name {
            bpdu-timeout-action {
                alarm;
                block;
            }
            cost cost;
            edge;
            mode (p2p | shared);
            no-root-port;
            priority interface-priority;
        }
        priority-hold-time seconds;
        vlan vlan-id {
            bridge-priority priority;
            forward-delay seconds;
            hello-time seconds;
            max-age seconds;
            interface interface-name {
                access-trunk
                bpdu-timeout-action {
                    alarm;
                    block;
                }
                cost cost;
                edge;
                mode (p2p | shared);
                no-root-port;
                priority interface-priority;
            }
        }
        traceoptions {
            file filename <files number> <size size> <world-readable | no-world-readable>;
            flag flag <flag-modifier> <disable>;
        }
    }
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
 [edit protocols],
 [edit routing-instances *routing-instance-name* protocols]

Release Information Statement introduced in Junos OS Release 9.0.
bpdu-block-on-edge statement added in Junos OS Release 9.4.
bpdu-timeout-action statement added in Junos OS Release 9.4.
 Support for logical systems added in Junos OS Release 9.6.

Description Configure VSTP parameters.

Options	The statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring VLAN Spanning-Tree Protocol on page 5077

[\[edit protocols layer2-control\] Hierarchy Level](#)

The following statement hierarchy can also be included at the [\[edit logical-systems logical-system-name\]](#) hierarchy level.

```
protocols {
  layer2-control {
    bpdu-block {
      disable-timeout seconds;
      interface [ interface-names ];
    }
    mac-rewrite {
      interface interface-name {
        protocol {
          cdp;
          stp;
          vtp;
        }
      }
    }
  }
  nonstop-bridging;
  traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag <disable>;
  }
}
```

Related Documentation	<ul style="list-style-type: none"> • <i>Notational Conventions Used in Junos OS Configuration Hierarchies</i> • [edit protocols] Hierarchy Level
------------------------------	--

interface (Layer 2 Protocol Tunneling)

Syntax	<code>interface <i>interface-name</i> { <code>protocol</code> (cdp stp vtp); }</code>
Hierarchy Level	[edit protocols <code>layer2-control</code> <code>mac-rewrite</code>]
Release Information	Statement introduced in Junos OS Release 9.1.
Description	<p>Configure an interface for Layer 2 protocol tunneling.</p> <p>The remaining statement is described separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Layer 2 Protocol Tunneling Through a Network Overview on page 5039

layer2-control

Syntax

```
layer2-control {
  bpd-block {
    interface interface-name;
    disable-timeout seconds;
  }
  mac-rewrite {
    interface interface-name {
      protocol (cdp | stp | vtp);
    }
  }
  nonstop-bridging;
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 8.4.
bpd-block statement added in Junos OS Release 9.4.

Description Configure Layer 2 control protocols to enable features such as Layer 2 protocol tunneling or nonstop bridging.

The remaining statements are described separately.



NOTE: For a detailed description of configuring the nonstop-bridging statement, see the *Junos OS High Availability Configuration Guide*. When this statement is configured on routing platforms with two Routing Engines, a master Routing Engine switches over gracefully to a backup Routing Engine and preserves Layer 2 Control Protocol (L2CP) information.

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

Related Documentation

- [Layer 2 Protocol Tunneling Through a Network Overview on page 5039](#)
- [Layer 2 Protocol Tunnel Interface on page 5046](#)
- [Layer 2 Protocol to be Tunneled on page 5047](#)
- [Configuring Layer 2 Protocol Tunneling on page 5086](#)
- *instance-type*

mac-rewrite

Syntax	<pre>mac-rewrite { interface <i>interface-name</i> { protocol (cdp stp vtp); } }</pre>
Hierarchy Level	[edit protocols layer2-control]
Release Information	Statement introduced in Junos OS Release 9.1.
Description	Enable rewriting of the MAC address for Layer 2 protocol tunneling. The remaining statements are described separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Layer 2 Protocol Tunneling Through a Network Overview on page 5039

protocol

Syntax	<pre>protocol (cdp stp vtp);</pre>
Hierarchy Level	[edit protocols layer2-control mac-rewrite interface <i>interface-name</i>],
Release Information	Statement introduced in Junos OS Release 9.1.
Description	Configure the protocol to be tunneled on an interface for Layer 2 protocol tunneling. To tunnel multiple protocols, include multiple protocol statements.
Options	cdp —Tunnel the Cisco discovery protocol. stp —Tunnel all versions of the spanning-tree protocol. vtp —Tunnel the VLAN trunk protocol.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Layer 2 Protocol Tunneling Through a Network Overview on page 5039• Layer 2 Protocol Tunnel Interface on page 5046• Layer 2 Protocol to be Tunneled on page 5047• Configuring Layer 2 Protocol Tunneling on page 5086

no-root-port

Syntax	no-root-port;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (mstp rstp vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols vstp vlan <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp rstp vstp) interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vstp vlan <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (mstp rstp vstp) interface <i>interface-name</i>],</p> <p>[edit protocols vstp vlan <i>vlan-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (mstp rstp vstp) interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vstp vlan <i>vlan-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.1.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p>
Description	Ensure the port is the spanning-tree designated port. If the port receives superior bridge protocol data unit (BPDU) packets, root protect moves this port to a root-prevented spanning-tree state.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Root Protection for Spanning-Tree Instance Interfaces Overview on page 5035 • Root Protect for a Spanning-Tree Instance Interface on page 5048 • Enabling Root Protect for a Spanning-Tree Instance Interface on page 5083

priority-hold-time

Syntax	<code>priority-hold-time seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols (mstp rstp)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp rstp)], [edit protocols (mstp rstp)], [edit routing-instances <i>routing-instance-name</i> protocols (mstp rstp)],
Release Information	Statement introduced in Junos OS Release 10.0.
Description	Specify the number of seconds to hold before switching to the primary priority when the first core domain comes up.
Options	seconds —Number of seconds to hold before switching to primary priority. Range: 1 through 255 Default: 2 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• VPLS Multihoming: Hold Time Before Switching to Primary Priority on page 5050

system-id

Syntax	<code>system-id system-id-value { ip-address(es); }</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (mstp rstp)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols vstp vlan <i>vlan-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (mstp rstp)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vstp vlan <i>vlan-id</i>],</p> <p>[edit protocols (mstp rstp)],</p> <p>[edit protocols vstp vlan <i>vlan-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (mstp rstp)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vstp vlan <i>vlan-id</i>]</p>
Release Information	Statement introduced in Junos OS Release 10.0.
Description	Determine the system identifier value for bridges in a VPLS multihomed Layer 2 ring with MPLS infrastructure.
Options	<p>system-id-value—System identifier in the format <i>nnnnnn:nnnnnn</i> where <i>n</i> = any digit from 0 through 9.</p> <p>Range: Any valid value</p> <p>Default: None</p> <p>ip-address(es)—Valid IP host addresses in the format <i>ip-address/32</i>.</p> <p>Range: Any valid IP address</p> <p>Default: None</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Overview on page 5036 • VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Topology on page 5037 • VPLS Multihoming: System Identifier for Bridges in the Ring on page 5050

vpls-flush-on-topology-change

Syntax	vpls-flush-on-topology-change;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (<i>mstp</i> <i>rstp</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <i>vstp</i> <i>vlan</i> <i>vlan-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<i>mstp</i> <i>rstp</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>vstp</i> <i>vlan</i> <i>vlan-id</i>],</p> <p>[edit protocols (<i>mstp</i> <i>rstp</i>)],</p> <p>[edit protocols <i>vstp</i> <i>vlan</i> <i>vlan-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<i>mstp</i> <i>rstp</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <i>vstp</i> <i>vlan</i> <i>vlan-id</i>]</p>
Release Information	Statement introduced in Junos OS Release 10.0.
Description	Determine the action the bridge should take when the topology of a multihomed Layer 2 ring with MPLS infrastructure changes: flush the media access control (MAC) cache or not. By default, the bridge does not flush the cache when the topology changes.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • VPLS Multihoming: Bridge Flush of MAC Cache on Topology Change on page 5051

[edit protocols layer2-control] Hierarchy Level

The following statement hierarchy can also be included at the [edit logical-systems *logical-system-name*] hierarchy level.

```

protocols {
  layer2-control {
    bpdu-block {
      disable-timeout seconds;
      interface [ interface-names ];
    }
    mac-rewrite {
      interface interface-name {
        protocol {
          cdp;
          stp;
          vtp;
        }
      }
    }
  }
  nonstop-bridging;
  traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
      no-world-readable>;
    flag flag <disable>;
  }
}

```



```
    }
  }
```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

bpdu-block

Syntax bpdu-block {
 [interface](#) *interface-name*;
 [disable-timeout](#) *seconds*;
 }

Hierarchy Level [edit protocols [layer2-control](#)]

Release Information Statement introduced in Junos OS Release 9.4.

Description Enable BPDU blocking on an interface.

The remaining statements are described separately.

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

Related Documentation

- [BPDU Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [BPDU Protection for Individual Spanning-Tree Instance Interfaces on page 5048](#)
- [Configuring BPDU Protection on Individual Interfaces on page 5084](#)

disable-timeout

Syntax	<code>disable-timeout <i>seconds</i>;</code>
Hierarchy Level	[edit protocols layer2-control bpd u-block]
Release Information	Statement introduced in Junos OS Release 9.4.
Description	Configure the timeout value to periodically check to see if an interface is still disabled with BPDU blocking. If this option is not configured, the interface is not periodically checked and remains disabled.
Options	<p><i>seconds</i>—Disable timeout value.</p> <p>Range: 10 through 3600</p> <p>Default: If this option is not configured, the interface is not periodically checked and remains disabled.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• BPDU Protection for Spanning-Tree Instance Interfaces Overview on page 5035• BPDU Protection for Individual Spanning-Tree Instance Interfaces on page 5048

interface (BPDU Blocking)

Syntax	<code>interface <i>interface-name</i>;</code>
Hierarchy Level	[edit protocols layer2-control bpd u-block]
Release Information	Statement introduced in Junos OS Release 9.4.
Description	Configure the interface to participate in BPDU blocking.
Options	<i>interface-name</i> —Name of a Gigabit Ethernet or 10-Gigabit Ethernet interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• BPDU Protection for Spanning-Tree Instance Interfaces Overview on page 5035• BPDU Protection for Individual Spanning-Tree Instance Interfaces on page 5048• Configuring BPDU Protection on Individual Interfaces on page 5084

interface (Layer 2 Protocol Tunneling)

Syntax	<code>interface <i>interface-name</i> { <code>protocol</code> (cdp stp vtp); }</code>
Hierarchy Level	[edit protocols <code>layer2-control mac-rewrite</code>]
Release Information	Statement introduced in Junos OS Release 9.1.
Description	Configure an interface for Layer 2 protocol tunneling. The remaining statement is described separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Layer 2 Protocol Tunneling Through a Network Overview on page 5039

Administration

- [Routine Monitoring on page 5125](#)
- [Operational Commands on page 5130](#)

Routine Monitoring

- [Checking the Status of Spanning-Tree Instance Interfaces on page 5125](#)
- [Clearing the Blocked Status of a Spanning-Tree Instance Interface on page 5126](#)
- [Checking for a MAC Rewrite Error Condition Blocking a Spanning-Tree Instance Interface on page 5127](#)
- [Clearing a MAC Rewrite Error Condition Blocking a Spanning-Tree Instance Interface on page 5127](#)
- [Tracing Spanning-Tree Operations on page 5128](#)

Checking the Status of Spanning-Tree Instance Interfaces

On an MX Series router with a spanning-tree protocol enabled, the detection of a possible bridging loop from spanning-tree protocol operation can raise a bridge protocol data unit (BPDU) error condition on the affected spanning-tree instance interface.

To check whether a spanning-tree instance interface is blocked due to a BPDU error condition:

1. To check the status of spanning-tree instance interface, use the **show interfaces** command.

```
user@host> show interfaces interface-name
```

2. You can determine the status of the interface as follows:

- If the **BPDU Error** field is **none**, the interface is enabled.
- If the **BPDU Error** field is **Detected** and the link is **down**, the interface is blocked.

**Related
Documentation**

- [Root Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [BPDU Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [BPDU Protection for Individual Spanning-Tree Instance Interfaces on page 5048](#)
- [BPDU Protection on All Edge Ports of the Bridge on page 5049](#)
- [Clearing the Blocked Status of a Spanning-Tree Instance Interface on page 5126](#)

Clearing the Blocked Status of a Spanning-Tree Instance Interface

To clear the blocked status of a spanning-tree instance interface:

- Use the **clear error bpdu** operational mode command.

```
user@host> clear error bpdu interface interface-name
```



NOTE: When you configure BPDU protection on individual interfaces (as opposed to on all the edge ports of the bridge), you can use the **disable-timeout seconds** option to specify that a blocked interface is automatically cleared after the specified time interval elapses (unless the interval is 0). For configuration details, see “[Configuring BPDU Protection on Individual Interfaces](#)” on page 5084.

**Related
Documentation**

- [Root Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [BPDU Protection for Spanning-Tree Instance Interfaces Overview on page 5035](#)
- [BPDU Protection for Individual Spanning-Tree Instance Interfaces on page 5048](#)
- [BPDU Protection on All Edge Ports of the Bridge on page 5049](#)
- [Checking the Status of Spanning-Tree Instance Interfaces on page 5125](#)

Checking for a MAC Rewrite Error Condition Blocking a Spanning-Tree Instance Interface

To check whether a spanning-tree instance interface is blocked due to a MAC rewrite error condition:

1. Use the **show interfaces** operational mode command:

```
user@host> show interfaces interface-name
```

2. You can determine the status of the interface as follows:

- If the value in the **Physical interface** includes **Enabled, Physical link is Up** and the value of the **BPDU Error** field is **None**, the interface is enabled
- If the value in the **Physical interface** field is **Enabled, Physical link is Down** and the value in the **BPDU Error** field is **Detected**, the interface is blocked.

Related Documentation

- [Layer 2 Protocol Tunneling Through a Network Overview on page 5039](#)
- [Configuring Layer 2 Protocol Tunneling on page 5086](#)
- [Clearing a MAC Rewrite Error Condition Blocking a Spanning-Tree Instance Interface on page 5127](#)

Clearing a MAC Rewrite Error Condition Blocking a Spanning-Tree Instance Interface

To clear the blocked status of a spanning-tree instance interface:

- Use the **clear error bpd** operational mode command.

```
user@host> clear error bpd interface interface-name
```

Related Documentation

- [Layer 2 Protocol Tunneling Through a Network Overview on page 5039](#)
- [Configuring Layer 2 Protocol Tunneling on page 5086](#)
- [Checking for a MAC Rewrite Error Condition Blocking a Spanning-Tree Instance Interface on page 5127](#)

Tracing Spanning-Tree Operations

You can enable global routing protocol tracing options at the **[edit routing-options] Hierarchy Level**. For general information about tracing and global tracing options, see the statement summary for the global *traceoptions* statement in the *Junos OS Routing Protocols Configuration Guide*.

In addition, you can enable STP-specific trace options at the following hierarchy levels:

- **[edit logical-systems logical-system-name protocols (mstp | rstp | vstp)]**
- **[edit logical-systems logical-system-name routing-instances routing-instance-name protocols (mstp | rstp | vstp)]**
- **[edit protocols (mstp | rstp | vstp)]**
- **[edit routing-instances routing-instance-name protocols (mstp | rstp | vstp)]**

The routing instance type can be either **virtual-switch** or **layer2-control**.

To enable tracing of spanning-tree protocol operations:

1. Enable configuration of the spanning-tree protocol whose operations are to be traced:

```
[edit]
user@host# edit ... protocols (mstp | rstp | vstp)
```

2. Enable configuration of spanning-tree protocol-specific trace options:

```
[edit ... protocols (mstp | rstp | vstp)]
user@host# edit traceoptions
```

3. Configure the files that contain trace logging information:

```
[edit ... protocols (mstp | rstp | vstp)]
user@host# set file filename <files number> <size bytes>
<world-readable | no-world-readable>
```

4. Configure spanning-tree protocol-specific options.

- a. To enable a spanning-tree protocol-specific option, include the
- flag**
- statement:

```
[edit ... protocols (mstp | rstp | vstp)]
user@host# set flag flag <flag-modifier> <disable>
```

You can specify the following spanning-tree protocol-specific **flag** options:

- **all**—Trace all operations.
- **all-failures**—Trace all failure conditions.
- **bpdu**—Trace BPDU reception and transmission.
- **bridge-detection-state-machine**—Trace the bridge detection state machine.
- **events**—Trace events of the protocol state machine.
- **port-information-state-machine**—Trace the port information state machine.
- **port-migration-state-machine**—Trace the port migration state machine.
- **port-receive-state-machine**—Trace the port receive state machine.
- **port-role-transit-state-machine**—Trace the port role transit state machine.
- **port-role-select-state-machine**—Trace the port role selection state machine.
- **port-transmit-state-machine**—Trace the port transmit state machine.
- **port-state-transit-state-machine**—Trace the port state transit state machine.
- **ppmd**—Trace the state and events for the ppm process.
- **state-machine-variables**—Trace when the state machine variables change.
- **timers**—Trace protocol timers.
- **topology-change-state-machine**—Trace the topology change state machine.



NOTE: Use the trace flag **all** with caution. This flag may cause the CPU to become very busy.

- b. To disable an individual spanning-tree protocol-specific option, include the
- disable**
- option with the
- flag**
- statement.

5. Verify the spanning-tree protocol-specific trace options.

```
[edit]
...
routing-options {
  traceoptions {
    ..global-trace-options-configuration...
  }
}
protocols {
  (mstp | rstp | vstp) {
```

```
traceoptions { # Spanning-tree protocol-specific.  
  file filename <files number> <size bytes> <world-readable | no-world-readable>;  
  flag flag <flag-modifier> <disable>;  
}  
}  
...  
}
```

- Related Documentation**
- [Spanning-Tree Protocols Supported on page 5032](#)
 - [Example: Tracing Spanning-Tree Protocol Operations on page 5065](#)

Operational Commands

show bridge mac-table

Syntax	<pre>show bridge mac-table <brief count detail extensive> <bridge-domain (all <i>bridge-domain-name</i>)> <global-count> <interface <i>interface-name</i>> <mac-address> <vlan-id (all-vlan <i>vlan-id</i>)></pre>
Release Information	Command introduced in Junos OS Release 8.4.
Description	(MX Series routers only) Display Layer 2 MAC address information.
Options	<p>none—Display all learned Layer 2 MAC address information.</p> <p>brief count detail extensive—(Optional) Display the specified level of output.</p> <p>bridge-domain (all <i>bridge-domain-name</i>)—(Optional) Display learned Layer 2 MAC addresses for all bridging domains or for the specified bridging domain.</p> <p>global-count—(Optional) Display the total number of learned Layer 2 MAC addresses on the system.</p> <p>instance <i>instance-name</i>—(Optional) Display learned Layer 2 MAC addresses for the specified routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Display learned Layer 2 MAC addresses for the specified interface.</p> <p>mac-address—(Optional) Display the specified learned Layer 2 MAC address information.</p> <p>vlan-id (all-vlan <i>vlan-id</i>)—(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.</p>
Additional Information	When Layer 2 protocol tunneling is enabled, the tunneling MAC address 01:00:0c:cd:cd:d0 is installed in the MAC table. When the Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP) is configured for Layer 2 protocol tunneling on an interface, the corresponding protocol MAC address is installed in the MAC table.
Required Privilege Level	view
List of Sample Output	show bridge mac-table on page 5132 show bridge mac-table brief on page 5133 show brief mac-table count on page 5133 show bridge mac-table detail on page 5133
Output Fields	Table 378 on page 5132 describes the output fields for the show bridge mac-table command. Output fields are listed in the approximate order in which they appear.

Table 378: show bridge mac-table Output fields

Field Name	Field Description
Routing instance	Name of the routing instance.
Bridging domain	Name of the bridging domain.
MAC address	MAC address or addresses learned on a logical interface.
MAC flags	Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> • S—Static MAC address is configured. • D—Dynamic MAC address is configured. • SE—MAC accounting is enabled. • NM—Non-configured MAC.
Logical interface	Name of the logical interface.
MAC count	Number of MAC addresses learned on the specific routing instance or interface.
Learning interface	Name of the logical interface on which the MAC address was learned.
Learning VLAN	VLAN ID of the routing instance or bridge domain in which the MAC address was learned.
Layer 2 flags	Debugging flags signifying that the MAC address is present in various lists.
Epoch	Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging.
Sequence number	Sequence number assigned to this MAC address. Used for debugging.
Learning mask	Mask of the Packet Forwarding Engines where this MAC address was learned. Used for debugging.
IPC generation	Creation time of the logical interface when this MAC address was learned. Used for debugging.

Sample Output

show bridge mac-table

```

user@host> show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC,
          SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vs1
Bridging domain : vlan100, VLAN : 100
  Learning MAC          MAC          Logical
  VLAN    address      flags      interface
    00:00:00:19:1c:db   D         ge-11/0/3.0
    00:00:00:59:3a:2f   D         xe-10/2/0.100

```

show bridge mac-table brief

```

user@host> show bridge mac-table brief
MAC flags (S -static MAC, D -dynamic MAC,
          SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vs1
  Bridging domain : vlan100, VLAN : 100
    Learning MAC          MAC          Logical
    VLAN      address      flags      interface
      00:00:00:19:1c:db    D          ge-11/0/3.0
      00:00:00:59:3a:2f    D          xe-10/2/0.100

```

show brief mac-table count

```

user@host> show bridge mac-table count
2 MAC address learned in routing instance vs1 bridge domain vlan100

MAC address count per interface within routing instance:
  Logical interface      MAC count
  ge-11/0/3.0            1
  ge-11/1/4.100          0
  ge-11/1/1.100          0
  ge-11/1/0.100          0
  xe-10/2/0.100          1
  xe-10/0/0.100          0

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID          MAC count
  0                       2

0 MAC address learned in routing instance vs1 bridge domain vlan200

MAC address count per interface within routing instance:
  Logical interface      MAC count
  ge-11/1/0.200          0
  ge-11/1/1.200          0
  ge-11/1/4.200          0
  xe-10/0/0.200          0
  xe-10/2/0.200          0

MAC address count per learn VLAN within routing instance:
  Learn VLAN ID          MAC count
  0                       0

```

show bridge mac-table detail

```

user@host> show bridge mac-table detail
MAC address: 00:00:00:19:1c:db
  Routing instance: vs1
  Bridging domain: vlan100
  Learning interface: ge-11/0/3.0   Learning VLAN: 0
  Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
  Epoch: 4                          Sequence number: 0
  Learning mask: 0x800              IPC generation: 0

MAC address: 00:00:00:59:3a:2f
  Routing instance: vs1
  Bridging domain: vlan100
  Learning interface: xe-10/2/0.100   Learning VLAN: 0
  Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel

```

Epoch: 7
Learning mask: 0x400

Sequence number: 0
IPC generation: 0

show mac-rewrite interface

Syntax	show mac-rewrite interface <brief detail> <interface-name>
Release Information	Command introduced in Junos OS Release 9.1.
Description	(MX Series routers only) Display Layer 2 protocol tunneling information.
Options	brief detail —(Optional) Display the specified level of output. interface <i>interface-name</i> —(Optional) Display Layer 2 protocol tunneling information for the specified interface.
Required Privilege Level	view
List of Sample Output	show mac-rewrite interface on page 5135
Output Fields	Table 379 on page 5135 lists the output fields for the show mac-rewrite interface command. Output fields are listed in the approximate order in which they appear.

Table 379: show mac-rewrite interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface that has Layer 2 protocol tunneling configured on it.	brief detail
Protocols	Layer 2 protocols being tunneled on this interface: Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP)	brief detail

Sample Output

show mac-rewrite interface

```

user@host> show mac-rewrite interface
Interface      Protocols

ge-1/0/1      STP VTP CDP

```

show spanning-tree bridge

Syntax	show spanning-tree bridge <brief detail> <msti <i>msti-id</i> > <routing-instance <i>routing-instance-name</i> > <vlan-id <i>vlan-id</i> >
Syntax (QFX Series)	show spanning-tree bridge <brief detail> <msti <i>msti-id</i> > <vlan-id <i>vlan-id</i> >
Release Information	Command introduced in Junos OS Release 8.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the configured or calculated Spanning Tree Protocol (STP) parameters.
Options	<p>none—(Optional) Display brief STP bridge information for all multiple spanning-tree instances (MSTIs).</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>msti <i>msti-id</i>—(Optional) Display STP bridge information for the specified MSTI.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) Display STP bridge information for the specified routing instance.</p> <p>vlan-id <i>vlan-id</i>—(Optional) Display STP bridge information for the specified VLAN.</p>
Required Privilege Level	view
List of Sample Output	show spanning-tree bridge routing-instance on page 5137 show spanning-tree bridge msti on page 5138 show spanning-tree bridge vlan-id (MSTP) on page 5139 show spanning-tree bridge (RSTP) on page 5139 show spanning-tree bridge vlan-id (RSTP) on page 5140
Output Fields	Table 380 on page 5136 lists the output fields for the show spanning-tree bridge command. Output fields are listed in the approximate order in which they appear.

Table 380: show spanning-tree bridge Output Fields

Field Name	Field Description
Routing instance name	Name of the routing instance under which the bridge is configured.
Enabled protocol	Spanning Tree Protocol type enabled.
Root ID	Bridge ID of the elected spanning-tree root bridge. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.

Table 380: show spanning-tree bridge Output Fields (*continued*)

Field Name	Field Description
Root cost	Calculated cost to reach the root bridge from the bridge where the command is entered.
Root port	Interface that is the current elected root port for this bridge.
CIST regional root	Bridge ID of the elected MSTP regional root bridge.
CIST internal root cost	Calculated cost to reach the regional root bridge from the bridge where the command is entered.
Hello time	Configured number of seconds between transmissions of configuration bridge protocol data units (BPDUs).
Maximum age	Configured maximum expected arrival time of hello bridge protocol data units (BPDUs).
Forward delay	How long an STP bridge port remains in the listening and learning states before transitioning to the forwarding state.
Hop count	Configured maximum number of hops a BPDUs can be forwarded in the MSTP region.
Message age	Number of elapsed seconds since the most recent BPDUs was received.
Number of topology changes	Total number of STP topology changes detected since the routing device last booted.
Time since last topology change	Number of elapsed seconds since the most recent topology change.
Bridge ID (Local)	Locally configured bridge ID. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.
Extended system ID	System identifier.
MSTI regional root	Bridge ID of the elected MSTP regional root bridge.

Sample Output

show spanning-tree bridge routing-instance

```

user@host> show spanning-tree bridge routing-instance vs1 detail
STP bridge parameters
Routing instance name      : vs1
Enabled protocol          : MSTP

STP bridge parameters for CIST
Root ID                    : 32768.00:13:c3:9e:c8:80
Root cost                  : 0

```

```
Root port : ge-10/2/0
CIST regional root : 32768.00:13:c3:9e:c8:80
CIST internal root cost : 22000
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 18
Message age : 0
Number of topology changes : 1
Time since last topology change : 1191 seconds
Local parameters
  Bridge ID : 32768.00:90:69:0b:7f:d1
  Extended system ID : 1

STP bridge parameters for MSTI 1
MSTI regional root : 32769.00:13:c3:9e:c8:80
Root cost : 22000
Root port : ge-10/2/0
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 18
Number of topology changes : 1
Time since last topology change : 1191 seconds
Local parameters
  Bridge ID : 32769.00:90:69:0b:7f:d1
  Extended system ID : 1

STP bridge parameters for MSTI 2
MSTI regional root : 32770.00:13:c3:9e:c8:80
Root cost : 22000
Root port : ge-10/2/0
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 18
Number of topology changes : 1
Time since last topology change : 1191 seconds
Local parameters
  Bridge ID : 32770.00:90:69:0b:7f:d1
  Extended system ID : 1
```

show spanning-tree bridge msti

```
user@host> show spanning-tree bridge msti 1 routing-instance vs1 detail
STP bridge parameters
Routing instance name : vs1
Enabled protocol : MSTP

STP bridge parameters for MSTI 1
MSTI regional root : 32769.00:13:c3:9e:c8:80
Root cost : 22000
Root port : xe-10/2/0
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 18
Number of topology changes : 1
Time since last topology change : 1191 seconds
Local parameters
```



```

Bridge ID                : 32769.00:90:69:0b:7f:d1
Extended system ID       : 1

```

show spanning-tree bridge vlan-id (MSTP)

```
user@host> show spanning-tree bridge vlan-id 1 101 routing-instance vs1 detail
```

STP bridge parameters

```

Routing instance name      : vs1
Enabled protocol           : MSTP

```

STP bridge parameters for CIST

```

Root ID                    : 32768.00:13:c3:9e:c8:80
Root cost                  : 0
Root port                  : xe-10/2/0
CIST regional root        : 32768.00:13:c3:9e:c8:80
CIST internal root cost   : 22000
Hello time                 : 2 seconds
Maximum age                : 20 seconds
Forward delay              : 15 seconds
Hop count                  : 18
Message age                : 0
Number of topology changes : 0

```

Local parameters

```

Bridge ID                  : 32768.00:90:69:0b:7f:d1
Extended system ID         : 1
Hello time                 : 2 seconds
Maximum age                : 20 seconds
Forward delay              : 15 seconds
Path cost method           : 32 bit
Maximum hop count          : 20

```

show spanning-tree bridge (RSTP)

```
user@host> show spanning-tree bridge
```

STP bridge parameters

```

Routing instance name      : GLOBAL
Enabled protocol           : RSTP
Root ID                    : 28672.00:90:69:0b:3f:d0
Hello time                 : 2 seconds
Maximum age                : 20 seconds
Forward delay              : 15 seconds
Message age                : 0
Number of topology changes : 58
Time since last topology change : 14127 seconds

```

Local parameters

```

Bridge ID                  : 28672.00:90:69:0b:3f:d0
Extended system ID         : 0

```

STP bridge parameters for bridge VLAN 10

```

Root ID                    : 28672.00:90:69:0b:3f:d0
Hello time                 : 2 seconds
Maximum age                : 20 seconds
Forward delay              : 15 seconds
Message age                : 0
Number of topology changes : 58
Time since last topology change : 14127 seconds

```

Local parameters

```

Bridge ID                  : 28672.00:90:69:0b:3f:d0
Extended system ID         : 0

```

STP bridge parameters for bridge VLAN 20

```
Root ID : 28672.00:90:69:0b:3f:d0
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Message age : 0
Number of topology changes : 58
Time since last topology change : 14127 seconds
Local parameters
  Bridge ID : 28672.00:90:69:0b:3f:d0
  Extended system ID : 0
```

show spanning-tree bridge vlan-id (RSTP)

```
user@host> show spanning-tree bridge vlan-id 10
STP bridge parameters
Routing instance name : GLOBAL
Enabled protocol : RSTP

STP bridge parameters for VLAN 10
Root ID : 28672.00:90:69:0b:3f:d0
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Message age : 0
Number of topology changes : 58
Time since last topology change : 14127 seconds
Local parameters
  Bridge ID : 28672.00:90:69:0b:3f:d0
  Extended system ID : 0
```

show spanning-tree interface

Syntax	show spanning-tree interface <brief detail> <msti <i>msti-id</i> > <routing-instance <i>routing-instance-name</i> > <vlan-id <i>vlan-id</i> >
Syntax (EX Series Switches and the QFX Series)	show spanning-tree interface <brief detail> <msti <i>msti-id</i> > <vlan-id <i>vlan-id</i> >
Release Information	Command introduced in Junos OS Release 8.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the configured or calculated interface-level STP parameters.
Options	<p>none—Display brief STP interface information.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>msti <i>msti-id</i>—(Optional) Display STP interface information for the specified MST instance.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) Display STP interface information for the specified routing instance.</p> <p>vlan-id <i>vlan-id</i>—(Optional) Display STP interface information for the specified VLAN.</p>
Required Privilege Level	view
List of Sample Output	show spanning-tree interface on page 5142 show spanning-tree interface (QFX Series) on page 5143 show spanning-tree interface detail on page 5143 show spanning-tree interface msti on page 5145 show spanning-tree interface vlan-id on page 5145 show spanning-tree interface (VSTP) on page 5146 show spanning-tree interface vlan-id (VSTP) on page 5146
Output Fields	Table 381 on page 5141 lists the output fields for the show spanning-tree interface command. Output fields are listed in the approximate order in which they appear.

Table 381: show spanning-tree Interface Output Fields

Field Name	Field Description
Interface name	Interface configured to participate in the STP, RSTP, VSTP, or MSTP instance.

Table 381: show spanning-tree Interface Output Fields (*continued*)

Field Name	Field Description
Port ID	Logical interface identifier configured to participate in the MSTP or VSTP instance.
Designated port ID	Port ID of the designated port for the LAN segment to which this interface is attached.
Designated bridge ID	Bridge ID of the designated bridge for the LAN segment to which this interface is attached.
Port Cost	Configured cost for the interface.
Port State	STP port state: forwarding (FWD), blocking (BLK), listening, learning, or disabled.
Port Role	MSTP, VSTP, or RSTP port role: designated (DESG), backup (BKUP), alternate (ALT), (ROOT), or Root Prevented (Root-Prev).
Link type	MSTP, VSTP, or RSTP link type. Shared or point-to-point (pt-pt) and edge or nonedge.
Alternate	Identifies the interface as an MSTP, VSTP, or RSTP alternate root port (Yes) or nonalternate root port (No).
Boundary Port	Identifies the interface as an MSTP regional boundary port (Yes) or nonboundary port (No).

Sample Output

show spanning-tree interface

```
user@host> show spanning-tree interface routing-instance vs1 detail
Spanning tree interface parameters for instance 0
```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae1	128:1	128:1	32768.0090690b47d1	1000	FWD	DESG
ge-2/1/2	128:2	128:2	32768.0090690b47d1	20000	FWD	DESG
ge-2/1/5	128:3	128:3	32768.0090690b47d1	29999	FWD	DESG
ge-2/2/1	128:4	128:26	32768.0013c39ec880	20000	FWD	ROOT
xe-9/2/0	128:5	128:5	32768.0090690b47d1	2000	FWD	DESG
xe-9/3/0	128:6	128:6	32768.0090690b47d1	2000	FWD	DESG

```
Spanning tree interface parameters for instance 1
```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae1	128:1	128:1	32769.0090690b47d1	1000	FWD	DESG
ge-2/1/2	128:2	128:2	32769.0090690b47d1	20000	FWD	DESG
ge-2/1/5	128:3	128:3	32769.0090690b47d1	29999	FWD	DESG
ge-2/2/1	128:4	128:26	32769.0013c39ec880	20000	FWD	ROOT
xe-9/2/0	128:5	128:5	32769.0090690b47d1	2000	FWD	DESG
xe-9/3/0	128:6	128:6	32769.0090690b47d1	2000	FWD	DESG

Spanning tree interface parameters for instance 2

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae1	128:1	128:1	32770.0090690b47d1	1000	FWD	DESG
ge-2/1/2	128:2	128:2	32770.0090690b47d1	20000	FWD	DESG
ge-2/1/5	128:3	128:3	32770.0090690b47d1	29999	FWD	DESG
ge-2/2/1	128:4	128:26	32770.0013c39ec880	20000	FWD	ROOT
xe-9/2/0	128:5	128:5	32770.0090690b47d1	2000	FWD	DESG
xe-9/3/0	128:6	128:6	32770.0090690b47d1	2000	FWD	DESG

show spanning-tree interface (QFX Series)

```
user@1f0> show spanning-tree interface routing-instance vs1 detail
Spanning tree interface parameters for instance 0
```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae1	128:1	128:1	32768.0090690b47d1	1000	FWD	DESG
ge-2/1/2	128:2	128:2	32768.0090690b47d1	20000	FWD	DESG
ge-2/1/5	128:3	128:3	32768.0090690b47d1	29999	FWD	DESG
ge-2/2/1	128:4	128:26	32768.0013c39ec880	20000	FWD	ROOT
xe-9/2/0	128:5	128:5	32768.0090690b47d1	2000	FWD	DESG
xe-9/3/0	128:6	128:6	32768.0090690b47d1	2000	FWD	DESG

Spanning tree interface parameters for instance 1

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae1	128:1	128:1	32769.0090690b47d1	1000	FWD	DESG
ge-2/1/2	128:2	128:2	32769.0090690b47d1	20000	FWD	DESG
ge-2/1/5	128:3	128:3	32769.0090690b47d1	29999	FWD	DESG
ge-2/2/1	128:4	128:26	32769.0013c39ec880	20000	FWD	ROOT
xe-9/2/0	128:5	128:5	32769.0090690b47d1	2000	FWD	DESG
xe-9/3/0	128:6	128:6	32769.0090690b47d1	2000	FWD	DESG

Spanning tree interface parameters for instance 2

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae1	128:1	128:1	32770.0090690b47d1	1000	FWD	DESG
ge-2/1/2	128:2	128:2	32770.0090690b47d1	20000	FWD	DESG
ge-2/1/5	128:3	128:3	32770.0090690b47d1	29999	FWD	DESG
ge-2/2/1	128:4	128:26	32770.0013c39ec880	20000	FWD	ROOT
xe-9/2/0	128:5	128:5	32770.0090690b47d1	2000	FWD	DESG
xe-9/3/0	128:6	128:6	32770.0090690b47d1	2000	FWD	DESG

show spanning-tree interface detail

```
user@host> show spanning-tree interface routing-instance vs1 detail
Spanning tree interface parameters for instance 0
```

```
Interface name           : ae1
Port identifier          : 128.1
Designated port ID      : 128.1
Port cost                : 1000
Port state               : Forwarding
Designated bridge ID     : 32768.00:90:69:0b:47:d1
Port role               : Designated
Link type                : Pt-Pt/NONEDGE
```

```
Boundary port                : No

Interface name                : ge-2/1/2
Port identifier               : 128.2
Designated port ID           : 128.2
Port cost                     : 20000
Port state                    : Forwarding
Designated bridge ID         : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                : No

Interface name                : ge-2/1/5
Port identifier               : 128.3
Designated port ID           : 128.3
Port cost                     : 29999
Port state                    : Forwarding
Designated bridge ID         : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                : No

Interface name                : ge-2/2/1
Port identifier               : 128.4
Designated port ID           : 128.26
Port cost                     : 20000
Port state                    : Forwarding
Designated bridge ID         : 32768.00:13:c3:9e:c8:80
Port role                     : Root
Link type                     : Pt-Pt/NONEDGE
Boundary port                : No

Interface name                : xe-9/2/0
Port identifier               : 128.5
Designated port ID           : 128.5
Port cost                     : 2000
Port state                    : Forwarding
Designated bridge ID         : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                : No

Interface name                : xe-9/3/0
Port identifier               : 128.6
Designated port ID           : 128.6
Port cost                     : 2000
Port state                    : Forwarding
Designated bridge ID         : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                : No
```

Spanning tree interface parameters for instance 1

```
Interface name                : ae1
Port identifier               : 128.1
Designated port ID           : 128.1
Port cost                     : 1000
Port state                    : Forwarding
Designated bridge ID         : 32768.00:90:69:0b:47:d1
```

```

Port role           : Designated
Link type           : Pt-Pt/NONEDGE
Boundary port       : No

Interface name       : ge-2/1/2
Port identifier      : 128.2
Designated port ID   : 128.2
Port cost            : 20000
Port state           : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role           : Designated
Link type           : Pt-Pt/NONEDGE
Boundary port       : No

Interface name       : ge-2/1/5
Port identifier      : 128.3
Designated port ID   : 128.3
Port cost            : 29999
Port state           : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role           : Designated
Link type           : Pt-Pt/NONEDGE
Boundary port       : No

Interface name       : ge-2/2/1
Port identifier      : 128.4
Designated port ID   : 128.26
Port cost            : 20000
Port state           : Forwarding
Designated bridge ID : 32768.00:13:c3:9e:c8:80
Port role           : Root
Link type           : Pt-Pt/NONEDGE
Boundary port       : No

...

```

show spanning-tree interface msti

```

user@host> show spanning-tree interface msti 1 routing-instance vs1 detail
Spanning tree interface parameters for instance 1

```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-7/0/0	128:1	128:1	32769.0090690b4fd1	2000	FWD	DESG
ge-5/1/0	128:2	128:2	32769.0090690b4fd1	20000	FWD	DESG
ge-5/1/1	128:3	128:3	32769.0090690b4fd1	20000	FWD	DESG
ae1	128:4	128:1	32769.0090690b47d1	10000	BLK	ALT
ge-5/1/4	128:5	128:3	32769.0090690b47d1	20000	BLK	ALT
xe-7/2/0	128:6	128:6	32769.0090690b47d1	2000	FWD	ROOT

show spanning-tree interface vlan-id

```

user@host> show spanning-tree interface vlan-id 101 routing-instance vs1 detail
Spanning tree interface parameters for instance 0

```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ge-11/0/5	128:1	128:1	32768.0090690b7fd1	20000	FWD	DESG
ge-11/0/6	128:2	128:1	32768.0090690b7fd1	20000	BLK	BKUP
ge-11/1/0	128:3	128:2	32768.0090690b4fd1	20000	BLK	ALT
ge-11/1/1	128:4	128:3	32768.0090690b4fd1	20000	BLK	ALT

ge-11/1/4	128:5	128:1	32768.0090690b47d1	20000	BLK	ALT
xe-10/0/0	128:6	128:5	32768.0090690b4fd1	2000	BLK	ALT
xe-10/2/0	128:7	128:4	32768.0090690b47d1	2000	FWD	ROOT

show spanning-tree interface (VSTP)

```
user@host> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0
```

Interface	Port ID	Designated port ID	Designated bridge ID	Cost	State	Role
ge-1/0/1	128:1	128:1	28672.0090690b3fe0	20000	FWD	DESG
ge-1/0/2	128:2	128:2	28672.0090690b3fe0	20000	FWD	DESG

```
Spanning tree interface parameters for VLAN 10
```

Interface	Port ID	Designated port ID	Designated bridge ID	Cost	State	Role
ge-1/0/1	128:1	128:1	28672.0090690b3fe0	20000	FWD	DESG
ge-1/0/2	128:2	128:2	28672.0090690b3fe0	20000	FWD	DESG

```
Spanning tree interface parameters for VLAN 20
```

Interface	Port ID	Designated port ID	Designated bridge ID	Cost	State	Role
ge-1/0/1	128:1	128:1	28672.0090690b3fe0	20000	FWD	DESG
ge-1/0/2	128:2	128:2	28672.0090690b3fe0	20000	FWD	DESG

show spanning-tree interface vlan-id (VSTP)

```
user@host> show spanning-tree interface vlan-id 10
```

```
Spanning tree interface parameters for VLAN 10
```

Interface	Port ID	Designated port ID	Designated bridge ID	Cost	State	Role
ge-1/0/1	128:1	128:1	28672.0090690b3fe0	20000	FWD	DESG
ge-1/0/2	128:2	128:2	28672.0090690b3fe0	20000	FWD	DESG

show spanning-tree mstp configuration

Syntax	show spanning-tree mstp configuration <brief detail> <routing-instance <i>routing-instance-name</i> >
Syntax (EX Series Switch and the QFX Series)	show spanning-tree mstp configuration <brief detail>
Release Information	Command introduced in Junos OS Release 8.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the MSTP configuration.
Options	none —Display MSTP configuration information. brief detail —(Optional) Display the specified level of output. routing-instance <i>routing-instance-name</i> —(Optional) Display MSTP configuration information for the specified routing instance.
Required Privilege Level	view
List of Sample Output	show spanning-tree mstp configuration detail on page 5148 show spanning-tree mstp configuration detail (QFX Series) on page 5148
Output Fields	Table 382 on page 5147 lists the output fields for the show spanning-tree mstp configuration command. Output fields are listed in the approximate order in which they appear.

Table 382: show spanning-tree mstp configuration Output Fields

Field Name	Field Description
Context id	Internally generated identifier.
Region name	MSTP region name carried in the MSTP BPDUs.
Revision	Revision number of the MSTP configuration.
Configuration digest	Numerical value derived from the VLAN-to-instance mapping table.
MSTI	MST instance identifier.
Member VLANs	VLAN identifiers associated with the MSTI.

Sample Output

show spanning-tree mstp configuration detail

```
user@host> show spanning-tree mstp configuration routing-instance vs1 detail
MSTP configuration information
Context identifier      : 1
Region name            : henry
Revision               : 3
Configuration digest    : 0x6da4b5c4fd587757eef35675365e1

MSTI      Member VLANs
  0 0-99,101-199,201-4094
  1 100
  2 200
```

show spanning-tree mstp configuration detail (QFX Series)

```
user@1f0> show spanning-tree mstp configuration routing-instance vs1 detail
MSTP configuration information
Context identifier      : 1
Region name            : henry
Revision               : 3
Configuration digest    : 0x6da4b5c4fd587757eef35675365e1

MSTI      Member VLANs
  0 0-99,101-199,201-4094
  1 100
  2 200
```

show spanning-tree statistics

Syntax	show spanning-tree statistics <brief detail> <interface <i>interface-name</i> > <routing-instance <i>routing-instance-name</i> >
Syntax (EX Series Switch and the QFX Series)	show spanning-tree statistics <brief detail> <interface <i>interface-name</i> vlan <i>vlan-id</i> >
Release Information	Command introduced in Junos OS Release 8.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for QFX Series switches.
Description	Display STP statistics.
Options	<p>none—Display brief STP statistics.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i>—(Optional) Display STP statistics for the specified interface.</p> <p>routing-instance <i>routing-instance-name</i>—(Optional) Display STP statistics for the specified routing instance.</p>
Required Privilege Level	view
List of Sample Output	show spanning-tree statistics routing-instance on page 5150 show spanning-tree statistics interface routing-instance detail on page 5150
Output Fields	Table 383 on page 5149 lists the output fields for the show spanning-tree statistics command. Output fields are listed in the approximate order in which they appear.

Table 383: show spanning-tree statistics Output Fields

Field Name	Field Description
Message type	Type of message being counted.
BPDUs sent	Total number of BPDUs sent.
BPDUs received	Total number of BPDUs received.
BPDUs sent in last interval	Number of BPDUs sent within a specified interval.
BPDUs received in last interval	Number of BPDUs received within a specified interval.
Interface	Interface for which the statistics are being displayed.

Table 383: show spanning-tree statistics Output Fields (*continued*)

Field Name	Field Description
Next BPDU transmission	Number of seconds until the next BPDU is scheduled to be sent.

Sample Output

show spanning-tree statistics routing-instance

```

user@host> show spanning-tree statistics routing-instance vs1 detail
Routing instance level STP statistics
Message type           : bpdus
BPDUs sent             : 1396
BPDUs received         : 1027
BPDUs sent in last interval : 5      (duration: 4 sec)
BPDUs received in last interval: 4    (duration: 4 sec)

```

show spanning-tree statistics interface routing-instance detail

```

user@host> show spanning-tree statistics interface ge-11/1/4 routing-instance vs1 detail
Interface  BPDUs sent  BPDUs received  Next BPDU
                                     transmission
ge-11/1/4      7           190           0

```

CHAPTER 20

VPNs

- [Multicast VPNs on page 5151](#)
- [VPLS on page 5299](#)

Multicast VPNs

- [Overview on page 5151](#)
- [Configuration on page 5162](#)
- [Troubleshooting on page 5298](#)

Overview

- [Multicast VPNs on page 5151](#)
- [Layer 3 VPNs on page 5158](#)

Multicast VPNs

- [MBGP Multicast VPN Sites on page 5151](#)
- [PIM Sparse Mode, PIM Dense Mode, Auto-RP, and BSR for MBGP MVPNs on page 5153](#)
- [Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 5153](#)
- [Multicast VPN Terminology on page 5158](#)
- [Supported Multicast VPN Standards on page 5158](#)

MBGP Multicast VPN Sites

The main characteristics of MBGP MVPNs are:

- They extend Layer 3 VPN service (RFC 4364) to support IP multicast for Layer 3 VPN service providers.
- They follow the same architecture as specified by RFC 4364 for unicast VPNs. Specifically, BGP is used as the provider edge (PE) router-to-PE router control plane for multicast VPN.

- They eliminate the requirement for the virtual router (VR) model (as specified in Internet draft draft-rosen-vpn-mcast, *Multicast in MPLS/BGP VPNs*) for multicast VPNs and the RFC 4364 model for unicast VPNs.
- They rely on RFC 4364-based unicast with extensions for intra-AS and inter-AS communication.

An MBGP MVPN defines two types of site sets, a sender site set and a receiver site set. These sites have the following properties:

- Hosts within the sender site set can originate multicast traffic for receivers in the receiver site set.
- Receivers outside the receiver site set should not be able to receive this traffic.
- Hosts within the receiver site set can receive multicast traffic originated by any host in the sender site set.
- Hosts within the receiver site set should not be able to receive multicast traffic originated by any host that is not in the sender site set.

A site can be in both the sender site set and the receiver site set, so hosts within such a site can both originate and receive multicast traffic. For example, the sender site set could be the same as the receiver site set, in which case all sites could both originate and receive multicast traffic from one another.

Sites within a given MBGP MVPN might be within the same organization or in different organizations, which means that an MBGP MVPN can be either an intranet or an extranet. A given site can be in more than one MBGP MVPN, so MBGP MVPNs might overlap. Not all sites of a given MBGP MVPN have to be connected to the same service provider, meaning that an MBGP MVPN can span multiple service providers. Feature parity for the MVPN extranet functionality or overlapping MVPNs on the Junos Trio chipset is supported in Junos OS Releases 11.1R2, 11.2R2, and 11.4.

Another way to look at an MBGP MVPN is to say that an MBGP MVPN is defined by a set of administrative policies. These policies determine both the sender site set and the receiver site set. These policies are established by MBGP MVPN customers, but implemented by service providers using the existing BGP and MPLS VPN infrastructure.

**Related
Documentation**

- *Example: Allowing MBGP MVPN Remote Sources*
- *Example: Configuring a PIM-SSM Provider Tunnel for an MBGP MVPN*
- *Example: Configuring MBGP Multicast VPN Extranets*

PIM Sparse Mode, PIM Dense Mode, Auto-RP, and BSR for MBGP MVPNs

You can configure PIM sparse mode, PIM dense mode, auto-RP, and bootstrap router (BSR) for MBGP MVPN networks:

- PIM sparse mode—Allows a router to use any unicast routing protocol and performs reverse-path forwarding (RPF) checks using the unicast routing table. PIM sparse mode includes an explicit join message, so routers determine where the interested receivers are and send join messages upstream to their neighbors, building trees from the receivers to the rendezvous point (RP).
- PIM dense mode—Allows a router to use any unicast routing protocol and performs reverse-path forwarding (RPF) checks using the unicast routing table. Packets are forwarded to all interfaces except the incoming interface. Unlike PIM sparse mode, where explicit joins are required for packets to be transmitted downstream, packets are flooded to all routers in the routing instance in PIM dense mode.
- Auto-RP—Uses PIM dense mode to propagate control messages and establish RP mapping. You can configure an auto-RP node in one of three different modes: discovery mode, announce mode, and mapping mode.
- BSR—Establishes RPs. A selected router in a network acts as a BSR, which selects a unique RP for different group ranges. BSR messages are flooded using a data tunnel between PE routers.

Related Documentation

- *Example: Allowing MBGP MVPN Remote Sources*
- *Example: Configuring a PIM-SSM Provider Tunnel for an MBGP MVPN*
- *Example: Configuring MBGP Multicast VPN Extranets*

Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN

Selective LSPs are also referred to as selective provider tunnels. Selective provider tunnels carry traffic from some multicast groups in a VPN and extend only to the PE routers that have receivers for these groups. You can configure a selective provider tunnel for group prefixes and source prefixes, or you can use wildcards for the group and source, as described in the Internet draft draft-rekhter-mvpn-wildcard-spmsi-01.txt, *Use of Wildcard in S-PMSI Auto-Discovery Routes*.

The following sections describe the scenarios and special considerations when you use wildcards for selective provider tunnels.

- [About S-PMSI on page 5154](#)
- [Scenarios for Using Wildcard S-PMSI on page 5155](#)
- [Types of Wildcard S-PMSI on page 5155](#)
- [Differences Between Wildcard S-PMSI and \(S,G\) S-PMSI on page 5155](#)
- [Wildcard \(*,*\) S-PMSI and PIM Dense Mode on page 5156](#)
- [Wildcard \(*,*\) S-PMSI and PIM-BSR on page 5156](#)
- [Wildcard Source and the 0.0.0.0/0 Source Prefix on page 5157](#)

About S-PMSI


The provider multicast service interface (PMSI) is a BGP tunnel attribute that contains the tunnel ID used by the PE router for transmitting traffic through the core of the provider network. A selective PMSI (S-PMSI) autodiscovery route advertises binding of a given MVPN customer multicast flow to a particular provider tunnel. The S-PMSI autodiscovery route advertised by the ingress PE router contains /32 IPv4 or /128 IPv6 addresses for the customer source and the customer group derived from the source-tree customer multicast route.

Figure 112 on page 5154 shows a simple MVPN topology. The ingress router, PE1, originates the S-PMSI autodiscovery route. The egress routers, PE2 and PE3, have join state as a result of receiving join messages from CE devices that are not shown in the topology. In response to the S-PMSI autodiscovery route advertisement sent by PE1, PE2, and PE3, elect whether or not to join the tunnel based on the join state. The selective provider tunnel configuration is configured in a VRF instance on PE1.



NOTE: The MVPN mode configuration (RPT-SPT or SPT-only) is configured on all three PE routers for all VRFs that make up the VPN. If you omit the MVPN mode configuration, the default mode is SPT-only.

Figure 112: Simple MVPN Topology



Scenarios for Using Wildcard S-PMSI

A wildcard S-PMSI has the source or the group (or both the source and the group) field set to the wildcard value of 0.0.0.0/0 and advertises binding of multiple customer multicast flows to a single provider tunnel in a single S-PMSI autodiscovery route.

The scenarios under which you might configure a wildcard S-PMSI are as follows:

- When the customer multicast flows are PIM-SM in ASM-mode flows. In this case, a PE router connected to an MVPN customer's site that contains the customer's RP (C-RP) could bind all the customer multicast flows traveling along a customer's RPT tree to a single provider tunnel.
- When a PE router is connected to an MVPN customer's site that contains multiple sources, all sending to the same group.
- When the customer multicast flows are PIM-bidirectional flows. In this case, a PE router could bind to a single provider tunnel all the customer multicast flows for the same group that have been originated within the sites of a given MVPN connected to that PE, and advertise such binding in a single S-PMSI autodiscovery route.
- When the customer multicast flows are PIM-SM in SSM-mode flows. In this case, a PE router could bind to a single provider tunnel all the customer multicast flows coming from a given source located in a site connected to that PE router.
- When you want to carry in the provider tunnel all the customer multicast flows originated within the sites of a given MVPN connected to a given PE router.

Types of Wildcard S-PMSI

The following types of wildcard S-PMSI are supported:

- A (*G) S-PMSI matches all customer multicast routes that have the group address. The customer source address in the customer multicast route can be any address, including 0.0.0.0/0 for shared-tree customer multicast routes. A (*, C-G) S-PMSI autodiscovery route is advertised with the source field set to 0 and the source address length set to 0. The multicast group address for the S-PMSI autodiscovery route is derived from the customer multicast joins.
- A (*,*) S-PMSI matches all customer multicast routes. Any customer source address and any customer group address in a customer multicast route can be bound to the (*,*) S-PMSI. The S-PMSI autodiscovery route is advertised with the source address and length set to 0 and the group address and length set 0. The remaining fields in the S-PMSI autodiscovery route follow the same rule as (C-S, C-G) S-PMSI, as described in section 12.1 of the BGP-MVPN draft (draft-ietf-l3vpn-2547bis-mcast-bgp-00.txt).

Differences Between Wildcard S-PMSI and (S,G) S-PMSI

For dynamic provider tunnels, each customer multicast stream is bound to a separate provider tunnel, and each tunnel is advertised by a separate S-PMSI autodiscovery route. For static LSPs, multiple customer multicast flows are bound to a single provider tunnel by having multiple S-PMSI autodiscovery routes advertise the same provider tunnel.

When you configure a wildcard (*;G) or (*;*) S-PMSI, one or more matching customer multicast routes share a single S-PMSI. All customer multicast routes that have a matching source and group address are bound to the same (*;G) or (*;*) S-PMSI and share the same tunnel. The (*;G) or (*;*) S-PMSI is established when the first matching remote customer multicast join message is received in the ingress PE router, and deleted when the last remote customer multicast join is withdrawn from the ingress PE router. Sharing a single S-PMSI autodiscovery route improves control plane scalability.

Wildcard (*;*) S-PMSI and PIM Dense Mode

For (S,G) and (*;G) S-PMSI autodiscovery routes in PIM dense mode (PIM-DM), all downstream PE routers receive PIM-DM traffic. If a downstream PE router does not have receivers that are interested in the group address, the PE router instantiates prune state and stops receiving traffic from the tunnel.

Now consider what happens for (*;*) S-PMSI autodiscovery routes. If the PIM-DM traffic is not bound by a longer matching (S,G) or (*;G) S-PMSI, it is bound to the (*;*) S-PMSI. As is always true for dense mode, PIM-DM traffic is flooded to downstream PE routers over the provider tunnel regardless of the customer multicast join state. Because there is no group information in the (*;*) S-PMSI autodiscovery route, egress PE routers join a (*;*) S-PMSI tunnel if there is any configuration on the egress PE router indicating interest in PIM-DM traffic.

Interest in PIM-DM traffic is indicated if the egress PE router has one of the following configurations in the VRF instance that corresponds to the instance that imports the S-PMSI autodiscovery route:

- At least one interface is configured in dense mode at the **[edit routing-instances instance-name protocols pim interface]** hierarchy level.
- At least one group is configured as a dense-mode group at the **[edit routing-instances instance-name protocols pim dense-groups group-address]** hierarchy level.

Wildcard (*;*) S-PMSI and PIM-BSR

For (S,G) and (*;G) S-PMSI autodiscovery routes in PIM bootstrap router (PIM-BSR) mode, an ingress PE router floods the PIM bootstrap message (BSM) packets over the provider tunnel to all egress PE routers. An egress PE router does not join the tunnel unless the message has the ALL-PIM-ROUTERS group. If the message has this group, the egress PE router joins the tunnel, regardless of the join state. The group field in the message determines the presence or absence of the ALL-PIM-ROUTERS address.

Now consider what would happen for (*;*) S-PMSI autodiscovery routes used with PIM-BSR mode. If the PIM BSM packets are not bound by a longer matching (S,G) or (*;G) S-PMSI, they are bound to the (*;*) S-PMSI. As is always true for PIM-BSR, BSM packets are flooded to downstream PE routers over the provider tunnel to the ALL-PIM-ROUTERS destination group. Because there is no group information in the (*;*) S-PMSI autodiscovery route, egress PE routers always join a (*;*) S-PMSI tunnel. Unlike PIM-DM, the egress PE routers might have no configuration suggesting use of PIM-BSR as the RP discovery mechanism in the VRF instance. To prevent all egress PE routers from always joining the (*;*) S-PMSI tunnel, the (*;*) wildcard group configuration must be ignored.

This means that if you configure PIM-BSR, a wildcard-group S-PMSI can be configured for all other group addresses. The (*,*) S-PMSI is not used for PIM-BSR traffic. Either a matching (*G) or (S,G) S-PMSI (where the group address is the ALL-PIM-ROUTERS group) or an inclusive provider tunnel is needed to transmit data over the provider core. For PIM-BSR, the longest-match lookup is (S,G), (*G), and the inclusive provider tunnel, in that order. If you do not configure an inclusive tunnel for the routing instance, you must configure a (*G) or (S,G) selective tunnel. Otherwise, the data is dropped. This is because PIM-BSR functions like PIM-DM, in that traffic is flooded to downstream PE routers over the provider tunnel regardless of the customer multicast join state. However, unlike PIM-DM, the egress PE routers might have no configuration to indicate interest or noninterest in PIM-BSR traffic.

Wildcard Source and the 0.0.0.0/0 Source Prefix

You can configure a 0.0.0.0/0 source prefix and a wildcard source under the same group prefix in a selective provider tunnel. For example, the configuration might look as follows:

```
routing-instances {
  vpn {
    provider-tunnel {
      selective {
        group 224.1.1.0/24 {
          source 0.0.0.0/0 {
            rsvp-te {
              label-switched-path-template {
                sptnl3;
              }
            }
          }
        }
        wildcard-source {
          rsvp-te {
            label-switched-path-template {
              sptnl2;
            }
            static-lsp point-to-multipoint-lsp-name;
          }
          threshold-rate kbps;
        }
      }
    }
  }
}
```

The functions of the **source 0.0.0.0/0** and **wildcard-source** configuration statements are different. The 0.0.0.0/0 source prefix only matches (C-S, C-G) customer multicast join messages and triggers (C-S, C-G) S-PMSI autodiscovery routes derived from the customer multicast address. Because all (C-S, C-G) join messages are matched by the 0.0.0.0/0 source prefix in the matching group, the wildcard source S-PMSI is used only for (*C-G) customer multicast join messages. In the absence of a configured 0.0.0.0/0 source prefix, the wildcard source matches (C-S, C-G) and (*C-G) customer multicast join messages. In the example, a join message for (10.0.1.0/24, 224.1.1.0/24) is bound to **sptnl3**. A join message for (*, 224.1.1.0/24) is bound to **sptnl2**.

- Related Documentation**
- [Configuring a Selective Provider Tunnel Using Wildcards on page 5240](#)
 - [Example: Configuring Selective Provider Tunnels Using Wildcards on page 5241](#)
 - [Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs on page 5225](#)
 - [Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs on page 5227](#)

Multicast VPN Terminology

I

- Inclusive tree** A single multicast distribution tree in the backbone that carries all the multicast traffic from a specified set of one or more multicast VPNs. An inclusive tree that carries the traffic of more than one multicast VPN is an aggregate inclusive tree. An inclusive tree contains as its members all the PE routers that attach to the receiver sites of any of the multicast VPNs using the tree.

S

- Selective tree** A single multicast distribution tree in the backbone that carries traffic belonging only to a specified set of one or more multicast groups, from one or more multicast VPNs. An aggregate selective tree carries traffic for multicast groups that belong to different multicast VPNs. By default, traffic from most multicast groups could be carried by an inclusive tree, whereas traffic from high-bandwidth groups should be carried by a selective tree.

Supported Multicast VPN Standards

Junos OS substantially supports the following Internet drafts, which define standards for multicast virtual private networks (VPNs).

- Internet draft draft-ietf-l3vpn-2547bis-mcast-10.txt, *Multicast in MPLS/BGP IP VPNs*
- Internet draft draft-ietf-l3vpn-2547bis-mcast-bgp-08.txt, *BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs*

- Related Documentation**
- *Supported Carrier-of-Carriers and Interprovider VPN Standards*
 - *Supported Layer 2 Circuit Standards*
 - *Supported Layer 2 VPN Standard*
 - *Supported Layer 3 VPN Standards*
 - [Supported VPLS Standards on page 5432](#)
 - *Supported MPLS Standards*
 - *Supported BGP Standards*
 - *Accessing Standards Documents on the Internet*

Layer 3 VPNs

- [Introduction to Configuring Layer 3 VPNs on page 5159](#)
- [Layer 3 VPN Platform Support on page 5161](#)

Introduction to Configuring Layer 3 VPNs

To configure Layer 3 virtual private network (VPN) functionality, you must enable VPN support on the provider edge (PE) router. You must also configure any provider (P) routers that service the VPN, and you must configure the customer edge (CE) routers so that their routes are distributed into the VPN.

To configure Layer 3 VPNs, you include the following statements:

```

description text;
instance-type vrf;
interface interface-name;
protocols {
  bgp {
    group group-name {
      peer-as as-number;
      neighbor ip-address;
    }
    multihop tth-value;
  }
  (ospf | ospf3) {
    area area {
      interface interface-name;
    }
    domain-id domain-id;
    domain-vpn-tag number;
    sham-link {
      local address;
    }
    sham-link-remote address <metric number>;
  }
  rip {
    rip-configuration;
  }
}
route-distinguisher (as-number:id | ip-address:id);
router-id address;
routing-options {
  autonomous-system autonomous-system {
    independent-domain;
    loops number;
  }
  forwarding-table {
    export [ policy-names ];
  }
  interface-routes {
    rib-group group-name;
  }
  martians {
    destination-prefix match-type <allow>;
  }
  maximum-paths {
    path-limit;
    log-interval interval;
    log-only;
  }
}

```

```
        threshold percentage;
    }
    maximum-prefixes {
        prefix-limit;
        log-interval interval;
        log-only;
        threshold percentage;
    }
    multipath {
        vpn-unequal-cost;
    }
    options {
        syslog (level level | upto level);
    }
    rib routing-table-name {
        martians {
            destination-prefix match-type <allow>;
        }
        multipath {
            vpn-unequal-cost;
        }
        static {
            defaults {
                static-options;
            }
            route destination-prefix {
                next-hop [next-hops];
                static-options;
            }
        }
    }
}
static {
    defaults {
        static-options;
    }
    route destination-prefix {
        policy [policy-names ];
        static-options;
    }
}
vrf-advertise-selective {
    family {
        inet-mvpn;
        inet6-mvpn;
    }
}
vrf-export [policy-names ];
vrf-import [policy-names ];
vrf-target (community | export community-name | import community-name);
vrf-table-label;
```

You can include these statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]

- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]**

For Layer 3 VPNs, only some of the statements in the **[edit routing-instances]** hierarchy are valid. For the full hierarchy, see the *Junos OS Routing Protocols Configuration Guide*.

In addition to these statements, you must enable a signaling protocol, IBGP sessions between the PE routers, and an interior gateway protocol (IGP) on the PE and P routers.

By default, Layer 3 VPNs are disabled.

Many of the configuration procedures for Layer 3 VPNs are common to all types of VPNs.

Related Documentation

- *Centralized Internet Access*
- *Configuring Hub-and-Spoke VPN Topologies: One Interface*
- *Configuring Hub-and-Spoke VPN Topologies: Two Interfaces*
- *Configuring Overlapping VPNs Using Automatic Route Export*
- *Configuring Overlapping VPNs Using Routing Table Groups*
- *Configuring a Full-Mesh VPN Topology with Route Reflectors*
- *Configuring a GRE Tunnel Interface Between PE Routers*
- *Configuring a GRE Tunnel Interface Between a PE and CE Router*
- *Configuring a Simple Full-Mesh VPN Topology*
- *Configuring an Application-Based Layer 3 VPN Topology*
- *Configuring an ES Tunnel Interface Between a PE and CE Router*
- *Configuring an LDP-over-RSVP VPN Topology*
- *Configuring an OSPF Domain ID for a Layer 3 VPN*
- *Distributed Internet Access*
- *Routing Internet Traffic Through a Separate NAT Device*
- *Routing VPN and Internet Traffic Through Different Interfaces*
- *Routing VPN and Internet Traffic Through the Same Interface Bidirectionally (VPN Has Private Addresses)*
- *Routing VPN and Internet Traffic Through the Same Interface Bidirectionally (VPN Has Public Addresses)*
- *Routing VPN and Outgoing Internet Traffic Through the Same Interface and Routing Return Internet Traffic Through a Different Interface*
- *Setting the Forwarding Class of the Ping Packets*

Layer 3 VPN Platform Support

Layer 3 VPNs are supported on most combinations of Juniper Networks routing and switching platforms and PICs capable of running the JUNOS Software.

MX Series routers configured to be in Ethernet services mode can support some of the Junos OS Layer 3 VPN features. For Layer 3 VPNs, Ethernet services mode supports configuring a loopback interface for a VPN routing and forwarding (VRF) instance. You can configure up to two VRF instances in Ethernet services mode. Each VRF instance can handle up to 10,000 routes. The **ping mpls l3vpn** operational mode command is also supported.

Configuration

- [Configuration Examples on page 5162](#)
- [Configuration Tasks on page 5221](#)
- [Configuration Statements on page 5250](#)

Configuration Examples

- [Example: Configuring MBGP Multicast VPNs on page 5162](#)
- [Example: Configuring PIM Join Load Balancing on Draft-Rosen Multicast VPN on page 5182](#)
- [Example: Configuring PIM Join Load Balancing On Next-Generation Multicast VPN on page 5190](#)
- [Example: Configuring PIM State Limits on page 5198](#)
- [Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs on page 5211](#)

Example: Configuring MBGP Multicast VPNs

This example provides a step-by-step procedure to configure multicast services across a multiprotocol BGP (MBGP) Layer 3 virtual private network.

- [Requirements on page 5162](#)
- [Overview and Topology on page 5163](#)
- [Configuration on page 5163](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.2 or later
- Five M Series, T Series, TX Series, or MX Series Juniper routers
- One host system capable of sending multicast traffic and supporting the Internet Group Management Protocol (IGMP)
- One host systems capable of receiving multicast traffic and supporting IGMP

Depending on the devices you are using, you might be required to configure static routes to:

- The multicast sender
- The Fast Ethernet interface to which the sender is connected on the multicast receiver
- The multicast receiver
- The Fast Ethernet interface to which the receiver is connected on the multicast sender

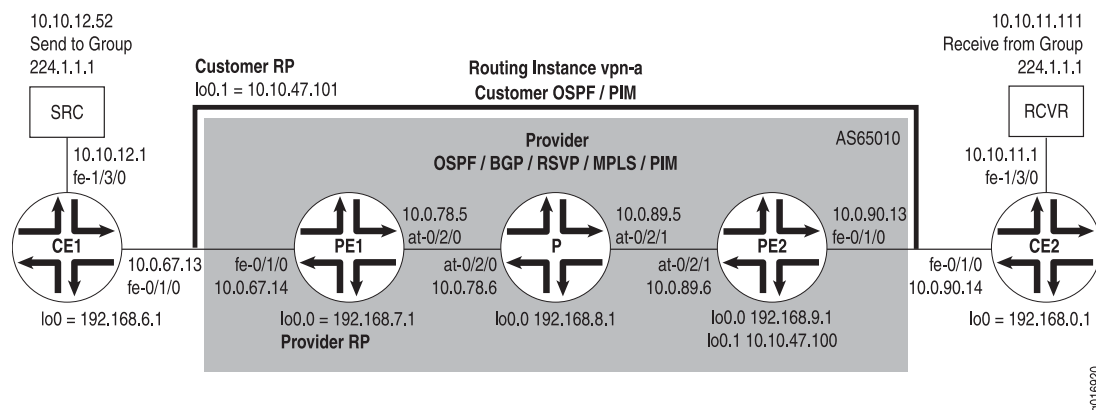
Overview and Topology

This example shows how to configure the following technologies:

- IPv4
- BGP
- OSPF
- RSVP
- MPLS
- PIM sparse mode
- Static RP

The topology of the network is shown in [Figure 113 on page 5163](#).

Figure 113: Multicast Over Layer 3 VPN Example Topology



Configuration



NOTE: In any configuration session, it is a good practice to periodically verify that the configuration can be committed using the `commit check` command.

In this example, the router being configured is identified using the following command prompts:

- **CE1** identifies the customer edge 1 (CE1) router
- **PE1** identifies the provider edge 1 (PE1) router
- **P** identifies the provider core (P) router
- **CE2** identifies the customer edge 2 (CE2) router
- **PE2** identifies the provider edge 2 (PE2) router

To configure MBGP multicast VPNs for the network shown in [Figure 113 on page 5163](#), perform the following steps:

- [Configuring Interfaces on page 5164](#)
- [Configuring OSPF on page 5165](#)
- [Configuring BGP on page 5166](#)
- [Configuring RSVP on page 5167](#)
- [Configuring MPLS on page 5168](#)
- [Configuring the VRF Routing Instance on page 5168](#)
- [Configuring PIM on page 5170](#)
- [Configuring the Provider Tunnel on page 5171](#)
- [Configuring the Rendezvous Point on page 5171](#)
- [Results on page 5172](#)

Configuring Interfaces

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

1. On each router, configure an IP address on the loopback logical interface 0 (**lo0.0**).

```
[edit interfaces]
```

```
user@CE1# set lo0 unit 0 family inet address 192.168.6.1/32 primary
```

```
user@PE1# set lo0 unit 0 family inet address 192.168.7.1/32 primary
```

```
user@P# set lo0 unit 0 family inet address 192.168.8.1/32 primary
```

```
user@PE2# set lo0 unit 0 family inet address 192.168.9.1/32 primary
```

```
user@CE2# set lo0 unit 0 family inet address 192.168.0.1/32 primary
```

Use the **show interfaces terse** command to verify that the IP address is correct on the loopback logical interface.

2. On the PE and CE routers, configure the IP address and protocol family on the Fast Ethernet interfaces. Specify the **inet** protocol family type.

```
[edit interfaces]
```

```
user@CE1# set fe-1/3/0 unit 0 family inet address 10.10.12.1/24
```

```
user@CE1# set fe-0/1/0 unit 0 family inet address 10.0.67.13/30
```

```
[edit interfaces]
```

```
user@PE1# set fe-0/1/0 unit 0 family inet address 10.0.67.14/30
```

```
[edit interfaces]
```

```
user@PE2# set fe-0/1/0 unit 0 family inet address 10.0.90.13/30
```

```
[edit interfaces]
```

```
user@CE2# set fe-0/1/0 unit 0 family inet address 10.0.90.14/30
user@CE2# set fe-1/3/0 unit 0 family inet address 10.10.11.1/24
```

Use the **show interfaces terse** command to verify that the IP address is correct on the Fast Ethernet interfaces.

3. On the PE and P routers, configure the ATM interfaces' VPI and maximum virtual circuits. If the default PIC type is different on directly connected ATM interfaces, configure the PIC type to be the same. Configure the logical interface VCI, protocol family, local IP address, and destination IP address.

```
[edit interfaces]
user@PE1# set at-0/2/0 atm-options pic-type atm1
user@PE1# set at-0/2/0 atm-options vpi 0 maximum-vcs 256
user@PE1# set at-0/2/0 unit 0 vci 0.128
user@PE1# set at-0/2/0 unit 0 family inet address 10.0.78.5/32 destination 10.0.78.6
```

```
[edit interfaces]
user@P# set at-0/2/0 atm-options pic-type atm1
user@P# set at-0/2/0 atm-options vpi 0 maximum-vcs 256
user@P# set at-0/2/0 unit 0 vci 0.128
user@P# set at-0/2/0 unit 0 family inet address 10.0.78.6/32 destination 10.0.78.5
user@P# set at-0/2/1 atm-options pic-type atm1
user@P# set at-0/2/1 atm-options vpi 0 maximum-vcs 256
user@P# set at-0/2/1 unit 0 vci 0.128
user@P# set at-0/2/1 unit 0 family inet address 10.0.89.5/32 destination 10.0.89.6
```

```
[edit interfaces]
user@PE2# set at-0/2/1 atm-options pic-type atm1
user@PE2# set at-0/2/1 atm-options vpi 0 maximum-vcs 256
user@PE2# set at-0/2/1 unit 0 vci 0.128
user@PE2# set at-0/2/1 unit 0 family inet address 10.0.89.6/32 destination 10.0.89.5
```

Use the **show configuration interfaces** command to verify that the ATM interfaces' VPI and maximum VCs are correct and that the logical interface VCI, protocol family, local IP address, and destination IP address are correct.

Configuring OSPF

- Step-by-Step Procedure**
1. On the P and PE routers, configure the provider instance of OSPF. Specify the **lo0.0** and ATM core-facing logical interfaces. The provider instance of OSPF on the PE router forms adjacencies with the OSPF neighbors on the other PE router and Router P.

```
user@PE1# set protocols ospf area 0.0.0.0 interface at-0/2/0.0
user@PE1# set protocols ospf area 0.0.0.0 interface lo0.0
```

```
user@P# set protocols ospf area 0.0.0.0 interface lo0.0
user@P# set protocols ospf area 0.0.0.0 interface all
user@P# set protocols ospf area 0.0.0.0 interface fxp0 disable
```

```
user@PE2# set protocols ospf area 0.0.0.0 interface lo0.0
user@PE2# set protocols ospf area 0.0.0.0 interface at-0/2/1.0
```

Use the **show ospf interfaces** command to verify that the **lo0.0** and ATM core-facing logical interfaces are configured for OSPF.

2. On the CE routers, configure the customer instance of OSPF. Specify the loopback and Fast Ethernet logical interfaces. The customer instance of OSPF on the CE routers form adjacencies with the neighbors within the VPN routing instance of OSPF on the PE routers.

```
user@CE1# set protocols ospf area 0.0.0.0 interface fe-0/1/0.0
user@CE1# set protocols ospf area 0.0.0.0 interface fe-1/3/0.0
user@CE1# set protocols ospf area 0.0.0.0 interface lo0.0
```

```
user@CE2# set protocols ospf area 0.0.0.0 interface fe-0/1/0.0
user@CE2# set protocols ospf area 0.0.0.0 interface fe-1/3/0.0
user@CE2# set protocols ospf area 0.0.0.0 interface lo0.0
```

Use the **show ospf interfaces** command to verify that the correct loopback and Fast Ethernet logical interfaces have been added to the OSPF protocol.

3. On the P and PE routers, configure OSPF traffic engineering support for the provider instance of OSPF.

The **shortcuts** statement enables the master instance of OSPF to use a label-switched path as the next hop.

```
user@PE1# set protocols ospf traffic-engineering shortcuts
```

```
user@P# set protocols ospf traffic-engineering shortcuts
```

```
user@PE2# set protocols ospf traffic-engineering shortcuts
```

Use the **show ospf overview** or **show configuration protocols ospf** command to verify that traffic engineering support is enabled.

Configuring BGP

Step-by-Step Procedure

1. On Router P, configure BGP for the VPN. The local address is the local **lo0.0** address. The neighbor addresses are the PE routers' **lo0.0** addresses.

The **unicast** statement enables the router to use BGP to advertise network layer reachability information (NLRI). The **signaling** statement enables the router to use BGP as the signaling protocol for the VPN.

```
user@P# set protocols bgp group group-mvpn type internal
user@P# set protocols bgp group group-mvpn local-address 192.168.8.1
user@P# set protocols bgp group group-mvpn family inet unicast
user@P# set protocols bgp group group-mvpn family inet-mvpn signaling
user@P# set protocols bgp group group-mvpn neighbor 192.168.9.1
user@P# set protocols bgp group group-mvpn neighbor 192.168.7.1
```

Use the **show configuration protocols bgp** command to verify that the router has been configured to use BGP to advertise NLRI.

2. On the PE and P routers, configure the BGP local autonomous system number.

```
user@PE1# set routing-options autonomous-system 0.65010
```

```
user@P# set routing-options autonomous-system 0.65010
```

```
user@PE2# set routing-options autonomous-system 0.65010
```

Use the **show configuration routing-options** command to verify that the BGP local autonomous system number is correct.

3. On the PE routers, configure BGP for the VPN. Configure the local address as the local **lo0.0** address. The neighbor addresses are the **lo0.0** addresses of Router P and the other PE router, PE2.

```
user@PE1# set protocols bgp group group-mvpn type internal
user@PE1# set protocols bgp group group-mvpn local-address 192.168.7.1
user@PE1# set protocols bgp group group-mvpn family inet-vpn unicast
user@PE1# set protocols bgp group group-mvpn family inet-mvpn signaling
user@PE1# set protocols bgp group group-mvpn neighbor 192.168.9.1
user@PE1# set protocols bgp group group-mvpn neighbor 192.168.8.1
```

```
user@PE2# set protocols bgp group group-mvpn type internal
user@PE2# set protocols bgp group group-mvpn local-address 192.168.9.1
user@PE2# set protocols bgp group group-mvpn family inet-vpn unicast
user@PE2# set protocols bgp group group-mvpn family inet-mvpn signaling
user@PE2# set protocols bgp group group-mvpn neighbor 192.168.7.1
user@PE2# set protocols bgp group group-mvpn neighbor 192.168.8.1
```

Use the **show bgp group** command to verify that the BGP configuration is correct.

4. On the PE routers, configure a policy to export the BGP routes into OSPF.

```
user@PE1# set policy-options policy-statement bgp-to-ospf from protocol bgp
user@PE1# set policy-options policy-statement bgp-to-ospf then accept
```

```
user@PE2# set policy-options policy-statement bgp-to-ospf from protocol bgp
user@PE2# set policy-options policy-statement bgp-to-ospf then accept
```

Use the **show policy bgp-to-ospf** command to verify that the policy is correct.

Configuring RSVP

Step-by-Step Procedure

1. On the PE routers, enable RSVP on the interfaces that participate in the LSP. Configure the Fast Ethernet and ATM logical interfaces.

```
user@PE1# set protocols rsvp interface fe-0/1/0.0
user@PE1# set protocols rsvp interface at-0/2/0.0
```

```
user@PE2# set protocols rsvp interface fe-0/1/0.0
user@PE2# set protocols rsvp interface at-0/2/1.0
```

2. On Router P, enable RSVP on the interfaces that participate in the LSP. Configure the ATM logical interfaces.

```
user@P# set protocols rsvp interface at-0/2/0.0
user@P# set protocols rsvp interface at-0/2/1.0
```

Use the **show configuration protocols rsvp** command to verify that the RSVP configuration is correct.

Configuring MPLS

Step-by-Step Procedure

1. On the PE routers, configure an MPLS LSP to the PE router that is the LSP egress point. Specify the IP address of the **lo0.0** interface on the router at the other end of the LSP. Configure MPLS on the ATM, Fast Ethernet, and **lo0.0** interfaces.

To help identify each LSP when troubleshooting, configure a different LSP name on each PE router. In this example, we use the name **to-pe2** as the name for the LSP configured on PE1 and **to-pe1** as the name for the LSP configured on PE2.

```
user@PE1# set protocols mpls label-switched-path to-pe2 to 192.168.9.1
user@PE1# set protocols mpls interface fe-0/1/0.0
user@PE1# set protocols mpls interface at-0/2/0.0
user@PE1# set protocols mpls interface lo0.0
```

```
user@PE2# set protocols mpls label-switched-path to-pe1 to 192.168.7.1
user@PE2# set protocols mpls interface fe-0/1/0.0
user@PE2# set protocols mpls interface at-0/2/1.0
user@PE2# set protocols mpls interface lo0.0
```

Use the **show configuration protocols mpls** and **show route label-switched-path to-pe1** commands to verify that the MPLS and LSP configuration is correct.

After the configuration is committed, use the **show mpls lsp name to-pe1** and **show mpls lsp name to-pe2** commands to verify that the LSP is operational.

2. On Router P, enable MPLS. Specify the ATM interfaces connected to the PE routers.

```
user@P# set protocols mpls interface at-0/2/0.0
user@P# set protocols mpls interface at-0/2/1.0
```

Use the **show mpls interface** command to verify that MPLS is enabled on the ATM interfaces.

3. On the PE and P routers, configure the protocol family on the ATM interfaces associated with the LSP. Specify the **mpls** protocol family type.

```
user@PE1# set interfaces at-0/2/0 unit 0 family mpls
```

```
user@P# set interfaces at-0/2/0 unit 0 family mpls
user@P# set interfaces at-0/2/1 unit 0 family mpls
```

```
user@PE2# set interfaces at-0/2/1 unit 0 family mpls
```

Use the **show mpls interface** command to verify that the MPLS protocol family is enabled on the ATM interfaces associated with the LSP.

Configuring the VRF Routing Instance

Step-by-Step Procedure

1. On the PE routers, configure a routing instance for the VPN and specify the **vrf** instance type. Add the Fast Ethernet and **lo0.1** customer-facing interfaces. Configure the VPN instance of OSPF and include the BGP-to-OSPF export policy.

```
user@PE1# set routing-instances vpn-a instance-type vrf
user@PE1# set routing-instances vpn-a interface lo0.1
user@PE1# set routing-instances vpn-a interface fe-0/1/0.0
```

```
user@PE1# set routing-instances vpn-a protocols ospf export bgp-to-ospf
user@PE1# set routing-instances vpn-a protocols ospf area 0.0.0.0 interface all
```

```
user@PE2# set routing-instances vpn-a instance-type vrf
user@PE2# set routing-instances vpn-a interface lo0.1
user@PE2# set routing-instances vpn-a interface fe-0/1/0.0
user@PE2# set routing-instances vpn-a protocols ospf export bgp-to-ospf
user@PE2# set routing-instances vpn-a protocols ospf area 0.0.0.0 interface all
```

Use the **show configuration routing-instances vpn-a** command to verify that the routing instance configuration is correct.

2. On the PE routers, configure a route distinguisher for the routing instance. A route distinguisher allows the router to distinguish between two identical IP prefixes used as VPN routes. Configure a different route distinguisher on each PE router. This example uses 65010:1 on PE1 and 65010:2 on PE2.

```
user@PE1# set routing-instances vpn-a route-distinguisher 65010:1
```

```
user@PE2# set routing-instances vpn-a route-distinguisher 65010:2
```

Use the **show configuration routing-instances vpn-a** command to verify that the route distinguisher is correct.

3. On the PE routers, configure default VRF import and export policies. Based on this configuration, BGP automatically generates local routes corresponding to the route target referenced in the VRF import policies. This example uses 2:1 as the route target.



NOTE: You must configure the same route target on each PE router for a given VPN routing instance.

```
user@PE1# set routing-instances vpn-a vrf-target target:2:1
```

```
user@PE2# set routing-instances vpn-a vrf-target target:2:1
```

Use the **show configuration routing-instances vpn-a** command to verify that the route target is correct.

4. On the PE routers, configure the VPN routing instance for multicast support.

```
user@PE1# set routing-instances vpn-a protocols mvpn
```

```
user@PE2# set routing-instances vpn-a protocols mvpn
```

Use the **show configuration routing-instance vpn-a** command to verify that the VPN routing instance has been configured for multicast support.

5. On the PE routers, configure an IP address on loopback logical interface 1 (lo0.1) used in the customer routing instance VPN.

```
user@PE1# set interfaces lo0 unit 1 family inet address 10.10.47.101/32
```

```
user@PE2# set interfaces lo0 unit 1 family inet address 10.10.47.100/32
```

Use the **show interfaces terse** command to verify that the IP address on the loopback interface is correct.

Configuring PIM

Step-by-Step Procedure

1. On the PE and P routers, enable the provider instance of PIM. Add the core-facing ATM interfaces. On the PE routers, also configure the **lo0.0** interface. Specify the mode as **sparse** and the version as **2**.

```
user@PE1# set protocols pim interface at-0/2/0.0 mode sparse
user@PE1# set protocols pim interface at-0/2/0.0 version 2
user@PE1# set protocols pim interface lo0.0 mode sparse
user@PE1# set protocols pim interface lo0.0 version 2
```

```
user@P# set protocols pim interface at-0/2/0.0 mode sparse
user@P# set protocols pim interface at-0/2/0.0 version 2
user@P# set protocols pim interface at-0/2/1.0 mode sparse
user@P# set protocols pim interface at-0/2/1.0 version 2
```

```
user@PE2# set protocols pim interface at-0/2/1.0 mode sparse
user@PE2# set protocols pim interface at-0/2/1.0 version 2
user@PE2# set protocols pim interface lo0.0 mode sparse
user@PE2# set protocols pim interface lo0.0 version 2
```

Use the **show pim interfaces** command to verify that PIM sparse-mode is enabled on the core-facing ATM interfaces.

2. On the PE routers, enable the VPN customer instance of PIM. Configure the **lo0.1** and the customer-facing Fast Ethernet interface. Specify the mode as **sparse** and the version as **2**.

```
user@PE1# set routing-instances vpn-a protocols pim interface lo0.1 mode sparse
user@PE1# set routing-instances vpn-a protocols pim interface lo0.1 version 2
user@PE1# set routing-instances vpn-a protocols pim interface fe-0/1/0.0 mode
sparse
user@PE1# set routing-instances vpn-a protocols pim interface fe-0/1/0.0 version
2
```

```
user@PE2# set routing-instances vpn-a protocols pim interface lo0.1 mode sparse
user@PE2# set routing-instances vpn-a protocols pim interface lo0.1 version 2
user@PE2# set routing-instances vpn-a protocols pim interface fe-0/1/0.0 mode
sparse
user@PE2# set routing-instances vpn-a protocols pim interface fe-0/1/0.0 version
2
```

Use the **show pim interfaces instance vpn-a** command to verify that PIM sparse-mode is enabled on the **lo0.1** interface and the customer-facing Fast Ethernet interface.

3. On the CE routers, enable the customer instance of PIM. In this example, we configure all interfaces. Specify the mode as **sparse** and the version as **2**.

```
user@CE1# set protocols pim interface all
```

```
user@CE2# set protocols pim interface all mode sparse
```



```
user@CE2# set protocols pim interface all version 2
```

Use the **show pim interfaces** command to verify that PIM sparse mode is enabled on all interfaces.

Configuring the Provider Tunnel

Step-by-Step Procedure

1. On Router PE1, configure the provider tunnel. Specify the multicast address to be used.

The **provider-tunnel** statement instructs the router to send multicast traffic across a tunnel. The **pim-asm** statement instructs the router to accept the multicast stream from any source.

```
user@PE1# set routing-instances vpn-a provider-tunnel pim-asm group-address 224.1.1.1
```

Use the **show configuration routing-instance vpn-a** command to verify that the multicast group address is correct on Router PE1.

2. On Router PE2, configure the provider tunnel. Specify the multicast address to be used.

```
user@PE2# set routing-instances vpn-a provider-tunnel pim-asm group-address 224.1.1.1
```

Use the **show configuration routing-instance vpn-a** command to verify that the multicast group address is correct on Router PE2.

Configuring the Rendezvous Point

Step-by-Step Procedure

1. Configure Router PE1 to be the rendezvous point for the provider instance of PIM. Specify the **lo0.0** address of Router PE1.

```
user@PE1# set protocols pim rp local address 192.168.7.1
```

Use the **show pim rps** command to verify that the correct local IP address is configured for the provider instance RP.

2. Configure the static rendezvous point on Router P and the PE2 router for the provider instance of PIM. Specify the **lo0.0** address of Router PE1. Specify the version as 2.

```
user@P# set protocols pim rp static address 192.168.7.1 version 2
```

```
user@PE2# set protocols pim rp static address 192.168.7.1 version 2
```

Use the **show pim rps** command to verify that the correct static IP address is configured for the provider instance RP.

3. Configure Router PE1 to be the rendezvous point for the customer instance of PIM. Specify the **lo0.1** address of Router PE1. Specify the multicast address to be used.

```
user@PE1# set routing-instances vpn-a protocols pim rp local address 10.10.47.101
user@PE1# set routing-instances vpn-a protocols pim rp local group-ranges 224.1.1.1/32
```

Use the **show pim rps instance vpn-a** command to verify that the correct local IP address is configured for the customer instance RP.

4. On Router PE2, configure the static rendezvous point for the customer instance of PIM. Specify the **lo0.1** address of Router PE1.

```
user@PE2# set routing-instances vpn-a protocols pim rp static address 10.10.47.101
```

Use the **show pim rps instance vpn-a** command to verify that the correct static IP address is configured for the customer instance RP.

5. On the CE routers, configure the static rendezvous point for the customer instance of PIM. Specify the **lo0.1** address of Router PE1.

```
user@CE1# set protocols pim rp static address 10.10.47.101 version 2
```

```
user@CE2# set protocols pim rp static address 10.10.47.101 version 2
```

Use the **show pim rps** command to verify that the correct static IP address is configured for the customer instance RP.

6. Use the **commit check** command to verify that the configuration can be successfully committed. If the configuration passes the check, commit the configuration.
7. Start the multicast sender device connected to CE1.
8. Start the multicast receiver device connected to CE2.
9. Verify that the receiver is receiving the multicast stream.
10. Use **show** commands to verify the routing, VPN, and multicast operation.

Results

The configuration and verification parts of this example have been completed. The following section is for your reference.

The relevant sample configuration for Router CE1 follows.

```
Router CE1 interfaces {
  lo0 {
    unit 0 {
      family inet {
        address 192.168.6.1/32 {
          primary;
        }
      }
    }
  }
  fe-0/1/0 {
    unit 0 {
      family inet {
        address 10.0.67.13/30;
      }
    }
  }
  fe-1/3/0 {
    unit 0 {
      family inet {
        address 10.10.12.1/24;
      }
    }
  }
}
```

```

    }
  }
}
protocols {
  ospf {
    area 0.0.0.0 {
      interface fe-0/1/0.0;
      interface lo0.0;
      interface fe-1/3/0.0;
    }
  }
  pim {
    rp {
      static {
        address 10.10.47.101 {
          version 2;
        }
      }
    }
    interface all;
  }
}

```

The relevant sample configuration for Router PE1 follows.

```

Router PE1  interfaces {
              lo0 {
                unit 0 {
                  family inet {
                    address 192.168.7.1/32 {
                      primary;
                    }
                  }
                }
              }
              fe-0/1/0 {
                unit 0 {
                  family inet {
                    address 10.0.67.14/30;
                  }
                }
              }
              at-0/2/0 {
                atm-options {
                  pic-type atm1;
                  vpi 0 {
                    maximum-vcs 256;
                  }
                }
                unit 0 {
                  vci 0.128;
                  family inet {
                    address 10.0.78.5/32 {
                      destination 10.0.78.6;
                    }
                  }
                }
              }
            }

```

```
        family mpls;
      }
    }
    lo0 {
      unit 1 {
        family inet {
          address 10.10.47.101/32;
        }
      }
    }
  }
  routing-options {
    autonomous-system 0.65010;
  }
  protocols {
    rsvp {
      interface fe-0/1/0.0;
      interface at-0/2/0.0;
    }
    mpls {
      label-switched-path to-pe2 {
        to 192.168.9.1;
      }
      interface fe-0/1/0.0;
      interface at-0/2/0.0;
      interface lo0.0;
    }
    bgp {
      group group-mvpn {
        type internal;
        local-address 192.168.7.1;
        family inet-vpn {
          unicast;
        }
        family inet-mvpn {
          signaling;
        }
        neighbor 192.168.9.1;
        neighbor 192.168.8.1;
      }
    }
    ospf {
      traffic-engineering {
        shortcuts;
      }
      area 0.0.0.0 {
        interface at-0/2/0.0;
        interface lo0.0;
      }
    }
    pim {
      rp {
        local {
          address 192.168.7.1;
        }
      }
    }
  }
}
```

```

        interface at-0/2/0.0 {
            mode sparse;
            version 2;
        }
        interface lo0.0 {
            mode sparse;
            version 2;
        }
    }
}
policy-options {
    policy-statement bgp-to-ospf {
        from protocol bgp;
        then accept;
    }
}
routing-instances {
    vpn-a {
        instance-type vrf;
        interface lo0.1;
        interface fe-0/1/0.0;
        route-distinguisher 65010:1;
        provider-tunnel {
            pim-asm {
                group-address 224.1.1.1;
            }
        }
        vrf-target target:2:1;
        protocols {
            ospf {
                export bgp-to-ospf;
                area 0.0.0.0 {
                    interface all;
                }
            }
            pim {
                rp {
                    local {
                        address 10.10.47.101;
                        group-ranges {
                            224.1.1.1/32;
                        }
                    }
                }
            }
            interface lo0.1 {
                mode sparse;
                version 2;
            }
            interface fe-0/1/0.0 {
                mode sparse;
                version 2;
            }
        }
        mvpn;
    }
}

```

```
}
```

The relevant sample configuration for Router P follows.

```
Router P  interfaces {
            lo0 {
              unit 0 {
                family inet {
                  address 192.168.8.1/32 {
                    primary;
                  }
                }
              }
            }
            at-0/2/0 {
              atm-options {
                pic-type atm1;
              }
              vpi 0 {
                maximum-vcs 256;
              }
            }
            unit 0 {
              vci 0.128;
              family inet {
                address 10.0.78.6/32 {
                  destination 10.0.78.5;
                }
              }
              family mpls;
            }
          }
          at-0/2/1 {
            atm-options {
              pic-type atm1;
            }
            vpi 0 {
              maximum-vcs 256;
            }
          }
          unit 0 {
            vci 0.128;
            family inet {
              address 10.0.89.5/32 {
                destination 10.0.89.6;
              }
            }
            family mpls;
          }
        }
        routing-options {
          autonomous-system 0.65010;
        }
        protocols {
          rsvp {
            interface at-0/2/0.0;
            interface at-0/2/1.0;
```

```

}
mpls {
  interface at-0/2/0.0;
  interface at-0/2/1.0;
}
bgp {
  group group-mvpn {
    type internal;
    local-address 192.168.8.1;
    family inet {
      unicast;
    }
    family inet-mvpn {
      signaling;
    }
    neighbor 192.168.9.1;
    neighbor 192.168.7.1;
  }
}
ospf {
  traffic-engineering {
    shortcuts;
  }
  area 0.0.0.0 {
    interface lo0.0;
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}
pim {
  rp {
    static {
      address 192.168.7.1 {
        version 2;
      }
    }
  }
  interface at-0/2/0.0 {
    mode sparse;
    version 2;
  }
  interface at-0/2/1.0 {
    mode sparse;
    version 2;
  }
}
}

```

The relevant sample configuration for Router PE2 follows.

```

Router PE2  interfaces {
              lo0 {
                unit 0 {
                  family inet {

```

```
        address 192.168.9.1/32 {
            primary;
        }
    }
}
fe-0/1/0 {
    unit 0 {
        family inet {
            address 10.0.90.13/30;
        }
    }
}
at-0/2/1 {
    atm-options {
        pic-type atm1;
        vpi 0 {
            maximum-vcs 256;
        }
    }
    unit 0 {
        vci 0.128;
        family inet {
            address 10.0.89.6/32 {
                destination 10.0.89.5;
            }
        }
        family mpls;
    }
}
lo0 {
    unit 1 {
        family inet {
            address 10.10.47.100/32;
        }
    }
}
}
routing-options {
    autonomous-system 0.65010;
}
protocols {
    rsvp {
        interface fe-0/1/0.0;
        interface at-0/2/1.0;
    }
    mpls {
        label-switched-path to-pe1 {
            to 192.168.7.1;
        }
        interface lo0.0;
        interface fe-0/1/0.0;
        interface at-0/2/1.0;
    }
}
bgp {
    group group-mvpn {
```



```

        type internal;
        local-address 192.168.9.1;
        family inet-vpn {
            unicast;
        }
        family inet-mvpn {
            signaling;
        }
        neighbor 192.168.7.1;
        neighbor 192.168.8.1;
    }
}
ospf {
    traffic-engineering {
        shortcuts;
    }
    area 0.0.0.0 {
        interface lo0.0;
        interface at-0/2/1.0;
    }
}
pim {
    rp {
        static {
            address 192.168.7.1 {
                version 2;
            }
        }
    }
    interface lo0.0 {
        mode sparse;
        version 2;
    }
    interface at-0/2/1.0 {
        mode sparse;
        version 2;
    }
}
}
policy-options {
    policy-statement bgp-to-ospf {
        from protocol bgp;
        then accept;
    }
}
routing-instances {
    vpn-a {
        instance-type vrf;
        interface fe-0/1/0.0;
        interface lo0.1;
        route-distinguisher 65010:2;
        provider-tunnel {
            pim-asm {
                group-address 224.1.1.1;
            }
        }
    }
}

```

```
vrf-target target:2:1;
protocols {
  ospf {
    export bgp-to-ospf;
    area 0.0.0.0 {
      interface all;
    }
  }
  pim {
    rp {
      static {
        address 10.10.47.101;
      }
    }
    interface fe-0/1/0.0 {
      mode sparse;
      version 2;
    }
    interface lo0.1 {
      mode sparse;
      version 2;
    }
  }
  mvpn;
}
}
```

The relevant sample configuration for Router CE2 follows.

```
Router CE2 interfaces {
  lo0 {
    unit 0 {
      family inet {
        address 192.168.0.1/32 {
          primary;
        }
      }
    }
  }
  fe-0/1/0 {
    unit 0 {
      family inet {
        address 10.0.90.14/30;
      }
    }
  }
  fe-1/3/0 {
    unit 0 {
      family inet {
        address 10.10.11.1/24;
      }
      family inet6 {
        address fe80::205:85ff:fe88:cdb/64;
      }
    }
  }
}
```

```

    }
  }
  protocols {
    ospf {
      area 0.0.0.0 {
        interface fe-0/1/0.0;
        interface lo0.0;
        interface fe-1/3/0.0;
      }
    }
    pim {
      rp {
        static {
          address 10.10.47.101 {
            version 2;
          }
        }
      }
      interface all {
        mode sparse;
        version 2;
      }
    }
  }
}

```

Related Documentation

- *Multicast over Layer 3 VPNs Overview*
- *Configuring BGP, MPLS, RSVP, and an IGP on the PE and Core Routers for Draft Rosen VPNs*
- *Configuring BGP, MPLS, RSVP, and an IGP on the PE and Core Routers for MBGP MVPNs*
- *Configuring Interfaces for Layer 3 VPNs*
- *Configuring Intra-AS Inclusive Point-to-Multipoint Traffic Engineering LSPs*
- *Configuring Intra-AS Selective Provider Tunnels*
- *Configuring MBGP MVPNs to Support IPv6 Multicast Traffic*
- *Configuring PIM and the VPN Group Address in a Routing Instance*
- *Configuring Provider Tunnels*
- *Configuring the Master PIM Instance on the PE Router for BGP-Based Multicast VPNs*
- *Configuring the Master PIM Instance on the PE Router in the Service Provider Network*
- *Creating a Routing Instance for a Multiprotocol BGP-Based Multicast VPN*
- *Configuring the Router's IPv4 Bootstrap Router Priority*
- *Creating a Unique Logical Loopback Interface for the Routing Instance for Draft Rosen VPNs*
- *Creating a Unique Logical Loopback Interface for the Routing Instance for MBGP MVPNs*
- *Dual PIM Draft-Rosen Multicast VPN Operation*
- *Enabling Multicast VPN in BGP*

- *MBGP Multicast VPN Extranets Configuration Guidelines*
- *Option: Configuring MSDP Within a Layer 3 VPN*
- *Option: Configuring Multicast Distribution Trees for Data*
- *Option: Configuring PIM Sparse Mode Graceful Restart for a Layer 3 VPN*
- *Option: Configuring Sender and Receiver Sites*
- *Option: Specifying Route Targets*
- *Understanding MBGP Multicast VPN Extranets*
- *Understanding Multiprotocol BGP-Based Multicast VPNs: Next-Generation*

Example: Configuring PIM Join Load Balancing on Draft-Rosen Multicast VPN

This example shows how to configure multipath routing for external and internal virtual private network (VPN) routes with unequal interior gateway protocol (IGP) metrics, and Protocol Independent Multicast (PIM) join load balancing on provider edge (PE) routers running Draft-Rosen multicast VPN (MVPN). This feature allows customer PIM (C-PIM) join messages to be load-balanced across external and internal BGP (EIBGP) upstream paths when the PE router has both external BGP (EBGP) and internal BGP (IBGP) paths toward the source or rendezvous point (RP).

- [Requirements on page 5182](#)
- [Overview and Topology on page 5183](#)
- [Configuration on page 5186](#)
- [Verification on page 5189](#)

Requirements

This example requires the following hardware and software components:

- Three routers that can be a combination of M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, or T Series Core Routers.
- Junos OS Release 12.1 or later running on all the devices.

Before you begin:

1. Configure the device interfaces.
2. Configure the following routing protocols on all PE routers:
 - OSPF
 - MPLS
 - LDP
 - PIM
 - BGP
3. Configure a multicast VPN.

Overview and Topology

Junos OS Release 12.1 and later support multipath configuration along with PIM join load balancing. This allows C-PIM join messages to be load-balanced across unequal EIBGP routes, if a PE router has EBGP and IBGP paths toward the source (or RP). In previous releases, only the active EBGP path was used to send the join messages. This feature is applicable to IPv4 C-PIM join messages.

During load balancing, if a PE router loses one or more EBGP paths toward the source (or RP), the C-PIM join messages that were previously using the EBGP path are moved to a multicast tunnel interface, and the reverse path forwarding (RPF) neighbor on the multicast tunnel interface is selected based on a hash mechanism.

On discovering the first EBGP path toward the source (or RP), only the new join messages get load-balanced across EIBGP paths, whereas the existing join messages on the multicast tunnel interface remain unaffected.

Though the primary goal for multipath PIM join load balancing is to utilize unequal EIBGP paths for multicast traffic, potential join loops can be avoided if a PE router chooses only the EBGP path when there are one or more join messages for different groups from a remote PE router. If the remote PE router's join message arrives after the PE router has already chosen IBGP as the upstream path, then the potential loops can be broken by changing the selected upstream path to EBGP.



NOTE: During a graceful Routing Engine switchover (GRES), the EIBGP path selection for C-PIM join messages can vary, because the upstream interface selection is performed again for the new Routing Engine based on the join messages it receives from the CE and PE neighbors. This can lead to disruption of multicast traffic depending on the number of join messages received and the load on the network at the time of the graceful restart. However, the nonstop active routing feature is not supported and has no impact on the multicast traffic in a Draft-Rosen MVPN scenario.

In this example, PE1 and PE2 are the upstream PE routers for which the multipath PIM join load-balancing feature is configured. Routers PE1 and PE2 have one EBGP path and one IBGP path each toward the source. The Source and Receiver attached to customer edge (CE) routers are Free BSD hosts.

On PE routers that have EIBGP paths toward the source (or RP), such as PE1 and PE2, PIM join load balancing is performed as follows:

1. The existing join-count-based load balancing is performed such that the algorithm first selects the least loaded C-PIM interface. If there is equal or no load on all the C-PIM interfaces, the join messages get distributed equally across the available upstream interfaces.

In [Figure 114 on page 5186](#), if the PE1 router receives PIM join messages from the CE2 router, and if there is equal or no load on both the EBGp and IBGP paths toward the source, the join messages get load-balanced on the EIBGP paths.

2. If the selected least loaded interface is a multicast tunnel interface, then there can be a potential join loop if the downstream list of the customer join (C-join) message already contains the multicast tunnel interface. In such a case, the least loaded interface among EBGp paths is selected as the upstream interface for the C-join message.

Assuming that the IBGP path is the least loaded, the PE1 router sends the join messages to PE2 using the IBGP path. If PIM join messages from the PE3 router arrive on PE1, then the downstream list of the C-join messages for PE3 already contains a multicast tunnel interface, which can lead to a potential join loop, because both the upstream and downstream interfaces are multicast tunnel interfaces. In this case, PE1 uses only the EBGp path to send the join messages.

3. If the selected least loaded interface is a multicast tunnel interface and the multicast tunnel interface is not present in the downstream list of the C-join messages, the loop prevention mechanism is not necessary. If any PE router has already advertised data multicast distribution tree (MDT) type, length, and values (TLVs), that PE router is selected as the upstream neighbor.

When the PE1 router sends the join messages to PE2 using the least loaded IBGP path, and if PE3 sends its join messages to PE2, no join loop is created.

4. If no data MDT TLV corresponds to the C-join message, the least loaded neighbor on a multicast tunnel interface is selected as the upstream interface.

On PE routers that have only IBGP paths toward the source (or RP), such as PE3, PIM join load balancing is performed as follows:

1. The PE router only finds a multicast tunnel interface as the RPF interface, and load balancing is done across the C-PIM neighbors on a multicast tunnel interface.

Router PE3 load-balances PIM join messages received from the CE4 router across the IBGP paths to the PE1 and PE2 routers.

2. If any PE router has already advertised data MDT TLVs corresponding to the C-join messages, that PE router is selected as the RPF neighbor.

For a particular C-multicast flow, at least one of the PE routers having EIBGP paths toward the source (or RP) must use only the EBGp path to avoid or break join loops. As a result of the loop avoidance mechanism, a PE router is constrained to choose among EIBGP paths when a multicast tunnel interface is already present in the downstream list.

In [Figure 114 on page 5186](#), assuming that the CE2 host is interested in receiving traffic from the Source and CE2 initiates multiple PIM join messages for different groups (Group 1 with group address 225.1.1.1, and Group 2 with group address 225.1.1.2), the join messages for both groups arrive on the PE1 router.

Router PE1 then equally distributes the join messages between the EIBGP paths toward the Source. Assuming that Group 1 join messages are sent to the CE1 router directly using the EBGp path, and Group 2 join messages are sent to the PE2 router using the IBGP path, PE1 and PE2 become the RPF neighbors for Group 1 and Group 2 join messages, respectively.

When the CE3 router initiates Group 1 and Group 2 PIM join messages, the join messages for both groups arrive on the PE2 router. Router PE2 then equally distributes the join messages between the EIBGP paths toward the Source. Since PE2 is the RPF neighbor for Group 2 join messages, it sends the Group 2 join messages directly to the CE1 router using the EBGp path. Group 1 join messages are sent to the PE1 router using the IBGP path.

However, if the CE4 router initiates multiple Group 1 and Group 2 PIM join messages, there is no control over how these join messages received on the PE3 router get distributed to reach the Source. The selection of the RPF neighbor by PE3 can affect PIM join load balancing on EIBGP paths.


- If PE3 sends Group 1 join messages to PE1 and Group 2 join messages to PE2, there is no change in RPF neighbor. As a result, no join loops are created.
- If PE3 sends Group 1 join messages to PE2 and Group 2 join messages to PE1, there is a change in the RPF neighbor for the different groups resulting in the creation of join loops. To avoid potential join loops, PE1 and PE2 do not consider IBGP paths to send the join messages received from the PE3 router. Instead, the join messages are sent directly to the CE1 router using only the EBGp path.

The loop avoidance mechanism in a Draft-Rosen MVPN has the following limitations:

- Because the timing of arrival of join messages on remote PE routers determines the distribution of join messages, the distribution could be sub-optimal in terms of join count.
- Because join loops cannot be avoided and can occur due to the timing of join messages, the subsequent RPF interface change leads to loss of multicast traffic. This can be avoided by implementing the PIM make-before-break feature.

The PIM make-before-break feature is an approach to detect and break C-PIM join loops in a Draft-Rosen MVPN. The C-PIM join messages are sent to the new RPF neighbor after establishing the PIM neighbor relationship, but before updating the related multicast forwarding entry. Though the upstream RPF neighbor would have updated its multicast forwarding entry and started sending the multicast traffic downstream, the downstream router does not forward the multicast traffic (because of RPF check failure) until the multicast forwarding entry is updated with the new RPF neighbor. This helps to ensure that the multicast traffic is available on the new path before switching the RPF interface of the multicast forwarding entry.

Figure 114: PIM Join Load Balancing on Draft-Rosen MVPN



Configuration

CLI Quick Configuration

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

```

PE1 set routing-instances vpn1 instance-type vrf
    set routing-instances vpn1 interface ge-5/0/4.0
    set routing-instances vpn1 interface ge-5/2/0.0
    set routing-instances vpn1 interface lo0.1
    set routing-instances vpn1 route-distinguisher 1:1
    set routing-instances vpn1 vrf-target target:1:1
    set routing-instances vpn1 routing-options multipath vpn-unequal-cost
        equal-external-internal
    set routing-instances vpn1 protocols bgp export direct
    set routing-instances vpn1 protocols bgp group bgp type external
    set routing-instances vpn1 protocols bgp group bgp local-address 44.44.44.1
    set routing-instances vpn1 protocols bgp group bgp family inet unicast
    set routing-instances vpn1 protocols bgp group bgp neighbor 44.44.44.2 peer-as 3
    set routing-instances vpn1 protocols bgp group bgp1 type external
    set routing-instances vpn1 protocols bgp group bgp1 local-address 11.11.11.1
    set routing-instances vpn1 protocols bgp group bgp1 family inet unicast
    set routing-instances vpn1 protocols bgp group bgp1 neighbor 11.11.11.2 peer-as 4
    set routing-instances vpn1 protocols pim vpn-group-address 224.1.1.1
    set routing-instances vpn1 protocols pim rp static address 10.255.8.168
  
```



```
set routing-instances vpn1 protocols pim interface all
set routing-instances vpn1 protocols pim join-load-balance
```

```
PE2 set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-2/0/3.0
set routing-instances vpn1 interface ge-4/0/5.0
set routing-instances vpn1 interface lo0.1
set routing-instances vpn1 route-distinguisher 2:2
set routing-instances vpn1 vrf-target target:1:1
set routing-instances vpn1 routing-options multipath vpn-unequal-cost
    equal-external-internal
set routing-instances vpn1 protocols bgp export direct
set routing-instances vpn1 protocols bgp group bgp1 type external
set routing-instances vpn1 protocols bgp group bgp1 local-address 10.90.10.1
set routing-instances vpn1 protocols bgp group bgp1 family inet unicast
set routing-instances vpn1 protocols bgp group bgp1 neighbor 10.90.10.2 peer-as 45
set routing-instances vpn1 protocols bgp group bgp type external
set routing-instances vpn1 protocols bgp group bgp local-address 10.50.10.2
set routing-instances vpn1 protocols bgp group bgp family inet unicast
set routing-instances vpn1 protocols bgp group bgp neighbor 10.50.10.1 peer-as 4
set routing-instances vpn1 protocols pim vpn-group-address 224.1.1.1
set routing-instances vpn1 protocols pim rp static address 10.255.8.168
set routing-instances vpn1 protocols pim interface all
set routing-instances vpn1 protocols pim join-load-balance
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#). To configure the PE1 router:



NOTE: Repeat this procedure for every Juniper Networks router in the MVPN domain, after modifying the appropriate interface names, addresses, and any other parameters for each router.

1. Configure a VPN routing and forwarding (VRF) instance.


```
[edit routing-instances vpn1]
user@PE1# set instance-type vrf
user@PE1# set interface ge-5/0/4.0
user@PE1# set interface ge-5/2/0.0
user@PE1# set interface lo0.1
user@PE1# set route-distinguisher 1:1
user@PE1# set vrf-target target:1:1
```
2. Enable protocol-independent load balancing for the VRF instance.


```
[edit routing-instances vpn1]
user@PE1# set routing-options multipath vpn-unequal-cost equal-external-internal
```
3. Configure BGP groups and neighbors to enable PE to CE routing.


```
[edit routing-instances vpn1 protocols]
user@PE1# set bgp export direct
user@PE1# set bgp group bgp type external
user@PE1# set bgp group bgp local-address 44.44.44.1
```

```
user@PE1# set bgp group bgp family inet unicast
user@PE1# set bgp group bgp neighbor 44.44.44.2 peer-as 3
user@PE1# set bgp group bgp1 type external
user@PE1# set bgp group bgp1 local-address 11.11.11.1
user@PE1# set bgp group bgp1 family inet unicast
user@PE1# set bgp group bgp1 neighbor 11.11.11.2 peer-as 4
```

4. Configure PIM to enable PE to CE multicast routing.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim vpn-group-address 224.1.1.1
user@PE1# set pim rp static address 10.255.8.168
```

5. Enable PIM on all network interfaces.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim interface all
```

6. Enable PIM join load balancing for the VRF instance.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim join-load-balance
```

Results From configuration mode, confirm your configuration by entering the **show routing-instances** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
routing-instances {
  vpn1 {
    instance-type vrf;
    interface ge-5/0/4.0;
    interface ge-5/2/0.0;
    interface lo0.1;
    route-distinguisher 1:1;
    vrf-target target:1:1;
    routing-options {
      multipath {
        vpn-unequal-cost equal-external-internal;
      }
    }
  }
  protocols {
    bgp {
      export direct;
      group bgp {
        type external;
        local-address 44.44.44.1;
        family inet {
          unicast;
        }
        neighbor 44.44.44.2 {
          peer-as 3;
        }
      }
    }
    group bgp1 {
      type external;
      local-address 11.11.11.1;
      family inet {
```

```

        unicast;
    }
    neighbor 11.11.11.2 {
        peer-as 4;
    }
}
pim {
    vpn-group-address 224.1.1.1;
    rp {
        static {
            address 10.255.8.168;
        }
    }
    interface all;
    join-load-balance;
}
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying PIM Join Load Balancing for Different Groups of Join Messages on page 5189](#)

Verifying PIM Join Load Balancing for Different Groups of Join Messages

Purpose Verify PIM join load balancing for the different groups of join messages received on the PE1 router.

Action From operational mode, run the **show pim join instance extensive** command.

```

user@PE1> show pim join instance extensive
Instance: PIM.vpn1 Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 225.1.1.1
Source: *
RP: 10.255.8.168
Flags: sparse,rptree,wildcard
Upstream interface: ge-5/2/0.1
Upstream neighbor: 10.10.10.2
Upstream state: Join to RP
Downstream neighbors:
Interface: ge-5/0/4.0
10.40.10.2 State: Join Flags: SRW Timeout: 207

Group: 225.1.1.2
Source: *
RP: 10.255.8.168
Flags: sparse,rptree,wildcard
Upstream interface: mt-5/0/10.32768
Upstream neighbor: 19.19.19.19
Upstream state: Join to RP

```

```
Downstream neighbors:
  Interface: ge-5/0/4.0
    10.40.10.2 State: Join Flags: SRW Timeout: 207

Group: 225.1.1.3
  Source: *
  RP: 10.255.8.168
  Flags: sparse,rptree,wildcard
  Upstream interface: ge-5/2/0.1
  Upstream neighbor: 10.10.10.2
  Upstream state: Join to RP
  Downstream neighbors:
    Interface: ge-5/0/4.0
      10.40.10.2 State: Join Flags: SRW Timeout: 207

Group: 225.1.1.4
  Source: *
  RP: 10.255.8.168
  Flags: sparse,rptree,wildcard
  Upstream interface: mt-5/0/10.32768
  Upstream neighbor: 19.19.19.19
  Upstream state: Join to RP
  Downstream neighbors:
    Interface: ge-5/0/4.0
      10.40.10.2 State: Join Flags: SRW Timeout: 207
```

Meaning The output shows how the PE1 router has load-balanced the C-PIM join messages for four different groups.

- For Group 1 (group address: 225.1.1.1) and Group 3 (group address: 225.1.1.3) join messages, the PE1 router has selected the EBGp path toward the CE1 router to send the join messages.
- For Group 2 (group address: 225.1.1.2) and Group 4 (group address: 225.1.1.4) join messages, the PE1 router has selected the IBGP path toward the PE2 router to send the join messages.

Related Documentation

- [PIM Join Load Balancing on Multipath MVPN Routes Overview on page 3741](#)
- [Example: Configuring PIM Join Load Balancing On Next-Generation Multicast VPN on page 5190](#)

Example: Configuring PIM Join Load Balancing On Next-Generation Multicast VPN

This example shows how to configure multipath routing for external and internal virtual private network (VPN) routes with unequal interior gateway protocol (IGP) metrics and Protocol Independent Multicast (PIM) join load balancing on provider edge (PE) routers running next-generation multicast VPN (MVPN). This feature allows customer PIM (C-PIM) join messages to be load-balanced across available internal BGP (IBGP) upstream paths when there is no external BGP (EBGP) path present, and across available EBGp upstream paths when external and internal BGP (EIBGP) paths are present toward the source or rendezvous point (RP).

- [Requirements on page 5191](#)
- [Overview and Topology on page 5191](#)

- [Configuration on page 5194](#)
- [Verification on page 5197](#)

Requirements

This example uses the following hardware and software components:

- Three routers that can be a combination of M Series, MX Series, or T Series routers.
- Junos OS Release 12.1 running on all the devices.

Before you begin:

1. Configure the device interfaces.
2. Configure the following routing protocols on all PE routers:
 - OSPF
 - MPLS
 - LDP
 - PIM
 - BGP
3. Configure a multicast VPN.

Overview and Topology

Junos OS Release 12.1 and later support multipath configuration along with PIM join load balancing. This allows C-PIM join messages to be load-balanced across all available IBGP paths when there are only IBGP paths present, and across all available upstream EBGP paths when EIBGP paths are present toward the source (or RP). Unlike Draft-Rosen MVPN, next-generation MVPN does not utilize unequal EIBGP paths to send C-PIM join messages. This feature is applicable to IPv4 C-PIM join messages.

By default, only one active IBGP path is used to send the C-PIM join messages for a PE router having only IBGP paths toward the source (or RP). When there are EIBGP upstream paths present, only one active EBGP path is used to send the join messages.

In a next-generation MVPN, C-PIM join messages are translated into (or encoded as) BGP customer multicast (C-multicast) MVPN routes and advertised with the BGP MCAST-VPN address family toward the sender PE routers. A PE router originates a C-multicast MVPN route in response to receiving a C-PIM join message through its PE router to customer edge (CE) router interface. The two types of C-multicast MVPN routes are:

- Shared tree join route (C-*, C-G)
 - Originated by receiver PE routers.
 - Originated when a PE router receives a shared tree C-PIM join message through its PE-CE router interface.
- Source tree join route (C-S, C-G)

- Originated by receiver PE routers.
- Originated when a PE router receives a source tree C-PIM join message (C-S, C-G), or originated by the PE router that already has a shared tree join route and receives a source active autodiscovery route.

The upstream path in a next-generation MVPN is selected using the Bitwise-XOR hash algorithm as specified in Internet draft draft-ietf-l3vpn-2547bis-mcast, *Multicast in MPLS/BGP IP VPNs*. The hash algorithm is performed as follows:

1. The PE routers in the candidate set are numbered from lower to higher IP address, starting from **0**.
2. A bitwise exclusive-or of all the bytes is performed on the C-root (source) and the C-G (group) address.
3. The result is taken modulo n , where n is the number of PE routers in the candidate set. The result is **N**.
4. **N** represents the IP address of the upstream PE router as numbered in Step 1.

During load balancing, if a PE router with one or more upstream IBGP paths toward the source (or RP) discovers a new IBGP path toward the same source (or RP), the C-PIM join messages distributed among previously existing IBGP paths get redistributed due to the change in the candidate PE router set.

In this example, PE1, PE2, and PE3 are the PE routers that have the multipath PIM join load-balancing feature configured. Router PE1 has two EBGp paths and one IBGP upstream path, PE2 has one EBGp path and one IBGP upstream path, and PE3 has two IBGP upstream paths toward the Source. Router CE4 is the customer edge (CE) router attached to PE3. Source and Receiver are the Free BSD hosts.

On PE routers that have EIBGP paths toward the source (or RP), such as PE1 and PE2, PIM join load balancing is performed as follows:

1. The C-PIM join messages are sent using EBGp paths only. IBGP paths are not used to propagate the join messages.

In [Figure 115 on page 5193](#), the PE1 router distributes the join messages between the two EBGp paths to the CE1 router, and PE2 uses the EBGp path to CE1 to send the join messages.

2. If a PE router loses one or more EBGp paths toward the source (or RP), the RPF neighbor on the multicast tunnel interface is selected based on a hash mechanism.

On discovering the first EBGp path, only new join messages get load-balanced across available EBGp paths, whereas the existing join messages on the multicast tunnel interface are not redistributed.

If the EBGp path from the PE2 router to the CE1 router goes down, PE2 sends the join messages to PE1 using the IBGP path. When the EBGp path to CE1 is restored, only new join messages that arrive on PE2 use the restored EBGp path, whereas join messages already sent on the IBGP path are not redistributed.

On PE routers that have only IBGP paths toward the source (or RP), such as the PE3 router, PIM join load balancing is performed as follows:

1. The C-PIM join messages from CE routers get load-balanced only as BGP C-multicast data messages among IBGP paths.


In [Figure 115 on page 5193](#), assuming that the CE4 host is interested in receiving traffic from the Source, and CE4 initiates source join messages for different groups (Group 1 [C-S,C-G1] and Group 2 [C-S,C-G2]), the source join messages arrive on the PE3 router.

Router PE3 then uses the Bitwise-XOR hash algorithm to select the upstream PE router to send the C-multicast data for each group. The algorithm first numbers the upstream PE routers from lower to higher IP address starting from 0.

Assuming that Router PE1 router is numbered 0 and Router PE2 is 1, and the hash result for Group 1 and Group 2 join messages is 0 and 1, respectively, the PE3 router selects PE1 as the upstream PE router to send Group 1 join messages, and PE2 as the upstream PE router to send the Group 2 join messages to the Source.

2. The shared join messages for different groups [C-*,C-G] are also treated in a similar way to reach the destination.

Figure 115: PIM Join Load Balancing on Next-Generation MVPN



Configuration

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

```

PE1  set routing-instances vpn1 instance-type vrf
      set routing-instances vpn1 interface ge-3/0/1.0
      set routing-instances vpn1 interface ge-3/3/2.0
      set routing-instances vpn1 interface lo0.1
      set routing-instances vpn1 route-distinguisher 1:1
      set routing-instances vpn1 provider-tunnel rsvp-te label-switched-path-template
        default-template
      set routing-instances vpn1 vrf-target target:1:1
      set routing-instances vpn1 vrf-table-label
      set routing-instances vpn1 routing-options multipath vpn-unequal-cost
        equal-external-internal
      set routing-instances vpn1 protocols bgp export direct
      set routing-instances vpn1 protocols bgp group bgp type external
      set routing-instances vpn1 protocols bgp group bgp local-address 10.40.10.1
      set routing-instances vpn1 protocols bgp group bgp family inet unicast
      set routing-instances vpn1 protocols bgp group bgp neighbor 10.40.10.2 peer-as 3
      set routing-instances vpn1 protocols bgp group bgp1 type external
      set routing-instances vpn1 protocols bgp group bgp1 local-address 10.10.10.1
      set routing-instances vpn1 protocols bgp group bgp1 family inet unicast
      set routing-instances vpn1 protocols bgp group bgp1 neighbor 10.10.10.2 peer-as 3
      set routing-instances vpn1 protocols pim rp static address 10.255.10.119
      set routing-instances vpn1 protocols pim interface all
      set routing-instances vpn1 protocols pim join-load-balance
      set routing-instances vpn1 protocols mvpn mvpn-mode rpt-spt
      set routing-instances vpn1 protocols mvpn mvpn-join-load-balance bitwise-xor-hash

PE2  set routing-instances vpn1 instance-type vrf
      set routing-instances vpn1 interface ge-1/0/9.0
      set routing-instances vpn1 interface lo0.1
      set routing-instances vpn1 route-distinguisher 2:2
      set routing-instances vpn1 provider-tunnel rsvp-te label-switched-path-template
        default-template
      set routing-instances vpn1 vrf-target target:1:1
      set routing-instances vpn1 vrf-table-label
      set routing-instances vpn1 routing-options multipath vpn-unequal-cost
        equal-external-internal
      set routing-instances vpn1 protocols bgp export direct
      set routing-instances vpn1 protocols bgp group bgp local-address 10.50.10.2
      set routing-instances vpn1 protocols bgp group bgp family inet unicast
      set routing-instances vpn1 protocols bgp group bgp neighbor 10.50.10.1 peer-as 3
      set routing-instances vpn1 protocols pim rp static address 10.255.10.119
      set routing-instances vpn1 protocols pim interface all
      set routing-instances vpn1 protocols mvpn mvpn-mode rpt-spt
      set routing-instances vpn1 protocols mvpn mvpn-join-load-balance bitwise-xor-hash

PE3  set routing-instances vpn1 instance-type vrf
      set routing-instances vpn1 interface ge-0/0/8.0
      set routing-instances vpn1 interface lo0.1

```



```

set routing-instances vpn1 route-distinguisher 3:3
set routing-instances vpn1 provider-tunnel rsvp-te label-switched-path-template
  default-template
set routing-instances vpn1 vrf-target target:1:1
set routing-instances vpn1 vrf-table-label
set routing-instances vpn1 routing-options multipath vpn-unequal-cost
  equal-external-internal
set routing-instances vpn1 routing-options autonomous-system 1
set routing-instances vpn1 protocols bgp export direct
set routing-instances vpn1 protocols bgp group bgp type external
set routing-instances vpn1 protocols bgp group bgp local-address 10.80.10.1
set routing-instances vpn1 protocols bgp group bgp family inet unicast
set routing-instances vpn1 protocols bgp group bgp neighbor 10.80.10.2 peer-as 2
set routing-instances vpn1 protocols pim rp static address 10.255.10.119
set routing-instances vpn1 protocols pim interface all
set routing-instances vpn1 protocols mvpn mvpn-mode rpt-spt
set routing-instances vpn1 protocols mvpn mvpn-join-load-balance bitwise-xor-hash

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#). To configure the PE1 router:



NOTE: Repeat this procedure for every Juniper Networks router in the MVPN domain, after modifying the appropriate interface names, addresses, and any other parameters for each router.

1. Configure a VPN routing forwarding (VRF) routing instance.

```

[edit routing-instances vpn1]
user@PE1# set instance-type vrf
user@PE1# set interface ge-3/0/1.0
user@PE1# set interface ge-3/3/2.0
user@PE1# set interface lo0.1
user@PE1# set route-distinguisher 1:1
user@PE1# set provider-tunnel rsvp-te label-switched-path-template
  default-template
user@PE1# set vrf-target target:1:1
user@PE1# set vrf-table-label

```

2. Enable protocol-independent load balancing for the VRF instance.

```

[edit routing-instances vpn1]
user@PE1# set routing-options multipath vpn-unequal-cost equal-external-internal

```

3. Configure BGP groups and neighbors to enable PE to CE routing.

```

[edit routing-instances vpn1 protocols]
user@PE1# set bgp export direct
user@PE1# set bgp group bgp type external
user@PE1# set bgp group bgp local-address 10.40.10.1
user@PE1# set bgp group bgp family inet unicast
user@PE1# set bgp group bgp neighbor 10.40.10.2 peer-as 3
user@PE1# set bgp group bgp1 type external
user@PE1# set bgp group bgp1 local-address 10.10.10.1

```

```
user@PE1# set bgp group bgp1 family inet unicast
user@PE1# set bgp group bgp1 neighbor 10.10.10.2 peer-as 3
```

4. Configure PIM to enable PE to CE multicast routing.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim rp static address 10.255.10.119
```

5. Enable PIM on all network interfaces.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim interface all
```

6. Enable PIM join load balancing for the VRF instance.

```
[edit routing-instances vpn1 protocols]
user@PE1# set pim join-load-balance
```

7. Configure the mode for C-PIM join messages to use rendezvous-point trees, and switch to the shortest-path tree after the source is known.

```
[edit routing-instances vpn1 protocols]
user@PE1# set mvpn mvpn-mode rpt-spt
```

8. Configure the VRF instance to use the Bytewise-XOR hash algorithm.

```
[edit routing-instances vpn1 protocols]
user@PE1# set mvpn mvpn-join-load-balance bytewise-xor-hash
```

Results

From configuration mode, confirm your configuration by entering the **show routing-instances** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show routing-instances
routing-instances {
  vpn1 {
    instance-type vrf;
    interface ge-3/0/1.0;
    interface ge-3/3/2.0;
    interface lo0.1;
    route-distinguisher 1:1;
    provider-tunnel {
      rsvp-te {
        label-switched-path-template {
          default-template;
        }
      }
    }
    vrf-target target:1:1;
    vrf-table-label;
    routing-options {
      multipath {
        vpn-unequal-cost equal-external-internal;
      }
    }
    protocols {
      bgp {
```

```

export direct;
group bgp {
    type external;
    local-address 10.40.10.1;
    family inet {
        unicast;
    }
    neighbor 10.40.10.2 {
        peer-as 3;
    }
}
group bgp1 {
    type external;
    local-address 10.10.10.1;
    family inet {
        unicast;
    }
    neighbor 10.10.10.2 {
        peer-as 3;
    }
}
}
pim {
    rp {
        static {
            address 10.255.10.119;
        }
    }
    interface all;
    join-load-balance;
}
mvpn {
    mvpn-mode {
        rpt-spt;
    }
    mvpn-join-load-balance {
        bitwise-xor-hash;
    }
}
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying MVPN C-Multicast Route Information for Different Groups of Join Messages on page 5197](#)

Verifying MVPN C-Multicast Route Information for Different Groups of Join Messages

Purpose Verify MVPN C-multicast route information for different groups of join messages received on the PE3 router.

Action From operational mode, run the **show mvpn c-multicast** command.

```

user@PE3> show mvpn c-multicast
MVPN instance:
Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel
Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Family : INET

Instance : vpn1
MVPN Mode : RPT-SPT
C-mcast IPv4 (S:G)          Ptnl          St
0.0.0.0/0:225.1.1.1/32      RSVP-TE P2MP:10.255.10.2, 5834,10.255.10.2
4.4.4.2/32:225.1.1.1/32     RSVP-TE P2MP:10.255.10.2, 5834,10.255.10.2
0.0.0.0/0:225.1.1.2/32     RSVP-TE P2MP:10.255.10.14, 47575,10.255.10.14
4.4.4.2/32:225.1.1.2/32     RSVP-TE P2MP:10.255.10.14, 47575,10.255.10.14

```

Meaning The output shows how the PE3 router has load-balanced the C-multicast data for the different groups.

- For source join messages (S,G):
 - 4.4.4.2/32:225.1.1.1/32 (S,G1) toward the PE1 router (10.255.10.2 is the loopback address of Router PE1).
 - 4.4.4.2/32:225.1.1.2/32 (S,G2) toward the PE2 router (10.255.10.14 is the loopback address of Router PE2).
- For shared join messages (*,G):
 - 0.0.0.0/0:225.1.1.1/32 (*,G1) toward the PE1 router (10.255.10.2 is the loopback address of Router PE1).
 - 0.0.0.0/0:225.1.1.2/32 (*,G2) toward the PE2 router (10.255.10.14 is the loopback address of Router PE2).

Related Documentation

- [PIM Join Load Balancing on Multipath MVPN Routes Overview on page 3741](#)
- [Example: Configuring PIM Join Load Balancing on Draft-Rosen Multicast VPN on page 5182](#)

Example: Configuring PIM State Limits

- [Controlling PIM Resources for Multicast VPNs Overview on page 5198](#)
- [Example: Configuring PIM State Limits on page 5201](#)

Controlling PIM Resources for Multicast VPNs Overview

A service provider network must protect itself from potential attacks from misconfigured or misbehaving customer edge (CE) devices and their associated VPN routing and forwarding (VRF) routing instances. Misbehaving CE devices can potentially advertise a large number of multicast routes toward a provider edge (PE) device, thereby consuming memory on the PE device and using other system resources in the network that are reserved for routes belonging to other VPNs.

To protect against potential misbehaving CE devices and VRF routing instances for specific multicast VPNs (MVPNs), you can control the following Protocol Independent Multicast (PIM) resources:

- Limit the number of accepted PIM join messages for any-source groups (*G) and source-specific groups (S,G).

Note how the device counts the PIM join messages:

- Each (*G) counts as one group toward the limit.
- Each (S,G) counts as one group toward the limit.
- Limit the number of PIM register messages received for a specific VRF routing instance. Use this configuration if the device is configured as a rendezvous point (RP) or has the potential to become an RP. When a source in a multicast network becomes active, the source's designated router (DR) encapsulates multicast data packets into a PIM register message and sends them by means of unicast to the RP router.

Note how the device counts PIM register messages:

- Each unique (S,G) join received by the RP counts as one group toward the configured register messages limit.
- Periodic register messages sent by the DR for existing or already known (S,G) entries do not count toward the configured register messages limit.
- Register messages are accepted until either the PIM register limit or the PIM join limit (if configured) is exceeded. Once either limit is reached, any new requests are dropped.
- Limit the number of group-to-RP mappings allowed in a specific VRF routing instance. Use this configuration if the device is configured as an RP or has the potential to become an RP. This configuration can apply to devices configured for automatic RP announce and discovery (Auto-RP) or as a PIM bootstrap router. Every multicast device within a PIM domain must be able to map a particular multicast group address to the same RP. Both Auto-RP and the bootstrap router functionality are the mechanisms used to learn the set of group-to-RP mappings. Auto-RP is typically used in a PIM dense-mode deployment, and the bootstrap router is typically used in a PIM sparse-mode deployment.



NOTE: The group-to-RP mappings limit does not apply to static RP or embedded RP configurations.

Some important things to note about how the device counts group-to-RP mappings:

- One group prefix mapped to five RPs counts as five group-to-RP mappings.
- Five distinct group prefixes mapped to one RP count as five group-to-RP mappings.

Once the configured limits are reached, no new PIM join messages, PIM register messages, or group-to-RP mappings are accepted unless one of the following occurs:

- You clear the current PIM join states by using the **clear pim join** command. If you use this command on an RP configured for PIM register message limits, the register limit count is also restarted because the PIM join messages are unknown by the RP.



NOTE: On the RP, you can also use the **clear pim register** command to clear all of the PIM registers. This command is useful if the current PIM register count is greater than the newly configured PIM register limit. After you clear the PIM registers, new PIM register messages are received up to the configured limit.

- The traffic responsible for the excess PIM join messages and PIM register messages stops and is no longer present.



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

You restart the PIM routing process on the device. This restart clears all of the configured limits but disrupts routing and therefore requires a maintenance window for the change.

System Log Messages for PIM Resources

You can optionally configure a system log warning threshold for each of the PIM resources. With this configuration, you can generate and review system log messages to detect if an excessive number of PIM join messages, PIM register messages, or group-to-RP mappings have been received on the device. The system log warning thresholds are configured per PIM resource and are a percentage of the configured maximum limits of the PIM join messages, PIM register messages, and group-to-RP mappings. You can further specify a log interval for each configured PIM resource, which is the amount of time (in seconds) between the log messages.

The log messages convey when the configured limits have been exceeded, when the configured warning thresholds have been exceeded, and when the configured limits drop below the configured warning threshold. [Table 384 on page 5200](#) describes the different types of PIM system messages that you might see depending on your system log warning and log interval configurations.

Table 384: PIM System Log Messages

System Log Message	Definition
RPD_PIM_SG_THRESHOLD_EXCEED	Records when the (S,G)/(*G) routes exceed the configured warning threshold.
RPD_PIM_REG_THRESH_EXCEED	Records when the PIM registers exceed the configured warning threshold.
RPD_PIM_GRP_RP_MAP_THRES_EXCEED	Records when the group-to-RP mappings exceed the configured warning threshold.

Table 384: PIM System Log Messages (*continued*)

System Log Message	Definition
RPD_PIM_SG_LIMIT_EXCEED	Records when the (S,G)/(*G) routes exceed the configured limit, or when the configured log interval has been met and the routes exceed the configured limit.
RPD_PIM_REGISTER_LIMIT_EXCEED	Records when the PIM registers exceed the configured limit, or when the configured log interval has been met and the registers exceed the configured limit.
RPD_PIM_GRP_RP_MAP_LIMIT_EXCEED	Records when the group-to-RP mappings exceed the configured limit, or when the configured log interval has been met and the mapping exceeds the configured limit.
RPD_PIM_SG_LIMIT_BELOW	Records when the (S,G)/(*G) routes drop below the configured limit and the configured log interval.
RPD_PIM_REGISTER_LIMIT_BELOW	Records when the PIM registers drop below the configured limit and the configured log interval.
RPD_PIM_GRP_RP_MAP_LIMIT_BELOW	Records when the group-to-RP mappings drop below the configured limit and the configured log interval.

Example: Configuring PIM State Limits

This example shows how to set limits on the Protocol Independent Multicast (PIM) state information so that a service provider network can protect itself from potential attacks from misconfigured or misbehaving customer edge (CE) devices and their associated VPN routing and forwarding (VRF) routing instances.

- [Requirements on page 5201](#)
- [Overview on page 5201](#)
- [Configuration on page 5202](#)
- [Verification on page 5209](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, a multiprotocol BGP-based multicast VPN (next-generation MBGP MVPN) is configured with limits on the PIM state resources.

The **sglimit maximum** statement sets a limit for the number of accepted (*G) and (S,G) PIM join states received for the vpn-l routing instance.

The **rp register-limit maximum** statement configures a limit for the number of PIM register messages received for the vpn-l routing instance. You configure this statement on the rendezvous point (RP) or on all the devices that might become the RP.


The **group-rp-mapping maximum** statement configures a limit for the number of group-to-RP mappings allowed in the vpn-1 routing instance.

For each configured PIM resource, the **threshold** statement sets a percentage of the maximum limit at which to start generating warning messages in the PIM log file.

For each configured PIM resource, the **log-interval** statement is an amount of time (in seconds) between system log message generation.

Figure 116 on page 5202 shows the topology used in this example.

Figure 116: PIM State Limits Topology



“CLI Quick Configuration” on page 5202 shows the configuration for all of the devices in Figure 116 on page 5202. The section “Step-by-Step Procedure” on page 5205 describes the steps on Device PE1.

Configuration

CLI Quick Configuration	To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.
Device CE1	<pre> set interfaces ge-1/2/0 unit 1 family inet address 10.1.1.1/30 set interfaces ge-1/2/0 unit 1 family mpls set interfaces lo0 unit 1 family inet address 1.1.1.1/32 set protocols ospf area 0.0.0.0 interface lo0.1 passive set protocols ospf area 0.0.0.0 interface ge-1/2/0.1 set protocols pim rp static address 100.1.1.2 set protocols pim interface all set routing-options router-id 1.1.1.1 </pre>
Device PE1	<pre> set interfaces ge-1/2/0 unit 2 family inet address 10.1.1.2/30 set interfaces ge-1/2/0 unit 2 family mpls set interfaces ge-1/2/1 unit 5 family inet address 10.1.1.5/30 set interfaces ge-1/2/1 unit 5 family mpls set interfaces vt-1/2/0 unit 2 family inet set interfaces lo0 unit 2 family inet address 1.1.1.2/32 set interfaces lo0 unit 102 family inet address 100.1.1.2/32 set protocols mpls interface ge-1/2/1.5 set protocols bgp group ibgp type internal set protocols bgp group ibgp local-address 1.1.1.2 set protocols bgp group ibgp family inet-vpn any </pre>


```

set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 1.1.1.4
set protocols bgp group ibgp neighbor 1.1.1.5
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/1.5
set protocols ldp interface ge-1/2/1.5
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface ge-1/2/0.2
set routing-instances vpn-1 interface vt-1/2/0.2
set routing-instances vpn-1 interface lo0.102
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 provider-tunnel ldp-p2mp
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.102 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/0.2
set routing-instances vpn-1 protocols pim sglimit family inet maximum 100
set routing-instances vpn-1 protocols pim sglimit family inet threshold 70
set routing-instances vpn-1 protocols pim sglimit family inet log-interval 10
set routing-instances vpn-1 protocols pim rp register-limit family inet maximum 100
set routing-instances vpn-1 protocols pim rp register-limit family inet threshold 80
set routing-instances vpn-1 protocols pim rp register-limit family inet log-interval 10
set routing-instances vpn-1 protocols pim rp group-rp-mapping family inet maximum 100
set routing-instances vpn-1 protocols pim rp group-rp-mapping family inet threshold 80
set routing-instances vpn-1 protocols pim rp group-rp-mapping family inet log-interval
  10
set routing-instances vpn-1 protocols pim rp static address 100.1.1.2
set routing-instances vpn-1 protocols pim interface ge-1/2/0.2 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 1.1.1.2
set routing-options autonomous-system 1001

```

```

Device P
set interfaces ge-1/2/0 unit 6 family inet address 10.1.1.6/30
set interfaces ge-1/2/0 unit 6 family mpls
set interfaces ge-1/2/1 unit 9 family inet address 10.1.1.9/30
set interfaces ge-1/2/1 unit 9 family mpls
set interfaces ge-1/2/2 unit 13 family inet address 10.1.1.13/30
set interfaces ge-1/2/2 unit 13 family mpls
set interfaces lo0 unit 3 family inet address 1.1.1.3/32
set protocols mpls interface ge-1/2/0.6
set protocols mpls interface ge-1/2/1.9
set protocols mpls interface ge-1/2/2.13
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.6
set protocols ospf area 0.0.0.0 interface ge-1/2/1.9
set protocols ospf area 0.0.0.0 interface ge-1/2/2.13
set protocols ldp interface ge-1/2/0.6
set protocols ldp interface ge-1/2/1.9
set protocols ldp interface ge-1/2/2.13
set protocols ldp p2mp
set routing-options router-id 1.1.1.3

```

Device PE2

```
set interfaces ge-1/2/0 unit 10 family inet address 10.1.1.10/30
set interfaces ge-1/2/0 unit 10 family mpls
set interfaces ge-1/2/1 unit 17 family inet address 10.1.1.17/30
set interfaces ge-1/2/1 unit 17 family mpls
set interfaces vt-1/2/0 unit 4 family inet
set interfaces lo0 unit 4 family inet address 1.1.1.4/32
set interfaces lo0 unit 104 family inet address 100.1.1.4/32
set protocols mpls interface ge-1/2/0.10
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 1.1.1.4
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 1.1.1.2
set protocols bgp group ibgp neighbor 1.1.1.5
set protocols ospf area 0.0.0.0 interface lo0.4 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.10
set protocols ldp interface ge-1/2/0.10
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface vt-1/2/0.4
set routing-instances vpn-1 interface ge-1/2/1.17
set routing-instances vpn-1 interface lo0.104
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.104 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/1.17
set routing-instances vpn-1 protocols pim rp group-rp-mapping family inet maximum 100
set routing-instances vpn-1 protocols pim rp group-rp-mapping family inet threshold 80
set routing-instances vpn-1 protocols pim rp group-rp-mapping family inet log-interval
    10
set routing-instances vpn-1 protocols pim rp static address 100.1.1.2
set routing-instances vpn-1 protocols pim interface ge-1/2/1.17 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 1.1.1.4
set routing-options autonomous-system 1001
```

Device PE3

```
set interfaces ge-1/2/0 unit 14 family inet address 10.1.1.14/30
set interfaces ge-1/2/0 unit 14 family mpls
set interfaces ge-1/2/1 unit 21 family inet address 10.1.1.21/30
set interfaces ge-1/2/1 unit 21 family mpls
set interfaces vt-1/2/0 unit 5 family inet
set interfaces lo0 unit 5 family inet address 1.1.1.5/32
set interfaces lo0 unit 105 family inet address 100.1.1.5/32
set protocols mpls interface ge-1/2/0.14
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 1.1.1.5
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 1.1.1.2
set protocols bgp group ibgp neighbor 1.1.1.4
set protocols ospf area 0.0.0.0 interface lo0.5 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.14
```

```

set protocols ldp interface ge-1/2/0.14
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface vt-1/2/0.5
set routing-instances vpn-1 interface ge-1/2/1.21
set routing-instances vpn-1 interface lo0.105
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.105 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/1.21
set routing-instances vpn-1 protocols pim rp static address 100.1.1.2
set routing-instances vpn-1 protocols pim interface ge-1/2/1.21 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 1.1.1.5
set routing-options autonomous-system 1001

```

Device CE2

```

set interfaces ge-1/2/0 unit 18 family inet address 10.1.1.18/30
set interfaces ge-1/2/0 unit 18 family mpls
set interfaces lo0 unit 6 family inet address 1.1.1.6/32
set protocols sap listen 224.1.1.1
set protocols ospf area 0.0.0.0 interface lo0.6 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.18
set protocols pim rp static address 100.1.1.2
set protocols pim interface all
set routing-options router-id 1.1.1.6

```

Device CE3

```

set interfaces ge-1/2/0 unit 22 family inet address 10.1.1.22/30
set interfaces ge-1/2/0 unit 22 family mpls
set interfaces lo0 unit 7 family inet address 1.1.1.7/32
set protocols ospf area 0.0.0.0 interface lo0.7 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.22
set protocols pim rp static address 100.1.1.2
set protocols pim interface all
set routing-options router-id 1.1.1.7

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode” on page 4704](#) in the *CLI User Guide*.

To configure PIM state limits:

1. Configure the network interfaces.

```

[edit interfaces]
user@PE1# set ge-1/2/0 unit 2 family inet address 10.1.1.2/30
user@PE1# set ge-1/2/0 unit 2 family mpls

user@PE1# set ge-1/2/1 unit 5 family inet address 10.1.1.5/30
user@PE1# set ge-1/2/1 unit 5 family mpls

user@PE1# set vt-1/2/0 unit 2 family inet

```

```
user@PE1# set lo0 unit 2 family inet address 1.1.1.2/32
user@PE1# set lo0 unit 102 family inet address 100.1.1.2/32
```

2. Configure MPLS on the core-facing interface.

```
[edit protocols mpls]
user@PE1# set interface ge-1/2/1.5
```

3. Configure internal BGP (IBGP) on the main router.

The IBGP neighbors are the other PE devices.

```
[edit protocols bgp group ibgp]
user@PE1# set type internal
user@PE1# set local-address 1.1.1.2
user@PE1# set family inet-vpn any
user@PE1# set family inet-mvpn signaling
user@PE1# set neighbor 1.1.1.4
user@PE1# set neighbor 1.1.1.5
```

4. Configure OSPF on the main router.

```
[edit protocols ospf area 0.0.0.0]
user@PE1# set interface lo0.2 passive
user@PE1# set interface ge-1/2/1.5
```

5. Configure a signaling protocol (RSVP or LDP) on the main router.

```
[edit protocols ldp]
user@PE1# set interface ge-1/2/1.5
user@PE1# set p2mp
```

6. Configure the BGP export policy.

```
[edit policy-options policy-statement parent_vpn_routes]
user@PE1# set from protocol bgp
user@PE1# set then accept
```

7. Configure the routing instance.

The customer-facing interfaces and the BGP export policy are referenced in the routing instance.

```
[edit routing-instances vpn-1]
user@PE1# set instance-type vrf
```

```
user@PE1# set interface ge-1/2/0.2
user@PE1# set interface vt-1/2/0.2
user@PE1# set interface lo0.102
```

```
user@PE1# set route-distinguisher 100:100
user@PE1# set provider-tunnel ldp-p2mp
user@PE1# set vrf-target target:1:1
```

```
user@PE1# set protocols ospf export parent_vpn_routes
user@PE1# set protocols ospf area 0.0.0.0 interface lo0.102 passive
user@PE1# set protocols ospf area 0.0.0.0 interface ge-1/2/0.2
```

```
user@PE1# set protocols pim rp static address 100.1.1.2
```

```
user@PE1# set protocols pim interface ge-1/2/0.2 mode sparse
```

```
user@PE1# set protocols mvpn
```

8. Configure the PIM state limits.

```
[edit routing-instances vpn-1 protocols pim]
user@PE1# set sglimit family inet maximum 100
user@PE1# set sglimit family inet threshold 70
user@PE1# set sglimit family inet log-interval 10
```

```
user@PE1# set rp register-limit family inet maximum 100
user@PE1# set rp register-limit family inet threshold 80
user@PE1# set rp register-limit family inet log-interval 10
```

```
user@PE1# set rp group-rp-mapping family inet maximum 100
user@PE1# set rp group-rp-mapping family inet threshold 80
user@PE1# set rp group-rp-mapping family inet log-interval 10
```

9. Configure the router ID and AS number.

```
[edit routing-options]
user@PE1# set router-id 1.1.1.2
user@PE1# set autonomous-system 1001
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, **show routing-instances**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@PE1# show interfaces
ge-1/2/0 {
  unit 2 {
    family inet {
      address 10.1.1.2/30;
    }
    family mpls;
  }
}
ge-1/2/1 {
  unit 5 {
    family inet {
      address 10.1.1.5/30;
    }
    family mpls;
  }
}
vt-1/2/0 {
  unit 2 {
    family inet;
  }
}
lo0 {
  unit 2 {
    family inet {
```

```
        address 1.1.1.2/32;
    }
}
unit 102 {
    family inet {
        address 100.1.1.2/32;
    }
}
}

user@PE1# show protocols
mpls {
    interface ge-1/2/1.5;
}
bgp {
    group ibgp {
        type internal;
        local-address 1.1.1.2;
        family inet-vpn {
            any;
        }
        family inet-mvpn {
            signaling;
        }
        neighbor 1.1.1.4;
        neighbor 1.1.1.5;
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.2 {
            passive;
        }
        interface ge-1/2/1.5;
    }
}
ldp {
    interface ge-1/2/1.5;
    p2mp;
}

user@PE1# show policy-options
policy-statement parent_vpn_routes {
    from protocol bgp;
    then accept;
}

user@PE1# show routing-instances
vpn-1 {
    instance-type vrf;
    interface ge-1/2/0.2;
    interface vt-1/2/0.2;
    interface lo0.102;
    route-distinguisher 100:100;
    provider-tunnel {
        ldp-p2mp;
    }
}
```

```

vrf-target target:1:1;
protocols {
  ospf {
    export parent_vpn_routes;
    area 0.0.0.0 {
      interface lo0.102 {
        passive;
      }
      interface ge-1/2/0.2;
    }
  }
  pim {
    sglimit {
      family inet {
        maximum 100;
        threshold 70;
        log-interval 10;
      }
    }
    rp {
      register-limit {
        family inet {
          maximum 100;
          threshold 80;
          log-interval 10;
        }
      }
      group-rp-mapping {
        family inet {
          maximum 100;
          threshold 80;
          log-interval 10;
        }
      }
      static {
        address 100.1.1.2;
      }
    }
    interface ge-1/2/0.2 {
      mode sparse;
    }
  }
  mvpn;
}

```

```

user@PE1# show routing-options
router-id 1.1.1.2;
autonomous-system 1001;

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Monitoring the PIM State Information

Purpose Verify that the counters are set as expected and are not exceeding the configured limits.

Action From operational mode, enter the **show pim statistics** command.

```
user@PE1> show pim statistics instance vpn-1
PIM Message type      Received      Sent  Rx errors
V2 Hello              393          390      0
...
V4 (S,G) Maximum                        100
V4 (S,G) Accepted                       0
V4 (S,G) Threshold                      70
V4 (S,G) Log Interval                   10
V4 (grp-prefix, RP) Maximum             100
V4 (grp-prefix, RP) Accepted             0
V4 (grp-prefix, RP) Threshold            80
V4 (grp-prefix, RP) Log Interval         10
V4 Register Maximum                     100
V4 Register Accepted                     0
V4 Register Threshold                    80
V4 Register Log Interval                 10
```

Meaning The V4 (S,G) Maximum field shows the maximum number of (S,G) IPv4 multicast routes accepted for the VPN routing instance. If this number is met, additional (S,G) entries are not accepted.

The V4 (S,G) Accepted field shows the number of accepted (S,G) IPv4 multicast routes.

The V4 (S,G) Threshold field shows the threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv4 multicast routes accepted by the device).

The V4 (S,G) Log Interval field shows the time (in seconds) between consecutive log messages.

The V4 (grp-prefix, RP) Maximum field shows the maximum number of group-to-rendezvous point (RP) IPv4 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.

The V4 (grp-prefix, RP) Accepted field shows the number of accepted group-to-RP IPv4 multicast mappings.

The V4 (grp-prefix, RP) Threshold field shows the threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv4 multicast mappings accepted by the device).

The V4 (grp-prefix, RP) Log Interval field shows the time (in seconds) between consecutive log messages.

The V4 Register Maximum field shows the maximum number of IPv4 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted. You configure the register limits on the RP.

The V4 Register Accepted field shows the number of accepted IPv4 PIM registers.

The V4 Register Threshold field shows the threshold at which a warning message is logged (percentage of the maximum number of IPv4 PIM registers accepted by the device).

The V4 Register Log Interval field shows the time (in seconds) between consecutive log messages.

Related Documentation

- [Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces on page 3676](#)
- [Examples: Configuring the Multicast Forwarding Cache](#)
- [Example: Configuring MSDP with Active Source Limits and Mesh Groups](#)

Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs

- [Understanding Redundant Virtual Tunnel Interfaces in MBGP MVPNs on page 5211](#)
- [Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs on page 5212](#)

Understanding Redundant Virtual Tunnel Interfaces in MBGP MVPNs

In multiprotocol BGP (MBGP) multicast VPNs (MVPNs), VT interfaces are needed for multicast traffic on routing devices that function as combined provider edge (PE) and provider core (P) routers to optimize bandwidth usage on core links. VT interfaces prevent traffic replication when a P router also acts as a PE router (an exit point for multicast traffic).

Starting in Junos OS Release 12.3, you can configure up to eight VT interfaces in a routing instance, thus providing Tunnel PIC redundancy inside the same multicast VPN routing instance. When the active VT interface fails, the secondary one takes over, and you can continue managing multicast traffic with no duplication.

Redundant VT interfaces are supported with RSVP point-to-multipoint provider tunnels as well as multicast LDP provider tunnels. This feature also works for extranets.

You can configure one of the VT interfaces to be the primary interface. If a VT interface is configured as the primary, it becomes the next hop that is used for traffic coming in from the core on the label-switched path (LSP) into the routing instance. When a VT interface is configured to be primary and the VT interface is used for both unicast and multicast traffic, only the multicast traffic is affected.

If no VT interface is configured to be the primary or if the primary VT interface is unusable, one of the usable configured VT interfaces is chosen to be the next hop that is used for traffic coming in from the core on the LSP into the routing instance. If the VT interface in use goes down for any reason, another usable configured VT interface in the routing instance is chosen. When the VT interface in use changes, all multicast routes in the instance also switch their reverse-path forwarding (RPF) interface to the new VT interface to allow the traffic to be received.

To realize the full benefit of redundancy, we recommend that when you configure multiple VT interfaces, at least one of the VT interfaces be on a different Tunnel PIC from the other VT interfaces. However, Junos OS does not enforce this.

Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs

This example shows how to configure redundant virtual tunnel (VT) interfaces in multiprotocol BGP (MBGP) multicast VPNs (MVPNs). To configure, include multiple VT interfaces in the routing instance and, optionally, apply the **primary** statement to one of the VT interfaces.

- [Requirements on page 5212](#)
- [Overview on page 5212](#)
- [Configuration on page 5212](#)
- [Verification on page 5219](#)

Requirements

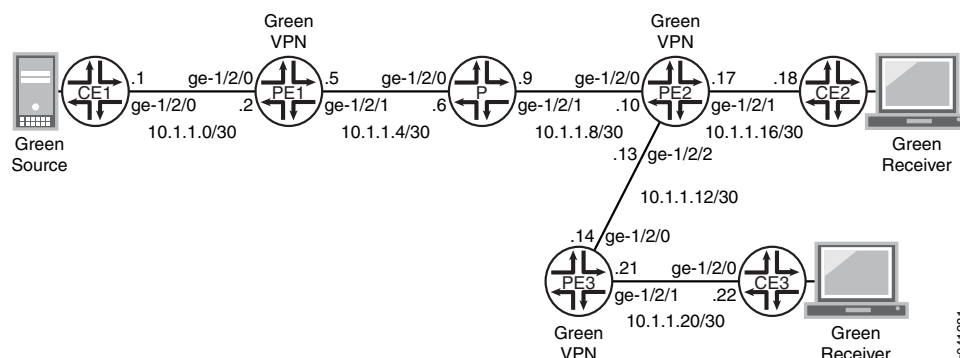
The routing device that has redundant VT interfaces configured must be running Junos OS Release 12.3 or later.

Overview

In this example, Device PE2 has redundant VT interfaces configured in a multicast LDP routing instance, and one of the VT interfaces is assigned to be the primary interface.

[Figure 117 on page 5212](#) shows the topology used in this example.

Figure 117: Multiple VT Interfaces in MBGP MVPN Topology



“CLI Quick Configuration” on page 5212 shows the configuration for the customer edge (CE), provider (P), and provider edge (PE) devices in [Figure 117 on page 5212](#). The section “Step-by-Step Procedure” on page 5215 describes the steps on Device PE2.

Configuration**CLI Quick Configuration**

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

Device CE1

```
set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.1/30
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.1/32
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.0
```

```

set protocols pim rp static address 100.1.1.2
set protocols pim interface all
set routing-options router-id 1.1.1.1

```

Device CE2

```

set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.18/30
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.6/32
set protocols sap listen 224.1.1.1
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.0
set protocols pim rp static address 100.1.1.2
set protocols pim interface all
set routing-options router-id 1.1.1.6

```

Device CE3

```

set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.22/30
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.7/32
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.0
set protocols pim rp static address 100.1.1.2
set protocols pim interface all
set routing-options router-id 1.1.1.7

```

Device P

```

set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.6/30
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 10.1.1.9/30
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.3/32
set protocols mpls interface ge-1/2/0.0
set protocols mpls interface ge-1/2/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.0
set protocols ospf area 0.0.0.0 interface ge-1/2/1.0
set protocols ldp interface ge-1/2/0.0
set protocols ldp interface ge-1/2/1.0
set protocols ldp p2mp
set routing-options router-id 1.1.1.3

```

Device PE1

```

set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.2/30
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 10.1.1.5/30
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces vt-1/2/0 unit 2 family inet
set interfaces lo0 unit 0 family inet address 1.1.1.2/32
set interfaces lo0 unit 1 family inet address 100.1.1.2/32
set protocols mpls interface ge-1/2/1.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 1.1.1.2
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 1.1.1.4
set protocols bgp group ibgp neighbor 1.1.1.5
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/1.0
set protocols ldp interface ge-1/2/1.0

```

```

set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface ge-1/2/0.0
set routing-instances vpn-1 interface vt-1/2/0.2 multicast
set routing-instances vpn-1 interface lo0.1
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 provider-tunnel ldp-p2mp
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.1 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/0.0
set routing-instances vpn-1 protocols pim rp static address 100.1.1.2
set routing-instances vpn-1 protocols pim interface ge-1/2/0.0 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 1.1.1.2
set routing-options autonomous-system 1001

```

Device PE2

```

set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.10/30
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/2 unit 0 family inet address 10.1.1.13/30
set interfaces ge-1/2/2 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 10.1.1.17/30
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces vt-1/1/0 unit 0 family inet
set interfaces vt-1/2/1 unit 0 family inet
set interfaces lo0 unit 0 family inet address 1.1.1.4/32
set interfaces lo0 unit 1 family inet address 100.1.1.4/32
set protocols mpls interface ge-1/2/0.0
set protocols mpls interface ge-1/2/2.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 1.1.1.4
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 1.1.1.2
set protocols bgp group ibgp neighbor 1.1.1.5
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.0
set protocols ospf area 0.0.0.0 interface ge-1/2/2.0
set protocols ldp interface ge-1/2/0.0
set protocols ldp interface ge-1/2/2.0
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface vt-1/1/0.0 multicast
set routing-instances vpn-1 interface vt-1/1/0.0 primary
set routing-instances vpn-1 interface vt-1/2/1.0 multicast
set routing-instances vpn-1 interface ge-1/2/1.0
set routing-instances vpn-1 interface lo0.1
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.1 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/1.0

```

```

set routing-instances vpn-1 protocols pim rp static address 100.1.1.2
set routing-instances vpn-1 protocols pim interface ge-1/2/1.0 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 1.1.1.4
set routing-options autonomous-system 1001

```

Device PE3

```

set interfaces ge-1/2/0 unit 0 family inet address 10.1.1.14/30
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 10.1.1.21/30
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces vt-1/2/0 unit 5 family inet
set interfaces lo0 unit 0 family inet address 1.1.1.5/32
set interfaces lo0 unit 1 family inet address 100.1.1.5/32
set protocols mpls interface ge-1/2/0.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 1.1.1.5
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 1.1.1.2
set protocols bgp group ibgp neighbor 1.1.1.4
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.0
set protocols ldp interface ge-1/2/0.0
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface vt-1/2/0.5 multicast
set routing-instances vpn-1 interface ge-1/2/1.0
set routing-instances vpn-1 interface lo0.1
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.1 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/1.0
set routing-instances vpn-1 protocols pim rp static address 100.1.1.2
set routing-instances vpn-1 protocols pim interface ge-1/2/1.0 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 1.1.1.5
set routing-options autonomous-system 1001

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see [“Using the CLI Editor in Configuration Mode”](#) on page 4704 in the *CLI User Guide*.

To configure redundant VT interfaces in an MBGP MVPN:

1. Configure the physical interfaces and loopback interfaces.

```

[edit interfaces]
user@PE2# set ge-1/2/0 unit 0 family inet address 10.1.1.10/30
user@PE2# set ge-1/2/0 unit 0 family mpls

user@PE2# set ge-1/2/2 unit 0 family inet address 10.1.1.13/30
user@PE2# set ge-1/2/2 unit 0 family mpls

```

```
user@PE2# set ge-1/2/1 unit 0 family inet address 10.1.1.17/30
user@PE2# set ge-1/2/1 unit 0 family mpls
```

```
user@PE2# set lo0 unit 0 family inet address 1.1.1.4/32
user@PE2# set lo0 unit 1 family inet address 100.1.1.4/32
```

2. Configure the VT interfaces.

Each VT interface is configurable under one routing instance.

```
[edit interfaces]
user@PE2# set vt-1/1/0 unit 0 family inet
user@PE2# set vt-1/2/1 unit 0 family inet
```

3. Configure MPLS on the physical interfaces.

```
[edit protocols mpls]
user@PE2# set interface ge-1/2/0.0
user@PE2# set interface ge-1/2/2.0
```

4. Configure BGP.

```
[edit protocols bgp group ibgp]
user@PE2# set type internal
user@PE2# set local-address 1.1.1.4
user@PE2# set family inet-vpn any
user@PE2# set family inet-mvpn signaling
user@PE2# set neighbor 1.1.1.2
user@PE2# set neighbor 1.1.1.5
```

5. Configure an interior gateway protocol.

```
[edit protocols ospf area 0.0.0.0]
user@PE2# set interface lo0.0 passive
user@PE2# set interface ge-1/2/0.0
user@PE2# set interface ge-1/2/2.0
```

6. Configure LDP.

```
[edit protocols ldp]
user@PE2# set interface ge-1/2/0.0
user@PE2# set interface ge-1/2/2.0
user@PE2# set p2mp
```

7. Configure the routing policy.

```
[edit policy-options policy-statement parent_vpn_routes]
user@PE2# set from protocol bgp
user@PE2# set then accept
```

8. Configure the routing instance.

```
[edit routing-instances vpn-1]
user@PE2# set instance-type vrf
user@PE2# set interface ge-1/2/1.0
user@PE2# set interface lo0.1
user@PE2# set route-distinguisher 100:100
user@PE2# set vrf-target target:1:1
user@PE2# set protocols ospf export parent_vpn_routes
user@PE2# set protocols ospf area 0.0.0.0 interface lo0.1 passive
user@PE2# set protocols ospf area 0.0.0.0 interface ge-1/2/1.0
```

```

user@PE2# set protocols pim rp static address 100.1.1.2
user@PE2# set protocols pim interface ge-1/2/1.0 mode sparse
user@PE2# set protocols mvpn

```

9. Configure redundant VT interfaces in the routing instance.

Make vt-1/1/0.0 the primary interface.

```

[edit routing-instances vpn-1]
user@PE2# set interface vt-1/1/0.0 multicast primary
user@PE2# set interface vt-1/2/1.0 multicast

```

10. Configure the router ID and autonomous system (AS) number.

```

[edit routing-options]
user@PE2# set router-id 1.1.1.4
user@PE2# set autonomous-system 1001

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, **show routing-instances**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

user@PE2# show interfaces
ge-1/2/0 {
  unit 0 {
    family inet {
      address 10.1.1.10/30;
    }
    family mpls;
  }
}
ge-1/2/2 {
  unit 0 {
    family inet {
      address 10.1.1.13/30;
    }
    family mpls;
  }
}
ge-1/2/1 {
  unit 0 {
    family inet {
      address 10.1.1.17/30;
    }
    family mpls;
  }
}
vt-1/1/0 {
  unit 0 {
    family inet;
  }
}
vt-1/2/1 {
  unit 0 {
    family inet;
  }
}

```

```
}
lo0 {
  unit 0 {
    family inet {
      address 1.1.1.4/32;
    }
  }
  unit 1 {
    family inet {
      address 100.1.1.4/32;
    }
  }
}

user@PE2# show protocols
mpls {
  interface ge-1/2/0.0;
  interface ge-1/2/2.0;
}
bgp {
  group ibgp {
    type internal;
    local-address 1.1.1.4;
    family inet-vpn {
      any;
    }
    family inet-mvpn {
      signaling;
    }
    neighbor 1.1.1.2;
    neighbor 1.1.1.5;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.0 {
      passive;
    }
    interface ge-1/2/0.0;
    interface ge-1/2/2.0;
  }
}
ldp {
  interface ge-1/2/0.0;
  interface ge-1/2/2.0;
  p2mp;
}

user@PE2# show policy-options
policy-statement parent_vpn_routes {
  from protocol bgp;
  then accept;
}

user@PE2# show routing-instances
vpn-1 {
  instance-type vrf;
```



```

interface vt-1/1/0.0 {
    multicast;
    primary;
}
interface vt-1/2/1.0 {
    multicast;
}
interface ge-1/2/1.0;
interface lo0.1;
route-distinguisher 100:100;
vrf-target target:1:1;
protocols {
    ospf {
        export parent_vpn_routes;
        area 0.0.0.0 {
            interface lo0.1 {
                passive;
            }
            interface ge-1/2/1.0;
        }
    }
    pim {
        rp {
            static {
                address 100.1.1.2;
            }
        }
        interface ge-1/2/1.0 {
            mode sparse;
        }
    }
}
mvpn;
}

user@PE2# show routing-options
router-id 1.1.1.4;
autonomous-system 1001;

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.



NOTE: The `show multicast route extensive instance instance-name` command also displays the VT interface in the multicast forwarding table when multicast traffic is transmitted across the VPN.

Checking the LSP Route

Purpose Verify that the expected LT interface is assigned to the LDP-learned route.

Action 1. From operational mode, enter the `show route table mpls` command.

```

user@PE2> show route table mpls
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

0          *[MPLS/0] 02:09:36, metric 1
            Receive
1          *[MPLS/0] 02:09:36, metric 1
            Receive
2          *[MPLS/0] 02:09:36, metric 1
            Receive
13         *[MPLS/0] 02:09:36, metric 1
            Receive
299776     *[LDP/9] 02:09:14, metric 1
            > via ge-1/2/0.0, Pop
299776(S=0) *[LDP/9] 02:09:14, metric 1
            > via ge-1/2/0.0, Pop
299792     *[LDP/9] 02:09:09, metric 1
            > via ge-1/2/2.0, Pop
299792(S=0) *[LDP/9] 02:09:09, metric 1
            > via ge-1/2/2.0, Pop
299808     *[LDP/9] 02:09:04, metric 1
            > via ge-1/2/0.0, Swap 299808
299824     *[VPN/170] 02:08:56
            > via ge-1/2/1.0, Pop
299840     *[VPN/170] 02:08:56
            > via ge-1/2/1.0, Pop
299856     *[VPN/170] 02:08:56
            receive table vpn-1.inet.0, Pop
299872     *[LDP/9] 02:08:54, metric 1
            > via vt-1/1/0.0, Pop
            > via ge-1/2/2.0, Swap 299872

```

- From configuration mode, change the primary VT interface by removing the **primary** statement from the vt-1/1/0.0 interface and adding it to the vt-1/2/1.0 interface.

```

[edit routing-instances vpn-1]
user@PE2# delete interface vt-1/1/0.0 primary
user@PE2# set interface vt-1/2/1.0 primary
user@PE2# commit

```

- From operational mode, enter the **show route table mpls** command.

```

user@PE2> show route table mpls
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

0          *[MPLS/0] 02:09:36, metric 1
            Receive
1          *[MPLS/0] 02:09:36, metric 1
            Receive
2          *[MPLS/0] 02:09:36, metric 1
            Receive
13         *[MPLS/0] 02:09:36, metric 1
            Receive
299776     *[LDP/9] 02:09:14, metric 1
            > via ge-1/2/0.0, Pop
299776(S=0) *[LDP/9] 02:09:14, metric 1
            > via ge-1/2/0.0, Pop
299792     *[LDP/9] 02:09:09, metric 1
            > via ge-1/2/2.0, Pop
299792(S=0) *[LDP/9] 02:09:09, metric 1
            > via ge-1/2/2.0, Pop

```

```

299808      *[LDP/9] 02:09:04, metric 1
              >   via ge-1/2/0.0, Swap 299808
299824      *[VPN/170] 02:08:56
              >   via ge-1/2/1.0, Pop
299840      *[VPN/170] 02:08:56
              >   via ge-1/2/1.0, Pop
299856      *[VPN/170] 02:08:56
              receive table vpn-1.inet.0, Pop
299872      *[LDP/9] 02:08:54, metric 1
              >   via vt-1/2/1.0, Pop
              via ge-1/2/2.0, Swap 299872

```

Meaning With the original configuration, the output shows the vt-1/1/0.0 interface. If you change the primary interface to vt-1/2/1.0, the output shows the vt-1/2/1.0 interface.

Related Documentation

- [Multiprotocol BGP MVPNs Overview](#)

Configuration Tasks

- [Introduction to Configuring MBGP MVPNs on page 5221](#)
- [Configuring Routing Instances for an MBGP MVPN on page 5223](#)
- [Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs on page 5225](#)
- [Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs on page 5227](#)
- [Limiting Routes to Be Advertised by an MVPN VRF Instance on page 5228](#)
- [Configuring VRF Route Targets for Routing Instances for an MBGP MVPN on page 5229](#)
- [Configuring NLRI Parameters for an MBGP MVPN on page 5232](#)
- [Configuring PIM Provider Tunnels for an MBGP MVPN on page 5233](#)
- [Configuring PIM-SSM GRE Selective Provider Tunnels on page 5233](#)
- [Configuring Point-to-Multipoint LSPs for an MBGP MVPN on page 5234](#)
- [Configuring a Selective Provider Tunnel Using Wildcards on page 5240](#)
- [Example: Configuring Selective Provider Tunnels Using Wildcards on page 5241](#)
- [Configuring Internet Multicast Using Ingress Replication Provider Tunnels on page 5243](#)
- [Configuring GRE Tunnels for Layer 3 VPNs on page 5246](#)

Introduction to Configuring MBGP MVPNs

You configure multiprotocol BGP-based (MBGP) multicast VPNs (MVPNs) at a number of different hierarchy levels within the Junos OS. However, a majority of MBGP MVPN statements are configured within a routing instance as follows:

```

description text;
instance-type vrf;
interface interface-name;
route-distinguisher (as-number:number | ip-address:number);
vrf-export [policy-names];
vrf-import [policy-names];
vrf-target (community | export community-name | import community-name);

```

```
protocols {
  mvpn {
    mvpn-mode (rpt-spt | spt-only);
    receiver-site;
    sender-site;
    route-target {
      export-target {
        target target-community;
        unicast;
      }
      import-target {
        target {
          target-value;
          receiver target-value;
          sender target-value;
        }
        unicast {
          receiver;
          sender;
        }
      }
    }
  }
}
provider-tunnel {
  ldp-p2mp;
  pim-ssm {
    group-address address;
  }
  rsvp-te {
    label-switched-path-template {
      (default-template | lsp-template-name);
    }
    static-lsp lsp-name;
  }
  selective {
    group mcast-prefix/prefix-length {
      source ip-prefix/prefix-length {
        ldp-p2mp;
        pim-ssm {
          group-range mcast-prefix;
        }
        rsvp-te {
          label-switched-path-template {
            (default-template | lsp-template-name);
          }
          static-lsp point-to-multipoint-lsp-name;
        }
        threshold-rate kbps;
      }
    }
    wildcard-source {
      ldp-p2mp;
      pim-ssm {
        group-range mcast-prefix;
      }
      rsvp-te {
```

```

        label-switched-path-template {
            (default-template | lsp-template-name);
        }
        static-lsp point-to-multipoint-lsp-name;
    }
    threshold-rate kbps;
}
tunnel-limit number;
wildcard-group-inet {
    wildcard-source {
        ldp-p2mp;
        pim-ssm {
            group-range multicast-prefix;
        }
        rsvp-te {
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
            static-lsp lsp-name;
        }
        threshold-rate number;
    }
}
wildcard-group-inet6 {
    wildcard-source {
        ldp-p2mp;
        pim-ssm {
            group-range multicast-prefix;
        }
        rsvp-te {
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
            static-lsp lsp-name;
        }
        threshold-rate number;
    }
}
threshold-rate number;
}

```

You can include these statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

Related Documentation

- *Multicast over Layer 3 VPNs*

Configuring Routing Instances for an MBGP MVPN

To configure MBGP MVPNs, include the **mvpn** statement:

```
mvpn {
  mvpn-mode (rpt-spt | spt-only);
  receiver-site;
  route-target {
    export-target {
      target target-community;
      unicast;
    }
    import-target {
      target {
        target-value;
        receiver target-value;
        sender target-value;
      }
      unicast {
        receiver;
        sender;
      }
    }
  }
  sender-site;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]

By default an MBGP MVPN routing instance is associated with both the multicast sender and the receiver sites. If you configure the **receiver-site** option, the routing instance is associated with only multicast receiver sites. Configuring the **sender-site** option associates the routing instance with only multicast sender sites.



NOTE: When you configure the routing instance for the MBGP MVPN, you must configure MPLS LSPs (either RSVP-signaled or LDP-signaled) between the PE routers of the routing instance to ensure VPN unicast connectivity. Point-to-multipoint LSPs are used for multicast data forwarding only.

Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs

For MBGP MVPNs (also referred to as next-generation Layer 3 multicast VPNs), the default mode of operation is shortest path tree only (SPT-only) mode. In SPT-only mode, the active multicast sources are learned through multicast VPN source-active routes. This mode of operation is described in section 14 of the BGP-MVPN draft (draft-ietf-l3vpn-2547bis-mcast-bgp-00.txt).

In contrast to SPT-only mode, rendezvous point tree (RPT)-SPT mode (also known as shared-tree data distribution) supports the native PIM model of transmitting (*G) messages from the receiver to the RP for intersite shared-tree join messages.

In SPT-only mode, when a PE router receives a (*, C-G) join message, the router looks for an active source transmitting data to the customer group. If the PE router has a source-active route for the customer group, the router creates a source tree customer multicast route and sends the route to the PE router connected to the VPN site with the source. The source is determined by MVPN's single-forwarder election. When a receiver sends a (*G) join message in a VPN site, the (*G) join message only travels as far as the PE router. After the join message is converted to a type 6 multicast route, which is equivalent to a (S,G) join message, the route is installed with the no-advertise community setting.



NOTE: The MVPN single-forwarder election follows the rule documented in section 9.1.1 of the BGP-MVPN draft (draft-ietf-l3vpn-2547bis-mcast-bgp-00.txt). The single-forwarder election winner is based on the following rules:

- If the active unicast route to the source is through the interface, then this route is used to determine the upstream multicast hop (UMH).
- If the active unicast route to the source is a VPN route, MVPN selects the UMH based on the highest IP address in the route import community for the VPN routes, and the local master loopback address for local VRF routes.

Single-forwarder election guarantees selection of a unique forwarder for a given customer source (C-S). The upstream PE router might differ for the source tree and the shared tree because the election is based on the customer source and C-RP, respectively. Although the single-forwarder election is sufficient for SPT-only mode, the alternative RPT-SPT mode involves procedures to prevent duplicate traffic from being sent on the shared tree and the source tree. These procedures might require administrator-configured parameters to reduce duplicate traffic and reduce blackholes during RPT to SPT switch and the reverse.

In SPT-only mode, when a source is active, PIM creates a register state for the source both on the DR and on the C-RP (or on a PE router that is running Multicast Source Discovery Protocol [MSDP] between itself and the C-RP). After the register states are created, MVPN creates a source-active route. These type 5 source-active routes are installed on all PE routers. When the egress PE router with the (*G) join message receives the source-active route, it has two routes that it can combine to produce the (S,G)

multicast route. The type 6 route informs the PE router that a receiver is interested in group G. The source active route informs the PE router that a source S is transmitting data to group G. MVPN combines this information to produce a multicast join message and advertises this to the ingress PE router, as determined by the single-forwarder election.



NOTE: If you are using a Trio FPC, the MVPN sender PE router can not have a local receiver.

For some service providers, the SPT-only implementation is not ideal because it creates a restriction on C-RP configuration. For a PE router to create customer multicast routes from (* C-G) join messages, the router must learn about active sources through MVPN type 5 source-active routes. These source-active routes can be originated only by a PE router. This means that a PE router in the MVPN must learn about all PIM register messages sent to the RP, which is possible only in the following cases:

- The C-RP is colocated on one of the PEs in the MVPN.
- MSDP is run between the C-RP and the VRF instance on one of the PE routers in the MVPN.

If this restriction is not acceptable, providers can use RPT-SPT mode instead of the default SPT-only mode. However, because SPT-only mode does not transmit (*G) routes between VPN sites, SPT-only mode has the following advantages over RPT-SPT mode:

- Simplified operations by exchanging and processing only source-tree customer multicast routes among PE routers
- Simplified operations by eliminating the need for the service provider to suppress MVPN transient duplicates during the switch from RPT to SPT
- Less control plane overhead in the service provider space by limiting the type of customer multicast routes exchanged, which results in more scalable deployments
- More stable traffic patterns in the backbone without the traffic shifts involved in the RPT-SPT mode
- Easier maintenance in the service provider space due to less state information

To configure SPT-only mode:

1. Explicitly configure SPT-only mode:

```
[edit routing-instances routing-instance-name protocols mvpn mvpn-mode]  
user@router# set spt-only
```

2. Include the **spt-only** statement for all VRFs that make up the VPN.

Related Documentation

- [Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs on page 5227](#)
- [Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 5153](#)

Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs

For MBGP MVPNs (also referred to as next-generation Layer 3 multicast VPNs), the default mode of operation supports only intersite shortest-path trees (SPTs) for customer PIM (C-PIM) join messages. It does not support rendezvous-point trees (RPTs) for C-PIM join messages. The default mode of operation provides advantages, but it requires either that the customer rendezvous point (C-RP) be located on a PE router or that the Multicast Source Discovery Protocol (MSDP) be used between the C-RP and a PE router so that the PE router can learn about active sources advertised by other PE routers.

If the default mode is not suitable for your environment, you can configure RPT-SPT mode (also known as *shared-tree data distribution*), as documented in section 13 of the BGP-MVPN draft (draft-ietf-l3vpn-2547bis-mcast-bgp-00.txt). RPT-SPT mode supports the native PIM model of transmitting (*G) messages from the receiver to the RP for intersite shared-tree join messages. This means that the type 6 (*G) routes get transmitted from one PE router to another. In RPT-SPT mode, the shared-tree multicast routes are advertised from an egress PE router to the upstream router connected to the VPN site with the C-RP. The single-forwarder election is performed for the C-RP rather than for the source. The egress PE router takes the upstream hop to advertise the (*G) and sends the type 6 route toward the upstream PE router. To send the data on the RPT, either inclusive or selective provider tunnels can be used. After the data starts flowing on the RPT, the last-hop router switches to SPT mode, unless you include the **spt-threshold infinity** statements in the configuration.



NOTE: The MVPN single-forwarder election follows the rule documented in section 9.1.1 of the BGP-MVPN draft (draft-ietf-l3vpn-2547bis-mcast-bgp-00.txt). The single-forwarder election winner is based on the following rules:

- If the active unicast route to the source is through the interface, then this route is used to determine the upstream multicast hop (UMH).
- If the active unicast route to the source is a VPN route, MVPN selects the UMH based on the highest IP address in the route import community for the VPN routes, and the local master loopback address for local VRF routes.

The switch to SPT mode is performed by PIM and not by MVPN type 5 and type 6 routes. After the last-hop router switches to SPT mode, the SPT (S,G) join messages follow the same rules as the SPT-only default mode.

The advantage of RPT-SPT mode is that it provides a method for PE routers to discover sources in the multicast VPN when the C-RP is located on the customer site instead of on a PE router. Because the shared C-tree is established between VPN sites, there is no need to run MSDP between the C-RP and the PE routers. RPT-SPT mode also enables egress PE routers to switch to receiving data from the PE connected to the source after the source information is learned, instead of receiving data from the RP.



CAUTION: When you configure RPT-SPT mode, receivers or sources directly attached to the PE router are not supported. As a workaround, place a CE router between any receiver or source and the PE router.

To configure RPT-SPT mode:

1. Enable shared-tree data distribution:

```
[edit routing-instances routing-instance-name protocols mvpn mvpn-mode]
user@router# set rpt-spt
```

2. Include the **rpt-spt** statement for all VRFs that make up the VPN.

Related Documentation

- [Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs on page 5225](#)
- [Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 5153](#)

Limiting Routes to Be Advertised by an MVPN VRF Instance

If a hub-and-spoke deployment uses one VPN routing and forwarding (VRF) routing instance for unicast routing and a separate VRF for MVPN routing, you need to limit the PE routers at the hub site to advertise only IPv4 MVPN routes, only IPv6 MVPN routes, or both. This is necessary to prevent the multicast VRF instance from advertising unicast VPN routes to other PE routers.



NOTE: This configuration does not prevent the exportation of VPN routes to other VRF instances on the same router if the **auto-export** statement is included in the **[edit routing-options]** hierarchy.

To configure a VRF routing instance with the name **green** to advertise MVPN routes from both the **inet** and **inet6** address families, perform the following steps:

1. Configure the VRF routing instance to advertise IPv4 routes.

```
user@host# set routing-instances green vrf-advertise-selective family inet-mvpn
```

2. Configure the VRF routing instance to advertise IPv6 routes.

```
user@host# set routing-instances green vrf-advertise-selective family inet6-mvpn
```

After the configuration is committed, only the MVPN routes for the specified address families are advertised from the VRF instance to remote PE routers. To remove the restriction on routes being advertised, delete the **vrf-advertise-selective** statement.



NOTE: You cannot include the `vrf-advertise-selective` statement and the `no-vrf-advertise` statement in the same VRF configuration. However, if you configure the `vrf-advertise-selective` statement without any of its options, the router has the same behavior as if you configured the `no-vrf-advertise` statement. VPN routes are prevented from being advertised from a VRF routing instance to the remote PE routers.

Related Documentation

- [family on page 5261](#)
- [inet-mvpn on page 5267](#)
- [inet6-mvpn on page 5268](#)
- `no-vrf-advertise`
- [vrf-advertise-selective on page 5295](#)

Configuring VRF Route Targets for Routing Instances for an MBGP MVPN

By default, the VPN routing and forwarding (VRF) import and export route targets (configured either using VRF import and export policies or using the **vrf-target** statement) are used for importing and exporting routes with the MBGP MVPN network layer reachability information (NLRI).

You can use the **export-target** and **import-target** statements to override the default VRF import and export route targets. Export and import targets can also be specified specifically for sender sites or receiver sites, or can be borrowed from a configured unicast route target. Note that a sender site export route target is always advertised when security association routes are exported.



NOTE: When you configure an MBGP MVPN routing instance, you should not configure a target value for an MBGP MVPN specific route target that is identical to a target value for a unicast route target configured in another routing instance.

Specifying route targets in the MBGP MVPN NLRI for sender and receiver sites is useful when there is a mix of sender only, receiver only, and sender and receiver sites. A sender site route target is used for exporting automatic discovery routes by a sender site and for importing automatic discovery routes by a receiver site. A receiver site route target is used for exporting routes by a receiver site and importing routes by a sender site. A sender and receiver site exports and imports routes with both route targets.

A provider edge (PE) router with sites in a specific MBGP MVPN must determine whether a received automatic discovery route is from a sender site or receiver site based on the following:

- If the PE router is configured to be only in a sender site, route targets are imported only from receiver sites. Imported automatic discovery routes must be from a receiver site.

- If the PE router is configured to be only in a receiver site, route targets are imported only from sender sites. Imported automatic discovery routes must be from a sender site.
- If a PE router is configured to be in both sender sites and receiver sites, these guidelines apply:
 - Along with an import route target, you can optionally configure whether the route target is from a receiver or a sender site.
 - If a configuration is not provided, an imported automatic discovery route is treated as belonging to both the sender site set and the receiver site set.

To configure a route target for the MBGP MVPN routing instance, include the **route-target** statement:

```
route-target {  
  export-target {  
    target target-community;  
    unicast;  
  }  
  import-target {  
    target {  
      target-value;  
      receiver target-value;  
      sender target-value;  
    }  
    unicast {  
      receiver;  
      sender;  
    }  
  }  
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols mvpn]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn]

The following sections describes how to configure the export target and the import target for an MBGP MVPN:

- [Configuring the Export Target for an MBGP MVPN on page 5230](#)
- [Configuring the Import Target for an MBGP MVPN on page 5231](#)

Configuring the Export Target for an MBGP MVPN

To configure an export target, include the **export-target** statement:

```
export-target {  
  target target-community;  
  unicast;  
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols mvpn route-target]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn route target]

Configure the **target** option to specify the export target community. Configure the **unicast** option to use the same target community that has been specified for unicast.

Configuring the Import Target for an MBGP MVPN

To configure an import target, include the **import-target** statement:

```
import-target {
  target target-value {
    receiver;
    sender;
  }
  unicast {
    receiver;
    sender;
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols mvpn route-target]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn route-target]

The following sections describe how to configure the import target and unicast parameters:

- [Configuring the Import Target Receiver and Sender for an MBGP MVPN on page 5231](#)
- [Configuring the Import Target Unicast Parameters for an MBGP MVPN on page 5232](#)

Configuring the Import Target Receiver and Sender for an MBGP MVPN

To configure the import target community, include the **target** statement and specify the target community. The target community must be in the format **target:x:y**. The x value is either an IP address or an AS number followed by an optional L to indicate a 4 byte AS number, and y is a number (for example, **target:123456L:100**)

```
target target-value {
  receiver;
  sender;
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols mvpn route-target import-target]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn route-target import-target]

You can specify the target community used when importing either receiver site sets or sender site sets by including one of the following statements:

- **receiver**—Specify the target community used when importing receiver site sets.
- **sender**—Specify the target community used when importing sender site sets.

Configuring the Import Target Unicast Parameters for an MBGP MVPN

To configure a unicast target community as the import target, include the **unicast** statement:

```
unicast {  
    receiver;  
    sender;  
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols mvpn route-target import-target]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn route-target import-target]

You can specify the unicast target community used when importing either receiver site sets or sender site sets by including one of the following statements:

- **receiver**—Specify the unicast target community used when importing receiver site sets.
- **sender**—Specify the unicast target community used when importing sender site sets.

Configuring NLRI Parameters for an MBGP MVPN

To enable VPN signaling where multiprotocol BGP carries multicast VPN NLRI for the IPv4 address family, include the **family inet-mvpn** statement:

```
inet-mvpn {  
    signaling {  
        accepted-prefix-limit {  
            maximum number;  
            teardown percentage {  
                idle-timeout (forever | minutes);  
            }  
        }  
        loops number;  
        prefix-limit {  
            maximum number;  
            teardown percentage {  
                idle-timeout (forever | minutes);  
            }  
        }  
    }  
}
```

To enable VPN signaling where multiprotocol BGP carries multicast VPN NLRI for the IPv6 address family, include the **family inet6-mvpn** statement:

```

inet6-mvpn {
  signaling {
    accepted-prefix-limit {
      maximum number;
      teardown percentage {
        idle-timeout (forever | minutes);
      }
    }
  }
  loops number
  prefix-limit {
    maximum number;
    teardown percentage {
      idle-timeout (forever | minutes);
    }
  }
}

```

Configuring PIM Provider Tunnels for an MBGP MVPN

To configure a Protocol Independent Multicast (PIM) sparse mode provider tunnel for a multicast VPN, include the **pim-asm** statement:

```

pim-asm {
  group-address address;
}

```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* provider-tunnel]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel]

To complete the PIM sparse mode provider tunnel configuration, you also need to specify the group address using the **group-address** option. The source address for a PIM sparse mode provider tunnel is configured to be the loopback address of the loopback interface in the **inet.0** routing table.

Configuring PIM-SSM GRE Selective Provider Tunnels

This topic describes how to configure a PIM-SSM GRE selective provider tunnel for an MBGP MVPN. A selective provider tunnel uses a point-to-multipoint LSP.

Creating a selective provider tunnel enables you to move high-rate traffic off the inclusive tunnel and deliver the multicast traffic only to receivers that request it. This improves bandwidth utilization.

To configure a PIM-SSM GRE selective provider tunnel for the 224.1.1.1/32 customer multicast group address, the 10.2.2.2/32 customer source address, and a virtual routing instance named **green**:

1. Configure the multicast group address range to be used for creating selective tunnels. The address prefix can be any valid nonreserved IPv4 multicast address range. Whether you configure a range of addresses or a single address, make sure that you configure enough group addresses for all the selective tunnels needed.

```
user@host# set routing-instances green provider-tunnel selective group 224.1.1.1/32
source 10.2.2.2/32 pim-ssm group-range 232.1.1.0/24
```

2. Configure the threshold rate in kilobits per second (Kbps) for triggering the creation of the selective tunnel. If you set the threshold rate to zero Kbps, the selective tunnel is created immediately, and the multicast traffic does not use an inclusive tunnel at all. Optionally, you can leave the threshold rate unconfigured and the result is the same as setting the threshold to zero.

```
user@host# set routing-instances green provider-tunnel selective group 224.1.1.1/32
source 10.2.2.2/32 threshold-rate 0
```

3. Configure the autonomous system number in the global routing options. This is required in MBGP MVPNs.

```
user@host# set routing-options autonomous-system 100
```

When configuring PIM-SSM GRE selective provider tunnels, keep the following in mind:

- Aggregation of multiple customer multicast routes to a single PIM S-PMSI is not supported.
- Provider tunnel multicast group addresses must be IPv4 addresses, even in configurations in which the customer multicast group and source are IPv6 addresses.

**Related
Documentation**

- [Multicast VPN Terminology on page 5158](#)
- [pim-ssm on page 5275](#)
- [group-range on page 5263](#)
- [threshold-rate on page 5290](#)

Configuring Point-to-Multipoint LSPs for an MBGP MVPN

The Junos OS supports point-to-multipoint label-switched paths (LSPs) for MBGP MVPNs. Point-to-multipoint LSPs for multicast VPNs are supported for intra-autonomous system (AS) environments (within an AS), but are not supported for inter-AS environments (between autonomous systems). A point-to-multipoint LSP is an RSVP-signaled LSP with a single source and multiple destinations.

You can configure point-to-multipoint LSPs for MBGP MVPNs as follows:

- Static point-to-multipoint LSPs—Configure static point-to-multipoint LSPs using the standard MPLS LSP statements specified at the **[edit protocols mpls]** hierarchy level. You manually configure each of the leaf nodes for the point-to-multipoint LSP.
- Dynamic point-to-multipoint LSPs using the default template—Configuring dynamic point-to-multipoint LSPs using the **default-template** option causes the leaf nodes to be discovered automatically. The leaf nodes are discovered through BGP intra-AS automatic discovery. The **default-template** option allows you to minimize the amount of configuration needed. However, it does not allow you to configure any of the standard MPLS options.
- Dynamic point-to-multipoint LSPs using a user-configured template—Configuring dynamic point-to-multipoint LSPs using a user-configured template also causes the

leaf nodes to be discovered automatically. By creating your own template for the point-to-multipoint LSPs, all of the standard MPLS features (such as bandwidth allocation and traffic engineering) can be configured.

Be aware of the following properties for the egress PE router in a point-to-multipoint LSP configured for a multicast VPN:

- Penultimate hop-popping is not used by point-to-multipoint LSPs for multicast VPNs. Only ultimate hop-popping is used.
- You must configure either the **vrf-table-label** statement or a virtual loopback tunnel interface on the egress PE router.
- If you configure the **vrf-table-label** statement on the egress PE router, and the egress PE router is also a transit router for the point-to-multipoint LSP, the penultimate hop router sends two copies of each packet over the link to the egress PE router.
- If you configure the **vrf-table-label** statement on the egress PE router, and the egress PE router is not a transit router for the point-to-multipoint LSP, the penultimate hop router can send just one copy of each packet over the link to the egress PE router.
- If you configure a virtual loopback tunnel interface on the egress PE router, and the egress PE router is also a transit router for the point-to-multipoint LSP, the penultimate hop router sends just one copy of each packet over the link to the egress PE router. A virtual loopback tunnel interface can perform two lookups on an incoming packet, one for the multicast MPLS lookup and one for the IP lookup.



NOTE: Junos OS Release 11.2 and earlier do not support point-to-multipoint LSPs with next-generation multicast VPNs on MX80 routers.

The following sections describe how to configure point-to-multipoint LSPs for MBGP MVPNs:

- [Configuring RSVP-Signaled Inclusive Point-to-Multipoint LSPs for an MBGP MVPN on page 5235](#)
- [Configuring Selective Provider Tunnels for an MBGP MVPN on page 5236](#)

Configuring RSVP-Signaled Inclusive Point-to-Multipoint LSPs for an MBGP MVPN

You can configure LDP-signaled or RSVP-signaled inclusive point-to-multipoint LSPs for MBGP MVPNs. Aggregation is not supported, so you need to configure an inclusive point-to-multipoint LSP for each sender PE router in each multicast VPN routing instance. The sender PE router is in the sender site set of the MBGP MVPN.

To configure a static RSVP-signaled inclusive point-to-multipoint LSP, include the **static-lsp** statement:

```
static-lsp lsp-name;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* provider-tunnel rsvp-te]**

- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel rsvp-te]

To configure dynamic inclusive point-to-multipoint LSPs, include the **label-switched-path-template** statement:

```
label-switched-path-template {  
  (default-template | lsp-template-name);  
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* provider-tunnel rsvp-te]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel rsvp-te]

You can configure either the **default-template** option or manually configure a point-to-multipoint LSP template and specify the template name.

Configuring Selective Provider Tunnels for an MBGP MVPN

You can configure LDP-signaled or RSVP-signaled selective point-to-multipoint LSPs (also referred to as selective provider tunnels) for MBGP MVPNs. Selective point-to-multipoint LSPs send traffic only to the receivers configured for the multicast VPNs, helping to minimize flooding in the service provider's network.

As with inclusive point-to-multipoint LSPs, you can configure both dynamic and static selective tunnels for the multicast VPN.

To configure selective point-to-multipoint provider tunnels, include the **selective** statement:

```
selective {  
  group multicast--prefix/prefix-length {  
    source ip--prefix/prefix-length {  
      ldp-p2mp;  
      pim-ssm {  
        group-range multicast-prefix;  
      }  
      rsvp-te {  
        label-switched-path-template {  
          (default-template | lsp-template-name);  
        }  
        static-lsp point-to-multipoint-lsp-name;  
      }  
      threshold-rate kbits;  
    }  
  }  
  wildcard-source {  
    ldp-p2mp;  
    pim-ssm {  
      group-range multicast-prefix;  
    }  
    rsvp-te {  
      label-switched-path-template {  
        (default-template | lsp-template-name);  
      }  
    }  
  }  
}
```

```

        }
        static-lsp point-to-multipoint-lsp-name;
    }
    threshold-rate kbits;
}
}
tunnel-limit number;
wildcard-group-inet {
    wildcard-source {
        ldp-p2mp;
        pim-ssm {
            group-range multicast-prefix;
        }
        rsvp-te {
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
            static-lsp lsp-name;
        }
        threshold-rate number;
    }
}
wildcard-group-inet6 {
    wildcard-source {
        ldp-p2mp;
        pim-ssm {
            group-range multicast-prefix;
        }
        rsvp-te {
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
            static-lsp lsp-name;
        }
        threshold-rate number;
    }
}
}

```

You can include these statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* provider-tunnel]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel]

The following sections describe how to configure selective point-to-multipoint LSPs for MBGP MVPNs:

- [Configuring the Multicast Group Address for an MBGP MVPN on page 5238](#)
- [Configuring the Multicast Source Address for an MBGP MVPN on page 5238](#)
- [Configuring Static Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 5238](#)
- [Configuring Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 5239](#)

- [Configuring the Threshold for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 5239](#)
- [Configuring the Tunnel Limit for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 5240](#)

Configuring the Multicast Group Address for an MBGP MVPN

To configure a point-to-multipoint LSP for an MBGP MVPN, you need to specify a multicast group address by including the **group** statement:

```
group address { ... }
```

You can include this statements at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* provider-tunnel selective]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective]**

The address must be a valid multicast group address. Multicast uses the Class D IP address range (224.0.0.0 through 239.255.255.255).

Configuring the Multicast Source Address for an MBGP MVPN

To configure a point-to-multipoint LSP for an MBGP MVPN, specify a multicast source address by including the **source** statement:

```
source address { ... }
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* provider-tunnel selective group address]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group address]**

Configuring Static Selective Point-to-Multipoint LSPs for an MBGP MVPN

You can configure a static selective point-to-multipoint LSP for an MBGP MVPN. You need to configure a static LSP using the standard MPLS LSP statements at the **[edit protocols mpls]** hierarchy level. You then include the static LSP in your selective point-to-multipoint LSP configuration by using the **static-lsp** statement. Once this functionality is enabled on the source PE router, the static point-to-multipoint LSP is created based on your configuration.

To configure a static selective point-to-multipoint LSP, include the **rsvp-te** and the **static-lsp** statements:

```
rsvp-te static-lsp lsp-name;
```

You can include these statements at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* provider-tunnel selective group address source *source-address*]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group address source *source-address*]**

Configuring Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN

You can configure a dynamic selective point-to-multipoint LSP for an MBGP MVPN. The leaf nodes for a dynamic point-to-multipoint LSP can be automatically discovered using leaf automatic discovery routes. Selective provider multicast service interface (S-PMSI) automatic discovery routes are also supported.

To configure a dynamic selective point-to-multipoint provider tunnel, include the **rsvp-te** and **label-switched-path-template** statements:

```
rsvp-te label-switched-path-template {
  (default-template | lsp-template-name);
}
```

You can include these statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* provider-tunnel selective group address source *source-address*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group address source *source-address*]

The **label-switched-path-template** statement includes the following options:

- **default-template**—Specify that point-to-multipoint LSPs are generated dynamically based on the default template. No user configuration is required for the LSPs. However, the automatically generated LSPs include none of the common LSP features, such as bandwidth allocation and traffic engineering.
- ***lsp-template-name***—Specify the name of an LSP template to be used for the point-to-multipoint LSP. You need to configure the LSP template to be used as a basis for the point-to-multipoint LSPs. You can configure any of the common LSP features for this template.

Configuring the Threshold for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN

To configure a selective point-to-multipoint LSP dynamically, you need to specify the data threshold (in kilobits per second) required before a new tunnel is created using the **threshold-rate** statement:

```
threshold-rate number;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* provider-tunnel selective group address source *source-address*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective group address source *source-address*]

Configuring the Tunnel Limit for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN

To configure a limit on the number of tunnels that can be generated for a dynamic point-to-multipoint LSP, include the **tunnel-limit** statement:

```
tunnel-limit number;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* provider-tunnel selective]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel selective]

Related Documentation

- *Example: Configuring Point-to-Multipoint LDPLSPs as the Data Plane for Intra-AS MBGP MVPNs*

Configuring a Selective Provider Tunnel Using Wildcards

When you configure a selective provider tunnel for MBGP MVPNs (also referred to as next-generation Layer 3 multicast VPNs), you can use wildcards for the multicast group and source address prefixes. Using wildcards enables a PE router to use a single route to advertise the binding of multiple multicast streams of a given MVPN customer to a single provider's tunnel, as described in

<http://tools.ietf.org/html/draft-rekhter-mvpn-wildcard-spmsi-00>.

Sharing a single route improves control plane scalability because it reduces the number of S-PMSI autodiscovery routes.

To configure a selective provider tunnel using wildcards:

1. Configure a wildcard group matching any group IPv4 address and a wildcard source for (*,*) join messages.

```
[edit routing-instances vpn provider-tunnel selective]
user@router# set wildcard-group-inet wildcard-source
```

2. Configure a wildcard group matching any group IPv6 address and a wildcard source for (*,*) join messages.

```
[edit routing-instances vpn provider-tunnel selective]
user@router# set wildcard-group-inet6 wildcard-source
```

3. Configure an IP prefix of a multicast group and a wildcard source for (*,G) join messages.

```
[edit routing-instances vpn provider-tunnel selective]
user@router# set group 224.0.0/24 wildcard-source
```

4. Map the IPv4 join messages to a selective provider tunnel.

```
[edit routing-instances vpn provider-tunnel selective wildcard-group-inet
wildcard-source]
user@router# set rsvp-te (Routing Instances Provider Tunnel Selective)
label-switched-path-template provider-tunnel
```

5. Map the IPv6 join messages to a selective provider tunnel.

```
[edit routing-instances vpna provider-tunnel selective wildcard-group-inet6
wildcard-source]
user@router# set rsvp-te (Routing Instances Provider Tunnel Selective)
label-switched-path-template provider-tunnel2
```

6. Map the (*,224.0.0/24) join messages to a selective provider tunnel.

```
[edit routing-instances vpna provider-tunnel selective group 224.0.0/24
wildcard-source]
user@router# set rsvp-te (Routing Instances Provider Tunnel Selective)
label-switched-path-template provider-tunnel3
```

Related Documentation

- [Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 5153](#)
- [Example: Configuring Selective Provider Tunnels Using Wildcards on page 5241](#)
- [Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs on page 5225](#)
- [Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs on page 5227](#)

Example: Configuring Selective Provider Tunnels Using Wildcards

With the (*,G) and (*,*) S-PMSI, a customer multicast join message can match more than one S-PMSI. In this case, a customer multicast join message is bound to the longest matching S-PMSI. The longest match is a (S,G) S-PMSI, followed by a (*,G) S-PMSI and a (*,*) S-PMSI, in that order.

Consider the following configuration:

```
routing-instances {
  vpna {
    provider-tunnel {
      selective {
        wildcard-group-inet {
          wildcard-source {
            rsvp-te {
              label-switched-path-template {
                sptnl1;
              }
            }
          }
        }
      }
    }
    group 224.1.1.0/24 {
      wildcard-source {
        rsvp-te {
          label-switched-path-template {
            sptnl2;
          }
        }
      }
    }
    source 10.1.1/24 {
      rsvp-te {
```

```
label-switched-path-template {  
    sptnl3;  
}  
}  
}  
}  
}  
}  
}
```

For this configuration, the longest-match rule works as follows:

- A customer multicast (10.1.1.1, 224.1.1.1) join message is bound to the sptnl3 S-PMSI autodiscovery route.
- A customer multicast (10.2.1.1, 224.1.1.1) join message is bound to the sptnl2 S-PMSI autodiscovery route.
- A customer multicast (10.1.1.1, 224.2.1.1) join message is bound to the sptnl1 S-PMSI autodiscovery route.

When more than one customer multicast route is bound to the same wildcard S-PMSI, only one S-PMSI autodiscovery route is created. An egress PE router always uses the same matching rules as the ingress PE router that advertises the S-PMSI autodiscovery route. This ensures consistent customer multicast mapping on the ingress and the egress PE routers.

**Related
Documentation**

- [Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 5153](#)
- [Configuring a Selective Provider Tunnel Using Wildcards on page 5240](#)

Configuring Internet Multicast Using Ingress Replication Provider Tunnels


The routing instance type `mpls-internet-multicast` uses ingress replication provider tunnels to carry IP multicast data between routers through an MPLS cloud, enabling a faster path for multicast traffic between sender and receiver routers in large-scale implementations.

The `mpls-internet-multicast` routing instance is a non-forwarding instance used only for control plane procedures; it does not support any interface configurations. Only one `mpls-internet-multicast` routing instance can be defined for a logical system. All multicast and unicast routes used for Internet multicast are associated only with the master instance (`inet.0`), not with the routing instance.

Each router participating in Internet multicast must be configured with BGP MPLS-based Internet multicast for control plane procedures and with ingress replication for the data plane provider tunnel, which forms a full mesh of MPLS point-to-point LSPs. The ingress replication tunnel can be selective or inclusive, matching the configuration of the provider tunnel in the routing instance.

The topology consists of routers on the edge of the IP multicast domain that have a set of IP interfaces and a set of MPLS core-facing interfaces, see [Figure 118 on page 5244](#). Internet multicast traffic is carried between the IP routers, through the MPLS cloud, using ingress replication tunnels for the data plane and a full-mesh IGBP session for the control plane.

Figure 118: Internet Multicast Topology



The `mpls-internet-multicast` routing instance type is configured for the default master instance on each router to support Internet multicast over MPLS. When using PIM as the multicast protocol, the `mpls-internet-multicast` configuration statement is also included at the `[edit protocols pim]` hierarchy level in the master instance. This creates a pseudo-interface that associates PIM with the `mpls-internet-multicast` routing instance.

When a new destination needs to be added to the ingress replication provider tunnel, the resulting behavior differs depending on how the ingress replication provider tunnel is configured:

- `create-new-ucast-tunnel`—When this statement is configured, a new unicast tunnel to the destination is created, and is deleted when the destination is no longer needed. Use this mode for RSVP LSPs using ingress replication.
- `label-switched-path-template`—When this statement is configured, an LSP template is used for the point-to-multipoint LSP for ingress replication.

Example: Configure Internet Multicast Using Ingress Replication Tunnels

This example configures VPN-B with the instance type **mpls-internet-multicast**. This example also uses PIM for the multicast protocol.

1. Configure the routing instance type for VPN-B as **mpls-internet-multicast**:

```
user@host# set routing-instances VPN-B instance-type mpls-internet-multicast
```

2. Configure the ingress replication provider tunnel to create a new unicast tunnel each time an application requests to add a destination:

```
user@host# set routing-instances VPN-B provider-tunnel ingress-replication
create-new-ucast-tunnel
```

3. Configure the point-to-point LSP to use the default template settings.

```
user@host# set routing-instances VPN-B provider-tunnel ingress-replication
label-switched-path label-switched-path-template default-template
```

4. Configure the ingress replication provider tunnel to be selective:

```
user@host# set routing-instances VPN-B provider-tunnel selective group
232.1.1.1/32 source 192.168.195.145/32 ingress-replication label-switched-path
```

5. Configure MVPN protocol in the routing instance:

```
user@host# set routing-instances VPN-B protocols mvpn
```

6. Commit the configuration:

```
user@host# commit
```

7. Use show command to verify the instance has been created:

```
user@host# run show mvpn instance VPN-B
MVPN instance:
Legend for provider tunnel I-P-tnl -- inclusive provider tunnel S-P-tnl --
selective provider tunnel
Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance : VPN-B
MVPN Mode : SPT-ONLY
Provider tunnel: I-P-tnl:INGRESS-REPLICATION:MPLS Label 18:10.255.245.6
Neighbor          I-P-tnl
10.255.245.2       INGRESS-REPLICATION:MPLS Label 22:10.255.245.2
10.255.245.7       INGRESS-REPLICATION:MPLS Label 19:10.255.245.7
C-mcast IPv4 (S:G) Ptnl          St
192.168.195.145/32:232.1.1.1/32 INGRESS-REPLICATION:MPLS Label
18:10.255.245.6          RM
```

8. Add the **mpls-internet-multicast** configuration statement under the **[edit protocols pim]** hierarchy level in the master instance:

```
user@host# set protocols pim mpls-internet-multicast
```

9. Commit the configuration:

```
user@host# commit
```

10. Use **show ingress-replication mvpn** command to verify configuration settings:

```
user@host# run show ingress-replication mvpn
Ingress Tunnel: mvpn:11
Application: MVPN
Unicast tunnels
Leaf Address      Tunnel-type      Mode      State
10.255.245.2      P2P LSP         New       Up
10.255.245.4      P2P LSP         New       Up
Ingress Tunnel: mvpn:2
```

Application: MVPN

Unicast tunnels

Leaf Address	Tunnel-type	Mode	State
10.255.245.2	P2P LSP	Existing	Up

11. Use this if you want to configure the ingress replication provider tunnel to be inclusive:

```
user@host# set routing-instances VPN-B provider-tunnel ingress-replication
create-new-ucast-tunnel
user@host# set routing-instances VPN-B provider-tunnel ingress-replication
label-switched-path label-switched-path-template default-template
```

12. Use show mvpn instance command to verify tunnel is inclusive:

```
user@host# run show mvpn instance VPN-B
MVPN instance:
Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance : VPN-A
MVPN Mode : SPT-ONLY
Provider tunnel: I-P-tnl:INGRESS-REPLICATION:MPLS Label 18:10.255.245.6
Neighbor          I-P-tnl
10.255.245.2      INGRESS-REPLICATION:MPLS Label 22:10.255.245.2
10.255.245.7      INGRESS-REPLICATION:MPLS Label 19:10.255.245.7
C-mcast IPv4 (S:G) Ptnl              St
192.168.195.145/32:232.1.1.1/32 INGRESS-REPLICATION:MPLS Label 18:10.255.245.6
RM
```

- Related Documentation
- [create-new-ucast-tunnel on page 5260](#)
 - [ingress-replication on page 5269](#)
 - [mpls-internet-multicast on page 5272](#)

Configuring GRE Tunnels for Layer 3 VPNs

Junos OS allows you to configure a generic routing encapsulation (GRE) tunnel between the PE and CE routers for a Layer 3 VPN. The GRE tunnel can have one or more hops. You can configure the tunnel from the PE router to a local CE router (as shown in [Figure 119 on page 5246](#)) or to a remote CE router (as shown in [Figure 120 on page 5247](#)).

Figure 119: GRE Tunnel Configured Between the Local CE Router and the PE Router

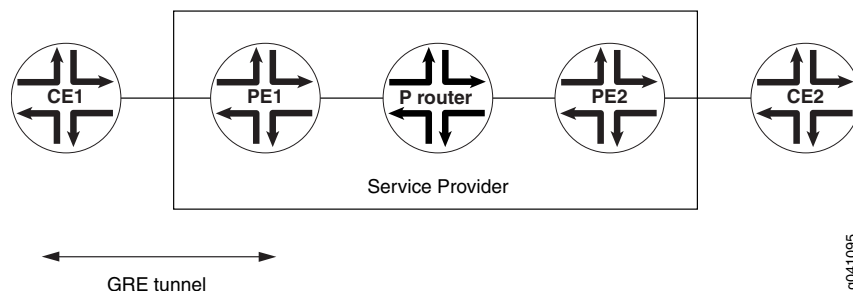



Figure 120: GRE Tunnel Configured Between the Remote CE Router and the PE Router



For more information about how to configure tunnel interfaces, see the *Junos Services Interfaces Configuration Release 11.2*.

You can configure the GRE tunnels manually or configure the Junos OS to instantiate GRE tunnels dynamically.

The following sections describe how to configure GRE tunnels manually and dynamically:

- [Configuring GRE Tunnels Manually Between PE and CE Routers on page 5247](#)
- [Configuring GRE Tunnels Dynamically on page 5248](#)

Configuring GRE Tunnels Manually Between PE and CE Routers

You can manually configure a GRE tunnel between a PE router and either a local CE router or a remote CE router for a Layer 3 VPN as explained in the following sections:

- [Configuring the GRE Tunnel Interface on the PE Router on page 5247](#)
- [Configuring the GRE Tunnel Interface on the CE Router on page 5248](#)

Configuring the GRE Tunnel Interface on the PE Router

You configure the GRE tunnel as a logical interface on the PE router. To configure the GRE tunnel interface, include the **unit** statement:

```
unit logical-unit-number {
  tunnel {
    source source-address;
    destination destination-address;
    routing-instance {
      destination routing-instance-name;
    }
  }
  family inet {
    address address;
  }
}
```

You can include this statement at the following hierarchy levels:

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

As part of the GRE tunnel interface configuration, you need to include the following statements:

- **source *source-address***—Specify the source or origin of the GRE tunnel, typically the PE router.
- **destination *destination-address***—Specify the destination or end point of the GRE tunnel. The destination can be a Provider router, the local CE router, or the remote CE router.

By default, the tunnel destination address is assumed to be in the default Internet routing table, **inet.0**. If the tunnel destination address is not in **inet.0**, you need to specify which routing table to search for the tunnel destination address by configuring the **routing-instance** statement. This is the case if the tunnel encapsulating interface is also configured under the routing instance.

- **destination *routing-instance-name***—Specify the name of the routing instance when configuring the GRE tunnel interface on the PE router.

To complete the GRE tunnel interface configuration, include the **interface** statement for the GRE interface under the appropriate routing instance:

```
interface interface-name;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

Configuring the GRE Tunnel Interface on the CE Router

You can configure either the local or the remote CE router to act as the endpoint for the GRE tunnel.

To configure the GRE tunnel interface on the CE router, include the **unit** statement:

```
unit logical-unit-number {  
  tunnel {  
    source address;  
    destination address;  
  }  
  family inet {  
    address address;  
  }  
}
```

You can include this statement at the following hierarchy levels:

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

Configuring GRE Tunnels Dynamically

When the router receives a VPN route to a BGP next hop address, but no MPLS path is available, a GRE tunnel can be dynamically generated to carry the VPN traffic across the

BGP network. The GRE tunnel is generated and then its routing information is copied into the inet.3 routing table. IPv4 routes are the only type of routes supported for dynamic GRE tunnels. Also, the routing platform must have a tunnel PIC.



NOTE: When configuring a dynamic GRE tunnel to a remote CE router, do not configure OSPF over the tunnel interface. It creates a routing loop forcing the router to take the GRE tunnel down. The router attempts to reestablish the GRE tunnel, but will be forced to take it down again when OSPF becomes active on the tunnel interface and discovers a route to the tunnel endpoint. This is not an issue when configuring static GRE tunnels to a remote CE router.

To generate GRE tunnels dynamically, include the **dynamic-tunnels** statement:

```
dynamic-tunnels tunnel-name {
  destination-networks prefix;
  source-address address;
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-options]
- [edit logical-systems *logical-system-name* routing-options]

Specify the IPv4 prefix range (for example, 10/8 or 11.1/16) for the destination network by including the **destination-networks** statement. Only tunnels within the specified IPv4 prefix range are allowed to be initiated.

```
destination-networks prefix;
```

You can include this statement at the following hierarchy levels:

- [edit routing-options dynamic-tunnels *tunnel-name*]
- [edit logical-systems *logical-system-name* routing-options dynamic-tunnels *tunnel-name*]

Specify the source address for the GRE tunnels by including the **source-address** statement. The source address specifies the address used as the source for the local tunnel endpoint. This could be any local address on the router (typically the router ID or the loopback address).

```
source-address address;
```

You can include this statement at the following hierarchy levels:

- [edit routing-options dynamic-tunnels *tunnel-name*]
- [edit logical-systems *logical-system-name* routing-options dynamic-tunnels *tunnel-name*]

Related Documentation

- [Example: Configuring a Two-Tiered Virtualized Data Center for Large Enterprise Networks on page 3481](#)

Configuration Statements

- [\[edit protocols bgp\] Hierarchy Level on page 5250](#)
- [Layer 2 Routing Instances Configuration Hierarchy on page 5255](#)

[edit protocols bgp] Hierarchy Level

Several statements in the **[edit protocols mpls]** hierarchy are valid at numerous locations within it. To make the complete hierarchy easier to read, the repeated statements are listed in “[Common BGP Family Options](#)” on page 369 and that section is referenced at the appropriate locations in “[Complete \[edit protocols bgp\] Hierarchy](#)” on page 369.

- [Common BGP Family Options on page 5250](#)
- [Complete \[edit protocols bgp\] Hierarchy on page 5250](#)

Common BGP Family Options

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit protocols bgp\] Hierarchy](#)” on page 369 instead of the statements being repeated.

- **[edit protocols bgp family inet (any | flow | labeled-unicast | multicast | unicast)]**
- **[edit protocols bgp family inet6 (any | labeled-unicast | multicast | unicast)]**
- **[edit protocols bgp family (inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) signaling]**
- **[edit protocols bgp family inet-vpn (any | flow | multicast | unicast)]**
- **[edit protocols bgp family inet6-vpn (any | multicast | unicast)]**
- **[edit protocols bgp family iso-vpn unicast]**

The common BGP family options are as follows:

```
accepted-prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
}
damping;
loops number;
prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
}
rib-group group-name;
topology name {
    community {
        target identifier;
    }
}
```

Complete [edit protocols bgp] Hierarchy

The statement hierarchy listed in this section can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.


```

protocols {
  bgp {
    disable;
    accept-remote-nexthop;
    advertise-external <conditional>;
    advertise-from-main-vpn-tables;
    advertise-inactive;
    (advertise-peer-as | no-advertise-peer-as);
    authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
    authentication-key key;
    authentication-key-chain key-chain;
    bfd-liveness-detection {
      authentication {
        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
          meticulous-keyed-sha-1 | simple-password);
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      holddown-interval milliseconds;
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      session-mode (automatic | multihop | single-hop);
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      version (1 | automatic);
    }
    cluster cluster-identifier;
    damping;
    description text-description;
    export [ policy-names ];
    family family-name {
      ... the family subhierarchies appear after the main [edit protocols bgp] hierarchy ...
    }
    graceful-restart {
      disable;
      restart-time seconds;
      stale-routes-time seconds;
    }
    group group-name {
      ... the group subhierarchy appears after the main [edit protocols bgp] hierarchy ...
    }
    hold-time seconds;
    idle-after-switch-over (seconds | forever);
    import [ policy-names ];
    include-mp-next-hop;
    ipsec-sa ipsec-sa;
    keep (all | none);
    local-address address;
    local-as autonomous-system <loops number> <alias> <private>;
  }
}

```

```
local-interface interface-name;
local-preference local-preference;
log-updown;
metric-out (metric | igp (delay-med-update | offset) | minimum-igp offset);
mtu-discovery;
multihop {
    no-nexthop-change;
    ttl ttl-value;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
outbound-route-filter {
    bgp-orf-cisco-mode;
    prefix-based {
        accept {
            inet;
            inet6;
        }
    }
}
passive;
path-selection {
    always-compare-med;
    as-path-ignore;
    cisco-non-deterministic;
    external-router-id;
    med-plus-igp {
        igp-multiplier number;
        med-multiplier number;
    }
}
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
vpn-apply-export;
}

bgp {
    family inet {
        (any | multicast) {
            ... statements in Common BGP Family Options on page 369 ...
        }
        flow {
            ... statements in Common BGP Family Options on page 369 PLUS ...
            no-validate [ validation-procedure-names ];
        }
        labeled-unicast {
            ... statements in Common BGP Family Options on page 369 PLUS ...
            add-path {
```

```

        receive;
        send {
            path-count number;
            prefix-policy [ policy-names ];
        }
    }
    aggregate-label {
        community community-name;
    }
    aigp [disable];
    explicit-null connected-only;
    per-group-label;
    per-prefix-label;
    resolve-vpn;
    rib (inet.3 | inet6.3);
    traffic-statistics {
        file filename <files number> <size maximum-file-size> <world-readable |
            no-world-readable>;
        interval seconds;
    }
}
unicast {
    ... statements in Common BGP Family Options on page 369 PLUS ...
    add-path {
        receive;
        send {
            path-count number;
            prefix-policy [ policy-names ];
        }
    }
    topology name {
        community target identifier;
    }
}
}

bgp {
    family inet6 {
        (any | multicast) {
            ... statements in Common BGP Family Options on page 369 ...
        }
        labeled-unicast {
            ... statements in Common BGP Family Options on page 369 PLUS ...
            add-path {
                receive;
                send {
                    path-count number;
                    prefix-policy [ policy-names ];
                }
            }
            aggregate-label {
                community community-name;
            }
            aigp [disable];
            explicit-null;

```

```
        per-group-label;
        traffic-statistics {
            file filename <files number> <size maximum-file-size> <world-readable |
            no-world-readable>;
            interval seconds;
        }
    }
    unicast {
        ... statements in Common BGP Family Options on page 369 PLUS ...
        topology name {
            community target identifier;
        }
    }
}

bgp {
    family (inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) {
        signaling {
            ... statements in Common BGP Family Options on page 369 ...
        }
    }
}

bgp {
    family inet-vpn {
        (any | multicast | unicast) {
            ... statements in Common BGP Family Options on page 369 PLUS ...
            aggregate-label <community community-name>;
        }
        flow {
            ... statements in Common BGP Family Options on page 369 ...
        }
    }
}

bgp {
    family inet6-vpn {
        (any | multicast | unicast) {
            ... statements in Common BGP Family Options on page 369 PLUS ...
            aggregate-label <community community-name>;
        }
    }
}

bgp {
    family iso-vpn {
        unicast {
            ... statements in Common BGP Family Options on page 369 PLUS ...
            aggregate-label <community community-name>;
        }
    }
}

bgp {
    family route-target {
```

```

accepted-prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
}
advertise-default;
external-paths number;
prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
}
proxy-generate <route-target-policy route-target-policy-name>;
}
}

bgp {
  group group-name {
    ... same statements as at the [edit protocols bgp] hierarchy level PLUS ...
    allow [ all ip-prefix</prefix-length> ];
    as-override;
    multipath <multiple-as>;
    neighbor address {
      ... the neighbor subhierarchy appears after the main [edit protocols bgp group
        group-name] hierarchy ...
    }
    type (external | internal);
    ... BUT NOT ...
    disable; # NOT valid at this level
    group group-name { ... } # NOT valid at this level
    path-selection { ... } # NOT valid at this level
  }

  group group-name {
    neighbor address {
      ... same statements as at the [edit protocols bgp] hierarchy level PLUS ...
      as-override;
      multipath <multiple-as>;
      ... BUT NOT ...
      disable; # NOT valid at this level
      group group-name { ... } # NOT valid at this level
      neighbor address { ... } # NOT valid at this level
      path-selection { ... } # NOT valid at this level
    }
  }
}
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

Layer 2 Routing Instances Configuration Hierarchy

Use the **vpls** routing instance type for point-to-multipoint LAN implementations between a set of sites in a VPN.

To configure routing instances for Layer 2 networks, include the following statements:

```
routing-instances {
  routing-instance-name {
    access {
      address-assignment {
        ... same statements as in the address-assignment subhierarchy in [edit access]
        Hierarchy Level ...
      }
      address-protection;
      description text;
      egress-protection {
        context-identifier context-id;
      }
      forwarding-options {
        ...forwarding-options...
      }
      instance-role role;
      instance-type type;
      interface interface-name;
      l2-domain-id-for-l3 id;
      l2vpn-id community;
      layer3-domain-identifier identifier;
      multicast-snooping-options {
        ... same statements as in [edit multicast-snooping-options] Hierarchy Level EXCEPT
        FOR ...
        traceoptions {...} # NOT valid at this level
      }
      no-irb-layer-2-copy;
      no-local-switching;
      no-vrf-advertise;
      no-vrf-propagate-ttl;
      pbb-options {
        default-bvlan bvlan;
        peer-instance instance;
        vlan-id vlan-id isid-list [ isid-numbers ]
      }
      protocols {
        ... the protocols subhierarchy appears after the main [edit routing-instances
        routing-instance-name] hierarchy ...
      }
      provider-tunnel {
        ... the provider-tunnel subhierarchy appears after the main [edit routing-instances
        routing-instance-name] hierarchy ...
      }
      route-distinguisher (as-number:number | ip-address:number);
      routing-interface interface;
      routing-options {
        ... the routing-options subhierarchy appears after the main [edit routing-instances
        routing-instance-name] hierarchy ...
      }
      service-groups {
        service-group-name {
          pbb-service-options {
            default-isid isid-number;
            isid isid-number vlan-id-list [ vlan-ids ];
          }
        }
      }
    }
  }
}
```

```

        mac-address mac-address;
    }
    service-type type;
}
}
services {
    mobile-ip {
        ... same statements as in [edit services mobile-ip] Hierarchy Level ...
    }
}
switch-options {
    ... same statements as in [edit switch-options] Hierarchy Level ...
}
vlan-id (id | all | none);
vlan-model one-to-one;
vlan-tags outer <tpid.>vlan-id inner <tpid.>vlan-id;
[edit vlans] Hierarchy Level on page 445 {
    ... same statements as in [edit vlans] Hierarchy Level ...
}
vrf-advertise-selective {
    family {
        inet-mvpn;
        inet6-mvpn;
    }
}
vrf-export [ policy-names ];
vrf-import [ policy-names ];
vrf-propagate-ttl;
vrf-table-label;
vrf-target {
    export community-name;
    import community-name;
}
protocols {
    ... protocols-configuration ...
}
routing-options {
    ... routing-options-configuration ...
}
bridge-domains {
    bridge-domain-name {
        domain-type bridge;
        interface interface-name;
        routing-interface routing-interface-name;
        vlan-id (Bridge Domain or VLAN) (none | all | number);
        vlan-tags outer number inner number;
        bridge-options {
            interface-mac-limit limit {
                packet-action drop;
            }
            interface interface-name {
                interface-mac-limit limit {
                    packet-action drop;
                }
            }
        }
        mac-statistics;
    }
}

```

```
        mac-table-size limit {  
            packet-action drop;  
        }  
        no-mac-learning;  
        static-mac mac-address;  
    }  
}  
}
```

With the exception of the **instance-type virtual-switch** statement (which configures a virtual-switch routing instance), you can include the statements at the following hierarchy levels:


- **[edit]**
- **[edit logical-systems *logical-system-name*]**

The **instance-type virtual-switch** statement is not supported at the **[edit logical-systems *logical-system-name*]** hierarchy level.

**Related
Documentation**

- *Routing Instances Overview*
- *Layer 2 Routing Instance Types*
- [Configuring a Layer 2 Virtual Switch on page 1847](#)
- *Configuring a Layer 2 Control Protocol Routing Instance*

advertise-from-main-vpn-tables

Syntax	advertise-from-main-vpn-tables;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp],
Release Information	Statement introduced in Junos OS Release 12.3.
Description	<p>Advertise VPN routes from the main VPN tables in the master routing instance (for example, <i>bgp.l3vpn.0</i>, <i>bgp.mvpn.0</i>) instead of advertising VPN routes from the tables in the VPN routing instances (for example, <i>instance-name.inet.0</i>, <i>instance-name.mvpn.0</i>).</p> <p>When this statement is enabled, before advertising a route for a VPN prefix, the path selection algorithm is run on all routes (local and received) that have the same route distinguisher (RD).</p> <div style="margin-top: 20px;">  <p>NOTE: Adding or removing this statement causes all BGP sessions that have VPN address families to be removed and then added again. On the other hand, having this statement in the configuration prevents BGP sessions from going down when route reflector (RR) or autonomous system border router (ASBR) functionality is enabled or disabled on a routing device that has VPN address families configured.</p> </div>
Default	If you do not include this statement, VPN routes are advertised from the tables in the VPN routing instances.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding Junos OS Routing Tables</i> • <i>Types of VPNs</i>

create-new-ucast-tunnel

Syntax	create-new-ucast-tunnel;
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> provider-tunnel ingress-replication], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source <i>source-address</i> ingress-replication]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	One of two modes for building unicast tunnels when ingress replication is configured for the provider tunnel. When this statement is configured, each time a new destination is added to the multicast distribution tree, a new unicast tunnel to the destination is created in the ingress replication tunnel. The new tunnel is deleted if the destination is no longer needed. Use this mode for RSVP LSPs using ingress replication.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs</i>• Configuring Routing Instances for an MBGP MVPN on page 5223• mpls-internet-multicast on page 5272• ingress-replication on page 5269

export-target

Syntax	<code>export-target { target <i>target-community</i>; unicast; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvpn route-target], [edit routing-instances <i>routing-instance-name</i> protocols mvpn route-target]
Release Information	Statement introduced in Junos OS Release 8.4.
Description	Enable you to override the Layer 3 VPN import and export route targets used for importing and exporting routes for the MBGP MVPN network layer reachability information (NLRI).
Options	target <i>target-community</i> —Specify the export target community. unicast —Use the same target community as specified for unicast.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Export Target for an MBGP MVPN on page 5230

family (VRF Advertisement)

Syntax	<code>family { inet-mvpn; inet6-mvpn; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> vrf-advertise-selective], [edit routing-instances <i>routing-instance-name</i> vrf-advertise-selective],
Release Information	Statement introduced in Junos OS Release 10.1.
Description	Explicitly enable IPv4 or IPv6 MVPN routes to be advertised from the VRF instance while preventing all other route types from being advertised. The options are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring PIM-SSM GRE Selective Provider Tunnels on page 5233 • inet-mvpn on page 5267 • inet6-mvpn on page 5268


group (Routing Instances)

Syntax	<pre>group address { source source-address { pim-ssm { group-range multicast-prefix; } ldp-p2mp; rsvp-te { label-switched-path-template { (default-template lsp-template-name); } static-lsp lsp-name; } threshold-rate number; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective]
Release Information	Statement introduced in Junos OS Release 8.5.
Description	Specify the IP address for the multicast group configured for point-to-multipoint label-switched paths (LSPs) and PIM-SSM GRE selective provider tunnels.
Options	address —Specify the IP address for the multicast group. This address must be a valid multicast group address. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Multicast Group Address for an MBGP MVPN on page 5238• Configuring PIM-SSM GRE Selective Provider Tunnels on page 5233

group-range (MBGP MVPN Tunnel)

Syntax	<code>group-range <i>multicast-prefix</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>group-address</i> source <i>source-address</i> pim-ssm],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>group-address</i> wildcard-source pim-ssm],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet wildcard-source pim-ssm],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet6 wildcard-source pim-ssm],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>group-address</i> source <i>source-address</i> pim-ssm],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>group-address</i> wildcard-source pim-ssm],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet wildcard-source pim-ssm],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet6 wildcard-source pim-ssm]</p>
Release Information	Statement introduced in Junos OS Release 10.1.
Description	Establish the multicast group address range to use for creating MBGP MVPN source-specific multicast selective PMSI tunnels.
Options	<p><i>multicast-prefix</i>—Multicast group address range to be used to create MBGP MVPN source-specific multicast selective PMSI tunnels.</p> <p>Range: Any valid, nonreserved IPv4 multicast address range</p> <p>Default: None</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring PIM-SSM GRE Selective Provider Tunnels on page 5233

group-rp-mapping

Syntax	<pre>group-rp-mapping { family (inet inet6) { log-interval seconds; maximum limit; threshold value; } log-interval seconds; maximum limit; threshold value; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols pim rp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp],</p> <p>[edit protocols pim rp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp]</p>
Release Information	Statement introduced in Junos OS Release 12.2.
Description	Configure a limit for the number of incoming group-to-RP mappings.
	<div>  <p>NOTE: The maximum limit settings that you configure with the <code>maximum</code> and the <code>family (inet inet6) maximum</code> statements are mutually exclusive. For example, if you configure a global maximum group-to-RP mapping limit, you cannot configure a limit at the family level for IPv4 or IPv6. If you attempt to configure a limit at both the global level and the family level, the device will not accept the configuration.</p> </div>
Options	<p>family (inet inet6)—(Optional) Specify either IPv4 or IPv6 messages to be counted towards the configured group-to-RP mapping limit.</p> <p>Default: Both IPv4 and IPv6 messages are counted towards the configured group-to-RP limit.</p> <p>The remaining statements are described separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring PIM State Limits on page 5198

import-target

Syntax	<pre> import-target { target { target-value; receiver target-value; sender target-value; } unicast { receiver; sender; } } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvpn route-target], [edit routing-instances <i>routing-instance-name</i> protocols mvpn route-target]
Release Information	Statement introduced in Junos OS Release 8.4.
Description	Enable you to override the Layer 3 VPN import and export route targets used for importing and exporting routes for the MBGP MVPN NLRI.
Options	The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Import Target for an MBGP MVPN on page 5231

inet-mvpn (BGP)

Syntax	<pre>inet-mvpn { signaling { accepted-prefix-limit { maximum <i>number</i>; teardown <i>percentage</i> { idle-timeout (forever <i>minutes</i>); } } damping; loops <i>number</i>; prefix-limit { maximum <i>number</i>; teardown <i>percentage</i> { idle-timeout (forever <i>minutes</i>); } } } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp family], [edit protocols bgp family], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family], [edit protocols bgp group <i>group-name</i> family]
Release Information	Statement introduced in Junos OS Release 8.4.
Description	Enable the inet-mvpn address family in BGP.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring NLRI Parameters for an MBGP MVPN on page 5232

inet-mvpn (VRF Advertisement)

Syntax	inet-mvpn;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> vrf-advertise-selective family], [edit routing-instances <i>routing-instance-name</i> vrf-advertise-selective family]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	Enable IPv4 MVPN routes to be advertised from the VRF instance.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Limiting Routes to Be Advertised by an MVPN VRF Instance on page 5228

inet6-mvpn (BGP)

Syntax	<pre>inet6-mvpn { signaling { accepted-prefix-limit { maximum <i>number</i>; teardown <i>percentage</i> { idle-timeout (forever <i>minutes</i>); } } } loops <i>number</i> prefix-limit { maximum <i>number</i>; teardown <i>percentage</i> { idle-timeout (forever <i>minutes</i>); } } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols bgp family], [edit protocols bgp family], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family], [edit protocols bgp group <i>group-name</i> family]
Release Information	Statement introduced in Junos OS Release 10.0.
Description	Enable the inet6-mvpn address family in BGP.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring NLRI Parameters for an MBGP MVPN on page 5232 • BGP Configuration Guidelines chapter in the <i>Routing Protocols Configuration Guide</i>

inet6-mvpn (VRF Advertisement)

Syntax	inet6-mvpn;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> vrf-advertise-selective family], [edit routing-instances <i>routing-instance-name</i> vrf-advertise-selective family],
Release Information	Statement introduced in Junos OS Release 10.1.
Description	Enable IPv6 MVPN routes to be advertised from the VRF instance.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Limiting Routes to Be Advertised by an MVPN VRF Instance on page 5228

ingress-replication

Syntax	<pre>ingress-replication { create-new-ucast-tunnel; label-switched-path { label-switched-path-template { (template-name default-template); } } }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source <i>source-address</i>]</p>
Release Information	Statement introduced in Junos OS Release 10.4.
Description	<p>A provider tunnel type used for passing multicast traffic between routers through the MPLS cloud, or between PE routers when using MVPN. The ingress replication provider tunnel uses MPLS point-to-point LSPs to create the multicast distribution tree.</p> <p>Optionally, you can specify a label-switched path template. If you configure ingress-replication label-switched-path and do not include label-switched-path-template, ingress replication works with existing LDP or RSVP tunnels. If you include label-switched-path-template, the tunnels must be RSVP.</p>
Options	<p>create-new-ucast-tunnel—A new unicast tunnel to the destination that is created and used for ingress replication. The unicast tunnel is deleted later if the destination is no longer included in the multicast distribution tree. A template must be specified when and only when create-new-ucast-tunnel is included in the configuration..</p> <p>template-name—Name of the point-to-point LSP used for the new unicast tunnel.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs • Configuring Routing Instances for an MBGP MVPN on page 5223 • create-new-ucast-tunnel on page 5260 • mpls-internet-multicast on page 5272

interface (Virtual Tunnel in Routing Instances)

Syntax	<pre>interface vt-<i>fpc/pic/port.unit-number</i> { multicast; primary; unicast; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced in Junos OS Release 9.4.
Description	<p>In a multiprotocol BGP (MBGP) multicast VPN (MVPN), configure a virtual tunnel (VT) interface.</p> <p>VT interfaces are needed for multicast traffic on routing devices that function as combined provider edge (PE) and provider core (P) routers to optimize bandwidth usage on core links. VT interfaces prevent traffic replication when a P router also acts as a PE router (an exit point for multicast traffic).</p> <p>In an MBGP MVPN extranet, if there is more than one VRF routing instance on a PE router that has receivers interested in receiving multicast traffic from the same source, VT interfaces must be configured on all instances.</p> <p>Starting in Junos OS Release 12.3, you can configure multiple VT interfaces in each routing instance. This provides redundancy. A VT interface can be used in only one routing instance.</p>
Options	<p><i>vt-fpc/pic/port.unit-number</i>—Name of the VT interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs on page 5212• Example: Configuring MBGP MVPN Extranets

label-switched-path-template

Syntax	label-switched-path-template { (default-template <i>lsp-template-name</i>); }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel rsvp-te], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel ingress-replication label-switched-path], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source <i>source-address</i> rsvp-te], [edit logical-systems <i>logical-system-name</i> routing-options dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry-name</i>], [edit routing-instances <i>routing-instance-name</i> provider-tunnel ingress-replication label-switched-path], [edit routing-instances <i>routing-instance-name</i> provider-tunnel rsvp-te], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source <i>source-address</i> rsvp-te], [edit routing-options dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry-name</i>]
Release Information	Statement introduced in Junos OS Release 8.5.
Description	Specify the LSP template. An LSP template is used as the basis for other dynamically generated LSPs. This feature can be used for a number of applications, including point-to-multipoint LSPs, flooding VPLS traffic, configuring ingress replication for IP multicast using MBGP MVPNs, and to enable RSVP automatic mesh. There is no default setting for the label-switched-path-template statement, so you must configure either the default-template using the default-template option, or you must specify the name of your preconfigured LSP template.
Options	default-template —Specify that the default LSP template be used for the dynamically generated LSPs. lsp-template-name —Specify the name of an LSP to be used as a template for the dynamically generated LSPs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs</i> • Configuring RSVP-Signaled Inclusive Point-to-Multipoint LSPs for an MBGP MVPN on page 5235 • Configuring Dynamic Point-to-Multipoint Flooding LSPs on page 5368 • <i>Configuring RSVP Automatic Mesh</i>

mpls-internet-multicast

Syntax	mpls-internet-multicast;
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> instance-type] [edit protocols pim]
Release Information	Statement introduced in Junos OS Release 11.1.
Description	<p>A nonforwarding routing instance type that supports Internet multicast over an MPLS network for the default master instance. No interfaces can be configured for it. Only one mpls-internet-multicast instance can be configured for each logical system.</p> <p>The mpls-internet-multicast configuration statement is also explicitly required under PIM in the master instance.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs</i>• ingress-replication on page 5269

multicast (Virtual Tunnel in Routing Instances)

Syntax	multicast;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> interface <i>vt-fpc/pic/port.unit-number</i>], [edit routing-instances <i>routing-instance-name</i> interface <i>vt-fpc/pic/port.unit-number</i>]
Release Information	Statement introduced in Junos OS Release 9.4.
Description	In a multiprotocol BGP (MBGP) multicast VPN (MVPN), configure the virtual tunnel (VT) interface to be used for multicast traffic only.
Default	If you omit this statement, the VT interface can be used for both multicast and unicast traffic.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs on page 5212• <i>Example: Configuring MBGP MVPN Extranets</i>

mvpn

Syntax	<pre> mvpn { mvpn-mode (rpt-spt spt-only); receiver-site; sender-site; route-target { export-target { target <i>target-community</i>; unicast; } import-target { target { <i>target-value</i>; receiver <i>target-value</i>; sender <i>target-value</i>; } unicast { receiver; sender; } } } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols]</p>
Release Information	Statement introduced in Junos OS Release 8.4.
Description	Enable next-generation multicast VPNs in a routing instance.
Options	<p>receiver-site—Allow sites with multicast receivers.</p> <p>sender-site—Allow sites with multicast senders.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Routing Instances for an MBGP MVPN on page 5223

mvpn-mode

Syntax	<code>mvpn-mode (rpt-spt spt-only);</code>
Hierarchy Level	[edit logical-systems <i>profile-name</i> routing-instances <i>instance-name</i> protocols mvpn], [edit routing-instances <i>instance-name</i> protocols mvpn]
Release Information	Statement introduced in Junos OS Release 10.0.
Description	Configure the mode for customer PIM (C-PIM) join messages. The remaining statements are explained separately.
Default	<code>spt-only</code>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs on page 5227• Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs on page 5225

pim-asm

Syntax	<code>pim-asm { group-address <i>address</i>; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel], [edit routing-instances <i>routing-instance-name</i> provider-tunnel]
Release Information	Statement introduced in Junos OS Release 8.3.
Description	Specify a Protocol Independent Multicast (PIM) sparse mode provider tunnel for an MBGP MVPN or for a draft-rosen MVPN. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring PIM Provider Tunnels for an MBGP MVPN on page 5233

pim-ssm (Selective Tunnel)

Syntax	<pre>pim-ssm { group-range <i>multicast-prefix</i>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>group-address</i> source <i>source-address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>group-address</i> wildcard-source],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet wildcard-source],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet6 wildcard-source],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>group-address</i> source <i>source-address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>group-address</i> wildcard-source],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet wildcard-source],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet6 wildcard-source]</p>
Release Information	Statement introduced in Junos OS Release 10.1.
Description	Establish the multicast group address range to use for creating MBGP MVPN source-specific multicast selective PMSI tunnels.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring PIM-SSM GRE Selective Provider Tunnels on page 5233

primary (Virtual Tunnel in Routing Instances)

Syntax	<code>primary;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> interface <i>vt-fpc/pic/port.unit-number</i>], [edit routing-instances <i>routing-instance-name</i> interface <i>vt-fpc/pic/port.unit-number</i>]
Release Information	Statement introduced in Junos OS Release 12.3.
Description	<p>In a multiprotocol BGP (MBGP) multicast VPN (MVPN), configure the virtual tunnel (VT) interface to be used as the primary interface for multicast traffic.</p> <p>Junos OS supports up to eight VT interfaces configured for multicast in a routing instance to provide redundancy for MBGP (next-generation) MVPNs. This support is for RSVP point-to-multipoint provider tunnels as well as multicast Label Distribution Protocol (MLDP) provider tunnels. This feature works for extranets as well.</p> <p>This statement allows you to configure one of the VT interfaces to be the primary interface, which is always used if it is operational. If a VT interface is configured as the primary, it becomes the nexthop that is used for traffic coming in from the core on the label-switched path (LSP) into the routing instance. When a VT interface is configured to be primary and the VT interface is used for both unicast and multicast traffic, only the multicast traffic is affected.</p> <p>If no VT interface is configured to be the primary or if the primary VT interface is unusable, one of the usable configured VT interfaces is chosen to be the nexthop that is used for traffic coming in from the core on the LSP into the routing instance. If the VT interface in use goes down for any reason, another usable configured VT interface in the routing instance is chosen. When the VT interface in use changes, all multicast routes in the instance also switch their reverse-path forwarding (RPF) interface to the new VT interface to allow the traffic to be received.</p> <p>To realize the full benefit of redundancy, we recommend that when you configure multiple VT interfaces, at least one of the VT interfaces be on a different Tunnel PIC from the other VT interfaces. However, Junos OS does not enforce this.</p>
Default	If you omit this statement, Junos OS chooses a VT interface to be the active interface for multicast traffic.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs on page 5212• Example: Configuring MBGP MVPN Extranets

provider-tunnel


```
Syntax provider-tunnel {
    ingress-replication {
        create-new-ucast-tunnel;
        label-switched-path-template {
            (default-template | lsp-template-name);
        }
    }
    ldp-p2mp;
    pim-asm {
        group-address address;
    }
    mdt {
        data-mdt-reuse;
        group-range multicast-prefix;
        threshold {
            group group-address {
                source source-address {
                    rate threshold-rate;
                }
            }
        }
        tunnel-limit limit;
    }
}
pim-ssm {
    group-address address;
}
}
rsvp-te {
    label-switched-path-template {
        (default-template | lsp-template-name);
    }
    static-lsp lsp-name;
}
selective {
    group multicast--prefix/prefix-length {
        source ip--prefix/prefix-length {
            ldp-p2mp;
            create-new-ucast-tunnel;
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
        }
    }
    pim-ssm {
        group-range multicast-prefix;
    }
    rsvp-te {
        label-switched-path-template {
            (default-template | lsp-template-name);
        }
        static-lsp point-to-multipoint-lsp-name;
    }
    threshold-rate kbps;
}
```

```
wildcard-source {
  pim-ssm {
    group-range multicast-prefix;
  }
  rsvp-te {
    label-switched-path-template {
      (default-template | lsp-template-name);
    }
    static-lsp point-to-multipoint-lsp-name;
  }
  threshold-rate kbits;
}
}
tunnel-limit number;
wildcard-group-inet {
  wildcard-source {
    pim-ssm {
      group-range multicast-prefix;
    }
    rsvp-te {
      label-switched-path-template {
        (default-template | lsp-template-name);
      }
      static-lsp lsp-name;
    }
  }
  threshold-rate number;
}
}
wildcard-group-inet6 {
  wildcard-source {
    pim-ssm {
      group-range multicast-prefix;
    }
    rsvp-te {
      label-switched-path-template {
        (default-template | lsp-template-name);
      }
      static-lsp lsp-name;
    }
  }
  threshold-rate number;
}
}
}
```

Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced in Junos OS Release 8.3. The selective statement and substatements added in Junos OS Release 8.5. The ingress-replication statement and substatements added in Junos OS Release 10.4.
Description	Configure virtual private LAN service (VPLS) flooding of unknown unicast, broadcast, and multicast traffic using point-to-multipoint LSPs. Also configure point-to-multipoint LSPs for MBGP MVPNs.

Options	The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Flooding Unknown Traffic Using Point-to-Multipoint LSPs on page 5365• Configuring RSVP-Signaled Inclusive Point-to-Multipoint LSPs for an MBGP MVPN on page 5235• <i>Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode</i>

register-limit

Syntax	<pre> register-limit { family (inet inet6) { log-interval <i>seconds</i>; maximum <i>limit</i>; threshold <i>value</i>; } log-interval <i>seconds</i>; maximum <i>limit</i>; threshold <i>value</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols pim rp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim rp],</p> <p>[edit protocols pim rp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rp]</p>
Release Information	Statement introduced in Junos OS Release 12.2.
Description	Configure a limit for the number of incoming (S,G) PIM registers.
	<div>  <p>NOTE: The maximum limit settings that you configure with the maximum and the family (inet inet6) maximum statements are mutually exclusive. For example, if you configure a global maximum PIM register message limit, you cannot configure a limit at the family level for IPv4 or IPv6. If you attempt to configure a limit at both the global level and the family level, the device will not accept the configuration.</p> </div>
Options	<p>family (inet inet6)—(Optional) Specify either IPv4 or IPv6 messages to be counted towards the configured register message limit.</p> <p>Default: Both IPv4 and IPv6 messages are counted towards the configured register message limit.</p> <p>The remaining statements are described separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring PIM State Limits on page 5198 • clear pim join on page 4017 • clear pim register on page 4020


route-target (Protocols MVPN)

Syntax	<pre> route-target { export-target { target <i>target-community</i>; unicast; } import-target { target { <i>target-value</i>; receiver <i>target-value</i>; sender <i>target-value</i>; } unicast { receiver; sender; } } } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvpn], [edit routing-instances <i>routing-instance-name</i> protocols mvpn]
Release Information	Statement introduced in Junos OS Release 8.4.
Description	Enable you to override the Layer 3 VPN import and export route targets used for importing and exporting routes for the MBGP MVPN NLRI.
Default	The multicast VPN routing instance uses the import and export route targets configured for the Layer 3 VPN.
Options	The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring VRF Route Targets for Routing Instances for an MBGP MVPN on page 5229

rpt-spt

Syntax	rpt-spt;
Hierarchy Level	[edit logical-systems <i>profile-name</i> routing-instances <i>instance-name</i> protocols mvpn mvpn-mode], [edit routing-instances <i>instance-name</i> protocols mvpn mvpn-mode]
Release Information	Statement introduced in Junos OS Release 10.0.
Description	Use rendezvous-point trees for customer PIM (C-PIM) join messages, and switch to the shortest-path tree after the source is known.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Shared-Tree Data Distribution Across Provider Cores for Providers of MBGP MVPNs on page 5227

rsvp-te (Routing Instances Provider Tunnel Selective)

Syntax	<pre> rsvp-te { label-switched-path-template { (default-template <i>lsp-template-name</i>); } static-lsp <i>lsp-name</i>; } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source <i>source-address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> wildcard-source], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet wildcard-source], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet6 wildcard-source], [edit routing-instances <i>routing-instance-name</i> provider-tunnel], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source <i>source-address</i>], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> wildcard-source], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet wildcard-source], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet6 wildcard-source] </pre>
Release Information	Statement introduced in Junos OS Release 8.5.
Description	<p>Configure the properties of the RSVP traffic-engineered point-to-multipoint LSP for MBGP MVPNs.</p> <p>The remaining statements are explained separately.</p>
	<div>  <p>NOTE: Junos OS Release 11.2 and earlier do not support point-to-multipoint LSPs with next-generation multicast VPNs on MX80 routers.</p> </div>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Selective Provider Tunnels for an MBGP MVPN on page 5236

selective

```
Syntax  selective {
    group multicast-prefix/prefix-length {
        source ip-prefix/prefix-length {
            ingress-replication {
                create-new-ucast-tunnel;
                label-switched-path-template {
                    (default-template | lsp-template-name);
                }
            }
        }
        ldp-p2mp;
        pim-ssm {
            group-range multicast-prefix;
        }
        rsvp-te {
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
            static-lsp point-to-multipoint-lsp-name;
        }
        threshold-rate kbits;
    }
    wildcard-source {
        ldp-p2mp;
        pim-ssm {
            group-range multicast-prefix;
        }
        rsvp-te {
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
        }
        static-lsp point-to-multipoint-lsp-name;
    }
    threshold-rate kbits;
}
tunnel-limit number;
wildcard-group-inet {
    wildcard-source {
        ldp-p2mp;
        pim-ssm {
            group-range multicast-prefix;
        }
        rsvp-te {
            label-switched-path-template {
                (default-template | lsp-template-name);
            }
        }
        static-lsp lsp-name;
    }
    threshold-rate number;
}
```


```

wildcard-group-inet6 {
  wildcard-source {
    ldp-p2mp;
    pim-ssm {
      group-range multicast-prefix;
    }
    rsvp-te {
      label-switched-path-template {
        (default-template | lsp-template-name);
      }
      static-lsp lsp-name;
    }
    threshold-rate number;
  }
}

```

Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel], [edit routing-instances <i>routing-instance-name</i> provider-tunnel]
Release Information	Statement introduced in Junos OS Release 8.5. The ingress-replication statement and substatements added in Junos OS Release 10.4.
Description	Configure selective point-to-multipoint LSPs for an MBGP MVPN. Selective point-to-multipoint LSPs send traffic only to the receivers configured for the MBGP MVPNs, helping to minimize flooding in the service provider's network. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Selective Provider Tunnels for an MBGP MVPN on page 5236 • Configuring PIM-SSM GRE Selective Provider Tunnels on page 5233

sglimit

Syntax	<pre>sglimit { family (inet inet6) { log-interval seconds; maximum limit; threshold value; } log-interval seconds; maximum limit; threshold value; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols pim],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],</p> <p>[edit protocols pim],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim]</p>
Release Information	Statement introduced in Junos OS Release 12.2.
Description	Configure a limit for the number of accepted (*G) and (S,G) PIM join states.
	<div>  <p>NOTE: The maximum limit settings that you configure with the <code>maximum</code> and the <code>family (inet inet6) maximum</code> statements are mutually exclusive. For example, if you configure a global maximum PIM join state limit, you cannot configure a limit at the family level for IPv4 or IPv6 joins. If you attempt to configure a limit at both the global level and the family level, the device will not accept the configuration.</p> </div>
Options	<p>family (inet inet6)—(Optional) Specify either IPv4 or IPv6 join states to be counted towards the configured join state limit.</p> <p>Default: Both IPv4 and IPv6 join states are counted towards the configured join state limit.</p> <p>The remaining statements are described separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring PIM State Limits on page 5198 • clear pim join on page 4017

source (Routing Instances Provider Tunnel Selective)

Syntax	<pre> source <i>source-address</i> { <i>ldp-p2mp</i>; pim-ssm { group-range <i>multicast-prefix</i>; } rsvp-te { label-switched-path-template { (default-template <i>lsp-template-name</i>); } static-lsp <i>lsp-name</i>; } threshold-rate <i>number</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i>]</p>
Release Information	Statement introduced in Junos OS Release 8.5.
Description	Specify the IP address for the multicast source. This statement is a part of the point-to-multipoint LSP and PIM-SSM GRE selective provider tunnel configuration for MBGP MVPNs.
Options	<p><i>source-address</i>—IP address for the multicast source.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Multicast Source Address for an MBGP MVPN on page 5238 • Configuring PIM-SSM GRE Selective Provider Tunnels on page 5233

spt-only

Syntax	spt-only;
Hierarchy Level	[edit logical-systems <i>profile-name</i> routing-instances <i>instance-name</i> protocols mvpn mvpn-mode], [edit routing-instances <i>instance-name</i> protocols mvpn mvpn-mode]
Release Information	Statement introduced in Junos OS Release 10.0.
Description	Set the MVPN mode to learn about active multicast sources using multicast VPN source-active routes. This is the default mode.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring SPT-Only Mode for Multiprotocol BGP-Based Multicast VPNs on page 5225

static-lsp

Syntax	static-lsp <i>lsp-name</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel rsvp-te], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source <i>source-address</i> rsvp-te], [edit routing-instances <i>routing-instance-name</i> provider-tunnel rsvp-te], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source <i>source-address</i> rsvp-te]
Release Information	Statement introduced in Junos OS Release 8.5.
Description	Specify the name of the static point-to-multipoint LSP used for an MBGP MVPN. Use this statement to specify the static LSP for both inclusive and selective point-to-multipoint LSPs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Selective Provider Tunnels for an MBGP MVPN on page 5236

target (Routing Instances MVPN)

Syntax	<code>target <i>target-value</i> { receiver <i>target-value</i>; sender <i>target-value</i>; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvpn route-target import-target], [edit routing-instances <i>routing-instance-name</i> protocols mvpn route-target import-target]
Release Information	Statement introduced in Junos OS Release 8.4.
Description	Specify the target value when importing sender and receiver site routes.
Options	<i>target-value</i> —Specify the target value when importing sender and receiver site routes. <i>receiver</i> —Specify the target community used when importing receiver site routes. <i>sender</i> —Specify the target community used when importing sender site routes.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Import Target Receiver and Sender for an MBGP MVPN on page 5231

threshold-rate

Syntax	<code>threshold-rate kbps;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group address source <i>source-address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>group-address</i> wildcard-source],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet wildcard-source],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet6 wildcard-source],</code> <code>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group address source <i>source-address</i>]</code> <code>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group address wildcard-source]</code> <code>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet wildcard-source],</code> <code>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet6 wildcard-source]</code>
Release Information	Statement introduced in Junos OS Release 8.5.
Description	Specify the data threshold required before a new tunnel is created for a dynamic selective point-to-multipoint LSP. This statement is part of the configuration for point-to-multipoint LSPs for MBGP MVPNs and PIM-SSM GRE or RSVP-TE selective provider tunnels.
Options	number —Specify the data threshold required before a new tunnel is created. Range: 0 through 1,000,000 kilobits per second. Specifying 0 is equivalent to not including the statement.
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Threshold for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 5239• Configuring PIM-SSM GRE Selective Provider Tunnels on page 5233• Configuring Intra-AS Selective Provider Tunnels

traceoptions (Protocols MVPN)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvpn],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols mvpn]</p>
Release Information	Statement introduced in Junos OS Release 8.4.
Description	Trace traffic flowing through an MBGP MVPN.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks (" ").</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named <i>trace-file</i> reaches this size, it is renamed <i>trace-file.0</i>. When <i>trace-file</i> again reaches its maximum size, <i>trace-file.0</i> is renamed <i>trace-file.1</i> and <i>trace-file</i> is renamed <i>trace-file.0</i>. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can specify any of the following flags:</p> <ul style="list-style-type: none"> • all—All multicast VPN tracing options • error—Error conditions • general—General events • nlri—Multicast VPN advertisements received or sent by means of the BGP • normal—Normal events • policy—Policy processing • route—Routing information • state—State transitions • task—Routing protocol task processing • timer—Routing protocol timer processing

- **topology**—Multicast VPN topology changes caused by reconfiguration or advertisements received from other provider edge (PE) routers using BGP

flag-modifier—(Optional) Modifier for the tracing flag. You can specify the following modifiers:

- **detail**—Provide detailed trace information
- **disable**—Disable the tracing flag
- **receive**—Trace received packets
- **send**—Trace sent packets

no-world-readable—Do not allow any user to read the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: *xk* to specify kilobytes, *xm* to specify megabytes, or *xg* to specify gigabytes

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—Allow any user to read the log file.

Required Privilege Level	routing—To view this statement in the configuration.
	routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Tracing MBGP MVPN Traffic and Operations on page 5299

tunnel-limit (Routing Instances Provider Tunnel Selective)

Syntax	tunnel-limit <i>number</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective]
Release Information	Statement introduced in Junos OS Release 8.5.
Description	Specify a limit on the number of selective tunnels that can be created for an LSP. This limit can be applied to the following types of selective tunnels: <ul style="list-style-type: none"> • Ingress replication tunnels • LDP-signaled LSP • LDP point-to-multipoint LSP • PIM-SSM provider tunnel • RSVP-signaled LSP • RSVP-signaled point-to-multipoint LSP
Options	<i>number</i> —Specify the tunnel limit. Range: 0 through 1024
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Tunnel Limit for Dynamic Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 5240 • selective on page 5284 • wildcard-source on page 5298

unicast (Route Target Community)

Syntax	<pre>unicast { receiver; sender; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols mvpn route-target import-target], [edit routing-instances <i>routing-instance-name</i> protocols mvpn route-target import-target]
Release Information	Statement introduced in Junos OS Release 8.4.
Description	Specify the same target community configured for unicast.
Options	receiver —Specify the unicast target community used when importing receiver site routes. sender —Specify the unicast target community used when importing sender site routes.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Import Target Unicast Parameters for an MBGP MVPN on page 5232

unicast (Virtual Tunnel in Routing Instances)

Syntax	<pre>unicast;</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> interface vt- <i>fpc/pic/port.unit-number</i>], [edit routing-instances <i>routing-instance-name</i> interface vt- <i>fpc/pic/port.unit-number</i>]
Release Information	Statement introduced in Junos OS Release 9.4.
Description	In a multiprotocol BGP (MBGP) multicast VPN (MVPN), configure the virtual tunnel (VT) interface to be used for unicast traffic only.
Default	If you omit this statement, the VT interface can be used for both multicast and unicast traffic.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs on page 5212• Example: Configuring MBGP MVPN Extranets

vrf-advertise-selective

Syntax	<pre>vrf-advertise-selective { family { inet-mvpn; inet6-mvpn; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced in Junos OS Release 10.1. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	<p>Explicitly enable IPv4 or IPv6 MVPN routes to be advertised from the VRF instance while preventing all other route types from being advertised.</p> <p>If you configure the vrf-advertise-selective statement without any of its options, the router or switch has the same behavior as if you configured the no-vrf-advertise statement. All VPN routes are prevented from being advertised from a VRF routing instance to the remote PE routers. This behavior is useful for hub-and-spoke configurations, enabling you to configure a PE router to not advertise VPN routes from the primary (hub) instance. Instead, these routes are advertised from the secondary (downstream) instance.</p> <p>The options are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Limiting Routes to Be Advertised by an MVPN VRF Instance on page 5228 • <i>no-vrf-advertise</i>

wildcard-group-inet

Syntax	<pre>wildcard-group-inet { wildcard-source { ldp-p2mp; pim-ssm { group-range <i>multicast-prefix</i>; } rsvp-te { label-switched-path-template { (default-template <i>lsp-template-name</i>); } static-lsp <i>lsp-name</i>; } threshold-rate <i>number</i>; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective]
Release Information	Statement introduced in Junos OS Release 10.0.
Description	Configure a wildcard group matching any group IPv4 address. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• wildcard-group-inet6 on page 5297• Example: Configuring Selective Provider Tunnels Using Wildcards on page 5241• Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 5153• Configuring a Selective Provider Tunnel Using Wildcards on page 5240

wildcard-group-inet6

Syntax	<pre>wildcard-group-inet6 { wildcard-source { ldp-p2mp; pim-ssm { group-range <i>multicast-prefix</i>; } rsvp-te { label-switched-path-template { (default-template <i>lsp-template-name</i>); } static-lsp <i>lsp-name</i>; } threshold-rate <i>number</i>; } }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective]</p>
Release Information	Statement introduced in Junos OS Release 10.0.
Description	<p>Configure a wildcard group matching any group IPv6 address.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • wildcard-group-inet on page 5296 • Example: Configuring Selective Provider Tunnels Using Wildcards on page 5241 • Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 5153 • Configuring a Selective Provider Tunnel Using Wildcards on page 5240

wildcard-source

Syntax	<pre>wildcard-source { ldp-p2mp; pim-ssm { group-range <i>multicast-prefix</i>; } rsvp-te { label-switched-path-template { (default-template <i>lsp-template-name</i>); } static-lsp <i>lsp-name</i>; } }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>group-prefix</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet6], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>group-prefix</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet6], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective wildcard-group-inet6]</pre>
Release Information	Statement introduced in Junos OS Release 10.0.
Description	<p>Configure a selective provider tunnel for a shared tree using a wildcard source.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• wildcard-group-inet on page 5296• wildcard-group-inet6 on page 5297• Example: Configuring Selective Provider Tunnels Using Wildcards on page 5241• Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN on page 5153• Configuring a Selective Provider Tunnel Using Wildcards on page 5240

Troubleshooting

- [Troubleshooting Multicast VPNs on page 5299](#)

Troubleshooting Multicast VPNs

- [Tracing MBGP MVPN Traffic and Operations on page 5299](#)

Tracing MBGP MVPN Traffic and Operations

To trace MBGP MVPN traffic, you can specify options with the **traceoptions** statement:

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

You can include this statement at the following hierarchy levels:

- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols mvpn]
- [edit routing-instances *routing-instance-name* protocols mvpn]

The following trace flags display the operations associated with multicast VPNs:

- **all**—All multicast VPN tracing options
- **error**—Error conditions
- **general**—General events
- **nlri**—Multicast VPN advertisements received or sent by means of BGP
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing
- **topology**—Multicast VPN topology changes caused by reconfiguration or advertisements received from other PE routers using BGP

VPLS

- [Overview on page 5299](#)
- [Configuration on page 5322](#)
- [Administration on page 5431](#)

Overview

- [VPLS on page 5300](#)

VPLS

- [Introduction to VPLS on page 5300](#)
- [VPLS Routing and Virtual Ports on page 5301](#)
- [VPLS and Aggregated Ethernet Interfaces on page 5302](#)
- [BGP Signaling for VPLS PE Routers Overview on page 5304](#)
- [BGP Route Reflectors for VPLS on page 5304](#)
- [VPLS Multihoming Overview on page 5305](#)
- [Enabling BGP Path Selection for Layer 2 VPNs and VPLS on page 5307](#)
- [VPLS Path Selection Process for PE Routers on page 5309](#)
- [BGP and VPLS Path Selection for Multihomed PE Routers on page 5311](#)
- [VPLS Multihoming Reactions to Network Failures on page 5313](#)
- [Interoperability between BGP Signaling and LDP Signaling in VPLS on page 5314](#)
- [VPLS Label Blocks Operation on page 5316](#)
- [PE Router Mesh Groups for VPLS Routing Instances on page 5320](#)
- [Understanding PIM Snooping for VPLS on page 5321](#)

Introduction to VPLS

VPLS is an Ethernet-based point-to-multipoint Layer 2 VPN. It allows you to connect geographically dispersed Ethernet local area networks (LAN) sites to each other across an MPLS backbone. For customers who implement VPLS, all sites appear to be in the same Ethernet LAN even though traffic travels across the service provider's network.

VPLS, in its implementation and configuration, has much in common with a Layer 2 VPN. In VPLS, a packet originating within a service provider customer's network is sent first to a customer edge (CE) device (for example, a router or Ethernet switch). It is then sent to a provider edge (PE) router within the service provider's network. The packet traverses the service provider's network over a MPLS label-switched path (LSP). It arrives at the egress PE router, which then forwards the traffic to the CE device at the destination customer site.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

The difference is that for VPLS, packets can traverse the service provider's network in point-to-multipoint fashion, meaning that a packet originating from a CE device can be broadcast to all the PE routers participating in a VPLS routing instance. In contrast, a Layer 2 VPN forwards packets in point-to-point fashion only.

The paths carrying VPLS traffic between each PE router participating in a routing instance are called pseudowires. The pseudowires are signaled using either BGP or LDP.

VPLS Routing and Virtual Ports

Because VPLS carries Ethernet traffic across a service provider network, it must mimic an Ethernet network in some ways. When a PE router configured with a VPLS routing instance receives a packet from a CE device, it first determines whether it has the destination of the VPLS packet in the appropriate routing table. If it does, it forwards the packet to the appropriate PE router or CE device. If it does not, it broadcasts the packet to all other PE routers and CE devices that are members of that VPLS routing instance. In both cases, the CE device receiving the packet must be different from the one sending the packet.



NOTE: In the VPLS documentation, the term *router* is used to refer to any device that provides routing functions.


When a PE router receives a packet from another PE router, it first determines whether it has the destination of the VPLS packet in the appropriate routing table. If it does, the PE router either forwards the packet or drops it depending on whether the destination is a local or remote CE device:

- If the destination is a local CE device, the PE router forwards the packet to it.
- If the destination is a remote CE device (connected to another PE router), the PE router discards the packet.

If the PE router cannot determine the destination of the VPLS packet, it floods the packet to all attached CE devices.

This process is illustrated in [Figure 121 on page 5301](#).

Figure 121: Flooding a Packet with an Unknown Destination to All PE Routers in the VPLS Instance



VPLS can be directly connected to an Ethernet switch. Layer 2 information gathered by an Ethernet switch (for example, media access control [MAC] addresses and interface ports) is included in the VPLS routing instance table. However, instead of all VPLS

interfaces being physical switch ports, the router allows remote traffic for a VPLS instance to be delivered across an MPLS LSP and arrive on a virtual port. The virtual port emulates a local, physical port. Traffic can be learned, forwarded, or flooded to the virtual port in almost the same way as traffic is sent to a local port.

The VPLS routing table learns MAC address and interface information for both physical and virtual ports. The main difference between a physical port and a virtual port is that the router captures additional information from the virtual port, an outgoing MPLS label used to reach the remote site and an incoming MPLS label for VPLS traffic received from the remote site. The virtual port is generated dynamically on a Tunnel Services Physical Interface Card (PIC) when you configure VPLS on the router.

You can also configure VPLS without a Tunnel Services PIC. To do so, you use a label-switched interface (LSI) to provide VPLS functionality. An LSI MPLS label is used as the inner label for VPLS. This label maps to a VPLS routing instance. On the PE router, the LSI label is stripped and then mapped to a logical LSI interface. The Layer 2 Ethernet frame is then forwarded using the LSI interface to the correct VPLS routing instance.

One restriction on flooding behavior in VPLS is that traffic received from remote PE routers is never forwarded to other PE routers. This restriction helps prevent loops in the core network. However, if a CE Ethernet switch has two or more connections to the same PE router, you must enable the Spanning Tree Protocol (STP) on the CE switch to prevent loops. STP is supported on MX Series routers and EX Series switches only.

The Junos OS allows standard Bridge Protocol Data Unit (BPDU) frames to pass through emulated Layer 2 connections, such as those configured with Layer 2 VPNs, Layer 2 circuits, and VPLS routing instances. However, CE Ethernet switches that generate proprietary BPDU frames might not be able to run STP across Juniper Networks routing platforms configured for these emulated Layer 2 connections.



NOTE: Under certain circumstances, VPLS provider routers might duplicate an Internet Control Message Protocol (ICMP) reply from a CE router when a PE router has to flood an ICMP request because the destination MAC address has not yet been learned. The duplicate ICMP reply can be triggered when a CE router with promiscuous mode enabled is connected to a PE router. The PE router automatically floods the promiscuous mode-enabled CE router, which then returns the ICMP request to the VPLS provider routers. The VPLS provider routers consider the ICMP request to be new and flood the request again, creating a duplicate ping reply.

VPLS and Aggregated Ethernet Interfaces

You can configure aggregated Ethernet interfaces between CE devices and PE routers for VPLS routing instances. Traffic is load-balanced across all of the links in the aggregated interface. If one or more links in the aggregated interface fails, the traffic is switched to the remaining links.

Forwarding is based on a lookup of the DA MAC address. For the remote site, if a packet needs to be forwarded over an LSP, the packet is encapsulated and forwarded through

the LSP. If the packet destination is a local site, it is forwarded over appropriate local site interface. For an aggregated Ethernet interface on the local site, packets are sent out of the load-balanced child interface. The Packet Forwarding Engine acquires the child link to transmit the data.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

When a received packet does not have a match to a MAC address in the forwarding database, the packet is forwarded over a set of interfaces determined from a lookup in the flooding database based on the incoming interface. This is denoted by a flood next hop. The flood next hop can include the aggregated Ethernet interface as the set of interfaces to flood the packet.

Each VPLS routing instance configured on a PE router has its own forwarding database entries that associate all of the MAC addresses the VPLS routing instance acquires with each corresponding port. A route is added to the kernel with a MAC address as the prefix and the next hop used to reach the destination. The route is an interface if the destination is local. For a remote destination, the route is a next hop for the remote site.

For local aggregated Ethernet interfaces on M Series and T Series routers, learning is based on the parent aggregated Ethernet logical interface. To age out MAC addresses for aggregated Ethernet interfaces, each Packet Forwarding Engine is queried to determine where the individual child interfaces are located. MAC addresses are aged out based on the age of the original interface.

For MX Series routers and EX Series switches, when a Dense Port Concentrator (DPC) learns a MAC address it causes the Routing Engine to age out the entry. This behavior applies to all logical interfaces. For an aggregated Ethernet logical interface, once all the member DPCs have aged out the entry, the entry is deleted from the Routing Engine.

**Related
Documentation**

- [Configuring Interfaces for VPLS Routing on page 5336](#)
- [Configuring Aggregated Ethernet Interfaces for VPLS on page 5341](#)

BGP Signaling for VPLS PE Routers Overview

BGP can autonomously signal pseudowires between the PE routers participating in the same virtual private LAN service (VPLS) network. As PE routers are added to and removed from the VPLS network, BGP can signal pseudowires to new PE routers and tear down old pseudowires to old PE routers. Each PE router only needs to be configured with the identity of the VPLS routing instance. Each PE router does not need to be configured with the identities of all of the PE routers that are or might become a part of the VPLS network.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

When you configure BGP for signaling in a VPLS network, customer sites can be either single-homed to a single PE router or multihomed to two or more PE routers. Multihoming provides redundancy for the connection between the customer site and the service provider's network.

You can either configure all of the PE routers in the VPLS network as a full mesh or you can use BGP route reflectors. For full mesh configurations, each PE router needs to be able to create a bidirectional pseudowire to each of the other PE routers participating in the VPLS network.

Related Documentation

- [VPLS Multihoming Overview on page 5305](#)
- [VPLS Path Selection Process for PE Routers on page 5309](#)

BGP Route Reflectors for VPLS

In large networks, it might be necessary to configure BGP route reflectors to reduce the control plane workload for the routers participating in the VPLS network. BGP route reflectors can help to reduce the workload of the network control plane in the following ways.

- Making it unnecessary to configure all of the VPLS PE routers in a full mesh.
- Limiting the total volume of BGP VPLS messages exchanged within the network by transmitting messages to interested routers only (instead of all of the BGP routers in the network)
- Reducing the network signaling load whenever another BGP router is added to or removed from the network

The basic solution to these problems is to deploy a small group of BGP route reflectors that are in a full mesh with one another. Each of the VPLS PE routers is configured to have a BGP session with one or more of the route reflectors, making it unnecessary to maintain a full mesh of BGP sessions between all of the PE routers.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

This type of configuration only affects the control plane of the VPLS network (how routers signal and tear down pseudowires to one another in the network). The actual data plane state and forwarding paths for the VPLS traffic are not modified by the route reflectors. Effectively, the VPLS pseudowires should take the same paths across the network whether or not you have configured route reflectors. For a description of how VPLS selects the best path to a PE router, see [“VPLS Path Selection Process for PE Routers” on page 5309](#).

The MAC addresses themselves are not exchanged or processed in any way by BGP. Each VPLS PE router performs all MAC address learning and aging individually. BGP's only function relative to VPLS is to exchange messages related to automatic discovery of PE routers being added to and removed from the VPLS network and the MPLS label exchange needed to signal a pseudowire from one PE router to another.

Related Documentation


- [VPLS Path Selection Process for PE Routers on page 5309](#)
- [Example: Configuring a Route Reflector](#)
- [Example: NG-VPLS Using Point-to-Multipoint LSPs](#)
- [Example: Next-Generation VPLS for Multicast with Multihoming](#)

VPLS Multihoming Overview

Virtual private LAN service (VPLS) multihoming enables you to connect a customer site to two or more PE routers to provide redundant connectivity. A redundant PE router can provide network service to the customer site as soon as a failure is detected. VPLS multihoming helps to maintain VPLS service and traffic forwarding to and from the multihomed site in the event of the following types of network failures:

- PE router to CE device link failure
- PE router failure
- MPLS-reachability failure between the local PE router and a remote PE router

Figure 122: CE Device Multihomed to Two PE Routers



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Figure 122 on page 5305 illustrates how a CE device could be multihomed to two PE routers. Device CE1 is multihomed to Routers PE1 and PE2. Device CE2 has two potential paths to reach Device CE1, but only one path is active at any one time. If Router PE1 were the designated VPLS edge (VE) device (also called a designated forwarder), BGP would signal a pseudowire from Router PE3 to Router PE1. If a failure occurred over this path, Router PE2 would be made the designated VE device, and BGP would re-signal the pseudowire from Router PE3 to Router PE2.

Multihomed PE routers advertise network layer reachability information (NLRI) for the multihomed site to the other PE routers in the VPLS network. The NLRI includes the site ID for the multihomed PE routers. For all of the PE routers multihomed to the same CE device, you need to configure the same site ID. The remote VPLS PE routers use the site ID to determine where to forward traffic addressed to the customer site. To avoid route collisions, the site ID shared by the multihomed PE routers must be different than the site IDs configured on the remote PE routers in the VPLS network.

Although you configure the same site ID for each of the PE routers multihomed to the same CE device, you can configure unique values for other parameters, such as the route distinguisher. These values help to determine which multihomed PE router is selected as the designated VE device to be used to reach the customer site.



BEST PRACTICE: We recommend that you configure unique route distinguishers for each multihomed PE router. Configuring unique route distinguishers helps with faster convergence when the connection to a primary multihomed PE router goes down. If you configure unique route distinguishers, the other PE routers in the VPLS network must maintain additional state for the multihomed PE routers.

Remote PE routers in the VPLS network need to determine which of the multihomed PE routers should forward traffic to reach the CE device. To make this determination, remote PE routers use the VPLS path-selection process to select one of the multihomed PE routers based on its NLRI advertisement. Because remote PE routers pick only one of the NLRI advertisements, it establishes a pseudowire to only one of the multihomed PE routers, the PE router that originated the winning advertisement. This prevents multiple paths from being created between sites in the network, preventing the formation of Layer 2 loops. If the selected PE router fails, all PE routers in the network automatically switch to the backup PE router and establish new pseudowires to it.



BEST PRACTICE: To prevent the formation of Layer 2 loops between the CE devices and the multihomed PE routers, we recommend that you employ the Spanning Tree Protocol (STP) on your CE devices. Layer 2 loops can form due to incorrect configuration. Temporary Layer 2 loops can also form during convergence after a change in the network topology.

The PE routers run the BGP path selection procedure on locally originated and received Layer 2 route advertisements to establish that the routes are suitable for advertisement to other peers, such as BGP route reflectors. If a PE router in a VPLS network is also a

route reflector, the path selection process for the multihomed site has no effect on the path selection process performed by this PE router for the purpose of reflecting Layer 2 routes. Layer 2 prefixes that have different route distinguishers are considered to have different NLRIs for route reflection. The VPLS path selection process enables the route reflector to reflect all routes that have different route distinguishers to the route reflector clients, even though only one of these routes is used to create the VPLS pseudowire to the multihomed site.

**Related
Documentation**

- [Configuring VPLS Multihoming \(FEC 128\) on page 5343](#)
- [Advantages of Using Autodiscovery for VPLS Multihoming](#)

Enabling BGP Path Selection for Layer 2 VPNs and VPLS

Layer 2 VPNs and VPLS share the same path selection process for determining the optimal path to reach all of the destinations shared within a single routing instance. For Layer 2 VPN and VPLS topologies that do not include multihomed PE routers, the path selection process is straightforward since there is just a single path from each PE router to each CE device. However, if multihoming is configured for one or more of the CE devices, the path selection process becomes more complex, since there can be two or more valid paths to reach each multihomed CE device.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

The Layer 2 VPN and VPLS path selection process uses the following path selection algorithms:

- On the Provider routers within the service providers network, the standard BGP path selection algorithm is used. Using the standard BGP path selection for Layer 2 VPN and VPLS routes allows service providers to leverage their existing Layer 3 VPN network infrastructure to also support Layer 2 VPNs and VPLS. The BGP path selection algorithm also helps to ensure that the service provider's network behaves predictably with regard to Layer 2 VPN and VPLS path selection. This is particularly important in networks employing route reflectors and multihoming.

When a Provider router receives multiple paths for the same destination prefix (for example, a multihomed CE device), one path is selected based on the BGP path selection algorithm and placed in the `bgp.l2vpn.0` routing table and the appropriate `instance.l2vpn.0` routing table. However, all of the available paths (including the backup paths for multihomed CE devices) are advertised to the intermediate BGP routers and the PE routers in the Layer 2 VPN or VPLS routing instances.

For more information about the BGP path selection process, see *Understanding BGP Path Selection*.

- Once a PE router receives all of the available paths to each CE device, it runs the designated forwarder path selection algorithm to select the preferred path to reach each CE device, independently of the results of the earlier BGP path selection algorithm run on the Provider router. The VPLS designated forwarder algorithm uses the D-bit, preference, and PE router identifier to determine which of the valid paths to each CE device to use. The PE router might select a path to reach a CE device which is different from the path selected by the BGP-based Provider routers. In this scenario, the following is the expected behavior for traffic sent to the multihomed CE device:
 1. If the path selected by the remote PE router is available, traffic will traverse the network to the multihomed CE device using the remote PE router's preferred path (again, ignoring the path selected by the BGP-based Provider routers).
 2. If the path selected by the remote PE router fails, the Provider routers switch the traffic destined for the multihomed CE device to the alternate path as soon as failure is detected. They then notify the remote PE router of the path failure. The remote PE router updates its routing table accordingly.

For more information about the VPLS designated forwarder path selection algorithm, see [“VPLS Path Selection Process for PE Routers” on page 5309](#). This algorithm is also described in the Internet draft [draft-ietf-l2vpn-vpls-multihoming-03.txt](#), *BGP based Multi-homing in Virtual Private LAN Service*.

**Related
Documentation**

- [Understanding BGP Path Selection](#)
- [VPLS Path Selection Process for PE Routers on page 5309](#)

VPLS Path Selection Process for PE Routers

The VPLS path selection process is used to select the best path between a remote PE router and a local PE router in a VPLS network. This path selection process is applied to routes received from both single-homed and multi-homed PE routers.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

When the VPLS path selection process is complete, a PE router is made the designated VPLS edge (VE) device. The designated VE device effectively acts as the endpoint for the VPLS pseudowire that is signaled from the remote PE router. Once a PE router is made the designated VE device, a pseudowire can be signaled between the remote PE router and the local PE router and then VPLS packets can begin to flow between the PE routers.

Routes from multihomed PE routers connected to the same customer site share the same site ID, but can have different route distinguishers and block offsets. You can alter the configurations of the route distinguishers and block offsets to make a router more likely or less likely to be selected as the designated VE device.

On each PE router in the VPLS network, the best path to the CE device is determined by completing the following VPLS path selection process on each route advertisement received:

1. If the advertisement has the down bit set to 0, the advertisement is discarded.
2. Select the path with a higher preference. The preference attribute is obtained from the site-preference configured using the **site-preference** statement at the **[edit routing-instances routing-instance-name protocols vpls site site-name]** hierarchy level. If the site preference is 0, the preference attribute is obtained from the local preference.
3. If the preference values are the same, select the path with the lower router ID.
4. If the router IDs are the same, the routes are from the same PE router and the advertisement is considered to be an update. The router ID corresponds to the value of the originator ID for the BGP attribute (if present). Otherwise, the IP address for the remote BGP peer is used.
5. If the block offset values are the same, the advertisement is considered to be an update.

Once the VPLS path selection process has been completed and the designated VE device has been selected, a pseudowire is signaled between the remote PE router and the local PE router.



NOTE: The VPLS path selection process works the same whether or not the route has been received from another PE router, a route reflector, or an autonomous system border router (ASBR).

When the remote PE router establishes or refreshes a pseudowire to the local PE router, it verifies that the prefix is in the range required for the site ID based on the block offset and label range advertised by the designated VE device. If the prefix is out of range, the pseudowire status is set to out of range.

The following cases outline the potential decisions that could be made when a PE router completes the VPLS path selection process for a Layer 2 advertisement in the VPLS network:

- The PE router originated one of the advertisements and selected its own advertisement as the best path.

This PE router has been selected as the designated VE device. Selection as the designated VE device triggers the creation of pseudowires to and from the other PE routers in the VPLS network. If the remote customer site is multihomed, the designated VE device triggers the creation of pseudowires to and from only the designated VE device for the remote site.

- The PE router originated one of the advertisements but did not select its own advertisement as the best path.

This PE router is a redundant PE router for a multihomed site, but it was not selected as the designated VE device. However, if this PE router has just transitioned from being the designated VE device (meaning it was receiving traffic from the remote PE routers addressed to the multihomed customer site), the PE router tears down all the pseudowires that it had to and from the other PE routers in the VPLS network.

- The PE router received the route advertisements and selected a best path. It did not originate any of these advertisements because it was not connected to the customer site.

If the best path to the customer site (the designated VE device) has not changed, nothing happens. If the best path has changed, this PE router brings up pseudowires to and from the newly designated VE device and tears down the pseudowires to and from the previously designated VE device.

If this PE router does not select a best path after running the VPLS path selection process, then the local PE router does not consider the remote site to exist.

When a VE device receives an advertisement for a Layer 2 NLRI that matches its own site ID but the site is not multihomed, the pseudowire between the VE device and the transmitting PE router transitions to a site collision state and is not considered to be up.

**Related
Documentation**

- [BGP Route Reflectors for VPLS on page 5304](#)

BGP and VPLS Path Selection for Multihomed PE Routers

The BGP and VPLS path selection procedures are used to select the best path between the remote PE router and one of the multihomed PE routers. As part of these path selection procedures, one of the multihomed PE routers is made the designated VE device. The designated VE device effectively acts as the endpoint for the VPLS pseudowire from the remote PE router. Once a multihomed PE router is made the designated VE device, a pseudowire can be created between the remote PE router and the multihomed PE router.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Routes from multihomed PE routers connected to the same customer site share the same site ID, but can have different route distinguishers and block offsets. On each PE router in the VPLS network, the best path to the multihomed PE router is determined by completing the following VE device-selection procedures on each route advertisement received from a multihomed PE router:

1. BGP designated VE device-selection procedure—Runs before the VPLS designated VE device-selection procedure. However, the BGP designated VE device-selection procedure is used only when the route distinguishers for the multihomed PE routers are identical. If the route distinguishers are unique, only the VPLS designated VE device-selection procedure is run.
2. VPLS designated VE device-selection procedure—Runs after the BGP designated VE device-selection procedure. However, if the route distinguishers for each multihomed PE router are unique, the advertisements are not considered relevant to the BGP designated VE device-selection procedure. As a consequence, only the VPLS designated VE device-selection procedure is used.

The BGP designated VE device-selection procedure is as follows:

1. If the advertisement has the down bit set to 0, the advertisement is discarded.
2. Select the path with a higher preference. The preference attribute is obtained from the site-preference configured using the **site-preference** statement at the [edit routing-instances *routing-instance-name* protocols vpls site *site-name*] hierarchy level. If the site-preference is 0, the preference attribute is obtained from the local-preference.
3. If the preference values are the same, select the path with the lower router-id.
4. If the router-ids are the same, the routes are from the same PE router and the advertisement is considered to be an update.

Once the BGP designated VE device-selection procedure is complete, the VPLS designated VE device-selection procedure begins. This procedure is carried out regardless of the outcome of the BGP designated VE device-selection procedure:

1. If the advertisement has the down bit set to 0, the advertisement is discarded.
2. Select the path with a higher preference. The preference attribute is obtained from the site-preference configured using the **site-preference** statement at the [edit routing-instances *routing-instance-name* protocols vpls site *site-name*] hierarchy level. If the site-preference is 0, the preference attribute is obtained from the local-preference.
3. If the preference values are the same, select the path with the lower router-id.
4. If the router-ids are the same, select the path with a lower route distinguisher.
5. If the route distinguishers are the same, select the path with the lower block offset value.
6. If the block offset values are the same, the advertisement is considered to be an update.

Once the BGP and VPLS path selection procedures have been completed and the designated VE devices have been selected, a pseudowire can be created between the remote PE router and the multihomed PE router.

When the remote PE router establishes or refreshes a pseudowire to the local PE router, it verifies that the prefix is in the range required for the site ID based on the block offset and label range advertised by the designated VE device. If the prefix is out of range, the pseudowire status is set to out of range.

The following cases outline the potential decisions that could be made when a PE router completes the BGP and VPLS path selection procedures for a Layer 2 advertisement in the VPLS network:

- The PE router originated one of the multihomed advertisements and selected its own advertisement as the best path.

This PE router has been selected as the designated VE device. Selection as the designated VE device triggers the creation of pseudowires to and from the other PE routers in the VPLS network. When the remote customer site is also multihomed, the designated VE device triggers the creation of pseudowires to and from only the designated VE device for the remote site.

- The PE router originated one of the multihomed advertisements but did not select its own advertisement as the best path.

This PE router is one of the redundant PE routers for the multihomed site; it was not selected as the designated VE device. However, if this PE router has just transitioned from being the designated VE device (meaning it was receiving traffic from the remote PE routers addressed to the multihomed customer site), the PE router tears down all the pseudowires that it had to and from the other PE routers in the VPLS network.

- The PE router receives the multihomed advertisements and selects a best path; it does not originate any of these advertisements because it is not connected to the multihomed customer site.

If the preferred path to the customer site (the designated VE device) has not changed, nothing happens. If the preferred path has changed, this PE router brings up pseudowires to and from the newly designated VE device and tears down the pseudowires to and from the previously designated VE device.

If this PE router does not select a best path after running the BGP and VPLS path selection process, the local PE router does not consider the remote site to exist.

When a VE device receives an advertisement for a Layer 2 NLRI which matches its own site ID but the site is not multihomed, the pseudowire between it and the transmitting PE router transitions to a site collision state and is not considered to be up.

VPLS Multihoming Reactions to Network Failures

VPLS multihoming is designed to protect customer sites from a loss of network connectivity in the event of the following types of network failures:

- Link failure between the CE device and the PE router—BGP on the PE router is notified when the link goes down. BGP sets the circuit status vector bit in the MP_REACH_NLRI to indicate that the circuit is down.

If all of the VPLS local attachment circuits are down, then BGP modifies the down bit in the VPLS advertisement Layer2-Extended-Community to indicate that the customer site is down. When the bit is modified, BGP advertises the route to all of the remote PE routers to notify them that the circuit (and site) is down. Each of the remote PE routers run the BGP and VPLS path selection procedures again and reroute the VPLS pseudowires as needed.

- MPLS connectivity failure to the remote PE router—On the multihomed PE router, BGP discovers that MPLS cannot connect to the BGP next hop in the service provider's network. BGP modifies the circuit status vector bit in the MP_REACH_NLRI to indicate that the LSP is down. Once the bit is modified, BGP readvertises the route to all of the remote PE routers to notify them that connectivity from the local site to the remote site is down.

The remote PE routers each run the BGP and VPLS path selection procedures again. With the LSP to the original multihomed PE router down, the remote PE routers designate the backup multihomed PE router as the VE device for the multihomed customer site. The pseudowires to and from the remote PE routers are then rerouted to the backup multihomed PE router.

- PE router failure—When either the multihomed PE router or the BGP process running on it fails, the remote PE routers detect the expiration of the holdtimer, bring down their peering sessions, and delete the Layer 2 advertisements from that multihomed PE router. The remote PE routers each run the BGP and VPLS path selection procedures again and reroute their pseudowires to the backup multihomed PE router.

Alternatively, the remote PE routers could discover that the BGP next hop, represented by the failed multihomed PE router, is unreachable. For this case, the remote PE routers mark the Layer 2 routes advertised by the multihomed PE router as unreachable. The

remote PE routers each run the BGP and VPLS path selection procedures again and reroute their pseudowires to the backup multihomed PE router.

The remote PE routers behave in the same manner if you reconfigure the local preference attribute of the primary multihomed PE router (effectively performing an administrative failover to the backup multihomed PE router). On the primary multihomed PE router, BGP advertises a Layer 2 update with the new local preference attribute to all of the remote PE routers. The remote PE routers each run the BGP and VPLS path selection procedures again and reroute their pseudowires to the backup multihomed PE router.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Interoperability between BGP Signaling and LDP Signaling in VPLS

You can configure a VPLS routing instance where some of the PE routers use BGP for signaling and some use LDP for signaling.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

The following concepts form the basis of the configuration needed to include both BGP-signaled and LDP-signaled PE routers in a VPLS routing instance:

- **PE router mesh group**—Consists of a set of routers participating in a VPLS routing instance that share the same signaling protocol, either BGP or LDP, and are also fully meshed. Each VPLS routing instance can have just one BGP mesh group. However, you can configure multiple LDP mesh groups for each routing instance.
- **Border router**—A PE router that must be reachable by all of the other PE routers participating in a VPLS routing instance, whether they are LDP-signaled or BGP-signaled. Bidirectional pseudowires are created between the border router and all of these PE routers. The border router is aware of the composition of each PE mesh group configured as a part of the VPLS routing instance. It can also have direct connections to local CE routers, allowing it to act as a typical PE router in a VPLS routing instance.

The following sections describe how the LDP-signaled and BGP-signaled PE routers function when configured to interoperate within a VPLS routing instance:


- [LDP-Signaled and BGP-Signaled PE Router Topology on page 5314](#)
- [Flooding Unknown Packets Across Mesh Groups on page 5316](#)
- [Unicast Packet Forwarding on page 5316](#)

LDP-Signaled and BGP-Signaled PE Router Topology

[Figure 123 on page 5315](#) illustrates a topology for a VPLS routing instance configured to support both BGP and LDP signaling. Router B is the border router. Routers PE1 and PE2 are in the LDP-signaled mesh group LDP-1. Routers PE3, PE4, and PE5 are in the

LDP-signaled mesh group LDP-2. Routers PE6, PE7, PE8, and router B (the border router) are in the BGP-signaled mesh group. The border router also acts as a standard VPLS router (having local connections to CE routers). All of the PE routers shown are within the same VPLS routing instance.

Figure 123: BGP and LDP Signaling for a VPLS Routing Instance



Two-way pseudowires are established between the PE routers in each mesh group and between each PE router in the VPLS routing instance and the border router. In [Figure 123 on page 5315](#), two-way pseudowires are established between routers PE1 and PE2 in mesh group LDP-1, routers PE3, PE4, and PE5 in mesh group LDP-2, and routers PE6, PE7, and PE8 in the BGP mesh group. Routers PE1 through PE8 also all have two-way pseudowires to the Border router. Based on this topology, the LDP-signaled routers are able to interoperate with the BGP-signaled routers. Both the LDP-signaled and BGP-signaled PE routers can logically function within a single VPLS routing instance.



NOTE: The following features are not supported for VPLS routing instances configured with both BGP and LDP signaling:

- Point-to-multipoint LSPs
- Integrated routing and bridging
- IGMP snooping

Flooding Unknown Packets Across Mesh Groups

Broadcast, multicast, and unicast packets of unknown origin received from a PE router are flooded to all local CE routers. They are also flooded to all of the PE routers in the VPLS routing instance except the PE routers that are a part of the originating PE router mesh group.

For example, if a multicast packet is received by the border router in [Figure 123 on page 5315](#), it is flooded to the two local CE routers. It is also flooded to routers PE1 and PE2 in the LDP-1 mesh group and to routers PE3, PE4, and PE5 in the LDP-2 mesh group. However, the packet is not flooded to routers PE6, PE7, and PE8 in the BGP mesh group.

Unicast Packet Forwarding

The PE border router is made aware of the composition of each PE router mesh group. From the data plane, each PE router mesh group is viewed as a virtual pseudowire LAN. The border router is configured to interconnect all of the PE router mesh groups belonging to a single VPLS routing instance. To interconnect the mesh groups, a common MAC table is created on the border router.

Unicast packets originating within a mesh group are dropped if the destination is another PE router within the same mesh group. However, if the destination MAC address of the unicast packet is a PE router located in a different mesh group, the packet is forwarded to that PE router.

VPLS Label Blocks Operation

A virtual private LAN service (VPLS) is a Layer 2 (L2) service that emulates a local area network (LAN) across a wide area network (WAN). VPLS labels are defined and exchanged in the Border Gateway Protocol (BGP) control plane. In the Junos OS implementation, label blocks are allocated and used in the VPLS control plane for two primary functions: autodiscovery and signaling.

- Autodiscovery—A method for automatically recognizing each provider edge (PE) router in a particular VPLS domain, using BGP update messages.
- Signaling—Each pair of PE routers in a VPLS domain sends and withdraws VPN labels to each other. The labels are used to establish and dismantle pseudowires between the routers. Signaling is also used to transmit certain characteristics of a pseudowire.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

The PE router uses BGP extended communities to identify the members of its VPLS. Once the PE router discovers its members, it is able to establish and tear down pseudowires between members by exchanging and withdrawing labels and transmitting certain characteristics of the pseudowires.

The PE router sends common update messages to all remote PE routers, using a distinct BGP update message, thereby reducing the control plane load. This is achieved by using VPLS label blocks.

Elements of Network Layer Reachability Information

VPLS BGP network layer reachability information (NLRI) is used to exchange VPLS membership and parameters. The elements of a VPLS BGP NLRI are defined in [Table 385 on page 5317](#).

Table 385: NLRI Elements

Element	Acronym	Description	Default Size (Octets)
Length		Total length of the NLRI size represented in bytes.	2
Route Distinguisher	RD	Unique identifier for each routing instance configured on a PE.	8
VPLS Edge ID	VE ID	Unique number to identify the edge site.	2
VE Block Offset	VBO	Value used to identify a label block from which a label value is selected to set up pseudowires for a remote site.	2
VE Block Size	VBS	Indicates the number of pseudowires that peers can have in a single block.	2
Label Base	LB	Starting value of the label in the advertised label block.	3

Requirements for NLRI Elements

Junos OS requires a unique route distinguisher (RD) for each routing instance configured on a PE router. A PE router might use the same RD across a VPLS (or VPN) domain or it might use different RDs. Using different RDs helps identify the originator of the VPLS NLRI.

The VPLS edge (VE) ID can be a unique VE ID, site ID, or customer edge (CE) ID. The VE ID is used by a VPLS PE router to index into label blocks used to derive the transmit and receive VPN labels needed for transport of VPLS traffic. The VE ID identifies a particular site, so it needs to be unique within the VPLS domain, except for some scenarios such as multihoming.

All PE routers have full mesh connectivity with each other to exchange labels and set up pseudowires. The VE block size (VBS) is a configurable value that represents the number of label blocks required to cover all the pseudowires for the remote peer.

A single label block contains 8 labels (1 octet) by default. The default VBS in Junos OS is 2 blocks (2 octets) for a total of 16 labels.

How Labels are Used in Label Blocks

Each PE router creates a mapping of the labels in the label block to the sites in a VPLS domain. A PE router advertising a label block with a block offset indicates which sites can use the labels to reach it. When a PE router is ready to advertise its membership to

a VPLS domain, it allocates a label block and advertises the VPLS NLRI. In this way, other PE routers in the same VPLS domain can learn of the existence of the VPLS and set up pseudowires to it if needed. The VPLS NLRI advertised for this purpose is referred to as the *default VPLS NLRI*. The label block in the default VPLS NLRI is referred to as the *default label block*.

Label Block Composition

A label block (set of labels) is used to reach a given site ID. A single label block contains 8 labels (1 octet) by default. The VBS is 2 octets by default in Junos OS.

The label block advertised is defined as a label base (LB) and a VE block size (VBS). It is a contiguous set of labels (LB, LB+1,...,LB+VBS-1). For example, when Router PE-A sends a VPLS update, it sends the same label block information to all other PE routers. Each PE router that receives the LB advertisement infers the label intended for Router PE-A by adding its own site ID to the label base.

In this manner, each receiving PE gets a unique label for PE-A for that VPLS. This simple method is enhanced by using a VE block offset (VBO).

A label block is defined as: <Label Base (LB), VE block offset (VBO), VE block size (VBS)> is the set {LB+VBO, LB+VBO+1,...,LB+VBO+VBS-1}.

Label Blocks in Junos OS

Instead of a single large label block to cover all VE IDs in a VPLS, the Junos OS implementation contains several label blocks, each with a different label base. This makes label block management easier, and also allows Router PE-A to seamlessly integrate a PE router joining a VPLS with a site ID not covered by the set of label blocks that Router PE-A has already advertised.

VPLS Label Block Structure

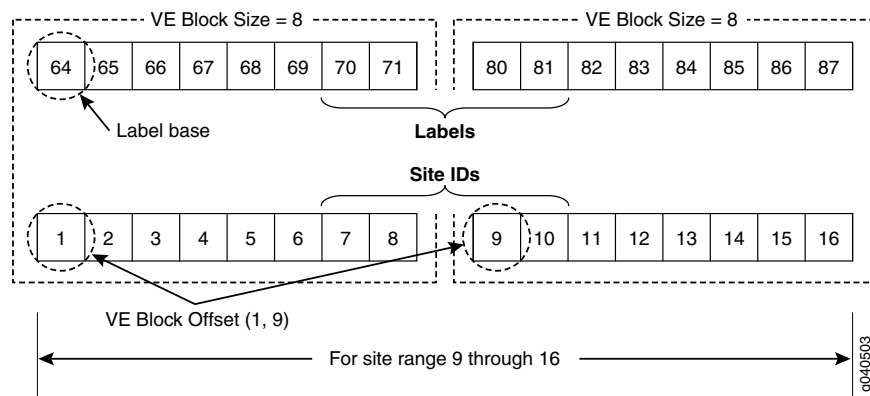
This section illustrates how a label block is uniquely identified.

A VPLS BGP NLRI with site ID V, VE block offset VBO, VE block size VBS, and label base LB communicates the following to its peers:

- Label block for V: Labels from LB to (LB + VBS -1).
- Remote VE set for V: from VBO to (VBO + VBS -1).

The label block advertised is a set of labels used to reach a given site ID. If there are several label blocks, the remote VE set helps to identify which label block to use. The example in [Figure 124 on page 5319](#) illustrates label blocks. There are two blocks and each block has eight labels. In this example, the label values are 64 to 71 and 80 to 87.

Figure 124: VPLS Label Block Structure



To create a one-to-one mapping of these 16 labels to 16 sites, assume the site IDs are the numbers 1 to 16, as shown in the illustration. The site block indicates which site ID can use which label in the label block. So, in the first block, site ID 1 uses 64, site ID 2 uses 65, and so forth. Finally, site ID 8 uses 71. The 9th site ID will use the second block instead of the first block.

The labels are calculated by comparing the values of $VBO \leq \text{Local site ID} < (VBO + VBS)$. Consequently, site ID 9 uses 80, site ID 10 uses 81, and so on.

To further illustrate the one-to-one mapping of labels to sites, assume a label block with site offset of 1 and a label base of 10. The combination of label base and block offset contained in the VPLS NLRI provides the mapping of labels to site IDs. The block offset is the starting site ID that can use the label block as advertised in the VPLS NLRI.

To advertise the default VPLS NLRI, a PE router picks a starting block offset that fits its own site ID and is such that the end block offset is a multiple of a single label block. In Junos OS a single label block is eight labels by default.

The end block offset is the last site ID that maps to the last label in the label block. The end offset for the first block is 8 which maps to label 17 and the second block is 16. For example, a site with ID 3 picks a block offset of 1 and advertises a label block of size 8 to cover sites with IDs 1 to 8. A site with ID 10 picks a block offset of 9 to cover sites with IDs 9 to 16.

The VPLS NLRI shown in [Figure 125 on page 5320](#) is for site ID 18. The label base contains value 262145. The block offset contains value 17. The illustration shows which site IDs correspond to which labels.

Figure 125: Label Mapping Example

VPLS NLRI for Site ID 18

Length
RD
VE ID - 18
VE Block Offset - 17
VE Block Size - 8
Label Base - 262145

Label Mapping for Site ID 18

Label Base = 262145

Label Block

Label	262145	262146	262147	262148	262149	262150	262151	262152
Site ID	17	18	19	20	21	22	23	24

Site Offset = 17

Site IDs

g040504

If a PE router configured with site ID 17 is in the same VPLS domain as a PE router configured with site ID 18, it receives the VPLS NLRI as shown in Figure 3. So it uses label 262145 to send traffic to site 18. Similarly, a PE router configured with site ID 19 uses label 262147 to send traffic to a PE router configured with site ID 18. However, only PE routers configured with site IDs 17 to 24 can use the label block shown to set up pseudowires.

Related Documentation

- *Example: Building a VPLS From Router 1 to Router 3 to Validate Label Blocks*

PE Router Mesh Groups for VPLS Routing Instances

A PE router mesh group consists of a set of routers participating in a VPLS routing instance that share the same signaling protocol, either BGP or LDP. Each VPLS routing instance can have just one BGP mesh group. However, you can configure multiple LDP mesh groups for each routing instance.

The Junos OS can support up to 16 mesh groups on MX Series routers and up to 128 on M Series and T Series routers. However, two mesh groups are created by default, one for the CE routers and one for the PE routers. Therefore, the maximum number of user-defined mesh groups is 14 for MX Series routers and 126 for M Series and T Series routers. PE router mesh groups are not supported on J Series routers.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

The Junos OS supports both forwarding equivalency class (FEC) 128 and FEC 129. FEC 129 uses VPLS autodiscovery to convey endpoint information. FEC 128 requires manually configured pseudowires.

The following describes the behavior of mesh groups in regards to BGP-signaled PE routers and LDP-signaled PE routers:

- BGP-signaled PE routers—Automatically discovered PE routers that use BGP for signaling are associated with the default VE mesh group. You cannot configure the Junos OS to associate these routers with a user-defined VE mesh group.

- LDP-signaled PE routers (FEC 128)—PE routers statically configured using FEC-128 LDP signaling are placed in a default mesh group. However, you can configure a VE mesh group and associate each LDP FEC-128 neighbor with it. Each configured VE mesh group contains a set of VEs that are in the same interior gateway protocol (IGP) routing instance and are fully meshed with each other in the control and data planes.
- LDP-signaled PE routers (FEC 129)—Configuration for a mesh group for FEC 129 is very similar to the configuration for FEC 128.

Note the following differences for FEC 129:

- Each user-defined mesh group must have a unique route distinguisher. Do not use the route distinguisher that is defined for the default mesh group at the **[edit routing-instances]** hierarchy level.
- Each user-defined mesh group must have its own import and export route target.
- Each user-defined mesh group can have a unique Layer 2 VPN ID. By default, all the mesh groups that are configured for the a VPLS routing-instance use the same Layer 2 VPN ID, the one that you configure at the **[edit routing-instances]** hierarchy level.

Related Documentation

- *Example: Configuring BGP Autodiscovery for LDP VPLS*

Understanding PIM Snooping for VPLS

There are two ways to direct PIM control packets:

- By the use of PIM snooping
- By the use of PIM proxying

PIM snooping configures a device to examine and operate only on PIM hello and join/prune packets. A PIM snooping device snoops PIM hello and join/prune packets on each interface to find interested multicast receivers and populates the multicast forwarding tree with this information. PIM snooping differs from PIM proxying in that both PIM hello and join/prune packets are transparently flooded in the VPLS as opposed to the flooding of only hello packets in the case of PIM proxying. PIM snooping is configured on PE routers connected through pseudowires. PIM snooping ensures that no new PIM packets are generated in the VPLS, with the exception of PIM messages sent through LDP on pseudowires.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

A device that supports PIM snooping snoops hello packets received on attachment circuits. It does not introduce latency in the VPLS core when it forwards PIM join/prune packets.

To configure PIM snooping on a PE router, use the **pim-snooping** statement at the **[edit routing-instances instance-name protocols]** hierarchy level:

```
routing-instances {
  customer {
    instance-type vpls;
    ...
    protocols {
      pim-snooping {
        traceoptions {
          file pim.log size 10m;
          flag all;
          flag timer disable;
        }
      }
    }
  }
}
```

“[Example: Configuring PIM Snooping for VPLS](#)” on [page 3746](#) explains the PIM snooping method. The use of the PIM proxying method is not discussed here and is outside the scope of this document. For more information about PIM proxying, see [PIM Snooping over VPLS](#).

**Related
Documentation**

- [PIM Snooping for VPLS Use Cases](#)
- [Example: Configuring PIM Snooping for VPLS on page 3746](#)

Configuration

- [Configuration Tasks on page 5322](#)
- [Configuration Statements on page 5376](#)

Configuration Tasks

- [Configuring an Ethernet Switch as the CE Device on page 5323](#)
- [Introduction to Configuring VPLS on page 5323](#)
- [Configuring VPLS Routing Instances on page 5324](#)
- [Configuring Interfaces for VPLS Routing on page 5336](#)
- [Configuring Static Pseudowires for VPLS on page 5342](#)
- [Configuring VPLS Multihoming \(FEC 128\) on page 5343](#)
- [Enabling BGP Path Selection for Layer 2 VPNs and VPLS on page 5346](#)
- [Configuring EXP-Based Traffic Classification for VPLS on page 5348](#)
- [Configuring VPLS Load Balancing on page 5348](#)
- [Configuring VPLS Fast Reroute Priority on page 5350](#)
- [Configuring VPLS Without a Tunnel Services PIC on page 5351](#)
- [Mapping VPLS Traffic to Specific LSPs on page 5352](#)
- [Configuring Firewall Filters and Policers for VPLS on page 5353](#)
- [Standard Firewall Filter Match Conditions for VPLS Traffic on page 5358](#)
- [Specifying the VT Interfaces Used by VPLS Routing Instances on page 5365](#)

- [Flooding Unknown Traffic Using Point-to-Multipoint LSPs on page 5365](#)
- [Configuring Interoperability Between BGP Signaling and LDP Signaling in VPLS on page 5369](#)
- [Tracing VPLS Traffic and Operations on page 5374](#)
- [Configuring the Label Block Size on page 5375](#)

Configuring an Ethernet Switch as the CE Device

For VPLS configurations, the CE device does not necessarily need to be a router. You can link the PE routers directly to Ethernet switches. However, there are a few configuration issues to be aware of:

- When you configure VPLS routing instances and establish two or more connections between a CE Ethernet switch and a PE router, you must enable the Spanning Tree Protocol (STP) on the switch to prevent loops.
- The Junos OS allows standard Bridge Protocol Data Unit (BPDU) frames to pass through emulated Layer 2 connections, such as those configured with Layer 2 VPNs, Layer 2 circuits, and VPLS instances. However, CE Ethernet switches that generate proprietary BPDU frames might not be able to run STP across Juniper Networks routing platforms configured for these emulated Layer 2 connections.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Introduction to Configuring VPLS

Virtual private LAN service (VPLS) allows you to provide a point-to-multipoint LAN between a set of sites in a virtual private network (VPN).

To configure VPLS functionality, you must enable VPLS support on the provider edge (PE) router. You must also configure PE routers to distribute routing information to the other PE routers in the VPLS and configure the circuits between the PE routers and the customer edge (CE) routers.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Each VPLS is configured under a routing instance of type **vpls**. A **vpls** routing instance can transparently carry Ethernet traffic across the service provider's network. As with other routing instances, all logical interfaces belonging to a VPLS routing instance are listed under that instance.

In addition to VPLS routing instance configuration, you must configure MPLS label-switched paths (LSPs) between the PE routers, IBGP sessions between the PE routers, and an interior gateway protocol (IGP) on the PE and provider (P) routers.

By default, VPLS is disabled.

Many configuration procedures for VPLS are identical to the procedures for Layer 2 VPNs and Layer 3 VPNs.

Configuring VPLS Routing Instances

To configure a VPLS routing instance, include the **vpls** statement:

```
vpls {  
  active-interface {  
    any;  
    primary interface-name;  
  }  
  connectivity-type (ce | irb | permanent);  
  encapsulation-type encapsulation-type;  
  interface-mac-limit limit;  
  label-block-size size;  
  mac-table-aging-time time;  
  mac-table-size size;  
  neighbor neighbor-id;  
  no-tunnel-services;  
  site site-name {  
    active-interface {  
      any;  
      primary interface-name;  
    }  
    interface interface-name {  
      interface-mac-limit limit;  
    }  
    mesh-group mesh-group-name;  
    multi-homing;  
    site-identifier identifier;  
    site-preference preference-value;  
  }  
  site-range number;  
  traceoptions {  
    file filename <files number> <size size> <world-readable | no-world-readable>;  
    flag flag <flag-modifier> <disable>;  
  }  
  tunnel-services {  
    devices device-names;  
    primary primary-device-name;  
  }  
  vpls-id vpls-id;  
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]



NOTE: You cannot configure a routing protocol (OSPF, RIP, IS-IS or BGP) inside a VPLS routing instance (`instance-type vpls`). The Junos CLI disallows this configuration.

The configuration for the VPLS routing instance statements is explained in the following sections:

- [Configuring BGP Signaling for VPLS on page 5325](#)
- [Configuring LDP Signaling for VPLS on page 5330](#)
- [Configuring VPLS Routing Instance and VPLS Interface Connectivity on page 5332](#)
- [Configuring the VPLS Encapsulation Type on page 5332](#)
- [Configuring the VPLS MAC Table Timeout Interval on page 5333](#)
- [Configuring the Size of the VPLS MAC Address Table on page 5333](#)
- [Limiting the Number of MAC Addresses Learned from an Interface on page 5334](#)
- [Removing Addresses from the MAC Address Database on page 5335](#)

Configuring BGP Signaling for VPLS

You can configure BGP signaling for the VPLS routing instance. BGP is used to signal the pseudowires linking each of the PE routers participating in the VPLS routing instance. The pseudowires carry VPLS traffic across the service provider's network between the VPLS sites.



NOTE: You cannot configure both BGP signaling and LDP signaling for the same VPLS routing instance. If you attempt to configure the statements that enable BGP signaling for the VPLS routing instance (the `site`, `site-identifier`, and `site-range` statements) and the statements that enable LDP signaling for the same instance (the `neighbor` and `vpls-id` statements), the commit operation fails.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Configure BGP signaling for the VPLS routing instance by completing the steps in the following sections:

- [Configuring the VPLS Site Name and Site Identifier on page 5326](#)
- [Configuring Automatic Site Identifiers for VPLS on page 5326](#)
- [Configuring the Site Range on page 5327](#)
- [Configuring the VPLS Site Interfaces on page 5329](#)
- [Configuring the VPLS Site Preference on page 5329](#)

Configuring the VPLS Site Name and Site Identifier

When you configure BGP signaling for the VPLS routing instance, on each PE router you must configure each VPLS site that has a connection to the PE router. All the Layer 2 circuits provisioned for a VPLS site are listed as the set of logical interfaces (using the **interface** statement) within the **site** statement.

You must configure a site name and site identifier for each VPLS site.

To configure the site name and the site identifier, include the **site** and the **site-identifier** statements:

```
site site-name {  
  interface interface-name {  
    interface-mac-limit limit;  
  }  
  site-identifier identifier;  
}
```

The numerical identifier can be any number from 1 through 65,534 that uniquely identifies the local VPLS site.

You can include these statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols vpls]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls]

Configuring Automatic Site Identifiers for VPLS

When you enable automatic site identifiers, the Junos OS automatically assigns site identifiers to VPLS sites. To configure automatic site identifiers for a VPLS routing instance, include the **automatic-site-id** statement:

```
automatic-site-id {  
  collision-detect-time seconds;  
  new-site-wait-time seconds;  
  reclaim-wait-time minimum seconds maximum seconds;  
  startup-wait-time seconds;  
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols vpls site *site-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls site *site-name*]

The **automatic-site-id** statement includes a number of options that control different delays in network layer reachability information (NLRI) advertisements. All of these options are configured with default values. See the statement summary for the **automatic-site-id** statement for more information.

The **automatic-site-id** statement includes the following options:

- **collision-detect-time**—The time in seconds to wait after a claim advertisement is sent to the other routers in a VPLS instance before a PE router can begin using a site identifier. If the PE router receives a competing claim advertisement for the same site identifier during this time period, it initiates the collision resolution procedure for site identifiers.
- **new-site-wait-time**—The time in seconds to wait to receive VPLS information for a newly configured routing instance or a new site. This time interval is also applied whenever the automatic site identifier feature is activated on a VPLS routing instance other than at startup. Effectively, this timer indicates how long to wait before an attempt is made to allocate a site identifier. This timer is also triggered whenever a VPLS routing instance is enabled.
- **reclaim-wait-time**—The time to wait before attempting to claim a site identifier after a collision. A collision occurs whenever an attempt is made to claim a site identifier by two separate VPLS sites.
- **startup-wait-time**—The time in seconds to wait at startup to receive all the VPLS information for the route targets configured on the other PE routers included in the VPLS routing instance.

Configuring the Site Range

When you enable BGP signaling for each VPLS routing instance, you can optionally configure the site range. The site range specifies an upper limit on the maximum site identifier that can be accepted to allow a pseudowire to be brought up. You must specify a value from 1 through 65,534. The default value is **65,534**. We recommend using the default. Pseudowires cannot be established to sites with site identifiers greater than the configured site range. If you issue the **show vpls connections** command, such sites are displayed as OR (out of range).

To configure the site range, include the **site-range** statement:

```
site-range number;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols vpls]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls]

There are networks that require that the site range be configured using a value smaller than the local site identifier, for example, a hub-and-spoke VPLS with multihomed sites. For this type of network, you need to allow pseudowires to be established between the spoke routers and the hub router. However, you also need to prevent pseudowires from being established between spoke routers directly. Due to the multihoming requirement of spoke sites, Layer 2 VPN NRIs need to be accepted from other spoke routers (at least

from spokes with the same site identifier as the locally configured sites) to determine the status of local spoke routers (active or not active) based on the local preference included in the NRLIs received from the other spoke routers.

This type of VPLS network can be implemented by, for example, numbering hub sites with identifiers 1 through 8 and spoke sites with identifiers 9 and larger. You can then configure a site range of 8 on each of the spoke sites. Although the spoke sites accept NRLIs and install them in the Layer 2 VPN routing tables (allowing the multihomed sites to determine the status of the local site), the spoke sites cannot establish pseudowires directly to the other spoke sites due to the configured site range.

The following configurations illustrate this concept. The configurations are for the VPLS routing instances on three routers, two spoke routers and one hub router:

Router 1—spoke:

```
routing-instance hub-and-spoke {
  no-local-switching;
  protocols {
    vpls {
      site-range 8;
      no-tunnel-services;
      site spoke-9 {
        site-identifier 9 {
          multi-homing;
          site-preference primary;
        }
      }
      site spoke-10 {
        site-identifier 10 {
          multi-homing;
          site-preference backup;
        }
      }
    }
  }
}
```

Router 2—spoke:

```
routing-instance hub-and-spoke {
  no-local-switching;
  protocols {
    vpls {
      site-range 8;
      no-tunnel-services;
      site spoke-9 {
        site-identifier 9 {
          multi-homing;
          site-preference backup;
        }
      }
      site spoke-10 {
        site-identifier 10 {
          multi-homing;
        }
      }
    }
  }
}
```

```

        site-preference primary;
    }
}
}
}

```

Hub—router 3:

```

routing-instance hub-and-spoke {
    no-local-switching;
    protocols {
        vpls {
            no-tunnel-services;
            site hub {
                site-identifier 1;
            }
        }
    }
}

```

Configuring the VPLS Site Interfaces

You must configure an interface for each of the pseudowires you specify for the VPLS site.

To configure an interface for the VPLS site, include the **interface** statement:

```

interface interface-name {
    interface-mac-limit limit;
}

```

You can include these statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols vpls site *site-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls]

You can also configure a limit on the number of MAC addresses that can be learned from the specified interface. For more information, see [“Limiting the Number of MAC Addresses Learned from an Interface” on page 5334](#).

Configuring the VPLS Site Preference

You can specify the local preference value advertised for a particular VPLS site. The site preference value is specified using the **site-preference** statement configured at the [edit routing-instances *routing-instance-name* protocols vpls site *site-name*] hierarchy level. By configuring the **site-preference** statement, a value configured for the **local-preference** statement at the [edit protocols bgp] hierarchy level is ignored by the VPLS routing instance. However, you can change the site preference value for VPLS routes exported to other routers by configuring an export policy. When a PE router receives multiple advertisements with the same VPLS edge (VE) device identifier, the advertisement with the highest local preference value is preferred.

To configure the VPLS site preference, include the **site-preference** statement:

```
site-preference preference-value {  
    backup;  
    primary;  
}
```

You can also specify either the **backup** option or the **primary** option for the **site-preference** statement. The **backup** option specifies the preference value as 1, the lowest possible value, ensuring that the VPLS site is the least likely to be selected. The **primary** option specifies the preference value as 65,535, the highest possible value, ensuring that the VPLS site is the most likely to be selected.

For a list of hierarchy levels at which you can include the **site-preference** statement, see the statement summary section for this statement.

Configuring LDP Signaling for VPLS

You can configure LDP as the signaling protocol for a VPLS routing instance. This functionality is described in RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*.

The Junos OS does not support all of RFC 4762. When enabling LDP signaling for a VPLS routing instance, network engineers should be aware that only the following values are supported:

- FEC—128 or 129
- Control bit—0
- Ethernet pseudowire type—0x0005
- Ethernet tagged mode pseudowire type—0x0004

To enable LDP signaling for the set of PE routers participating in the same VPLS routing instance, you need to use the **vpls-id** statement configured at the **[edit routing-instances routing-instance-name protocols vpls]** hierarchy level to configure the same VPLS identifier on each of the PE routers. The VPLS identifier must be globally unique. When each VPLS routing instance (domain) has a unique VPLS identifier, it is possible to configure multiple VPLS routing instances between a given pair of PE routers.

LDP signaling requires that you configure a full-mesh LDP session between the PE routers in the same VPLS routing instance. Neighboring PE routers are statically configured. Tunnels are created between the neighboring PE routers to aggregate traffic from one PE router to another. Pseudowires are then signaled to demultiplex traffic between VPLS routing instances. These PE routers exchange the pseudowire label, the MPLS label that acts as the VPLS pseudowire demultiplexer field, by using LDP forwarding equivalence classes (FECs). Tunnels based on both MPLS and generic routing encapsulation (GRE) are supported.



NOTE: You cannot configure both BGP signaling and LDP signaling for the same VPLS routing instance. If you attempt to configure the statements that enable BGP signaling for the VPLS routing instance (the `site`, `site-identifier`, and `site-range` statements), and the statements that enable LDP signaling for the same instance, `neighbor` and `vpls-id`, the commit operation fails.

To enable LDP signaling for the VPLS routing instance, complete the steps in the following sections:

- [Configuring LDP Signaling for the VPLS Routing Instance on page 5331](#)
- [Configuring LDP Signaling on the Router on page 5331](#)

Configuring LDP Signaling for the VPLS Routing Instance

To configure the VPLS routing instance to use LDP signaling, you must configure the same VPLS identifier on each PE router participating in the instance. Specify the VPLS identifier with the `vpls-id` statement:

```
vpls-id vpls-id;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols vpls]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls]

To configure the VPLS routing instance to use LDP signaling, you also must include the `neighbor` statement to specify each of the neighboring PE routers that are a part of this VPLS domain:

```
neighbor neighbor-id;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols vpls]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls]

Configuring LDP Signaling on the Router

To enable LDP signaling, you need to configure LDP on each PE router participating in the VPLS routing instance. A minimal configuration is to enable LDP on the loopback interface, which includes the router identifier (`router-id`), on the PE router using the `interface` statement:

```
interface interface-name;
```

You can include this statement at the following hierarchy levels:

- [edit protocols ldp]
- [edit logical-systems *logical-system-name* protocols ldp]

You can enable LDP on all the interfaces on the router using the **all** option for the **interfaces** statement. For more information about how to configure LDP, see the *Junos OS MPLS Applications Configuration Guide*.

Configuring VPLS Routing Instance and VPLS Interface Connectivity

You can configure the VPLS routing instance to take down or maintain its VPLS connections depending on the status of the interfaces configured for the VPLS routing instance. By default, the VPLS connection is taken down whenever a customer-facing interface configured for the VPLS routing instance fails. This behavior can be explicitly configured by specifying the **ce** option for the **connectivity-type** statement:

```
connectivity-type ce;
```

You can alternatively specify that the VPLS connection remain up so long as an Integrated Routing and Bridging (IRB) interface is configured for the VPLS routing instance by specifying the **irb** option for the **connectivity-type** statement:

```
connectivity-type irb;
```

To ensure that the VPLS connection remain up until explicitly taken down, specify the **permanent** option for the **connectivity-type** statement:

```
connectivity-type permanent;
```

This option is reserved for use in configuring Layer 2 Wholesale subscriber networks. See the *Broadband Subscriber Management Solutions Guide* for details about configuring a Layer 2 Wholesale network.

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols vpls]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls]

Configuring the VPLS Encapsulation Type

You can specify a VPLS encapsulation type for the pseudowires established between VPLS neighbors. The encapsulation type is carried in the LDP-signaling messages exchanged between VPLS neighbors when pseudowires are created. You might need to alter the encapsulation type depending on what other vendors' equipment is deployed within your network.

VPLS effectively provides a bridge between Ethernet networks. As a consequence, only two encapsulation types are available:

- **ethernet**—Ethernet
- **ethernet-vlan**—Ethernet virtual LAN (VLAN)

If you do not specify an encapsulation type for the VPLS routing instance or the VPLS neighbor, **ethernet** is used.

To specify an encapsulation type for the VPLS routing instance, include the **encapsulation-type** statement:

encapsulation-type (ethernet | ethernet-vlan);

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols vpls]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls]

You can also specify an encapsulation type for a specific VPLS neighbor by including the **encapsulation-type** statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols vpls neighbor *address*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls neighbor *address*]

Configuring the VPLS MAC Table Timeout Interval

You can modify the timeout interval for the VPLS table. We recommend you that configure longer values for small, stable VPLS networks and shorter values for large, dynamic VPLS networks. If the VPLS table does not receive any updates during the timeout interval, the router waits one additional interval before automatically clearing the MAC address entries from the VPLS table.

To modify the timeout interval for the VPLS table, include the **mac-table-aging-time** statement:

mac-table-aging-time *seconds*;

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols vpls]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls]



NOTE: The **mac-table-aging-time** statement is not available on MX Series routers.

Configuring the Size of the VPLS MAC Address Table

You can modify the size of the VPLS media access control (MAC) address table. The default table size is 512 MAC addresses, the minimum is 16 addresses, and the maximum is 65,536 addresses.



NOTE: T4000 routers with Type 5 FPCs support up to 262,143 MAC addresses per VPLS routing instance. To enable the improved VPLS MAC address learning limit (that is, 262,143 MAC addresses), you must include the **enhanced-mode** statement at the [edit chassis network-services] hierarchy level, reboot the router, and then modify the size of the VPLS MAC address table.

If the MAC table limit is reached, new MAC addresses can no longer be added to the table. Eventually the oldest MAC addresses are removed from the MAC address table automatically. This frees space in the table, allowing new entries to be added. However, as long as the table is full, new MAC addresses are dropped.

To change the VPLS MAC table size for each VPLS or VPN routing instance, include the **mac-table-size** statement:

```
mac-table-size size;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols vpls]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls]

When you include the **mac-table-size** statement, the affected interfaces include all interfaces within the VPLS routing instance, including the local interfaces, the LSI interfaces, and the VT interfaces.

Limiting the Number of MAC Addresses Learned from an Interface

You can configure a limit on the number of MAC addresses learned by a VPLS routing instance using the **mac-table-size** statement. If the MAC table limit is reached, new MAC addresses can no longer be added to the table. Eventually the oldest MAC addresses are removed from the MAC address table automatically. This frees space in the table, allowing new entries to be added. However, as long as the table is full, new MAC addresses are dropped.

Because this limit applies to each VPLS routing instance, the MAC addresses of a single interface can consume all the available space in the table, preventing the routing instance from acquiring addresses from other interfaces.

You can limit the number of MAC addresses learned from each interface configured for a VPLS routing instance. To do so, include the **interface-mac-limit** statement:

```
interface-mac-limit limit;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols vpls]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls]

The **interface-mac-limit** statement affects the local interfaces only (the interfaces facing CE devices).

Configuring the **interface-mac-limit** statement at the [edit routing-instances *routing-instance-name* protocols vpls] hierarchy level causes the same limit to be applied to all of the interfaces configured for that specific routing instance.

You can also limit the number of MAC addresses learned by a specific interface configured for a VPLS routing instance. This gives you the ability to limit particular interfaces that you expect might generate a lot of MAC addresses.

To limit the number of MAC addresses learned by a specific interface, include the **interface-mac-limit** statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols vpls site *site-name* interfaces *interface-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls site *site-name* interfaces *interface-name*]

The MAC limit configured for an individual interface at this hierarchy level overrides any value configured at the [edit routing-instances *routing-instance-name* protocols vpls] hierarchy level. Also, the MAC limit configured using the **mac-table-size** statement can override the limit configured using the **interface-mac-limit** statement.

The MAC address limit applies to customer-facing interfaces only.

Removing Addresses from the MAC Address Database

You can enable MAC flush processing for the VPLS routing instance or for the mesh group under a VPLS routing instance. MAC flush processing removes MAC addresses from the MAC address database that have been learned dynamically. With the dynamically learned MAC addresses removed, MAC address convergence requires less time to complete.

You can clear dynamically learned MAC addresses from the MAC address database by including the **mac-flush** statement:

mac-flush [*explicit-mac-flush-message-options*];

To clear dynamically learned MAC addresses globally across all devices participating in the routing instance, you can include the statement at the following hierarchy levels:

- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls]
- [edit routing-instances *routing-instance-name* protocols vpls]

To clear the MAC addresses on the routers in a specific mesh group, you can include the statement at the following hierarchy levels:

- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls mesh-group *mesh-group-name*]
- [edit routing-instances *routing-instance-name* protocols vpls mesh-group *mesh-group-name*]

For certain cases where MAC flush processing is not initiated by default, you can also specify **explicit-mac-flush-message-options** to additionally configure the router to send explicit MAC flush messages under specific conditions. For a list of the explicit MAC flush message options you can include with this statement, see the summary section for this statement.

Related Documentation

- [Configuring Improved VPLS MAC Address Learning on T4000 Routers with Type 5 FPCs](#)
- [enhanced-mode](#)

Configuring Interfaces for VPLS Routing

On each PE router and for each VPLS routing instance, specify which interfaces are intended for the VPLS traffic traveling between PE and CE routers. To specify the interface for VPLS traffic, include the **interface** statement in the routing instance configuration:

```
interface interface-name;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

You must also define each interface by including the following statements:

```
vlan-tagging;  
encapsulation encapsulation-type;  
unit logical-unit-number {  
    family vpls;  
    vlan-id vlan-id-number;  
}
```

You can include these statements at the following hierarchy levels:

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



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

The following sections provide enough information to enable you to configure interfaces for VPLS routing. For detailed information about configuring interfaces and the statements at the [edit interfaces] hierarchy level, see the *Junos® OS Network Interfaces*.

- [Configuring the VPLS Interface Name on page 5336](#)
- [Configuring VPLS Interface Encapsulation on page 5337](#)
- [Enabling VLAN Tagging on page 5339](#)
- [Configuring VLAN IDs for Logical Interfaces on page 5340](#)
- [Enabling VLANs for Hub and Spoke VPLS Networks on page 5340](#)
- [Configuring Aggregated Ethernet Interfaces for VPLS on page 5341](#)

Configuring the VPLS Interface Name

Specify both the physical and logical portions of the interface name, in the following format:

physical.logical

For example, in **ge-1/2/1.2**, **ge-1/2/1** is the physical portion of the interface name and **2** is the logical portion. If you do not specify the logical portion of the interface name, **0** is set by default.

A logical interface can be associated with only one routing instance.

If you enable a routing protocol on all instances by specifying **interfaces all** when configuring the master instance of the protocol at the **[edit protocols]** hierarchy level, and you configure a specific interface for VPLS routing at the **[edit routing-instances routing-instance-name]** hierarchy level, the latter interface statement takes precedence and the interface is used exclusively for VPLS.

If you explicitly configure the same interface name at both the **[edit protocols]** and **[edit routing-instances routing-instance-name]** hierarchy levels and then attempt to commit the configuration, the commit operation fails.

Configuring VPLS Interface Encapsulation

You need to specify an encapsulation type for each PE-router-to-CE-router interface configured for VPLS. This section describes the **encapsulation** statement configuration options available for VPLS. For a complete description of all of the options available for this statement, see the *Junos® OS Network Interfaces*.

To configure the encapsulation type on the physical interface, include the **encapsulation** statement:

encapsulation (ethernet-vpls | ether-vpls-over-atm-llc | extended-vlan-vpls | vlan-vpls);

You can include the **encapsulation** statement for physical interfaces at the following hierarchy levels:

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

You can configure the following physical interface encapsulations for VPLS routing instances:

- **ethernet-vpls**—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled. The PE router expects to receive Ethernet frames with VLAN tags that are not service-delimiting. The Ethernet frames are not meaningful to the PE router and cannot be used by the service provider to separate customer traffic.

On M Series routers (except the M320 router), the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

- **ether-vpls-over-atm-llc**—For ATM intelligent queuing (IQ) interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from

the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

- **extended-vlan-vpls**—Use extended virtual LAN (VLAN) VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901. On M Series routers (except the M320 router), the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.



NOTE: The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

- **vlan-vpls**—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN 802.1Q tagging and VPLS enabled. The PE router expects to receive Ethernet frames with VLAN tags that are service-delimiting. These VLAN tags can be used by the service provider to separate customer traffic. For example, LAN traffic from different customers can flow through the same service provider switch, which can then apply VLAN tags to distinguish one customer's traffic from the others. The traffic can then be forwarded to the PE router.

Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only. On M Series routers (except the M320 router), the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

To configure the encapsulation type for logical interfaces, include the **encapsulation** statement:

```
encapsulation (ether-vpls-over-atm-llc | vlan-vpls);
```

You can include the **encapsulation** statement for logical interfaces at the following hierarchy levels:

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

You can configure the following logical interface encapsulations for VPLS routing instances:

- **ether-vpls-over-atm-llc**—Use Ethernet VPLS over Asynchronous Transfer Mode (ATM) logical link control (LLC) encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3-encapsulated Ethernet frames with the frame check sequence (FCS) field removed. This encapsulation type is supported on ATM intelligent queuing (IQ) interfaces only.
- **vlan-vpls**—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN 802.1Q tagging and VPLS enabled. The PE router expects to receive Ethernet frames with VLAN tags that are service-delimiting. These VLAN tags can be used by the service provider to separate customer traffic. For example, LAN traffic from different customers

can flow through the same service provider switch, which can then apply VLAN tags to distinguish one customer's traffic from the others. The traffic can then be forwarded to the PE router.

Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only. On M Series routers (except the M320 router), the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.



NOTE: Label-switched interfaces (LSIs) do not support VLAN VPLS encapsulation. Therefore, you can only use VLAN VPLS encapsulation on a PE-router-to-CE-router interface and not a core-facing interface.

When you configure the physical interface encapsulation as **vlan-vpls**, you also need to configure the same interface encapsulation for the logical interface. You need to configure the **vlan-vpls** encapsulation on the logical interface because the **vlan-vpls** encapsulation allows you to configure a mixed mode, where some of the logical interfaces use regular Ethernet encapsulation (the default for logical interfaces) and some use **vlan-vpls**. For more information, see the *Junos® OS Network Interfaces*.

Enabling VLAN Tagging

Junos OS supports receiving and forwarding routed Ethernet frames with 802.1Q virtual local area network (VLAN) tags and running the Virtual Router Redundancy Protocol (VRRP) over 802.1Q-tagged interfaces. For VPLS to function properly, configure the router to receive and forward frames with 802.1Q VLAN tags by including the **vlan-tagging** statement at the **[edit interfaces interface-name]** hierarchy level:

```
[edit interfaces interface-name]
vlan-tagging;
```

Gigabit Ethernet interfaces can be partitioned. You can assign up to 4095 different logical interfaces, one for each VLAN, but you are limited to a maximum of 1024 VLANs on any single Gigabit Ethernet or 10-Gigabit Ethernet port. Fast Ethernet interfaces can also be partitioned, with a maximum of:

- 1024 logical interfaces for the 4-port FE PIC
- 1024 logical interfaces for the 2-port Fixed Interface Card (FIC) on an M7i router
- 16 logical interfaces for the M40e router

Table 386 on page 5339 lists VLAN ID ranges by interface type.

Table 386: VLAN ID Range by Interface Type

Interface Type	VLAN ID Range
Fast Ethernet	512 through 1023
Gigabit Ethernet	512 through 4094

Configuring VLAN IDs for Logical Interfaces

You can bind a VLAN identifier to a logical interface by including the **vlan-id** statement:

```
vlan-id number;
```

You can include this statement at the following hierarchy levels:

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

You can also configure a logical interface to forward packets and learn MAC addresses within each VPLS routing instance configured with a VLAN ID that matches a VLAN ID specified in a list using the **vlan-id-list** statement. VLAN IDs can be entered individually using a space to separate each ID, entered as an inclusive list separating the starting VLAN ID and ending VLAN ID with a hyphen, or a combination of both.

For example, to configure the VLAN IDs 20 and 45 and the range of VLAN IDs between 30 and 40, issue the following command from the CLI:

```
set interfaces ge-1/0/1 unit 1 vlan-id-list [20 30-40 45];
```

To configure a list of VLAN IDs for a logical interface, include the **vlan-id-list** statement:

```
vlan-id-list list-of-vlan-ids;
```

You can include this statement at the following hierarchy levels:

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

For more information about how to configure VLANs, see the *Junos® OS Network Interfaces*. For detailed information about how VLAN identifiers in a VPLS routing instance are processed and translated, see the *MX Series Layer 2 Configuration Guide*.

Enabling VLANs for Hub and Spoke VPLS Networks

For hub and spoke VPLS networks, you need to configure the **swap** option for the **output-vlan-map** statement on the hub facing interface of each spoke PE router. The **output-vlan-map** statement ensures that the vlan ID of the spoke PE router matches the VLAN ID of the hub PE router in the VPLS network. The following configuration example illustrates an interface configuration with the output-vlan-map statement included:

```
[edit interfaces xe-4/0/0]
vlan-tagging;
encapsulation flexible-ethernet-services;
unit 610 {
  encapsulation vlan-ccc;
  vlan-id 610;
  output-vlan-map swap;
}
```

Configuring Aggregated Ethernet Interfaces for VPLS

You can configure aggregated Ethernet interfaces between CE devices and PE routers for VPLS routing instances. Traffic is load-balanced across all of the links in the aggregated interface. If one or more links in the aggregated interface fails, the traffic is switched to the remaining links.

For more information about how aggregated Ethernet interfaces function in the context of VPLS, see [“VPLS and Aggregated Ethernet Interfaces” on page 5302](#).

To configure aggregated Ethernet interfaces for VPLS, configure the interface for the VPLS routing instance as follows:

```
interfaces aex {
  vlan-tagging;
  encapsulation encapsulation-type;
  unit logical-unit-number {
    vlan-id number;
  }
}
```

You can configure the following physical link-layer encapsulation types for the VPLS aggregated Ethernet interface:

- **ethernet-vpls**
- **extended-vlan-vpls**
- **flexible-ethernet-services**
- **vlan-vpls**

For the **interface** configuration statement, in **aex**, the **x** represents the interface instance number to complete the link association; **x** can be from 0 through 127, for a total of 128 aggregated interfaces.

For more information about how to configure aggregated Ethernet interfaces, see the *Junos® OS Network Interfaces*.

The aggregated Ethernet interface must also be configured for the VPLS routing instance as shown in the following example:

```
[edit]
routing-instances {
  green {
    instance-type vpls;
    interface ae0.0;
    route-distinguisher 10.255.234.34:1;
    vrf-target target:1111:1;
    protocols {
      vpls {
        site-range 10;
        site green3 {
          site-identifier 3;
        }
      }
    }
  }
}
```

```

    }
  }
}

```

Interface **ae0.0** represents the aggregated Ethernet interface in the routing instance configuration. The VPLS routing instance configuration is otherwise standard.

Configuring Static Pseudowires for VPLS

You can configure a VPLS domain using static pseudowires. A VPLS domain consists of a set of PE routers that act as a single virtual Ethernet bridge for the customer sites connected to these routers. By configuring static pseudowires for the VPLS domain, you do not need to configure the LDP or BGP protocols that would normally be used for signaling. However, if you configure static pseudowires, any changes to the VPLS network topology have to be managed manually.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Static pseudowires require that you configure a set of in and out labels for each pseudowire configured for the VPLS domain. You still need to configure a VPLS identifier and neighbor identifiers for a static VPLS domain. You can configure both static and dynamic neighbors within the same VPLS routing instance.

To configure a static pseudowire for a VPLS neighbor, include the **static** statement:

```

static (Protocols VPLS) {
    incoming-label label;
    outgoing-label label;
}

```

You must configure an incoming and outgoing label for the static pseudowire using the **incoming-label** and **outgoing-label** statements. These statements identify the static pseudowire's incoming traffic and destination.

To configure a static pseudowire for a VPLS neighbor, include the **static** statement at the **[edit routing-instances routing-instance-name protocols vpls neighbor address]** hierarchy level.

You can also configure the **static** statement for a backup neighbor (if you configure the neighbor as static the backup must also be static) by including it at the **[edit routing-instances routing-instance-name protocols vpls neighbor address backup-neighbor address]** hierarchy level and for a mesh group by including it at the **[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name neighbor address]** hierarchy level.

For a list of hierarchy levels at which you can include the **static** statement, see the statement summary section for this statement.

To enable static VPLS on a router, you need to either configure a virtual tunnel interface (requires the router to have a tunnel services PIC) or you can configure a label switching interface (LSI). To configure an LSI, include the **no-tunnel-services** statement at the **[edit**

`protocols vpls static-vpls]` hierarchy level. For more information, see [“Configuring VPLS Without a Tunnel Services PIC” on page 5351](#).



NOTE: Static pseudowires for VPLS using an LSI is supported on MX series routers and EX Series switches only. For M series and T series routers, a tunnel services PIC is required.

If you issue a **show vpls connections** command, static neighbors are displayed with **"SN"** next to their addresses in the command output.

Related Documentation

- [Configuring VPLS Without a Tunnel Services PIC on page 5351](#)

Configuring VPLS Multihoming (FEC 128)

VPLS multihoming allows you to connect a customer site to multiple PE routers to provide redundant connectivity while preventing the formation of Layer 2 loops in the service provider's network. A VPLS site multihomed to two or more PE routers provides redundant connectivity in the event of a PE router-to-CE device link failure or the failure of a PE router. For more information about VPLS multihoming, see [“VPLS Multihoming Overview” on page 5305](#).



NOTE: If you want to enable multihoming for a VPLS routing instance, you cannot also enable LDP signaling. You can only enable BGP signaling.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

The following sections describe how to configure VPLS multihoming. Some information is also provided on single-homed site configuration versus multihomed site configuration.

- [VPLS Multihomed Site Configuration on page 5343](#)
- [VPLS Single-Homed Site Configuration on page 5345](#)

VPLS Multihomed Site Configuration

The following describes the requirements for a VPLS multihomed site configuration:

- Assign the same site ID on all PE routers connected to the same CE devices.
- Assign the same route distinguisher on all PE routers connected to the same CE devices.
- Reference all interfaces assigned to the multihomed VPLS site on each PE router. Only one of these interfaces is used to send and receive traffic for this site at a time.
- Either designate a primary interface or allow the router to select the interface to be used as the primary interface.

If the router selects the interface, the interface used to connect the PE router to the site depends on the order in which interfaces are listed in the PE router's configuration. The first operational interface in the set of configured interfaces is chosen to be the designated interface. If this interface fails, the next interface in the list is selected to send and receive traffic for the site.

- Configure multihoming for the site.

The following configuration shows the statements you need to configure to enable VPLS multihoming:

```
[edit routing-instances routing-instance-name]  
instance-type vpls;  
interface interface-name;  
interface interface-name;  
protocols vpls {  
  site site-name {  
    active-interface {  
      any;  
      primary interface-name;  
    }  
    interface interface-name;  
    interface interface-name;  
    multi-homing;  
    site-identifier number;  
  }  
}  
route-distinguisher (as-number:id | ip-address:id);
```



NOTE: If you add a direct connection between CE devices that are multihomed to the same VPLS site on different PE routers, the traffic can loop and a loss of connectivity might occur. We do not recommend this topology.

Most of these statements are explained in more detail in the rest of this chapter. The following sections explain how to configure the statements that are specific to VPLS multihoming:

- [Specifying an Interface as the Active Interface on page 5344](#)
- [Configuring Multihoming on the PE Router on page 5345](#)

Specifying an Interface as the Active Interface

You need to specify one of the interfaces for the multihomed site as the primary interface. If there are multiple interfaces, the remaining interfaces are activated only when the primary interface goes down. If no active interfaces are configured at the site level, all traffic for a VPLS site travels through a single, non-multihomed PE router.

You must configure one of the following options for the **active-interface** statement:

- **any**—One configured interface is randomly designated as the active interface for the VPLS site.

- **primary**—Specify the name of the multihomed interface to be used as the primary interface by the VPLS site.

To specify a multihomed interface as the primary interface for the VPLS site, include the **active-interface** statement:

```
active-interface {
    any;
    primary interface-name;
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols vpls site *site-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls site *site-name*]

Configuring Multihoming on the PE Router

When a CE device is connected to the same VPLS site on more than one PE router, include the **multi-homing** statement on all associated PE routers. Configuration of this statement tracks BGP peers. If no BGP peer is available, VPLS deactivates all active interfaces for a site. To specify that the PE router is part of a multihomed VPLS site, include the **multi-homing** statement:

```
multi-homing;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols vpls site *site-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls site *site-name*]

Include the **multi-homing** statement on all PE routers associated with a particular VPLS site.

VPLS Single-Homed Site Configuration

All VPLS single-homed sites are connected to the same default VE device. All interfaces in a VPLS routing instance that are not configured as part of a multihomed site are assumed to be single-homed to the default VE device.

Enabling BGP Path Selection for Layer 2 VPNs and VPLS

Layer 2 VPNs and VPLS share the same path selection process for determining the optimal path to reach all of the destinations shared within a single routing instance. For Layer 2 VPN and VPLS topologies, the path selection process is straightforward if there is just a single path from each PE router to each CE device. However, the path selection process becomes more complex if the PE routers receive two or more valid paths to reach a specific CE device.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

The following network scenarios provide examples of what might cause a PE router to receive more than one valid path to reach a specific CE device:

- Multihoming—One or more CE devices within a routing instance are multihomed to two or more PE routers. Each multihomed CE device has at least two valid paths.
- Route reflectors—There are multiple route reflectors deployed within the same network and they are supporting PE routers within the same routing instance. Due to time delays in large complex networks, the route reflectors can separately receive a different valid path to reach a CE device at different times. When they readvertise these valid paths, a PE router could receive two or more separate but apparently valid paths to the same CE device.

By default, Juniper Networks routers use just the designated forwarder path selection algorithm to select the best path to reach each Layer 2 VPN or VPLS routing instance destination (for more information, see [“VPLS Path Selection Process for PE Routers” on page 5309](#)). However, you can also configure the routers in your network to use both the BGP path selection algorithm and the designated forwarder path selection algorithm as follows:

- On the Provider routers within the service providers network, the standard BGP path selection algorithm is used (for more information, see *Understanding BGP Path Selection*). Using the standard BGP path selection for Layer 2 VPN and VPLS routes allows a service provider to leverage the existing Layer 3 VPN network infrastructure to also support Layer 2 VPNs and VPLS. The BGP path selection algorithm also helps to ensure that the service provider’s network behaves predictably with regard to Layer 2 VPN and VPLS path selection. This is particularly important in networks employing route reflectors and multihoming.

When a Provider router receives multiple paths for the same destination prefix (for example, a multihomed CE device), one path is selected based on the BGP path selection algorithm and placed in the `bgp.l2vpn.0` routing table and the appropriate `instance.l2vpn.0` routing table.

- When a PE router receives all of the available paths to each CE device, it runs the designated forwarder path selection algorithm to select the preferred path to reach each CE device, independently of the results of the earlier BGP path selection algorithm run on the Provider router. The VPLS designated forwarder algorithm uses the D-bit, preference, and PE router identifier to determine which of the valid paths to each CE device to use. The PE router might select a path to reach a CE device which is different from the path selected by the BGP-based Provider routers. In this scenario, the following is the expected behavior for traffic sent to the multihomed CE device:
 - If the path selected by the remote PE router is available, traffic will traverse the network to the multihomed CE device using the remote PE router's preferred path (again, ignoring the path selected by the BGP-based Provider routers).
 - If the path selected by the remote PE router fails:
 1. The Provider routers switch the traffic destined for the multihomed CE device to the alternate path as soon as failure is detected.
 2. The Provider routers notify the remote PE routers of the path failure.
 3. The remote PE routers update their routing tables accordingly.

For more information about the VPLS designated forwarder path selection algorithm, see [“VPLS Path Selection Process for PE Routers” on page 5309](#). This algorithm is also described in the Internet draft [draft-kompella-l2vpn-vpls-multihoming-03.txt](#), *Multi-homing in BGP-based Virtual Private LAN Service*.

To enable the BGP path selection algorithm for Layer 2 VPN and VPLS routing instances, complete the following steps:

1. Run Junos OS Release 12.3 or later on all of the PE and Provider routers participating in Layer 2 VPN or VPLS routing instances.

Attempting to enable this functionality on a network with a mix of routers that both do and do not support this feature can result in anomalous behavior.

2. Specify a unique route distinguisher on each PE router participating in a Layer 2 VPN or VPLS routing instance.
3. Configure the **l2vpn-use-bgp-rules** statement on all of the PE and Provider routers participating in Layer 2 VPN or VPLS routing instances.

You can configure this statement at the **[edit protocols bgp path-selection]** hierarchy level to apply this behavior to all of the routing instances on the router or at the **[edit routing-instances routing-instance-name protocols bgp path-selection]** hierarchy level to apply this behavior to a specific routing instance.

Related Documentation

- [Understanding BGP Path Selection](#)
- [VPLS Path Selection Process for PE Routers on page 5309](#)
- [l2vpn-use-bgp-rules](#)
- [route-distinguisher](#)

Configuring EXP-Based Traffic Classification for VPLS

You can enable EXP classification on traffic entering core facing VPLS LSI interfaces on a VPLS routing instance by configuring either a logical tunnel interface (**lt-**) or the **no-tunnel-services** statement. By configuring either of these, a default EXP classifier is enabled on every core facing interface that includes **family mpls** in its configuration. This feature works on MX Series routers and EX Series switches only. You can configure an EXP classifier explicitly at the **[edit class-of-service]** hierarchy level. For more information about EXP classifiers, see the *Junos OS Class of Service Configuration Guide*.

To enable EXP classification on traffic entering core facing VPLS LSI interfaces on a VPLS routing instance, include the **no-tunnel-services** statement:

```
no-tunnel-services;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-instances *routing-instance-name* protocols vpls]**
- **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls]**

Configuring VPLS Load Balancing

By default, when there are multiple equal-cost paths to the same destination for the active route, the Junos OS uses a hash algorithm to select one of the next-hop addresses to install in the forwarding table. Whenever the set of next hops for a destination changes, the next-hop address is reselected using the hash algorithm.

You can configure the Junos OS so that, for the active route, all next-hop addresses for a destination are installed in the forwarding table. This feature is called per-packet load balancing. You can use load balancing to spread traffic across multiple paths between routers. You can also configure per-packet load balancing to optimize VPLS traffic flows across multiple paths.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

You can load-balance VPLS traffic based on Layer 2 media access control (MAC) information, IP information and MPLS labels, or MPLS labels only.



NOTE: For platform support information, see **family multiservice**.

To optimize VPLS traffic flows across multiple paths, include the **family multiservice** statement at the **[edit forwarding-options hash-key]** hierarchy level:

```
family multiservice {  
    destination-mac;  
    label-1;  
    label-2;  
}
```

```

payload {
  ip {
    layer-3 {
      (destination-ip-only | source-ip-only);
    }
    layer-3-only;
    layer-4;
  }
}
source-mac;
symetric-hash {
  complement;
}
}

```

You can configure one or more of the following options to load-balance using the specified packet information:

- **destination-mac**—Include the destination-address MAC information in the hash key for Layer 2 load balancing.
- **source-mac**—Include the source-address MAC information in the hash key.
- **label-1**—Include the first MPLS label in the hash key. Used for including a one-label packet for per-flow load balancing of IPv4 VPLS traffic based on IP information and MPLS labels.
- **label-2**—Include the second MPLS label in the hash key. If both **label-1** and **label-2** are specified, the entire first label and the first 16 bits of the second label are hashed.
- **payload**—Include the packet's IP payload in the hash key.
- **ip**—Include the IP address of the IPv4 or IPv6 payload in the hash key.
- **layer-3-only**—Include only the Layer 3 information from the packet's IP payload in the hash key.
- **layer-3**—Include Layer 3 information from the packet's IP payload in the hash key.
- **destination-address-only**—Include only the destination IP address in the payload in the hash key.



NOTE: You can include either the **source-address-only** or the **destination-address-only** statement, not both. They are mutually exclusive.

- **source-address-only**—Include only the source IP address in the payload in the hash key.



NOTE: You can include either the **source-address-only** or the **destination-address-only** statement, not both. They are mutually exclusive.

- **layer-4**—Include Layer 4 information from the packet's IP payload in the hash key.

- **symmetric-hash**—Configure the symmetric hash or symmetric hash complement for configuring symmetrical load balancing on an 802.3ad Link Aggregation Group.
- **complement**—Include the complement of the symmetric hash in the hash key.

For more information about how to configure per-packet load balancing, see the *Routing Policy Configuration Guide*.

**Related
Documentation**

- *Configuring VPLS Load Balancing Based on IP and MPLS Information*
- *Configuring VPLS Load Balancing on MX Series 3D Universal Edge Routers*

Configuring VPLS Fast Reroute Priority

When a path is rerouted after a link failure by using the MPLS fast reroute feature, the affected next hops are repaired by switching them from the active label switched path (LSP) to the standby LSP. To specify the order in which the next hops are repaired and traffic convergence is restored for VPLS routing instances after a fast reroute event, you can use the **fast-reroute-priority** statement to configure **high**, **medium**, or **low** fast reroute priority for a VPLS routing instance. By default, the fast reroute priority for a VPLS routing instance is **low**.

Next hops are repaired and known unicast, unknown unicast, broadcast, and multicast traffic for VPLS routing instances are restored in the following order, based on the fast reroute priority configuration:

1. Next hops for high-priority VPLS routing instances are repaired.
2. Next hops for medium-priority VPLS routing instances are repaired.
3. Next hops for low-priority VPLS routing instances are repaired.

Because next hops for VPLS routing instances configured with **high** fast reroute priority are repaired first, the traffic traversing high-priority VPLS instances is restored faster than the traffic for VPLS instances configured with **medium** or **low** fast reroute priority. The ability to prioritize specific VPLS routing instances for faster convergence and traffic restoration enables service providers to offer differentiated service levels to their customers.

Within a particular fast reroute priority level (**high**, **medium**, or **low**), no particular order for traffic restoration of VPLS routing instances is followed.



NOTE: VPLS fast reroute priority is not supported on J Series routers.

To configure **high**, **medium**, or **low** fast reroute priority for a VPLS routing instance, include the **fast-reroute-priority** statement:

fast-reroute-priority (high | medium | low);

You can include this statement at the following hierarchy levels:

- **[edit forwarding-options]**

- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options]
- [edit routing-instances *routing-instance-name* forwarding-options]

You can configure fast reroute priority only for routing instances with the **instance-type** set to **vpls**. If you attempt to configure fast reroute priority for a routing instance with an **instance-type** other than **vpls**, the router displays a warning message and the configuration fails.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

The following example snippet shows configuration of **high** fast reroute priority for a VPLS routing instance named **test-vpls**:

```
test-vpls {
  instance-type vpls;
  forwarding-options {
    fast-reroute-priority high;
  }
}
```

To display the fast reroute priority setting configured for a VPLS routing instance, use the **show route instance detail** operational command. For information about using this command, see the *Junos OS Operational Mode Commands*.

Configuring VPLS Without a Tunnel Services PIC

VPLS normally uses a dynamic virtual tunnel logical interface on a Tunnel Services PIC to model traffic from a remote site (a site on a remote PE router that is in a VPLS domain). All traffic coming from a remote site is treated as coming in over the virtual port representing this remote site, for the purposes of Ethernet flooding, forwarding, and learning. An MPLS lookup based on the inner VPN label is done on a PE router. The label is stripped and the Layer 2 Ethernet frame contained within is forwarded to a Tunnel Services PIC. The PIC loops back the packet and then a lookup based on Ethernet MAC addresses is completed. This approach requires that the router have a Tunnel Services PIC and that the PE router complete two protocol lookups.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

You can configure VPLS without a Tunnel Services PIC by configuring the **no-tunnel-services** statement. This statement creates a label-switched interface (LSI) to provide VPLS functionality. An LSI MPLS label is used as the inner label for VPLS. This label maps to a VPLS routing instance. On the PE router, the LSI label is stripped and then mapped to a logical LSI interface. The Layer 2 Ethernet frame is then forwarded using the LSI interface to the correct VPLS routing instance.

By default, VPLS requires a Tunnel Services PIC. To configure VPLS on a router without a Tunnel Services PIC and create an LSI, include the **no-tunnel-services** statement:

no-tunnel-services;

For a list of the hierarchy levels at which you can include this statement, see the summary section for this statement.

To configure a VPLS routing instance on a router without a tunnel services PIC, include the **no-tunnel-services** statement at the **[edit routing-instances routing-instance-name protocols vpls]** hierarchy level. To configure static VPLS on a router without a tunnel services PIC, include the **no-tunnel-services** statement at the **[edit protocols vpls static-vpls]** hierarchy level.

When you configure VPLS without a Tunnel Services PIC by including the **no-tunnel-services** statement, the following limitations apply:

- An Enhanced FPC is required.
- ATM1 interfaces are not supported.
- Aggregated SONET/SDH interfaces are not supported as core-facing interfaces.
- Channelized interfaces are not supported as core-facing interfaces.
- GRE-encapsulated interfaces are not supported as core-facing interfaces.

**Related
Documentation**

- [Configuring Static Pseudowires for VPLS on page 5342](#)

Mapping VPLS Traffic to Specific LSPs

You can map VPLS traffic to specific LSPs by configuring forwarding table policies. This procedure is optional but can be useful. The following example illustrates how you can map lower priority VPLS routing instances to slower LSPs while mapping other higher priority VPLS routing instances to faster LSPs. In this example configuration, **a-to-b1** and **a-to-c1** are high-priority LSPs between the PE routers, while **a-to-b2** and **a-to-c2** are low-priority LSPs between the PE routers.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

To map VPLS traffic, include the **policy-statement vpls-priority** statement:

```
policy-statement vpls-priority {
  term a {
    from {
      rib mpls.0;
      community company-1;
    }
    then {
      install-nexthop lsp [ a-to-b1 a-to-c1 ];
      accept;
    }
  }
}
```

```

}
term b {
  from {
    rib mpls.0;
    community company-2;
  }
  then {
    install-nexthop lsp-regex [ "^a-to-b2$" "^a-to-c2$" ];
    accept;
  }
}
}
community company-1 members target:11111:1;
community company-2 members target:11111:2;

```

You can include the **policy-statement vpls-priority** statement at the following hierarchy levels:

- [edit policy-options]
- [edit logical-systems *logical-system-name* policy-options]

Include the **export** statement to apply the **vpls-priority** policy to the forwarding table:

```
export vpls-priority;
```

You can include this statement at the following hierarchy levels:

- [edit routing-options forwarding-table]
- [edit logical-systems *logical-system-name* routing-options forwarding-table]

For more information about how to configure routing policies, see the *Routing Policy Configuration Guide*.

Configuring Firewall Filters and Policers for VPLS

You can configure both firewall filters and policers for VPLS. Firewall filters allow you to filter packets based on their components and to perform an action on packets that match the filter. Policers allow you to limit the amount of traffic that passes into or out of an interface.

VPLS filters and policers act on a Layer 2 frame that includes the media access control (MAC) header (after any VLAN rewrite or other rules are applied), but does not include the cyclical redundancy check (CRC) field.

You can apply VPLS filters and policers on the PE router to customer-facing interfaces only.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

The following sections explain how to configure filters and policers for VPLS:

- [Configuring a VPLS Filter on page 5354](#)
- [Configuring a VPLS Policer on page 5357](#)

Configuring a VPLS Filter

To configure a filter for VPLS, include the **filter** statement at the **[edit firewall family vpls]** hierarchy level:

```
[edit firewall family vpls]
filter filter-name {
  interface-specific;
  term term-name {
    from {
      match-conditions;
    }
    then {
      actions;
    }
  }
}
```

For more information about how to configure firewall filters, see the *Junos OS Firewall Filters and Traffic Policers Configuration Guide*. For information on how to configure a VPLS filter match condition, see “[Standard Firewall Filter Match Conditions for VPLS Traffic](#)” on page 4727.

To configure a filter for VPLS traffic, complete the following tasks:

- [Configuring an Interface-Specific Counter for VPLS on page 5354](#)
- [Configuring an Action for the VPLS Filter on page 5355](#)
- [Configuring VPLS FTFs on page 5355](#)
- [Changing Precedence for Spanning-Tree BPDU Packets on page 5355](#)
- [Applying a VPLS Filter to an Interface on page 5355](#)
- [Applying a VPLS Filter to a VPLS Routing Instance on page 5356](#)
- [Configuring a Filter for Flooded Traffic on page 5356](#)

Configuring an Interface-Specific Counter for VPLS

When you configure a firewall filter for VPLS and apply it to multiple interfaces, you can specify individual counters specific to each interface. This allows you to collect separate statistics on the traffic transiting each interface.

To generate an interface-specific counter for VPLS, you configure the **interface-specific** statement. A separate instantiation of the filter is generated. This filter instance has a different name (based on the interface name) and collects statistics on the interface specified only.

To configure interface-specific counters, include the **interface-specific** statement at the **[edit firewall family vpls filter filter-name]** hierarchy level:

```
[edit firewall family vpls filter filter-name]
```


interface-specific;



NOTE: The counter name is restricted to 24 bytes. If the renamed counter exceeds this maximum length, it might be rejected.

For more information about the **interface-specific** statement and an example of how to configure it, see the *Junos OS Firewall Filters and Traffic Policers Configuration Guide*.

Configuring an Action for the VPLS Filter

You can configure the following actions for a VPLS filter at the [edit firewall family vpls filter *filter-name* term *term-name* then] hierarchy level: **accept**, **count**, **discard**, **forwarding-class**, **loss-priority**, **next**, **policer**.

Configuring VPLS FTFs

Forwarding table filters (FTFs) are filters configured for forwarding tables. For VPLS, they are attached to the destination MAC (DMAC) forwarding table of the VPLS routing instance. You define VPLS FTFs in the same manner as any other type of FTF. You can only apply a VPLS FTF as an input filter.

To specify a VPLS FTF, include the **filter input** statement at the [edit routing-instance *routing-instance-name* forwarding-options family vpls] hierarchy level:

```
[edit routing-instance routing-instance-name forwarding-options family vpls]
filter input filter-name;
```

For the statement summaries of these statements, see the *Routing Policy Configuration Guide*.

Changing Precedence for Spanning-Tree BPDU Packets

Spanning tree BPDU packets are automatically set to a high precedence. The queue number on these packets is set to 3. On M Series routers (except the M320 router) by default, a queue value of 3 indicates high precedence. To enable this higher precedence on BPDU packets, an instance-specific BPDU precedence filter named **default_bpdu_filter** is automatically attached to the VPLS DMAC table. This filter places a high precedence on all packets sent to **01:80:c2:00:00:00/24**.

You can overwrite this filter by configuring a VPLS FTF filter and applying it to the VPLS routing instance. For more information, see [“Configuring VPLS FTFs” on page 5355](#) and [“Applying a VPLS Filter to a VPLS Routing Instance” on page 5356](#).

Applying a VPLS Filter to an Interface

To apply a VPLS filter to an interface, include the **filter** statement:

```
filter {
  input input-filter-name;
  output output-filter-name;
  group index;
}
```

You can include this statement at the following hierarchy levels:

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

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



NOTE: For output interface filters, MAC addresses are learned after the filter action is completed. When an output interface filter's action is **discard**, the packet is dropped before the MAC address is learned. However, an input interface filter learns the MAC address before discarding the packet.

For the statement summaries for these statements, see the *Junos® OS Network Interfaces*.

Applying a VPLS Filter to a VPLS Routing Instance

You can apply a VPLS filter to a VPLS routing instance. The filter checks traffic passing through the specified routing instance.

Input routing instance filters learn the MAC address before the filter action is completed, so if the filter action is **discard**, the MAC address is learned before the packet is dropped.

To apply a VPLS filter to packets arriving at a VPLS routing instance and specify the filter, include the **filter input** statement at the [edit routing-instances *routing-instance-name* forwarding-options family vpls] hierarchy level:

```
[edit routing-instances routing-instance-name forwarding-options family vpls]
  filter input input-filter-name;
```

Configuring a Filter for Flooded Traffic

You can configure a VPLS filter to filter flooded packets. CE routers typically flood the following types of packets to PE routers in VPLS routing instances:

- Layer 2 broadcast packets
- Layer 2 multicast packets
- Layer 2 unicast packets with an unknown destination MAC address
- Layer 2 packets with a MAC entry in the DMAC routing table

You can configure filters to manage how these flooded packets are distributed to the other PE routers in the VPLS routing instance.

To apply a flooding filter to packets arriving at the PE router in the VPLS routing instance, and specify the filter, include the **flood input** statement:

```
flood input filter-name;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* forwarding-options family vpls]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options family vpls]

Configuring a VPLS Policer

You can configure a policer for VPLS traffic. The VPLS policer configuration is similar to the configuration of any other type of policer.

VPLS policers have the following characteristics:

- You cannot police the default VPLS routes stored in the flood table from PE router-sourced flood traffic.
- When specifying policing bandwidth, the VPLS policer considers all Layer 2 bytes in a packet to determine the packet length.

To configure a VPLS policer, include the **policer** statement at the [edit firewall] hierarchy level:

```
[edit firewall]
policer policer-name {
  bandwidth-limit limit;
  burst-size-limit limit;
  then action;
}
```

For the statement summaries of these statements and more information about how to configure policers, see the *Junos OS Firewall Filters and Traffic Policers Configuration Guide*.

To apply a VPLS policer to an interface, include the **policer** statement:

```
policer {
  input input-policer-name;
  output output-policer-name;
}
```

You can include this statement at the following hierarchy levels:

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

In the **input** statement, list the name of the VPLS policer to be evaluated when packets are received on the interface. In the **output** statement, list the name of the VPLS policer to be evaluated when packets are transmitted on the interface. This type of VPLS policer can only apply to unicast packets. For information about how to filter flood packets, see [“Configuring a Filter for Flooded Traffic” on page 5356](#).

For the statement summaries for these statements, see the *Junos® OS Network Interfaces*.

Standard Firewall Filter Match Conditions for VPLS Traffic

In the **from** statement in the VPLS filter term, you specify conditions that the packet must match for the action in the **then** statement to be taken. All conditions in the **from** statement must match for the action to be taken. The order in which you specify match conditions is not important, because a packet must match all the conditions in a term for a match to occur.

If you specify no match conditions in a term, that term matches all packets.

An individual condition in a **from** statement can contain a list of values. For example, you can specify numeric ranges. You can also specify multiple source addresses or destination addresses. When a condition defines a list of values, a match occurs if one of the values in the list matches the packet.

Individual conditions in a **from** statement can be negated. When you negate a condition, you are defining an explicit mismatch. For example, the negated match condition for **forwarding-class** is **forwarding-class-except**. If a packet matches a negated condition, it is immediately considered not to match the **from** statement, and the next term in the filter is evaluated, if there is one. If there are no more terms, the packet is discarded.

You can configure a standard firewall filter with match conditions for Virtual Private LAN Service (VPLS) traffic (**family vpls**). [Table 346 on page 4728](#) describes the *match-conditions* you can configure at the `[edit firewall family vpls filter filter-name term term-name from]` hierarchy level.



NOTE: Not all match conditions for VPLS traffic are supported on all routing platforms or switching platforms. A number of match conditions for VPLS traffic are supported only on MX Series 3D Universal Edge Routers.

In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Table 387: Standard Firewall Filter Match Conditions for VPLS Traffic

Match Condition	Description
destination-mac-address <i>address</i>	Match the destination media access control (MAC) address of a VPLS packet.

Table 387: Standard Firewall Filter Match Conditions for VPLS Traffic (*continued*)

Match Condition	Description
destination-port <i>number</i>	<p>(MX Series routers and EX Series switches only) Match the UDP or TCP destination port field.</p> <p>You cannot specify both the port and destination-port match conditions in the same term.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed): afs (1483), bgp (179), biff (512), bootpc (68), bootps (67), cmd (514), cvspserver (2401), dhcp (67), domain (53), eklogin (2105), ekshell (2106), exec (512), finger (79), ftp (21), ftp-data (20), http (80), https (443), ident (113), imap (143), kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544), ldap (389), ldp (646), login (513), mobileip-agent (434), mobilip-mn (435), msdp (639), netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123), pop3 (110), pptp (1723), printer (515), radacct (1813), radius (1812), rip (520), rkinit (2108), smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514), tacacs (49), tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525), who (513), or xmcp (177).</p>
destination-port-except <i>number</i>	<p>(MX Series routers and EX Series switches only) Do not match on the TCP or UDP destination port field. You cannot specify both the port and destination-port match conditions in the same term.</p>
destination-prefix-list <i>name</i>	<p>(MX Series routers and EX Series switches only) Match destination prefixes in the specified list. Specify the name of a prefix list defined at the [edit policy-options prefix-list <i>prefix-list-name</i>] hierarchy level.</p> <p>NOTE: VPLS prefix lists support only IPv4 addresses. IPv6 addresses included in a VPLS prefix list will be discarded.</p>
destination-prefix-list <i>name</i> except	<p>(MX Series routers and EX Series switches only) Do not match destination prefixes in the specified list. For more information, see the destination-prefix-list match condition.</p>
dscp <i>number</i>	<p>(MX Series routers and EX Series switches only) Match the Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant 6 bits of this byte form the DSCP. For more information, see the <i>Junos OS Class of Service Configuration Guide</i>.</p> <p>You can specify a numeric value from 0 through 63. To specify the value in hexadecimal form, include 0x as a prefix. To specify the value in binary form, include b as a prefix.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> RFC 3246, <i>An Expedited Forwarding PHB (Per-Hop Behavior)</i>, defines one code point: ef (46). RFC 2597, <i>Assured Forwarding PHB Group</i>, defines 4 classes, with 3 drop precedences in each class, for a total of 12 code points: <p>af11 (10), af12 (12), af13 (14),</p> <p>af21 (18), af22 (20), af23 (22),</p> <p>af31 (26), af32 (28), af33 (30),</p> <p>af41 (34), af42 (36), af43 (38)</p>
dscp-except <i>number</i>	<p>(MX Series routers and EX Series switches only) Do not match on the DSCP. For details, see the dscp match condition.</p>

Table 387: Standard Firewall Filter Match Conditions for VPLS Traffic (*continued*)

Match Condition	Description
ether-type values	<p>Match the 2-octet IEEE 802.3 Length/EtherType field to the specified value or list of values.</p> <p>You can specify decimal or hexadecimal values from 0 through 65535 (0xFFFF). A value from 0 through 1500 (0x05DC) specifies the length of an Ethernet Version 1 frame. A value from 1536 (0x0600) through 65535 specifies the EtherType (nature of the MAC client protocol) of an Ethernet Version 2 frame.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the hexadecimal values are also listed): aarp (0x80F3), appletalk (0x809B), arp (0x0806), ipv4 (0x0800), ipv6 (0x86DD), mpls-multicast (0x8848), mpls-unicast (0x8847), oam (0x8902), ppp (0x880B), pppoe-discovery (0x8863), pppoe-session (0x8864), or sna (0x80D5).</p>
ether-type-except values	<p>Do not match the 2-octet Length/EtherType field to the specified value or list of values.</p> <p>For details about specifying the values, see the ether-type match condition.</p>
forwarding-class class	Match the forwarding class. Specify assured-forwarding , best-effort , expedited-forwarding , or network-control .
forwarding-class-except class	Do not match the forwarding class. For details, see the forwarding-class match condition.
icmp-code message-code	<p>Match the ICMP message code field.</p> <p>If you configure this match condition, we recommend that you also configure the next-header icmp or next-header icmp6 match condition in the same term.</p> <p>If you configure this match condition, you must also configure the icmp-type message-type match condition in the same term. An ICMP message code provides more specific information than an ICMP message type, but the meaning of an ICMP message code is dependent on the associated ICMP message type.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> parameter-problem: ip6-header-bad (0), unrecognized-next-header (1), unrecognized-option (2) time-exceeded: ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0) destination-unreachable: address-unreachable (3), administratively-prohibited (1), no-route-to-destination (0), port-unreachable (4)
icmp-code-except message-code	Do not match the ICMP message code field. For details, see the icmp-code match condition.

Table 387: Standard Firewall Filter Match Conditions for VPLS Traffic (*continued*)

Match Condition	Description
icmp-code <i>number</i>	<p>(MX Series routers and EX Series switches only) Match the ICMP message code field.</p> <p>If you configure this match condition, we recommend that you also configure the ip-protocol icmp or ip-protocol icmp6 match condition in the same term.</p> <p>If you configure this match condition, you must also configure the icmp-type <i>message-type</i> match condition in the same term. An ICMP message code provides more specific information than an ICMP message type, but the meaning of an ICMP message code is dependent on the associated ICMP message type.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> parameter-problem: ip6-header-bad (0), unrecognized-next-header (1), unrecognized-option (2) time-exceeded: ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0) destination-unreachable: address-unreachable (3), administratively-prohibited (1), no-route-to-destination (0), port-unreachable (4)
icmp-code-except <i>number</i>	<p>(MX Series routers and EX Series switches only) Do not match on the ICMP code field. For details, see the icmp-code match condition.</p>
icmp-type <i>number</i>	<p>(MX Series routers and EX Series switches only) Match the ICMP message type field.</p> <p>If you configure this match condition, we recommend that you also configure the ip-protocol icmp, ip-protocol icmp6, or ip-protocol icmpv6 match condition in the same term.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): destination-unreachable (1), echo-reply (129), echo-request (128), membership-query (130), membership-report (131), membership-termination (132), neighbor-advertisement (136), neighbor-solicit (135), node-information-reply (140), node-information-request (139), packet-too-big (2), parameter-problem (4), redirect (137), router-advertisement (134), router-renumbering (138), router-solicit (133), or time-exceeded (3).</p>
icmp-type-except <i>number</i>	<p>(MX Series routers and EX Series switches only) Do not match the ICMP message type field. For details, see the icmp-type match condition.</p>
interface <i>interface-name</i>	<p>Interface on which the packet was received. You can configure a match condition that matches packets based on the interface on which they were received.</p> <p>NOTE: If you configure this match condition with an interface that does not exist, the term does not match any packet.</p>
interface-group <i>group-number</i>	<p>Match the logical interface on which the packet was received to the specified interface group or set of interface groups. For group-number, specify a single value or a range of values from 0 through 255.</p> <p>To assign a logical interface to an interface group group-number, specify the group-number at the [interfaces <i>interface-name</i> unit <i>number</i> family <i>family</i> filter group] hierarchy level.</p> <p>For more information, see “Filtering Packets Received on a Set of Interface Groups Overview” on page 4514.</p> <p>NOTE: This match condition is not supported on T4000 Type 5 FPCs.</p>

Table 387: Standard Firewall Filter Match Conditions for VPLS Traffic (*continued*)

Match Condition	Description
interface-group-except <i>group-name</i>	Do not match the logical interface on which the packet was received to the specified interface group or set of interface groups. For details, see the interface-group match condition. NOTE: This match condition is not supported on T4000 Type 5 FPCs.
interface-set <i>interface-set-name</i>	Match the interface on which the packet was received to the specified interface set. To define an interface set, include the interface-set statement at the [edit firewall] hierarchy level. For more information, see “Filtering Packets Received on an Interface Set Overview” on page 4515 .
ip-address <i>address</i>	(MX Series routers and EX Series switches only) 32-bit address that supports the standard syntax for IPv4 addresses.
ip-destination-address <i>address</i>	(MX Series routers and EX Series switches only) 32-bit address that is the final destination node address for the packet.
ip-precedence <i>ip-precedence-field</i>	(MX Series routers and EX Series switches only) IP precedence field. In place of the numeric field value, you can specify one of the following text synonyms (the field values are also listed): critical-ecp (0xa0), flash (0x60), flash-override (0x80), immediate (0x40), internet-control (0xc0), net-control (0xe0), priority (0x20), or routine (0x00).
ip-precedence-except <i>ip-precedence-field</i>	(MX Series routers and EX Series switches only) Do not match on the IP precedence field.
ip-protocol <i>number</i>	(MX Series routers and EX Series switches only) IP protocol field.
ip-protocol-except <i>number</i>	(MX Series routers and EX Series switches only) Do not match on the IP protocol field.
ip-source-address <i>address</i>	(MX Series routers and EX Series switches only) IP address of the source node sending the packet.
learn-vlan-1p-priority <i>number</i>	(MX Series routers and EX Series switches only) Match on the IEEE 802.1p learned VLAN priority bits in the provider VLAN tag (the only tag in a single-tag frame with 802.1Q VLAN tags or the outer tag in a dual-tag frame with 802.1Q VLAN tags). Specify a single value or multiple values from 0 through 7. Compare with the user-vlan-1p-priority match condition.
learn-vlan-1p-priority-except <i>number</i>	(MX Series routers and EX Series switches only) Do not match on the IEEE 802.1p learned VLAN priority bits. For details, see the learn-vlan-1p-priority match condition.
learn-vlan-dei	(MX Series routers and EX Series switches only) Match the user VLAN ID drop eligibility indicator (DEI) bit.
learn-vlan-dei-except	(MX Series routers and EX Series switches only) Do not match the user VLAN ID DEI bit.
learn-vlan-id <i>number</i>	(MX Series routers and EX Series switches only) VLAN identifier used for MAC learning.
learn-vlan-id-except <i>number</i>	(MX Series routers and EX Series switches only) Do not match on the VLAN identifier used for MAC learning.

Table 387: Standard Firewall Filter Match Conditions for VPLS Traffic (*continued*)

Match Condition	Description
loss-priority <i>level</i>	<p>Packet loss priority (PLP) level. Specify a single level or multiple levels: low, medium-low, medium-high, or high.</p> <p>Supported on M120 and M320 routers; M7i and M10i routers with the Enhanced CFEB (CFEB-E); and MX Series routers.</p> <p>For IP traffic on M320, MX Series, and T Series routers with Enhanced II Flexible PIC Concentrators (FPCs) and EX Series switches, you must include the tri-color statement at the [edit class-of-service] hierarchy level to commit a PLP configuration with any of the four levels specified. If the tri-color statement is not enabled, you can only configure the high and low levels. This applies to all protocol families.</p> <p>For information about the tri-color statement and about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>Junos OS Class of Service Configuration Guide</i>.</p>
loss-priority-except <i>level</i>	<p>Do not match on the packet loss priority level. Specify a single level or multiple levels: low, medium-low, medium-high, or high.</p> <p>For information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see the <i>Junos OS Class of Service Configuration Guide</i>.</p>
port <i>number</i>	(MX Series routers and EX Series switches only) TCP or UDP source or destination port. You cannot specify both the port match condition and either the destination-port or source-port match condition in the same term.
port-except <i>number</i>	(MX Series routers and EX Series switches only) Do not match on the TCP or UDP source or destination port. You cannot specify both the port match condition and either the destination-port or source-port match condition in the same term.
prefix-list <i>name</i>	<p>(MX Series routers and EX Series switches only) Match the destination or source prefixes in the specified list. Specify the name of a prefix list defined at the [edit policy-options prefix-list prefix-list-name] hierarchy level.</p> <p>NOTE: VPLS prefix lists support only IPv4 addresses. IPv6 addresses included in a VPLS prefix list will be discarded.</p>
prefix-list <i>name</i> except	(MX Series routers and EX Series switches only) Do not match the destination or source prefixes in the specified list. For more information, see the destination-prefix-list match condition.
source-mac-address <i>address</i>	Source MAC address of a VPLS packet.
source-port <i>number</i>	(MX Series routers and EX Series switches only) TCP or UDP source port field. You cannot specify the port and source-port match conditions in the same term.
source-port-except <i>number</i>	(MX Series routers and EX Series switches only) Do not match on the TCP or UDP source port field. You cannot specify the port and source-port match conditions in the same term.
source-prefix-list <i>name</i>	<p>(MX Series routers and EX Series switches only) Match the source prefixes in the specified prefix list. Specify a prefix list name defined at the [edit policy-options prefix-list prefix-list-name] hierarchy level.</p> <p>NOTE: VPLS prefix lists support only IPv4 addresses. IPv6 addresses included in a VPLS prefix list will be discarded.</p>

Table 387: Standard Firewall Filter Match Conditions for VPLS Traffic (*continued*)

Match Condition	Description
source-prefix-list <i>name</i> except	(MX Series routers and EX Series switches only) Do not match the source prefixes in the specified prefix list. For more information, see the source-prefix-list match condition.
tcp-flags <i>flags</i>	<p>Match one or more of the low-order 6 bits in the 8-bit TCP flags field in the TCP header.</p> <p>To specify individual bit fields, you can specify the following text synonyms or hexadecimal values:</p> <ul style="list-style-type: none"> • fin (0x01) • syn (0x02) • rst (0x04) • push (0x08) • ack (0x10) • urgent (0x20) <p>In a TCP session, the SYN flag is set only in the initial packet sent, while the ACK flag is set in all packets sent after the initial packet.</p> <p>You can string together multiple flags using the bit-field logical operators.</p> <p>If you configure this match condition for IPv6 traffic, we recommend that you also configure the next-header tcp match condition in the same term to specify that the TCP protocol is being used on the port.</p>
traffic-type <i>type-name</i>	(MX Series routers and EX Series switches only) Traffic type. Specify broadcast , multicast , unknown-unicast , or known-unicast .
traffic-type-except <i>type-name</i>	(MX Series routers and EX Series switches only) Do not match on the traffic type. Specify broadcast , multicast , unknown-unicast , or known-unicast .
user-vlan-1p-priority <i>number</i>	<p>(MX Series routers and EX Series switches only) Match on the IEEE 802.1p user priority bits in the customer VLAN tag (the inner tag in a dual-tag frame with 802.1Q VLAN tags). Specify a single value or multiple values from 0 through 7.</p> <p>Compare with the learn-vlan-1p-priority match condition.</p>
user-vlan-1p-priority-except <i>number</i>	(MX Series routers and EX Series switches only) Do not match on the IEEE 802.1p user priority bits. For details, see the user-vlan-1p-priority match condition.
user-vlan-id <i>number</i>	(MX Series routers and EX Series switches only) Match the first VLAN identifier that is part of the payload.
user-vlan-id-except <i>number</i>	(MX Series routers and EX Series switches only) Do not match on the first VLAN identifier that is part of the payload.
vlan-ether-type <i>value</i>	VLAN Ethernet type field of a VPLS packet.
vlan-ether-type-except <i>value</i>	Do not match on the VLAN Ethernet type field of a VPLS packet.

Related Documentation

- [Guidelines for Configuring Standard Firewall Filters on page 4478](#)
- [Standard Firewall Filter Terminating Actions on page 4742](#)

- [Standard Firewall Filter Nonterminating Actions on page 4744](#)

Specifying the VT Interfaces Used by VPLS Routing Instances

By default, the Junos OS automatically selects one of the virtual tunnel (VT) interfaces available to the router for de-encapsulating traffic from a remote site. The Junos OS cycles through the currently available VT interfaces, regularly updating the list of available VT interfaces as new remote sites are discovered and new connections are brought up. However, you can also explicitly configure which VT interfaces will receive the VPLS traffic.

By including the **tunnel-services** statement at the **[edit routing-instances routing-instance-name protocols vpls]** hierarchy level, you can specify that traffic for particular VPLS routing instances be forwarded to specific VT interfaces. Doing so allows you to load-balance VPLS traffic among all the available VT interfaces on the router.

The **tunnel-services** statement includes the following options:

- **devices**—Specifies the VT interfaces acceptable for use by the VPLS routing instance. If you do not configure this option, all VT interfaces available to the router can be used for de-encapsulating traffic for this instance.
- **primary**—Specifies the primary VT interface to be used by the VPLS routing instance. The VT interface specified is used to de-encapsulate all VPLS traffic from the MPLS core network for this routing instance. If the VT interface specified is unavailable, then one of the other acceptable VT interfaces (specified in the **devices** option) is used for handling the VPLS traffic. If you do not configure this option, any acceptable VT interface can be used to de-encapsulate VPLS traffic from the core.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

To specify that traffic for a particular VPLS routing instance be forwarded to specific VT interfaces, include the **tunnel-services** statement:

```
tunnel-services {
  devices device-names;
  primary primary-device-name;
}
```

These statements can be configured at the following hierarchy levels:

- **[edit routing-instances routing-instance-name protocols vpls]**
- **[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls]**

Flooding Unknown Traffic Using Point-to-Multipoint LSPs

For a VPLS routing instance, you can flood unknown unicast, broadcast, and multicast traffic using point-to-multipoint (also called *P2MP*) LSPs. By default, VPLS relies upon

ingress replication to flood unknown traffic to the members of a VPLS routing instance. This can cause replication of data at routing nodes shared by multiple VPLS members, as shown in [Figure 126 on page 5366](#). The flood data is tripled between PE router PE1 and provider router P1 and doubled between provider routers P1 and P2. By configuring point-to-multipoint LSPs to handle flood traffic, the VPLS routing instance can avoid this type of traffic replication in the network, as shown in [Figure 127 on page 5366](#).

Figure 126: Flooding Unknown VPLS Traffic Using Ingress Replication





Figure 127: Flooding Unknown VPLS Traffic Using a Point-to-Multipoint LSP



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

The point-to-multipoint LSP used for VPLS flooding can be either static or dynamic. In either case, for each VPLS routing instance, the PE router creates a dedicated point-to-multipoint LSP. All of the neighbors of the VPLS routing instance are added to the point-to-multipoint LSP when the feature is enabled. If there are n PE routers in the VPLS routing instance, n point-to-multipoint LSPs are created in the network where each PE router is the root of the point-to-multipoint tree and includes the rest of the $n - 1$ PE routers as leaf nodes. If you configured static point-to-multipoint LSPs for flooding, any additional VPLS neighbors added to the routing instance later are not automatically added to the point-to-multipoint LSP. You will need to manually add the new VPLS neighbors to the static point-to-multipoint flooding LSP. If you configure dynamic point-to-multipoint LSPs, whenever VPLS discovers a new neighbor through BGP, a sub-LSP for this neighbor is added to the point-to-multipoint LSP for the routing instance.

This feature can be enabled incrementally on any PE router that is part of a specific VPLS routing instance. The PE routers can then use point-to-multipoint LSPs to flood traffic, whereas other PE routers in the same VPLS routing instance can still use ingress replication

to flood traffic. However, when this feature is enabled on any PE router, you must ensure that all PE routers in the VPLS routing instance that participate in the flooding of traffic over point-to-multipoint LSPs are upgraded to Junos OS Release 8.3 or later to support this feature.

To flood unknown unicast, broadcast, and multicast traffic using point-to-multipoint LSPs, configure the **rsvp-te** statement as follows:

```
rsvp-te {
  label-switched-path-template {
    (default-template | lsp-template-name);
  }
  static-lsp lsp-name;
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instance *routing-instance-name* provider-tunnel]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel]

You can configure either a static point-to-multipoint LSP for VPLS flooding or a dynamic point-to-multipoint LSP.



NOTE: You cannot specify both the **static** and **label-switched-path-template** statements at the same time.

The following sections describe how to configure static and dynamic point-to-multipoint LSPs for flooding unknown traffic in a VPLS routing instance:

- [Configuring Static Point-to-Multipoint Flooding LSPs on page 5367](#)
- [Configuring Dynamic Point-to-Multipoint Flooding LSPs on page 5368](#)

Configuring Static Point-to-Multipoint Flooding LSPs

The **static-lsp** option creates a static flooding point-to-multipoint LSP that includes all of the neighbors in the VPLS routing instance. Flood traffic is sent to all of the VPLS neighbors using the generated point-to-multipoint LSP. VPLS neighbors added to the routing instance later are not automatically added to the point-to-multipoint LSP. You will need to manually add the new VPLS neighbors to the static point-to-multipoint flooding LSP. By configuring static point-to-multipoint LSPs for flooding, you have more control over which path each sub-LSP follows.

To configure a static flooding point-to-multipoint LSP, specify the name of the static flooding point-to-multipoint LSP by including the **static-lsp** statement:

```
static-lsp lsp-name;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* provider-tunnel rsvp-te]

- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel rsvp-te]

Configuring Dynamic Point-to-Multipoint Flooding LSPs

To configure a dynamic point-to-multipoint flooding LSP, include the **label-switched-path-template** statement option at the [edit routing-instances *routing-instance-name* provider-tunnel rsvp-te] hierarchy level:

```
[edit routing-instances routing-instance-name provider-tunnel rsvp-te]
label-switched-path-template {
  (default-template | lsp-template-name);
}
```

You can automatically generate the point-to-multipoint LSP to be used for flooding unknown traffic or you can manually configure the point-to-multipoint LSP:

- [Configuring Dynamic Point-to-Multipoint Flooding LSPs with the Default Template on page 5368](#)
- [Configuring Dynamic Point-to-Multipoint Flooding LSPs with a Preconfigured Template on page 5368](#)

Configuring Dynamic Point-to-Multipoint Flooding LSPs with the Default Template

The **default-template** option, specified at the [edit routing-instances *routing-instance-name* provider-tunnel rsvp-te label-switched-path-template] hierarchy level, causes the point-to-multipoint LSPs to be created with default parameters. The default parameters are for a minimally configured point-to-multipoint LSP. The name of this point-to-multipoint LSP is also generated automatically and is based on the following model:

id:vppls:router-id:routing-instance-name

The following **show** command output for **show mpls lsp p2mp ingress** illustrates how a point-to-multipoint flood LSP name could appear if you configure the **label-switched-path-template** statement with the **default-template** option:

```
user@host> show mpls lsp p2mp ingress
Ingress LSP: 2 sessions P2MP name: static, P2MP branch count: 3
To          From          State Rt ActivePath      P      LSPname
10.255.14.181 10.255.14.172 Up    0
10.255.14.177 10.255.14.172 Up    0 path2          *      vpn02-vpn11
10.255.14.174 10.255.14.172 Up    0 path3          *      vpn02-vpn07
10.255.14.174 10.255.14.172 Up    0 path3          *      vpn02-vpn04
P2MP name: 9:vppls:10.255.14.172:green, P2MP branch count: 2
To          From          State Rt ActivePath      P      LSPname
10.255.14.177 10.255.14.172 Up    0
11:vppls:10.255.14.172:green
10.255.14.174 10.255.14.172 Up    0
10:vppls:10.255.14.172:green
Total 5 displayed, Up 5, Down 0
```

The dynamically generated point-to-multipoint LSP name is **9:vppls:10.255.14.172:green**.

Configuring Dynamic Point-to-Multipoint Flooding LSPs with a Preconfigured Template

You can configure a point-to-multipoint flooding LSP template for the VPLS routing instance. The template allows you to specify the properties of the dynamic

point-to-multipoint LSPs that are used to flood traffic for the VPLS routing instance. You can specify all of the standard options available for a point-to-multipoint LSP within this template. These properties are inherited by the dynamic point-to-multipoint flood LSPs.

To configure a point-to-multipoint LSP template for flooding VPLS traffic, specify all of the properties you want to include in a point-to-multipoint LSP configuration. To specify this LSP as a point-to-multipoint flooding template, include the **p2mp** and **template** statements:

```
p2mp;  
template;
```

You can include these statements at the following hierarchy levels:

- [edit protocols mpls label-switched-path *p2mp-lsp-template-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *p2mp-lsp-template-name*]

For more information about how to configure the **p2mp** statement and point-to-multipoint LSPs, see the *Junos OS MPLS Applications Configuration Guide*.

Once you have configured the point-to-multipoint LSP template, specify the name of the point-to-multipoint LSP template with the **label-switched-path-template** statement:

```
label-switched-path-template p2mp-lsp-template-name;
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* provider-tunnel rsvp-te]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* provider-tunnel rsvp-te]

Configuring Interoperability Between BGP Signaling and LDP Signaling in VPLS

A single VPLS routing instance can encompass one set of PE routers that use BGP for signaling and another set of PE routers that use LDP for signaling. Within each set, all of the PE routers are fully meshed in both the control and data planes and have a bidirectional pseudowire to each of the other routers in the set. However, the BGP-signaled routers cannot be directly connected to the LDP-signaled routers. To be able to manage the two separate sets of PE routers in a single VPLS routing instance, a border PE router must be configured to interconnect the two sets of routers.

The VPLS RFCs and Internet drafts require that all of the PE routers participating in a single VPLS routing instance must be fully meshed in the data plane. In the control plane, each fully meshed set of PE routers in a VPLS routing instance is called a PE router mesh group. The border PE router must be reachable by and have bidirectional pseudowires to all of the PE routers that are a part of the VPLS routing instance, both the LDP-signaled and BGP-signaled routers.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

For LDP BGP interworking to function, LDP-signaled routers can be configured with forwarding equivalence class (FEC) 128 or FEC 129.

The following sections describe how to configure BGP LDP interworking for VPLS:

- [LDP BGP Interworking Platform Support on page 5370](#)
- [Configuring FEC 128 VPLS Mesh Groups for LDP BGP Interworking on page 5370](#)
- [Configuring FEC 129 VPLS Mesh Groups for LDP BGP Interworking on page 5371](#)
- [Configuring Switching Between Pseudowires Using VPLS Mesh Groups on page 5371](#)
- [Configuring Integrated Routing and Bridging Support for LDP BGP Interworking with VPLS on page 5372](#)
- [Configuring Inter-AS VPLS with MAC Processing at the ASBR on page 5372](#)

LDP BGP Interworking Platform Support

LDP BGP interworking is supported on the following Juniper Networks routers and routing platforms:

- M7i
- M10i
- M40e
- M120
- M320
- MX Series routers
- T Series routers
- TX Matrix routers
- EX Series switches

Configuring FEC 128 VPLS Mesh Groups for LDP BGP Interworking

To configure FEC 128 LDP BGP interworking for VPLS, include the **mesh-group** statement in the VPLS routing instance configuration of the PE border router:

```
mesh-group mesh-group-name {  
  local-switching;  
  mac-flush [ explicit-mac-flush-message-options ];  
  neighbor address;  
  peer-as all;  
  vpls-id number;  
}
```

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols vpls]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls]

Using the **neighbor** statement, configure each PE router that is a part of the mesh group. You must separate the LDP-signaled routers and the BGP-signaled routers into their own respective mesh groups. The LDP-signaled routers can be divided into multiple mesh groups. The BGP-signaled routers must be configured within a single mesh group for each routing instance.

Configuring FEC 129 VPLS Mesh Groups for LDP BGP Interworking

Configuration for a mesh group for FEC 129 is very similar to the configuration for FEC 128.

Note the following differences for FEC 129:

- Each user-defined mesh group must have a unique route distinguisher. Do not use the route distinguisher that is defined for the default mesh group at the **[edit routing-instances]** hierarchy level.
- Each user-defined mesh group must have its own import and export route target.
- Each user-defined mesh group can have a unique Layer 2 VPN ID. By default, all the mesh groups that are configured for the a VPLS routing-instance use the same Layer 2 VPN ID, the one that you configure at the **[edit routing-instances]** hierarchy level.

Configuring Switching Between Pseudowires Using VPLS Mesh Groups

To configure switching between Layer 2 circuit pseudowires using VPLS mesh groups, you can do either of the following:

- Configure a mesh group for each Layer 2 circuit pseudowire terminating at a VPLS routing instance. The Junos OS can support up to 16 mesh groups on MX Series routers and up to 128 on M Series and T Series routers. However, two mesh groups are created by default, one for the CE routers and one for the PE routers. Therefore, the maximum number of user-defined mesh groups is 14 for MX Series routers and 126 for M Series and T Series routers. PE router mesh groups are not supported on J Series routers.
- Configure a single mesh group, terminate all the Layer 2 circuit pseudowires into it, and enable local switching between the pseudowires by including the **local-switching** statement at the **[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name]** hierarchy level. By default, you cannot configure local switching for mesh groups (except for the CE mesh group) because all of the VPLS PE routers must be configured in a full mesh. However, local switching is useful if you are terminating Layer 2 circuit pseudowires in a mesh group configured for an LDP signaled VPLS routing instance.



NOTE: Do not include the **local-switching** statement on PE routers configured in a full mesh VPLS network.

To terminate multiple pseudowires at a single VPLS mesh group, include the **local-switching** statement:

local-switching;

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols vpls mesh-group *mesh-group-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls mesh-group *mesh-group-name*]

Configuring Integrated Routing and Bridging Support for LDP BGP Interworking with VPLS

Beginning with Junos OS Release 9.4, you can configure an integrated routing and bridging (IRB) interface on a router that functions as an autonomous system border router (ASBR) in an inter-AS VPLS environment between BGP-signaled VPLS and LDP-signaled VPLS. Previously, IRB interfaces were supported only on Provider Edge (PE) routers.

To configure a IRB support for LDP BGP Interworking with VPLS, include the **routing-interface *interface-name*** statement.

You can include this statement at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-routers *logical-router-name* routing-instances *routing-instance-name*]

Configuring Inter-AS VPLS with MAC Processing at the ASBR

Inter-AS VPLS with MAC processing at the ASBR enables you to interconnect customer sites that are located in different ASs. In addition, you can configure the ASs with different signaling protocols. You can configure one of the ASs with BGP-signaled VPLS and the other with LDP-signaled VPLS. For more information about how to configure LDP-signaled and BGP signaled VPLS, see [“Configuring Interoperability Between BGP Signaling and LDP Signaling in VPLS” on page 5369](#).

For inter-AS VPLS to function properly, you need to configure IBGP peering between the PE routers, including the ASBRs in each AS, just as you do for a typical VPLS configuration. You also need to configure EBGP peering between the ASBRs in the separate ASs. The EBGP peering is needed between the ASBRs only. The link between the ASBR routers does not have to be Ethernet. You can also connect a CE router directly to one of the ASBRs, meaning you do not have to have a PE router between the ASBR and the CE router.

The configuration for the connection between the ASBRs makes inter-AS VPLS with MAC operations unique. The other elements of the configuration are described in other sections of this manual. An extensive configuration example for inter-AS VPLS with MAC operations is provided in the *Junos OS Feature Guides*.

The following sections describe how to configure inter-AS VPLS with MAC operations:

- [Inter-AS VPLS with MAC Operations Configuration Summary on page 5373](#)
- [Configuring the ASBRs for Inter-AS VPLS on page 5373](#)

Inter-AS VPLS with MAC Operations Configuration Summary

This section provides a summary of all of the elements which must be configured to enable inter-AS VPLS with MAC operations. These procedures are described in detail later in this chapter and in other parts of the *Junos OS VPNs Configuration Guide*.

The following lists all of major elements of an inter-AS VPLS with MAC operations configuration:

- Configure IBGP between all of the routers within each AS, including the ASBRs.
- Configure EBGP between the ASBRs in the separated ASs. The EBGP configuration includes the configuration that interconnects the ASs.
- Configure a full mesh of LSPs between the ASBRs.
- Configure a VPLS routing instance encompassing the ASBR routers. The ASBRs are VPLS peers and are linked by a single pseudowire. Multihoming between ASs is not supported. A full mesh of pseudowires is needed between the ASBR routers in all of the interconnected ASs.
- Configure the VPLS routing instances using either BGP signaling or LDP signaling. LDP BGP interworking is supported for inter-AS VPLS with MAC operations, so it is possible to interconnect the BGP-signaled VPLS routing instances with the LDP-signaled VPLS routing instances.
- Configure a single VPLS mesh group for all of the ASBRs interconnected using inter-AS VPLS.

Configuring the ASBRs for Inter-AS VPLS

This section describes the configuration on the ASBRs needed to enable inter-AS VPLS with MAC operations.

On each ASBR, you need to configure a VPLS mesh group within the VPLS routing instance which needs to include all of the PE routers within the AS, in addition to the ASBR. You need to configure the same mesh group for each of the ASs you want to interconnect using inter-AS VPLS. The mesh group name should be identical on each AS. You also must include the **peer-as all** statement. This statement enables the router to establish a single pseudowire to each of the other ASBRs.

To configure the mesh group on each ASBR, include the **mesh-group** and **peer-as all** statements:

```
mesh-group mesh-group-name {
  peer-as all;
}
```

You can include these statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols vpls]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls]

- Related Documentation**
- *Example: Configuring BGP Autodiscovery for LDP VPLS*
 - *Example: Configuring BGP Autodiscovery for LDP VPLS with User-Defined Mesh Groups*

Tracing VPLS Traffic and Operations

To trace VPLS traffic, include the **traceoptions** statement:

```
traceoptions {  
    file filename <files number> <size size> <world-readable | no-world-readable>;  
    flag flag <flag-modifier> <disable>;  
}
```

You can include this statement at the following hierarchy levels:

- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls]
- [edit routing-instances *routing-instance-name* protocols vpls]

The following trace flags display the operations associated with VPLS:

- **all**—All VPLS tracing options
- **connections**—VPLS connections (events and state changes)
- **error**—Error conditions
- **nlri**—VPLS advertisements received or sent using BGP
- **route**—Trace-routing information
- **topology**—VPLS topology changes caused by reconsideration or advertisements received from other PE routers using BGP

Configuring the Label Block Size

VPLS MPLS packets have a two-label stack. The outer label is used for normal MPLS forwarding in the service provider's network. If BGP is used to establish VPLS, the inner label is allocated by a PE router as part of a label block. One inner label is needed for each remote VPLS site. Four sizes are supported. We recommend using the default size of 8, unless the network design requires a different size for optimal label usage, to allow the router to support a larger number of VPLS instances.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

If you allocate a large number of small label blocks to increase efficiency, you also increase the number of routes in the VPLS domain. This has an impact on the control plane overhead.

Changing the configured label block size causes all existing pseudowires to be deleted. For example, if you configure the label block size to be 4 and then change the size to 8, all existing label blocks of size 4 are deleted, which means that all existing pseudowires are deleted. The new label block of size 8 is created, and new pseudowires are established.

Four label block sizes are supported: 2, 4, 8, and 16. Consider the following scenarios:

- 2—Allocate the label blocks in increments of 2. For a VPLS domain that has only two sites with no future expansion plans.
- 4—Allocate the label blocks in increments of 4.
- 8 (default)—Allocate the label blocks in increments of 8.
- 16—Allocate the label blocks in increments of 16. A label block size of 16 enables you to minimize the number of routes in the VPLS domain. Use this setting only if the number of routes is the most important concern.

Configure the label block size:


```
[edit routing-instances instance-name protocols vpls]
user@router# set label-block-size 2
```

Related Documentation


- [Configuring VPLS Routing Instances on page 5324](#)

Configuration Statements

active-interface (VPLS Multihoming)

Syntax	<pre>active-interface { any; primary interface-name; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls multi-homing],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls multi-homing site <i>site-name</i>],</p> <p>[edit routing-instances <i>instance-name</i> protocols vpls multi-homing],</p> <p>[edit routing-instances <i>instance-name</i> protocols vpls multi-homing site <i>site-name</i>]</p>
Release Information	Statement introduced in Junos OS Release 7.5.
Description	Specify a multihomed interface as the primary interface for the VPLS site. If there are multiple interfaces, the remaining interfaces are activated only when the primary interface goes down. If no active interfaces are configured at the site level, it is assumed that all traffic for a VPLS site travels through a single, nonmultihomed PE router.
	<div>  <p>NOTE: In the VPLS documentation, the word <i>router</i> in terms such as <i>PE router</i> is used to refer to any device that provides routing functions.</p> </div>
	The remaining statements are explained separately.
	For FEC 128, use the [edit routing-instances <i>instance-name</i> protocols vpls multi-homing] hierarchy level.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Specifying an Interface as the Active Interface on page 5344

any (VPLS Multihoming)

Syntax	any;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls multi-homing active-interface],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls multi-homing site <i>site-name</i> active-interface],</p> <p>[edit routing-instances <i>instance-name</i> protocols vpls multi-homing active-interface],</p> <p>[edit routing-instances <i>instance-name</i> protocols vpls multi-homing site <i>site-name</i> active-interface]</p>
Release Information	Statement introduced in Junos OS Release 7.5.
Description	Specify that any multihomed interface can be used as the primary interface by the VPLS site. Depending on the order in which interfaces are listed in the PE router's configuration, the first operational interface in the set of configured interfaces is chosen to be the primary interface.
	<div>  <p>NOTE: In the VPLS documentation, the word <i>router</i> in terms such as <i>PE router</i> is used to refer to any device that provides routing functions.</p> </div> <p>For FEC 128, use the [edit routing-instances <i>instance-name</i> protocols vpls multi-homing active-interface] hierarchy level.</p>
Default	This is the default behavior.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Specifying an Interface as the Active Interface on page 5344 • primary on page 5416

automatic-site-id

Syntax	<pre>automatic-site-id { collision-detect-time <i>seconds</i>; new-site-wait-time <i>seconds</i>; reclaim-wait-time minimum <i>seconds</i> maximum <i>seconds</i>; startup-wait-time <i>seconds</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i>]
Release Information	Statement introduced in Junos OS Release 9.1.
Description	Enable automatic site identifiers for VPLS routing instances.

When you configure **automatic-site-id** for the first time, you must deactivate and then activate **protocol vpls**. However, if you already have **automatic-site-id** configured, you do not need to deactivate and then activate **protocol vpls**.

Options **collision-detect-time**—The time in seconds to wait after a claim advertisement is sent to the other routers in a VPLS instance before a PE router can begin using a site identifier. If the PE router receives a competing claim advertisement for the same site identifier during this time period, it initiates the collision resolution procedure for site identifiers.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

new-site-wait-time—The time in seconds to wait to receive VPLS information for a newly configured routing instance or a new site. This time interval is also applied whenever the automatic site identifier feature is activated on a VPLS routing instance other than at startup. Effectively, this timer indicates how long to wait before an attempt is made to allocate a site identifier. This timer is also triggered whenever a VPLS routing instance is enabled.

reclaim-wait-time—The time in seconds to wait to receive VPLS information for a newly configured routing instance or a new site. This time interval is also applied whenever the automatic site identifier feature is activated on a VPLS routing instance other than at startup. Effectively, this timer indicates how long to wait before an attempt is made to allocate a site identifier. This timer is also triggered whenever a VPLS routing instance is enabled. You can configure two values for this option: the **minimum** wait time and the **maximum** wait time.

startup-wait-time—The time in seconds to wait at startup to receive all the VPLS information for the route targets configured on the other PE routers included in the VPLS routing instance.

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

Related Documentation • [Configuring Automatic Site Identifiers for VPLS on page 5326](#)

best-site

Syntax best-site;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls site *site-name*],
[edit routing-instances *routing-instance-name* protocols vpls site *site-name*]

Release Information Statement introduced in Junos OS Release 12.2.

Description Enables the VPLS multihoming best site functionality, allowing the site on which it has been enabled to be the preferred site for this PE router. This statement must be configured on all PE routers within the optimized VPLS routing instance.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.


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

Related Documentation • [Example: VPLS Multihoming, Improved Convergence Time](#)

bfd-liveness-detection (VPLS)

Syntax	<pre> bfd-liveness-detection { detection-time { threshold <i>milliseconds</i>; } minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; multiplier <i>number</i>; no-adaptation; transmit-interval { threshold <i>milliseconds</i>; minimum-interval <i>milliseconds</i>; } version (1 automatic); } </pre>
Hierarchy Level	<p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls oam],</p>
Release Information	Statement introduced in Junos OS Release 10.0.
Description	<p>Configure bidirectional failure detection timers.</p> <p>The BFD failure detection timers are adaptive and can be adjusted to be more or less aggressive. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the clear bfd adaptation command to return BFD interval timers to their configured values. The clear bfd adaptation command is hitless, meaning that the command does not affect traffic flow on the routing device.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BFD for Static Routes on page 3197

connectivity-type

Syntax	connectivity-type (ce irb permanent);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls], [edit routing-instances <i>routing-instance-name</i> protocols vpls]
Release Information	Statement introduced in Junos OS Release 9.1. irb option introduced in Junos OS Release 9.3. permanent option introduced in Junos OS Release 10.4.
Description	Specify when a VPLS connection is taken down depending on whether or not the interface for the VPLS routing instance is customer-facing or integrated routing and bridging (IRB).
Default	ce
Options	<p>ce—Require that for the VPLS connection to be up, the customer-facing interface for the VPLS routing instance must also be up. If the customer-facing interface fails, the VPLS connection is taken down.</p> <p>irb—Allow a VPLS connection to remain up so long as an IRB interface is configured for the VPLS routing instance.</p> <p>permanent—Allow a VPLS connection to remain up until specifically taken down. This option is reserved for use in configuring Layer 2 Wholesale subscriber networks. See the <i>Broadband Subscriber Management Solutions Guide</i> for details about configuring a Layer 2 Wholesale network.</p>
<div>  <p>NOTE: To specifically take down a VPLS routing instance that is using the permanent option, all associated static logical interfaces must also be down.</p> </div>	
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring VPLS Routing Instance and VPLS Interface Connectivity on page 5332 • Configuring Separate NNI Routing Instances for Layer 2 Wholesale Service Retailers

encapsulation (Physical Interface)

Syntax	encapsulation (atm-ccc-cell-relay atm-pvc cisco-hdlc cisco-hdlc-ccc cisco-hdlc-tcc ethernet-bridge ethernet-ccc ethernet-over-atm ethernet-tcc ethernet-vpls ethernet-vpls-fr ether-vpls-over-atm-llc ethernet-vpls-ppp extended-frame-relay-ccc extended-frame-relay-ether-type-tcc extended-frame-relay-tcc extended-vlan-bridge extended-vlan-ccc extended-vlan-tcc extended-vlan-vpls flexible-ethernet-services flexible-frame-relay frame-relay frame-relay-ccc frame-relay-ether-type frame-relay-ether-type-tcc frame-relay-port-ccc frame-relay-tcc generic-services multilink-frame-relay-uni-nni ppp ppp-ccc ppp-tcc vlan-ccc vlan-vci-ccc vlan-vpls);
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces rlsq <i>number:number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches (flexible-ethernet-services , ethernet-ccc , and ethernet-tcc options only).
Description	Specify the physical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.
Default	ppp —Use serial PPP encapsulation.
Options	<p>atm-ccc-cell-relay—Use ATM cell-relay encapsulation.</p> <p>atm-pvc—Use ATM PVC encapsulation.</p> <p>cisco-hdlc—Use Cisco-compatible High-Level Data Link Control (HDLC) framing.</p> <p>cisco-hdlc-ccc—Use Cisco-compatible HDLC framing on CCC circuits.</p> <p>cisco-hdlc-tcc—Use Cisco-compatible HDLC framing on TCC circuits for connecting different media.</p> <p>ethernet-bridge—Use Ethernet bridge encapsulation on Ethernet interfaces that have bridging enabled and that must accept all packets.</p> <p>ethernet-ccc—Use Ethernet CCC encapsulation on Ethernet interfaces that must accept packets carrying standard Tag Protocol ID (TPID) values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, CCC is not supported.</p> <p>ethernet-over-atm—For interfaces that carry IPv4 traffic, use Ethernet over ATM encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces. As defined in RFC 2684, <i>Multiprotocol Encapsulation over ATM Adaptation Layer 5</i>, this encapsulation type allows ATM interfaces to connect to devices that support only bridge protocol data units (BPDUs). Junos OS does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet LLC/SNAP frames with IP or ARP in the payload, and drops the rest. For packets destined to the Ethernet LAN, a route lookup is done using the destination</p>

IP address. If the route lookup yields a full address match, the packet is encapsulated with an LLC/SNAP and MAC header, and the packet is forwarded to the ATM interface.

ethernet-tcc—For interfaces that carry IPv4 traffic, use Ethernet TCC encapsulation on interfaces that must accept packets carrying standard TPID values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC is not supported.

ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard TPID values. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

ethernet-vpls-fr—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

ethernet-vpls-ppp—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer layer 2 PPP connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use it to forward the packet into a given VPLS instance.

ether-vpls-over-atm-llc—For ATM intelligent queuing (IQ) interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

extended-frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC.

extended-frame-relay-ether-type-tcc—Use extended Frame Relay ether type TCC for Cisco-compatible Frame Relay for DLCIs 1 through 1022. This encapsulation type is used for circuits with different media on either side of the connection.

extended-frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect different media. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC.

extended-vlan-bridge—Use extended VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q VLAN tagging and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

extended-vlan-ccc—Use extended VLAN encapsulation on CCC circuits with Gigabit Ethernet and 4-port Fast Ethernet interfaces that must accept packets carrying 802.1Q values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, extended VLAN CCC is not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC is not supported.

extended-vlan-tcc—For interfaces that carry IPv4 traffic, use extended VLAN encapsulation on TCC circuits with Gigabit Ethernet interfaces on which you want to use 802.1Q tagging. For 4-port Gigabit Ethernet PICs, extended VLAN TCC is not supported.

extended-vlan-vpls—Use extended VLAN VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.



NOTE: The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

flexible-ethernet-services—For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), use flexible Ethernet services encapsulation when you want to configure multiple per-unit Ethernet encapsulations. Aggregated Ethernet bundles can use this encapsulation type. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs), and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

flexible-frame-relay—For IQ interfaces only, use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value from 1 through 1022.

frame-relay—Use Frame Relay encapsulation.

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with the Cisco Frame Relay.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media.

frame-relay-port-ccc—Use Frame Relay port CCC encapsulation to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect different media.

generic-services—Use generic services encapsulation for services with a hierarchical scheduler.

multilink-frame-relay-uni-nni—Use MLFR UNI NNI encapsulation. This encapsulation is used on link services, voice services interfaces functioning as FRF.16 bundles, and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

ppp—Use serial PPP encapsulation.

ppp-ccc—Use serial PPP encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

ppp-tcc—Use serial PPP encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.

vlan-ccc—Use Ethernet VLAN encapsulation on CCC circuits.

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only. All logical interfaces configured on the Ethernet interface must also have the encapsulation type set to **vlan-vci-ccc**.

vlan-vpls—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN tagging and VPLS enabled. Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.



NOTE: Label-switched interfaces (LSIs) do not support VLAN VPLS encapsulation. Therefore, you can only use VLAN VPLS encapsulation on a PE-router-to-CE-router interface and not a core-facing interface.

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

**Related
Documentation**

- *Configuring Interface Encapsulation on Physical Interfaces*
- *Configuring CCC Encapsulation for Layer 2 VPNs*
- *Configuring Layer 2 Switching Cross-Connects Using CCC*
- *Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits*
- *Configuring ATM Interface Encapsulation*
- *Configuring ATM-to-Ethernet Interworking*
- *Configuring VLAN Encapsulation*
- *Configuring Extended VLAN Encapsulation*
- *Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces*
- *Configuring Interfaces for Layer 2 Circuits*
- *Configuring Interface Encapsulation on PTX Series Packet Transport Switches*
- *Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)*
- *Configuring MPLS LSP Tunnel Cross-Connects Using CCC*
- *Configuring TCC*
- [Configuring VPLS Interface Encapsulation on page 5337](#)
- [Configuring Interfaces for VPLS Routing on page 5336](#)
- *Defining the Encapsulation for Switching Cross-Connects*
- *Understanding Encapsulation on an Interface*

encapsulation-type (Layer 2 VPNs)

Syntax	<code>encapsulation-type (atm-aal5 atm-cell atm-cell-port-mode atm-cell-vc-mode atm-cell-vp-mode cesop cisco-hdlc ethernet ethernet-vlan frame-relay frame-relay-port-mode interworking ppp satop-e1 satop-e3 satop-t1 satop-t3);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls],</p> <p>[edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 11.1 for EX Series switches.</p>
Description	Specify the type of Layer 2 traffic originating from the CE device. Only the ethernet and ethernet-vlan encapsulation types are supported for VPLS. Not all encapsulation types are supported on the switches. See the switch CLI.
Options	<p>atm-aal5—ATM Adaptation Layer (AAL/5)</p> <p>atm-cell—ATM cell relay</p> <p>atm-cell-port-mode—ATM cell relay port promiscuous mode</p> <p>atm-cell-vc-mode—ATM VC cell relay nonpromiscuous mode</p> <p>atm-cell-vp-mode—ATM virtual path (VP) cell relay promiscuous mode</p> <p>cesop—CESOP-based Layer 2 VPN</p> <p>cisco-hdlc—Cisco Systems-compatible HDLC</p> <p>ethernet—Ethernet</p> <p>ethernet-vlan—Ethernet VLAN</p> <p>frame-relay—Frame Relay</p> <p>frame-relay-port-mode—Frame Relay port mode</p> <p>interworking—Layer 2.5 interworking VPN</p> <p>ppp—PPP</p> <p>satsop-e1—SATSOP-E1-based Layer 2 VPN</p>

satsop-e3—SATSOP-E3-based Layer 2 VPN

satsop-t1—SATSOP-T1-based Layer 2 VPN

satsop-t3—SATSOP-T3-based Layer 2 VPN

Default: For VPLS networks, the default encapsulation type is **ethernet**.

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

Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Local Site on PE Routers in Layer 2 VPNs</i>• Configuring VPLS Routing Instances on page 5324• <i>Configuring Interfaces for Layer 2 Circuits</i>• <i>Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)</i>
------------------------------	---

family multiservice

```
Syntax  family multiservice {
        destination-mac;
        label-1;
        label-2;
        payload {
            ip {
                layer-3 {
                    (source-ip-only | destination-ip-only);
                }
                layer-3-only;
                layer-4;
            }
        }
        source-mac;
        symmetric-hash {
            complement;
        }
    }
```

Hierarchy Level [edit [forwarding-options](#) hash-key]

Release Information Statement introduced in Junos OS Release 8.0.
ip, **label-1**, **label-2**, **layer-3-only**, and **payload** options introduced in Junos OS Release 9.4.
layer-3, **layer-4**, **source-ip-only**, and **destination-ip-only** options introduced in Junos OS Release 9.5.
symmetric-hash and **complement** options introduced in Junos OS Release 9.6.

Description Configure load balancing based on Layer 2 media access control information. On MX Series routers, configure VPLS load balancing. On M120 and M320 routers only, configure VPLS load balancing based on MPLS labels and IP information. For IPv4 traffic, only the IP source and destination addresses are included in the hash key. For MPLS and IPv4 traffic, one or two MPLS labels and IPv4 source and destination addresses are included. For MPLS Ethernet pseudowires, only one or two MPLS labels are included in the hash key.

Options You can configure one or more options to load-balance using the packet information that you specify.

destination-mac—Include the destination-address MAC information in the hash key for Layer 2 load balancing.

label-1 (M120 and M320 routers only)—Include the first MPLS label in the hash key. Used for including a one-label packet for per-flow load balancing of IPv4 VPLS traffic based on IP information and MPLS labels.

label-2 (M120 and M320 routers only)—Include the second MPLS label in the hash key. If both **label-1** and **label-2** are specified, the entire first label and the first 16 bits of the second label are hashed.

payload (MX Series, M120, and M320 routers only)—Include the packet's IP payload in the hash key.

- **ip** (MX Series, M120, and M320 routers only)—Include the IP address of the IPv4 or IPv6 payload in the hash key.
- **layer-3** (MX Series routers only)—Use this to include Layer 3 information from the packet's IP payload in the hash key.
 - **destination-ip-only** (MX Series routers only)—Use this to include only the destination IP address in the payload in the hash key.
 - **source-ip-only** (MX Series routers only)—Use this to include only the source IP address in the payload in the hash key.



NOTE: You can include either the **source-ip-only** or the **destination-ip-only** statement, not both. They are mutually exclusive.

- **layer-3-only** (M120, and M320 routers only)—Include only the Layer 3 information from the packet's IP payload in the hash key.
- **layer-4** (MX Series routers only)—Include Layer 4 information from the packet's IP payload in the hash key.



NOTE: On MX Series routers only, you can configure either Layer 3 or Layer 4 load balancing, or both at the same time.



NOTE: On I chip platforms, an unknown Layer 4 header is excluded from load-balance hashing to avoid undesired packet reordering.

source-mac—Include the source-address MAC information in the hash key.

symmetric-hash (MX Series routers only)—Configure the symmetric hash or symmetric hash complement for configuring symmetrical load balancing on an 802.3ad Link Aggregation Group.

- **complement** —Include the complement of the symmetric hash in the hash key.


Required Privilege Level

interface—To view this statement in the configuration.



interface-control—To add this statement to the configuration.

- Related Documentation**
- *Configuring Load Balancing Based on MAC Addresses*
 - *Configuring VPLS Load Balancing Based on IP and MPLS Information*
 - *Configuring VPLS Load Balancing on MX Series 3D Universal Edge Routers*
 - [Configuring VPLS Load Balancing on page 5348](#)

fast-reroute-priority

Syntax	fast-reroute-priority (high low medium);
Hierarchy Level	[edit forwarding-options] [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options], [edit routing-instances <i>routing-instance-name</i> forwarding-options]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Description	Specify the fast reroute priority for a VPLS routing instance. You can configure high , medium , or low fast reroute priority to prioritize specific VPLS routing instances for faster convergence and traffic restoration. Because the router repairs next hops for high-priority VPLS routing instances first, the traffic traversing a VPLS routing instance configured with high fast reroute priority is restored faster than the traffic for VPLS routing instances configured with medium or low fast reroute priority.
	<div> NOTE: In the VPLS documentation, the word <i>router</i> in terms such as <i>PE router</i> is used to refer to any device that provides routing functions.</div>
Default	low
Options	high —Set the fast reroute priority for a VPLS routing instance to high. During a fast reroute event, the router repairs next hops for high-priority VPLS routing instances first. low —Set the fast reroute priority for a VPLS routing instance to low, which is the default. During a fast reroute event, the router repairs next hops for low-priority VPLS routing instances last. medium —Set the fast reroute priority for a VPLS routing instance to medium. During a fast reroute event, the router repairs next hops for medium-priority VPLS instances after high-priority VPLS routing instances but before low-priority VPLS routing instances.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Configuring VPLS Fast Reroute Priority

identifier (VPLS Multihoming for FEC 129)

Syntax	<code>identifier <i>identifier</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls multi-homing site <i>site-name</i>], [edit routing-instances <i>instance-name</i> protocols vpls multi-homing site <i>site-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3.
Description	Configure a Layer 2 VPN or VPLS multihoming identifier (MHID). An identifier needs to be configured for each multihomed site. Multihoming site identifiers are specific to a VPLS domain. They need not be unique on a provider edge (PE) router when multiple VPLS instances are present. The network layer reachability information (NLRI) advertisements sent to a CE device are identified as candidates for designated forwarder selection because the advertisements have the same multihoming identifier. Thus, you should assign the same identifier on all VPLS PE routers that are multihomed to the same customer site.
	<div>  <p>NOTE: The route distinguisher must be unique among PE routers participating in a multihomed site, so that the RD:MHID combination is unique across multiple VPLS domains. For example, one PE router might have a route distinguisher of 1.1.1.4:1, and another PE router in the same site might have a route distinguisher of 1.1.1.2:1. The first number can be, for example, the loopback interface address that identifies the PE router. The second number is the multihoming identifier.</p> </div>
	<div>  <p>NOTE: In the VPLS documentation, the word <i>router</i> in terms such as <i>PE router</i> is used to refer to any device that provides routing functions.</p> </div>
Options	<i>identifier</i> —Number that identifies the multihomed site. Range: 1 through 65535
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring VPLS Multihoming (FEC 129)</i>

interface (Routing Instances)

Syntax	<code>interface <i>interface-name</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Interface over which the VPN traffic travels between the PE router or switch and customer edge (CE) router or switch. You configure the interface on the PE router or switch. If the value vrf is specified for the instance-type statement included in the routing instance configuration, this statement is required.
Options	<i>interface-name</i> —Name of the interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>instance-type</i>• <i>Configuring Routing Instances on PE Routers in VPNs</i>• <i>Example: Configuring MPLS-Based Layer 3 VPNs on EX Series Switches</i>• interface (VPLS Routing Instances) on page 5396

interface (VPLS Multihoming for FEC 129)

Syntax	interface <i>interface-name</i> { preference <i>preference-value</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls multi-homing site <i>site-name</i>], [edit routing-instances <i>instance-name</i> protocols vpls multi-homing site <i>site-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3.
Description	Configure the interface that connects this site to the VPN. The remaining statement is explained separately.
Options	<i>interface-name</i> —Name of the interface (for example, ge-0/1/0.1).
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring VPLS Multihoming (FEC 129)</i>

interface (VPLS Routing Instances)

Syntax	<code>interface <i>interface-name</i> { interface-mac-limit <i>limit</i>; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls], [edit routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the interface for a pseudowire to the VPLS customer site. To complete the configuration of interfaces for a VPLS routing instance, you must also configure the interfaces specified for a VPLS site at the [edit routing-instances <i>routing-instance-name</i>] hierarchy level as described in <i>Configuring Routing Instances on PE Routers in VPNs</i> .
Options	<i>interface-name</i> —Specify the name of the interface used by the VPLS site. The remaining statement is explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the VPLS Site Interfaces on page 5329• Configuring Routing Instances on PE Routers in VPNs• interface (Routing Instances) on page 5394

interface-mac-limit

Syntax	<code>interface-mac-limit <i>limit</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i> interfaces <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i> interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the maximum number of media access control (MAC) addresses that can be learned by the VPLS routing instance. You can configure the same limit for all interfaces configured for a routing instance. You can also configure a limit for a specific interface.
Options	limit —Specify the number of MAC addresses that can be learned from each interface. Range: 16 through 65,536 MAC addresses Default: 512 addresses
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Limiting the Number of MAC Addresses Learned from an Interface on page 5334 • mac-table-size on page 5404

l2vpn-id

Syntax	<code>l2vpn-id (<i>as-number:id</i> <i>ip-address:id</i>);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i>], [edit routing-instances <i>instance-name</i>]
Release Information	Statement introduced in Junos OS Release 10.4R2.
Description	Specify a globally unique Layer 2 VPN community identifier for the instance.
Options	<p>as-number:id—Autonomous system number (<i>l2vpn-id:as-number:2-byte-number</i>. For example: <i>l2vpn-id l2vpn-id:100:200</i>. The AS number can be in the range from 1 through 65,535.</p> <p>ip-address:id—IP address (<i>l2vpn-id:ip-address:2-byte-number</i>. For example: <i>l2vpn-id l2vpn-id:10.1.1.1:2</i>. The IP address can be any globally unique unicast address.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring BGP Autodiscovery for LDP VPLS • Example: Configuring BGP Autodiscovery for LDP VPLS with User-Defined Mesh Groups

label-block-size

Syntax	label-block-size <i>size</i> ;
Hierarchy Level	[edit logical-systems <i>profile-name</i> routing-instances <i>instance-name</i> protocols vpls], [edit routing-instances <i>instance-name</i> protocols vpls]
Release Information	Statement introduced in Junos OS Release 10.0.
Description	Configure the label block size for VPLS labels.
Default	8
Options	<ul style="list-style-type: none">• 2—Allocate the label blocks in increments of 2. For a VPLS domain that has only two sites with no future expansion plans.• 4—Allocate the label blocks in increments of 4.• 8 (default)—Allocate the label blocks in increments of 8.• 16—Allocate the label blocks in increments of 16. A label block size of 16 enables you to minimize the number of routes in the VPLS domain. Use this setting only if the number of routes is the most important concern.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Label Block Size on page 5375

label-switched-path-template

Syntax	label-switched-path-template { (default-template <i>lsp-template-name</i>); }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel rsvp-te], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel ingress-replication label-switched-path], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source <i>source-address</i> rsvp-te], [edit logical-systems <i>logical-system-name</i> routing-options dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry-name</i>], [edit routing-instances <i>routing-instance-name</i> provider-tunnel ingress-replication label-switched-path], [edit routing-instances <i>routing-instance-name</i> provider-tunnel rsvp-te], [edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source <i>source-address</i> rsvp-te], [edit routing-options dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry-name</i>]
Release Information	Statement introduced in Junos OS Release 8.5.
Description	Specify the LSP template. An LSP template is used as the basis for other dynamically generated LSPs. This feature can be used for a number of applications, including point-to-multipoint LSPs, flooding VPLS traffic, configuring ingress replication for IP multicast using MBGP MVPNs, and to enable RSVP automatic mesh. There is no default setting for the label-switched-path-template statement, so you must configure either the default-template using the default-template option, or you must specify the name of your preconfigured LSP template.
Options	default-template —Specify that the default LSP template be used for the dynamically generated LSPs. lsp-template-name —Specify the name of an LSP to be used as a template for the dynamically generated LSPs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs • Configuring RSVP-Signaled Inclusive Point-to-Multipoint LSPs for an MBGP MVPN on page 5235 • Configuring Dynamic Point-to-Multipoint Flooding LSPs on page 5368 • Configuring RSVP Automatic Mesh

local-switching (VPLS)

Syntax	local-switching;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>]
Release Information	Statement introduced in Junos OS Release 9.3.
Description	Allows you to terminate multiple Layer 2 circuit pseudowires at a single VPLS mesh group.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Switching Between Pseudowires Using VPLS Mesh Groups on page 5371

mac-flush

Syntax	<code>mac-flush [<i>explicit-mac-flush-message-options</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls], [edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>]
Release Information	Statement introduced in Junos OS Release 10.0.
Description	Enable media access control (MAC) flush processing for the virtual private LAN service (VPLS) routing instance or for the mesh group under a VPLS routing instance. MAC flush processing removes MAC addresses from the MAC address database that have been learned dynamically. With the dynamically learned MAC addresses removed, MAC address convergence requires less time to complete.

For certain cases where MAC flush processing is not initiated by default, you can also specify *explicit-mac-flush-message-options* that additionally configure the router to send explicit MAC flush messages. To configure the router to send explicit MAC flush messages under specific conditions, include *explicit-mac-flush-message-options* with the statement.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

In certain cases, BGP updates sent by the provider edge (PE) device are delayed for 1 to 5 seconds.

This happens when all of the following conditions are true:

- BGP-based VPLS multihoming sites are configured.
- The **mac-flush** statement is included in the configuration.
- a non-minimum designated-forwarder site (site-x, for example) transitions to non-designated-forwarder status

The BGP update being delayed corresponds to the explicit-MAC flush notification message sent by site-x's PE device (PE2, for example). This BGP update message is not deferred if the designated-forwarder status is lost due to a locally-triggered event (for example, a local attachment-circuit interface going down). In other words, BGP update messages are deferred (in Device PE2) only when the designated-forwarder state is lost due to external events taking place in remote PE devices that also hold site-x (for example, in PE1). Suppose, for example, that Device PE1 is the default designated-forwarder with site-x's local interface in the DOWN state. Device PE2 defers BGP update message after Device PE1's local interface comes back to the UP state.

Options *explicit-mac-flush-message-options*—(Optional) You can specify one or more of the following explicit MAC flush message options:

- **any-interface**—(Optional) Send a MAC flush message when any customer-facing attachment circuit interface goes down.
- **any-spoke**—(Optional) Send a MAC FLUSH-FROM-ME flush message to all provider edge (PE) routers in the core when one of the spoke pseudowires between the multitenant unit switch and the other network-facing provider edge (NPE) router goes down, causing the multitenant unit switch to switch to this NPE router.



.....

NOTE: This option has a similar effect in a VPLS multihoming environment with multiple multitenant unit switches connected to NPE routers, where both multitenant unit switches have pseudowires that terminate in a mesh group with local-switching configured. If the **any-spoke** option is enabled, then both PE routers send MAC FLUSH-FROM-ME flush messages to all PEs in the core.

.....

- **propagate**—(Optional) Propagate MAC flush to the core.

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


Related Documentation

- [Configuring VPLS Routing Instances on page 5324](#)
- [Configuring Interoperability Between BGP Signaling and LDP Signaling in VPLS on page 5369](#)


mac-table-aging-time

Syntax	<code>mac-table-aging-time <i>time</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls], [edit routing-instances <i>routing-instance-name</i> protocols vpls]
Release Information	Statement introduced in Junos OS Release 7.4.
Description	Modify the timeout interval for the VPLS table.
Options	<i>time</i> —Specify the number of seconds to wait between VPLS table clearings. Range: 10 through 1,000,000 seconds Default: 300 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the VPLS MAC Table Timeout Interval on page 5333

mac-table-size

Syntax	<code>mac-table-size size;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls], [edit routing-instances <i>routing-instance-name</i> protocols vpls]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Modify the size of the VPLS MAC address table.
Options	<p>size—Specify the size of the MAC address table.</p> <p>Range:</p> <ul style="list-style-type: none">• (M Series and T Series routers only) 16 through 65,536 MAC addresses• (MX Series routers only) 16 through 1,048,575 MAC addresses• (T4000 routers with Type 5 FPCs only) 16 through 262,143 MAC addresses <div><p>NOTE: Before modifying the size of the VPLS MAC address table (to 262,143 addresses), you must enable network services mode by including the <code>enhanced-mode</code> statement at the [edit chassis network-services] hierarchy level and then reboot the router.</p></div>
Default:	512 MAC addresses
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Size of the VPLS MAC Address Table on page 5333• <i>Configuring Improved VPLS MAC Address Learning on T4000 Routers with Type 5 FPCs</i>• <i>enhanced-mode</i>

mesh-group (Protocols VPLS)

Syntax	<pre> mesh-group <i>mesh-group-name</i> { l2vpn-id (<i>as-number:id</i> <i>ip-address:id</i>); local-switching; mac-flush [<i>explicit-mac-flush-message-options</i>]; neighbor <i>address</i> {...} peer-as all; pseudowire-status-tlv; route-distinguisher (<i>as-number:id</i> <i>ip-address:id</i>); vpls-id <i>number</i>; vrf-export [<i>policy-names</i>]; vrf-import [<i>policy-names</i>]; vrf-target { <i>community</i>; import <i>community-name</i>; export <i>community-name</i>; } } </pre>
Hierarchy Level	<pre> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls], [edit routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i>] </pre>
Release Information	<p>Statement introduced in Junos OS Release 9.0.</p> <p>local-switching, mac-tlv-receive, mac-tlv-send, and peer-as options introduced in Junos OS Release 9.3.</p> <p>pseudowire-status-tlv and mac-flush options introduced in Junos OS Release 10.0.</p> <p>route-distinguisher option introduced in Junos OS Release 11.2.</p>
Description	<p>Specify the virtual private LAN service (VPLS) mesh group. The statement options allow you to specify each provider edge (PE) router that is a member of the mesh group. This statement is also used in the configuration of inter-autonomous system (AS) VPLS with media access control (MAC) operations.</p>
	<div>  <p>NOTE: In the VPLS documentation, the word <i>router</i> in terms such as <i>PE router</i> is used to refer to any device that provides routing functions.</p> </div>
Options	<p><i>mesh-group-name</i>—Name of the VPLS mesh group.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Configuring VPLS Routing Instances on page 5324](#)
 - [Configuring Interoperability Between BGP Signaling and LDP Signaling in VPLS on page 5369](#)
 - [Configuring Inter-AS VPLS with MAC Processing at the ASBR on page 5372](#)

multi-homing (VPLS Multihoming for FEC 128)

Syntax	multi-homing;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i>]
Release Information	Statement introduced in Junos OS Release 7.5.
Description	Specify the PE router as being a part of a multihomed site. Include this statement on all PE routers associated with a particular site. Configuration of this statement tracks BGP peers. If no BGP peer is available, all active interfaces for a site are deactivated.



.....

NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

.....

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	• Configuring Multihoming on the PE Router on page 5345

multi-homing (VPLS Multihoming for FEC 129)

```
Syntax  multi-homing {
        peer-active;
        site site-name {
            active-interface interface-name {
                any;
                primary interface-name;
            }
            identifier identifier;
            interface interface-name {
                preference preference-value;
            }
            peer-active;
            preference (preference-value | backup | primary);
        }
    }
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *instance-name* protocols vpls],
[edit routing-instances *instance-name* protocols vpls]

Release Information Statement introduced in Junos OS Release 12.3.

Description For VPLS autodiscovery (FEC 129), specify the parameters for multihoming to two or more provider edge (PE) routers.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

The remaining statements are explained separately.

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

Related Documentation

- *Example: Configuring VPLS Multihoming (FEC 129)*

neighbor (Protocols VPLS)

Syntax	<pre> neighbor <i>neighbor-id</i> { associate-profile { <i>dynamic-profile-name</i>; profile-variable-set <i>profile-variable-set-name</i>; } backup-neighbor {...} community <i>community-name</i>; connection-protection; encapsulation-type (Layer 2 Circuits) <i>type</i>; ignore-encapsulation-mismatch; oam { bfd-liveness-detection { detection-time { threshold <i>milliseconds</i>; } minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; multiplier <i>number</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } version <i>bfd-protocol-version</i>; } ping-interval; } pseudowire-status-tlv; psn-tunnel-endpoint <i>address</i>; revert-time <i>seconds</i>; static { incoming-label <i>label</i>; outgoing-label <i>label</i>; } switchover-delay <i>milliseconds</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls],</p> <p>[edit routing-instances <i>instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>]</p>
Release Information	Statement introduced in Junos OS Release 8.4. The pseudowire-status-tlv option was added in Junos OS Release 10.0.
Description	Specify each of the PE routers participating in the VPLS domain. Configuring this statement enables LDP for signaling VPLS.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Options *neighbor-id*—Specify the neighbor identifier for each PE router participating in the VPLS domain.


The remaining statements are explained separately.

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

Related Documentation

- [Configuring LDP Signaling for VPLS on page 5330](#)

no-tunnel-services

Syntax	no-tunnel-services;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols vpls static-vpls], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls], [edit protocols vpls static-vpls], [edit routing-instances <i>routing-instance-name</i> protocols vpls]
Release Information	Statement introduced in Junos OS Release 7.6. Support for static VPLS added in Junos OS Release 10.2.
Description	<p>Configure VPLS on a router without a Tunnel Services PIC. Configuring the no-tunnel-services statement creates a label-switched interface (LSI) to provide VPLS functionality. An LSI MPLS label is used as the inner label for VPLS. This label maps to a VPLS routing instance. On the PE router, the LSI label is stripped and then mapped to a logical LSI interface. The Layer 2 Ethernet frame is then forwarded using the LSI interface to the correct VPLS routing instance.</p> <p>Label-switched interfaces configured with the no-tunnel-services statement are not supported with GRE tunnels.</p>
	<div><p>NOTE: In the VPLS documentation, the word <i>router</i> in terms such as <i>PE router</i> is used to refer to any device that provides routing functions.</p></div>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring VPLS Without a Tunnel Services PIC on page 5351• Configuring Static Pseudowires for VPLS on page 5342• Configuring EXP-Based Traffic Classification for VPLS on page 5348


peer-active (VPLS Multihoming for FEC 129)

Syntax	<code>peer-active;</code>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls multi-homing], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls multi-homing site <i>site-name</i>], [edit routing-instances <i>instance-name</i> protocols vpls multi-homing], [edit routing-instances <i>instance-name</i> protocols vpls multi-homing site <i>site-name</i>]</pre>
Release Information	Statement introduced in Junos OS Release 12.3.
Description	Keep customer edge (CE) interfaces in the up state when all BGP peers go down.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Consider a scenario in which two provider edge (PE) routers are sharing two multihomed sites under one routing instance, with two CE devices, CE1 and CE2.



If the BGP peering session drops between Router PE1 and Router PE2, each one would consider itself to be the designated forwarder (DF) for Device CE1 and Device CE2. This creates a loop through the two CE devices, in which traffic loops from one CE device to the other then back to the first.


Junos OS overcomes this scenario by dropping all multihomed CE interface traffic on all multihoming PE routers when the BGP session drops between the PE routers. This functionality is enabled by default for all sites in a routing instance.

The **peer-active** statement disables the default functionality, so that PE routers keep their multihomed CE interfaces in the up state, even though the BGP peering session is down.

If you configure this statement in the **multi-homing** hierarchy, the default functionality is disabled for all sites. If you configure this statement for a site, the default functionality is disabled only for that particular site.

Default	If you omit this statement, Junos OS drops all multihomed CE interface traffic on all multihoming PE routers when the BGP session drops between the PE routers.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring VPLS Multihoming (FEC 129)

peer-as (VPLS)

Syntax	<pre>peer-as { all; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>]
Release Information	Statement introduced in Junos OS Release 9.3.
Description	Enable the autonomous system border router (ASBR) to establish a single pseudowire to each of the other ASBRs interconnected using inter-AS VPLS with MAC processing at the ASBR.
<div> NOTE: In the VPLS documentation, the word <i>router</i> in terms such as <i>PE router</i> is used to refer to any device that provides routing functions.</div>	
Options	all —This option is required. All peer routers, the ASBRs, are placed within the same VPLS mesh group.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Inter-AS VPLS with MAC Processing at the ASBR on page 5372

ping-interval

Syntax	ping-interval;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> oam],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols l2vpn oam],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls neighbor <i>address</i> oam],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>address</i> oam],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls oam],</p> <p>[edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> oam],</p> <p>[edit routing-instances <i>instance-name</i> protocols l2vpn oam],</p> <p>[edit routing-instances <i>instance-name</i> protocols vpls neighbor <i>address</i> oam],</p> <p>[edit routing-instances <i>instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>address</i> oam],</p> <p>[edit routing-instances <i>instance-name</i> protocols vpls oam]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.0.</p> <p>Support for FEC 129 VPLS added in Junos OS Release 12.2.</p>
Description	Configure the time interval between ping messages for bidirectional forwarding detection (BFD) sessions enabled over pseudowires inside a VPN.
Options	<p><i>seconds</i>—Time interval between ping messages.</p> <p>Range: 30 through 3600</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Configuring BFD for VCCV for Layer 2 VPNs, Layer 2 Circuits, and VPLS in the Junos OS VPNs Configuration Guide</i>


preference (Interface-Level Preference for VPLS Multihoming for FEC 129)

Syntax	<code>preference <i>preference-value</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls multi-homing site <i>site-name</i> interface <i>interface-name</i>], [edit routing-instances <i>instance-name</i> protocols vpls multi-homing site <i>site-name</i> interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3.
Description	<p>Specify a preference for the interface to become the designated forwarder (DF) for a multihomed VPLS site. This preference statement can be useful when you want the interface preference for a site to change dynamically so that the DF election can be influenced depending on the interface state. Among the list of interface preferences, Junos OS advertises the best preference as the VPLS site's preference value. For example, if the site has three interfaces configured with preference values 12, 10, and 9, respectively, 12 is advertised as the site preference. If that interface goes down, 10 is advertised as the site preference.</p> <p>If you configure interface-level preference, you cannot configure site-level preference.</p>
Options	<p><i>preference-value</i>—Preference value for the interface.</p> <p>Range: 1 through 65535</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring VPLS Multihoming (FEC 129)</i>

preference (Site-Level Preference for VPLS Multihoming for FEC 129)

Syntax	<code>preference (<i>preference-value</i> backup primary);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls multi-homing site <i>site-name</i>], [edit routing-instances <i>instance-name</i> protocols vpls multi-homing site <i>site-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3.
Description	Influence the designated forwarder (DF) selection for a multihomed VPLS site. Configure the preference in terms of keywords primary and backup , or configure the preference value explicitly.
Default	If this statement is omitted, the default preference value for the site is 100.
Options	<i>preference-value</i> —Preference value for the DF. Range: 1 through 65535 backup —Less likely to become the DF. primary —Most likely to become the DF.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring VPLS Multihoming (FEC 129)</i>

primary (VPLS Multihoming)

Syntax	<code>primary interface-name;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls multi-homing active-interface]</code> , <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls multi-homing site <i>site-name</i> active-interface]</code> , <code>[edit routing-instances <i>instance-name</i> protocols vpls multi-homing active-interface]</code> , <code>[edit routing-instances <i>instance-name</i> protocols vpls multi-homing site <i>site-name</i> active-interface]</code>
Release Information	Statement introduced in Junos OS Release 7.5.
Description	<p>Specify the name of the multihomed interface to be used as the primary interface by the VPLS site.</p> <p>For FEC 128, use the <code>[edit routing-instances <i>instance-name</i> protocols vpls multi-homing active-interface]</code> hierarchy level.</p>
Default	If you omit this statement, depending on the order in which interfaces are listed in the PE router's configuration, the first operational interface in the set of configured interfaces is chosen to be the primary interface.
<div> NOTE: In the VPLS documentation, the word <i>router</i> in terms such as <i>PE router</i> is used to refer to any device that provides routing functions.</div>	
Options	<code>interface-name</code> —Name of the interface (for example, <code>ge-0/1/0.1</code>).
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Specifying an Interface as the Active Interface on page 5344• any on page 5377

rsvp-te (Routing Instances Provider Tunnel)

Syntax	<pre> rsvp-te { label-switched-path-template { (default-template <i>lsp-template-name</i>); } static-lsp <i>lsp-name</i>; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel], [edit routing-instances <i>routing-instance-name</i> provider-tunnel]
Release Information	Statement introduced in Junos OS Release 8.3.
Description	Configure VPLS unknown unicast, broadcast, and multicast traffic flooding using point-to-multipoint LSPs.
Options	<p>static-lsp <i>lsp-name</i>—Create a static point-to-multipoint LSP and automatically include all of the neighbors in the VPLS routing instance.</p> <p>The remaining option is explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Flooding Unknown Traffic Using Point-to-Multipoint LSPs on page 5365

site (VPLS Multihoming for FEC 128)

Syntax	<pre>site <i>site-name</i> { active-interface (<i>any</i> <i>primary interface-name</i>); best-site; interface <i>interface-name</i> { interface-mac-limit <i>limit</i>; } mesh-group <i>mesh-group-name</i>; multi-homing; site-identifier <i>identifier</i>; site-preference <i>preference-value</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls], [edit routing-instances <i>routing-instance-name</i> protocols vpls]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the site name and site identifier for a site. Allows you to configure a remote site ID for remote sites.
Options	<p><i>site-name</i>—Name of the site.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the VPLS Site Name and Site Identifier on page 5326

site (VPLS Multihoming for FEC 129)

Syntax

```

site site-name {
    active-interface interface-name {
        any;
        primary interface-name;
    }
    identifier identifier;
    interface interface-name {
        preference preference-value;
    }
    peer-active;
    preference (preference-value | backup | primary);
}

```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *instance-name* protocols vpls multi-homing],
[edit routing-instances *instance-name* protocols vpls multi-homing]

Release Information Statement introduced in Junos OS Release 12.3.

Description For VPLS autodiscovery (FEC 129), specify the parameters for a VPLS site that is multihomed to two or more provider edge (PE) routers.



NOTE: In the VPLS documentation, the word *router* in terms such as *PE router* is used to refer to any device that provides routing functions.

Options *site-name*—Name of the site.

The remaining statements are explained separately.

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


Related Documentation

- *Example: Configuring VPLS Multihoming (FEC 129)*

site-identifier (VPLS)

Syntax	site-identifier <i>identifier</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the numerical identifier for the local VPLS site.
Options	<i>identifier</i> —Specify the numerical identifier for the local VPLS site. The identifier must be an unsigned 16-bit number greater than zero.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the VPLS Site Name and Site Identifier on page 5326

site-preference (VPLS)

Syntax	<pre>site-preference <i>preference-value</i> { backup; primary; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn site <i>site-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn site <i>site-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls site <i>site-name</i>]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the preference value advertised for a particular Layer 2 VPN or VPLS site. When a PE router receives multiple advertisements with the same VE identifier, the advertisement with the highest local preference value is preferred.
<div>  <p>NOTE: In the VPLS documentation, the word <i>router</i> in terms such as <i>PE router</i> is used to refer to any device that provides routing functions.</p> </div>	
Options	<p><i>preference-value</i>—Specify the preference value advertised for a Layer 2 VPN or VPLS site.</p> <p>Range: 1 through 65,535</p> <p>backup—Set the preference value to 1.</p> <p>primary—Set the preference value to 65,535.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the VPLS Site Preference on page 5329

site-range

Syntax	<code>site-range <i>number</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls], [edit routing-instances <i>routing-instance-name</i> protocols vpls]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify an upper limit on the maximum site identifier that can be accepted to allow a pseudowire to be brought up. Pseudowires cannot be established to sites with site identifiers greater than the configured site range. If you issue the show vpls connections command, such sites are displayed as OR (out of range). You must specify a value from 1 through 65,534. We recommend using the default.
Default	65,534
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Site Range on page 5327

static (Protocols VPLS)

Syntax	static { incoming-label <i>label</i> ; outgoing-label <i>label</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i> backup-neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i> backup-neighbor <i>address</i>]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	Specifies a static pseudowire for a VPLS domain. By configuring static pseudowires for the VPLS domain, you do not need to configure the LDP or BGP protocols that would normally be used for signaling. Static pseudowires require that you configure a set of in and out labels for each pseudowire configured for the VPLS domain. You can configure both static and dynamic neighbors within the same VPLS routing instance. You can also configure a static pseudowire for a backup neighbor (if you configure the neighbor as static the backup must also be static) and for a mesh group.
Options	incoming-label <i>label</i> —You must configure an incoming label for the static pseudowire. Range: 29,696 through 41,983 and 1,000,000 through 1,048,575 outgoing-label <i>label</i> —You must configure an outgoing label for the static pseudowire. Range: 16 through 1,048,575
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> See Configuring Static Pseudowires for VPLS on page 5342.

template

Syntax	template;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>p2mp-lsp-template-name</i>], [edit protocols mpls label-switched-path <i>p2mp-lsp-template-name</i>]
Release Information	Statement introduced in Junos OS Release 8.3.
Description	Specify a template for the dynamically generated point-to-multipoint LSPs used for VPLS flooding.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Dynamic Point-to-Multipoint Flooding LSPs with a Preconfigured Template on page 5368

traceoptions (Protocols VPLS)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls], [edit routing-instances <i>routing-instance-name</i> protocols vpls]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Trace traffic flowing through a VPLS routing instance.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks (" ").</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named <i>trace-file</i> reaches this size, it is renamed <i>trace-file.0</i>. When <i>trace-file</i> again reaches its maximum size, <i>trace-file.0</i> is renamed <i>trace-file.1</i> and <i>trace-file</i> is renamed <i>trace-file.0</i>. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can specify the following tracing flags:</p> <ul style="list-style-type: none"> • all—All VPLS tracing options • connections—VPLS connections (events and state changes) • error—Error conditions • nlri—VPLS advertisements received or sent by means of the BGP • route—Routing information • topology—VPLS topology changes caused by reconfiguration or advertisements received from other provider edge (PE) routers using BGP <p>flag-modifier—(Optional) Modifier for the tracing flag. You can specify the following modifiers:</p> <ul style="list-style-type: none"> • detail—Provide detailed trace information.

- **disable**—Disable the tracing flag.
- **receive**—Trace received packets.
- **send**—Trace sent packets.

no-world-readable—Do not allow any user to read the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify kilobytes, **xm** to specify megabytes, or **xg** to specify gigabytes


Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—Allow any user to read the log file.

Required Privilege Level	routing—To view this statement in the configuration.
	routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Tracing VPLS Traffic and Operations on page 5374

tunnel-services (Routing Instances VPLS)

Syntax	<pre>tunnel-services { devices <i>device-names</i>; primary <i>primary-device-name</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls], [edit routing-instances <i>routing-instance-name</i> protocols vpls]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify that traffic for particular VPLS routing instances be forwarded to specific virtual tunnel (VT) interfaces, allowing you to load-balance VPLS traffic among all the available VT interfaces on the router.
<div>  <p>NOTE: In the VPLS documentation, the word <i>router</i> in terms such as <i>PE router</i> is used to refer to any device that provides routing functions.</p> </div>	
Options	<p>devices <i>device-names</i>—Specify the VT interfaces acceptable for use by the VPLS routing instance. If you do not configure this option, all VT interfaces available to the router can be used for de-encapsulating traffic for this instance.</p> <p>primary <i>primary-device-name</i>—Specify the primary VT interface to be used by the VPLS routing instance. The VT interface specified is used to de-encapsulate all VPLS traffic from the MPLS core network for this routing instance. If the VT interface specified is unavailable, then one of the other acceptable VT interfaces is used for handling the VPLS traffic. If you do not configure this option, any acceptable VT interface can be used to de-encapsulate VPLS traffic from the core.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Specifying the VT Interfaces Used by VPLS Routing Instances on page 5365

vlan-id

Syntax	<code>vlan-id <i>number</i>;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For Fast Ethernet and Gigabit Ethernet interfaces only, bind an 802.1Q VLAN tag ID to a logical interface.
Options	<i>number</i> —A valid VLAN identifier. Range: For 4-port Fast Ethernet PICs configured to handle VPLS traffic, 512 through 1023. For 1-port and 10-port Gigabit Ethernet PICs configured to handle VPLS traffic, 512 through 4094.
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Interfaces for VPLS Routing on page 5336

vlan-id-list (Interface in VPLS)

Syntax	<code>vlan-id-list [<i>numbers number-number</i>];</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</code>
Release Information	Statement introduced for VPLS in Junos OS Release 10.2.
Description	Configure a logical interface to forward packets and learn MAC addresses within each VPLS routing instance configured with a VLAN ID that matches a VLAN ID specified in the list. VLAN IDs can be entered individually using a space to separate each ID, entered as an inclusive list separating the starting VLAN ID and ending VLAN ID with a hyphen, or a combination of both.
Options	<i>number number</i> —Individual VLAN IDs separated by a space. <i>number-number</i> —Starting VLAN ID and ending VLAN ID in an inclusive range. Range: 1 through 4095
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Interfaces for VPLS Routing on page 5336• Configuring VLAN IDs for Logical Interfaces on page 5340

vlan-tagging

Syntax	vlan-tagging;
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Description	For Fast Ethernet and Gigabit Ethernet interfaces and aggregated Ethernet interfaces configured for VPLS, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch</i> • <i>Example: Configuring BGP Autodiscovery for LDP VPLS</i> • <i>Configuring a Layer 3 Subinterface (CLI Procedure)</i> • Configuring Tagged Aggregated Ethernet Interfaces on page 2302 • Configuring Interfaces for VPLS Routing on page 5336 • Enabling VLAN Tagging on page 5339 • <i>802.1Q VLANs Overview</i> • <i>vlan-id</i>

vpls (Interfaces)

Syntax	vpls;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the VPLS protocol family information for the logical interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Interfaces for VPLS Routing on page 5336

vpls (Routing Instance)

```
Syntax  vpls {
    active-interface {
        any;
        primary interface-name;
    }
    connectivity-type (ce | irb);
    interface interface-name;
    interface-mac-limit limit;
    label-block-size size;
    mac-flush [ explicit-mac-flush-message-options ];
    mac-table-aging-time time;
    mac-table-size size;
    mesh-group mesh-group-name {
        l2vpn-id (as-number:id | ip-address:id);
        local-switching;
        mac-flush [ explicit-mac-flush-message-options ];
        neighbor address {...}
        peer-as all;
        pseudowire-status-tlv;
        route-distinguisher (as-number:id | ip-address:id);
        vpls-id number;
        vrf-export [ policy-names ];
        vrf-import [ policy-names ];
        vrf-target {
            community;
            import community-name;
            export community-name;
        }
    }
    no-tunnel-services;
    site site-name {
        active-interface interface-name {
            any;
            primary preference-value;
        }
        best-site;
        interface interface-name {
            interface-mac-limit limit;
        }
        mesh-group mesh-group-name;
        multi-homing;
        site-identifier identifier;
        site-preference preference-value;
    }
    site-range number;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
    tunnel-services {
        devices device-names;
        primary primary-device-name;
    }
}
```

}

Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. The mac-flush option was added in Junos OS Release 10.0.
Description	Configure a virtual private LAN service (VPLS) routing instance. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring VPLS Routing Instances on page 5324

vpls-id

Syntax	vpls-id <i>vpls-id</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols l2vpn], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols l2vpn mesh-group <i>mesh-group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>], [edit routing-instances <i>instance-name</i> protocols l2vpn], [edit routing-instances <i>instance-name</i> protocols l2vpn mesh-group <i>mesh-group-name</i>], [edit routing-instances <i>instance-name</i> protocols vpls], [edit routing-instances <i>instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>]
Release Information	Statement introduced in Junos OS Release 8.4.
Description	Identify the virtual circuit identifier used for the VPLS routing instance or mesh group. This statement is a part of the configuration to enable LDP signaling for VPLS.
Options	<i>vpls-id</i> —Specify a valid identifier for the VPLS routing instance.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring LDP Signaling for VPLS on page 5330

Administration

- [VPLS Reference on page 5432](#)

VPLS Reference

- [Supported VPLS Standards on page 5432](#)

Supported VPLS Standards

Junos OS substantially supports the following Internet RFCs and draft, which define standards for virtual private LAN service (VPLS).

- RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*
- RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*
FEC 128, FEC 129, control bit 0, the Ethernet pseudowire type 0x0005, and the Ethernet tagged mode pseudowire type 0x0004 are supported.
- Internet draft draft-kompella-l2vpn-vpls-multihoming, *Multi-homing in BGP-based Virtual Private LAN Service*

Related Documentation

- [Supported Carrier-of-Carriers and Interprovider VPN Standards](#)
- [Supported Layer 2 Circuit Standards](#)
- [Supported Layer 2 VPN Standard](#)
- [Supported Layer 3 VPN Standards](#)
- [Supported Multicast VPN Standards on page 5158](#)
- [Accessing Standards Documents on the Internet](#)